
Blackbox optimization of unimodal functions

Abstract

We provide an intuitive new algorithm for black-box stochastic optimization of unimodal functions, a function class that we observe empirically can capture hyperparameter-tuning loss surfaces. Our method’s convergence guarantee automatically adapts to Lipschitz constants and other problem difficulty parameters, recovering and extending prior results. We complement our theoretical development with experimental validation on hyperparameter tuning tasks.

1 BLACK-BOX OPTIMIZATION

This paper considers the problem of *black-box stochastic optimization*. Specifically, we are interested in solving a maximization problem of the form:

$$\max_x F(x) = \mathbb{E}_z[f(x, z)]$$

Here, F is the objective function that takes the form $F(x) = \mathbb{E}_z[f(x, z)]$, for some function f and some random variable z . We have access to a black-box evaluation oracle (sometimes called a *stochastic zeroth-order oracle*), which returns a random sample $f(x, z)$ given any x . This is the only information we have about f ; we do not know the function definition or its gradients.

The ability to solve this problem forms a fundamental primitive in more complex systems. For example, it can be used to optimize design of new materials [Terayama et al., 2021], to improve user interfaces, or search for hyperparameters in machine learning [Golovin et al., 2017, Feurer and Hutter, 2019, Hazan et al., 2018]. Although the techniques we present in this work are very general, we focus on the task of optimizing hyperparameters as an important motivating application.

Black-box optimization has been studied from a variety of perspectives, which is reflected in the variety of names that

refer to essentially similar concepts. Depending on the author, it may be referred to as “zero-order” or “derivative-free” optimization [Rios and Sahinidis, 2013, Duchi et al., 2015, Jamieson et al., 2012], “bandit” optimization [Agarwal et al., 2011, Agrawal, 1995, Kleinberg et al., 2008, Auer et al., 2007, Shamir, 2013, Jun et al., 2017], sequential experimental design [Chernoff, 1959], or Bayesian optimization [Shahriari et al., 2015, Srinivas et al., 2009, Snoek et al., 2012]. While the details of the methods used to solve the problem are diverse, the overall idea behind all methods is necessarily similar: first, assume some sort of “structure” on either the distribution of z or the shape of f (or both) that constrains the number of possibilities for F - popular examples include convexity, smoothness, a Bayesian prior. Then, use new evaluations $f(x, z)$ to further constrain the possibilities until one can reliably identify a maximizing point.

An algorithm’s usefulness can thus be measured on two axes: the degree that whatever structural assumptions it makes are in fact reflected in practice, and the number of samples required by the algorithm to optimize F subject to the assumptions. The goal is to find assumptions expressive enough to capture real problems while still admitting efficient optimization. In this paper, we focus on the assumption of *unimodality*. That is (in the 1D case), the function F has only one local maximum. This assumption is motivated via empirical observation of hyperparameter-tuning problems. For example, in Figure 1 we plot the influence of the learning rate hyperparameter on the accuracy of an AlexNet type architecture trained on the CIFAR100 image classification task [Krizhevsky, 2009]. It is evident that this relationship is at least approximately unimodal, and we describe further evidence for unimodality in Appendix E. While our algorithms simply assume unimodality and so may behave poorly when this condition does not hold, we observe empirically in Section 5 that in practice our approach works well.

Although unimodality is much less well-studied in the literature than its popular and more stringent cousin convexity, we




Figure 1: Accuracy of CNN on the CIFAR100 validation set as we vary the learning rate from 1e-6 to 1e-1 on a uniform grid in the log scale. We repeat the experiment 6 times for each of the 1000 learning rate values and show the mean accuracy here, with confidence intervals.

are not the first to consider unimodality. For example, Yu and Mannor [2011], Combes and Proutiere [2014], Combes et al. [2020] all provide algorithms for F that satisfy particular further assumptions in addition to unimodality. We develop a new algorithm that provides theoretical improvements over these works in terms of robustness to these further assumptions: our method simultaneously enjoys provably good performance for these prior function classes as well as other natural classes.

We then conduct an empirical evaluation of our method on a benchmark hyperparameter tuning task involving tuning hyperparameters across a number of different sub-tasks [Arango et al., 2021]. We compare our algorithm to the simple random search baseline as well as the more advanced Gaussian process upper-confidence bound algorithm, which is a standard approach to hyperparameter tuning. Our method not only obtains slightly improved average performance across the different sub-tasks, it also obtains significantly better performance on tasks with less than 6 hyperparameters to tune, forming a very promising initial study of our approach.

In summary, we make the following contributions:

- In Section 3, we propose a novel algorithm for black-box optimization of 1D unimodal functions, with theoretical analysis on a variety of function classes. We simultaneously match the convergence rate of Kleinberg et al. [2008] for Lipschitz (Theorem 1) and smooth functions (Theorem 2), and that of Combes et al. [2020] for functions that meet certain growth conditions (Theorem 6). Our proposed algorithm does *not* need prior knowledge of problem parameters such as Lipschitz or smoothness constants to achieve these results.

- In Section 4, we extend our 1D algorithm to higher dimensions. This extension provably converges to a local maximum.
- In Section 5, we empirically show that our multi-dimensional algorithm is generally competitive with the state-of-the-art Bayesian optimization methods on hyperparameter tuning tasks. For modest dimension counts (i.e. $d \approx 5$), our method significantly outperforms Bayesian optimization methods.

2 DEFINITIONS AND SETTING

In this paper we consider exclusively loss functions of the form $f : [0, 1]^d \times \mathcal{Z} \rightarrow [0, 1]$, where \mathcal{Z} is some arbitrary set. Given a \mathcal{Z} -valued random variable Z , we let $F(x) : [0, 1]^d \rightarrow [0, 1]$ be given by $F(x) = \mathbb{E}[f(x, Z)]$. Our algorithms have access to a stochastic value oracle: given a point x , we may generate a new i.i.d. sample z and compute $f(x, z)$. We call such a process a “query” or a “sample” interchangeably. Our goal is to find a point \hat{x} such that $F(\hat{x})$ is as large as possible using at most B samples for some given budget B .

A 1-D function is unimodal if it has a single local maximum. We extend this notion to arbitrary dimensions by imposing the definition on each coordinate axis.

Definition 1 (Unimodality). *A function $F : [0, 1]^d \rightarrow [0, 1]$ is unimodal if for all $(x_1, \dots, x_d) \in [0, 1]^d$, and all $i \in \{1, \dots, d\}$, there exists an $x_i^* \in [0, 1]$ such that for all $a, b \in [0, 1]$ with either $x_i^* \geq a \geq b$ or $x_i^* \leq a \leq b$, we have*

$$F(x_1, \dots, x_{i-1}, a, \dots, x_d) \geq F(x_1, \dots, x_{i-1}, b, \dots, x_d)$$

We also define a notion of *local optimality* to non-differentiable functions.

Definition 2. *A point $x = (x_1, \dots, x_d) \in [0, 1]^d$ is an τ -approximate local optimum point of $F : [0, 1]^d \rightarrow [0, 1]$ if for all $i \in \{1, \dots, d\}$, for all $x_i^* \in [0, 1]$,*

$$F(x_1, \dots, x_{i-1}, x_i^*, \dots, x_d) \leq F(x) + \tau$$

In other words, an approximate local optimum point is a point at which only minimal progress can be made by changing only one coordinate.

3 ALGORITHM AND ANALYSIS FOR 1-D FUNCTIONS

We first build an algorithm for the special-case of 1-D unimodal functions. We extend it in Section 4 to the high dimensional case by employing coordinate descent. For this section, let $F : [0, 1] \rightarrow [0, 1]$ and $x_* = \operatorname{argmax}_{x \in [0, 1]} F(x)$.

3.1 PREVIOUS WORK

Golden section search method [Kiefer, 1953] is a minimax optimal algorithm for deterministic 1-D unimodal functions. The key observation is the following:

If $F(a) \geq F(b)$ for some $a, b \in [0, 1]$, then $a \geq b$ implies $x_ \geq b$, and $a \leq b$ implies $x_* \leq b$.*

This suggests a natural strategy: maintain a *candidate interval* $[l, r]$ such that $x_* \in [l, r]$ and iteratively narrow it down. Combes et al. [2020], Yu and Mannor [2011] generalized the golden section search stochastic settings. They first commit to a small number of possible values $y_1, \dots, y_K \in [l, r]$, such that each y_i is far from the boundary of the current candidate interval $[l, r]$. Then, they repeatedly sample $f(y_i, z_t)$ until they can estimate $F(y_1), \dots, F(y_K)$ with high accuracy and identify y_i and y_j where $F(y_i) \geq F(y_j)$.

However, this approach is intuitively wasteful because it puts too many eggs in few baskets. For example, it fails to make progress when $F(y_i)$ is *equal* for all i , and all the repeat samples of $f(y_i, z_t)$ are wasted. To address this issue, Combes et al. [2020] imposes additional structural conditions on F that prevent the derivative from being too close to zero. **Our approach does not require such conditions, but it automatically enjoys a faster convergence rate when the conditions are satisfied.**

Algorithm 1: One Elimination Round of Unimodal Optimization

Require: Confidence parameter δ , interval $\Delta = [l, r]$, threshold $\tau \geq 0$, spacing GAP such that $\frac{u-l}{\text{GAP}}$ is an integer, confidence scaling constant $h > 0$ to be set by Lemma 2.

- 1: Define the points $l = x^{(1)} < x^{(2)} < \dots < x^{(N)} = r$ for $N = 1 + \frac{l-r}{\text{GAP}}$ equally spaced in the interval Δ .
 - 2: Compute $f(x^{(k)}, z^{(k)})$ for i.i.d. $z^{(1)}, \dots, z^{(N)}$.
 - 3: Given any $i, j \in \{1, \dots, N\}$, define $m_{ij} = \sum_{k=i}^j f(x^{(k)}, z^{(k)}) / (j - i + 1)$, $s_{ij} = h\sqrt{\log(2N/\delta)/\sqrt{j-i+1}}$.
 - 4: Define upper confidence bound $U(i, j) = m_{ij} + s_{ij}$ for each $i, j \in \{1, \dots, N\}$.
 - 5: Define lower confidence bound $L(i, j) = m_{ij} - s_{ij}$ for each $i, j \in \{1, \dots, N\}$.
 - 6: Let $S_l = \max\{x^{(i)} \mid \exists i \leq j \leq k \leq l \text{ s.t. } U(i, j) < L(k, l) - \tau\} \cup \{x^{(1)}\}$.
 - 7: Let $S_r = \min\{x^{(l)} \mid \exists i \leq j \leq k \leq l \text{ s.t. } U(k, l) < L(i, j) - \tau\} \cup \{x^{(N)}\}$.
 - 8: Let $I_{\text{best}} = \text{argmax}_{[x^{(i)}, x^{(j)}]} L(i, j)$.
 - 9: Return new interval $[S_l, S_r] \subset \Delta$ and the best interval I_{best} .
-

Our approach is inspired by prior work on continuum armed-bandits [Auer et al., 2007] that tries to identify an interval containing x_* . We extend their technique to unimodality so that the intervals can be used for elimination.

3.2 ALGORITHM

Our algorithm samples a large number of points only *once* (Line 5 of Algorithm 1). As a result, our algorithm cannot build good enough estimates of F at a single point to identify a pair of points (a, b) with $F(a) \geq F(b)$. Instead, **our algorithm identifies a pair of intervals** $I_a = [a_1, a_2]$ and $I_b = [b_1, b_2]$ for which there is some (unknown) $a \in I_a, b \in I_b$ satisfying $F(a) \geq F(b)$. Thus, we can either eliminate all points less than b_1 if $b_2 < a_1$, or all points greater than b_2 if $a_2 < b_1$. (Line 5-5).

Note that since $b \in [b_1, b_2]$, we do not eliminate as much as we would be able to if we had found the exact point b . That is, our algorithm operates with a potentially larger candidate interval than the existing algorithms at first, but casts a wider net for finding sub-intervals to eliminate. Intuitively speaking, our algorithm explores the geometry of the function F more broadly.

In order to find these intervals, we start by generating a grid of uniformly spaced points $x^{(1)}, x^{(2)}, \dots, x^{(N)}$ with $|x^{(i)} - x^{(i+1)}| = \Delta$ for some grid-spacing parameter Δ . Then, we evaluate $f(x^{(i)}, z^{(i)})$ for each $i \in \{1, \dots, N\}$ for independent samples z_1, \dots, z_N . Now, for each interval $I = [x_i, x_j]$, we can form the intersection $S_I = I \cap \{x^{(1)}, \dots, x^{(N)}\}$ and estimate the average $F_I = \frac{1}{|S_I|} \sum_{x^{(i)} \in S_I} F(x^{(i)})$ by $\frac{1}{|S_I|} \sum_{x^{(i)} \in S_I} f(x^{(i)}, z^{(i)})$. Moreover, for each I we can also form *confidence intervals* of width $O(\sqrt{\log(N^2/\delta)/\sqrt{|S_I|}})$ valid with probability $1 - \delta/N^2$. Since there are $O(N^2)$ possible sets S_I , the confidence intervals are simultaneously valid with probability $1 - \delta$. Finally, if any two intervals I_a and I_b are such that the confidence interval for I_a is larger and disjoint from that for I_b , then there must be $a \in I_a$ and $b \in I_b$ with $F(a) \geq F(b)$.

We formally specify this algorithm in Algorithm 1 and provide analysis in Lemma 1. The algorithm includes a "threshold" parameter τ , which is used for extensions to multiple dimensions in Section 4. For the remainder of this section, we will assume $\tau = 0$.

Lemma 1. *There is a universal constant h that can be provided to Algorithm 1 such that with probability at least $1 - \delta$, if $x_* \in \Delta$ and $[S_l, S_r]$ is the output of Algorithm 1 with any GAP and τ , then $x_* \in [S_l, S_r]$, and $\frac{1}{j-i+1} \sum_{k=i}^j F(x^{(k)}) \in [L(i, j), U(i, j)]$ for all i, j .*

Now, our full 1-D algorithm combines this interval shrinking routine with a doubling argument: we repeatedly call Algorithm 1 with Δ set to the previously returned $[S_l, S_r]$ and GAP set to the previous GAP divided by two. Intuitively, if Algorithm 1 does *not* significantly shrink the interval Δ , then this approach re-runs the algorithm with double the number of samples. Otherwise, if Δ shrinks by a constant factor, we are making sure to still take a reasonable number of samples in this new smaller interval.

The final ingredient in our algorithm is how to select the

final output \hat{x} . One plausible approach would be to pick a random element of the interval Δ . However, we opt for a more refined method: we pick a random element of the interval that has the largest lower confidence bound. This enables us to make claims about the quality of our final output even in the case that the algorithm never eliminates any intervals at all (i.e. $\Delta = [0, 1]$ always). The formal description is provided in Algorithm 2.

Algorithm 2: 1D Unimodal Optimization

Require: Confidence parameter δ , threshold $\tau \geq 0$, budget B .

- 1: Initialize $t = 1$, $\Delta_1 = [0, 1]$, $\text{GAP}_1 = 0.5$, $N = N_1 = 3$.
- 2: **while** $N \leq B$ **do**
- 3: Call Algorithm 1 with input $6\delta/\pi^2 t^2$, Δ_t , τ , GAP_t to obtain outputs $[S_l, S_r]$, I_{best} .
- 4: Let $\Delta_{t+1} = [S_l, S_r]$.
- 5: Set $\text{GAP}_{t+1} = \text{GAP}_t/2$.
- 6: Set $N_{t+1} = 1 + \frac{S_r - S_l}{\text{GAP}}$ //Budget to be consumed by next iteration.
- 7: Set $N = N + N_t$. //Total budget consumed at end of next iteration.
- 8: Set $t = t + 1$.
- 9: **end while**
- 10: Return a random element of I_{best} .

3.3 CONVERGENCE ANALYSIS

In order to prove convergence rates, we will need to make some additional assumptions about the shape of F beyond just unimodality. To see this, consider the function $F(0.7) = 1$ and $F(x) = 0$ for $x \neq 0.7$. Such an F is certainly unimodal, and yet clearly no algorithm can obtain any non-trivial bound on $F(x_*) - F(\hat{x})$. As our first assumption, we consider Lipschitz functions.

Definition 3 (Lipschitz). $F(x)$ is L -Lipschitz (i.e. $|F(x) - F(y)| \leq L|x - y|$ for all $x, y \in [0, 1]$).

Theorem 1. Suppose F is L -Lipschitz. Let \hat{x} be the output of Algorithm 2 with total sample budget B and input failure probability δ . Then, for any $\delta < 1/2$, there is a constant C and an event E that occurs with probability at least $1 - \delta$ such that:

$$\begin{aligned} \mathbb{E}[F(\hat{x})|E] &\geq F(x_*) - 3(LC^2 \text{GAP}_T \log(B/\delta))^{1/3} \\ &\geq F(x_*) - \frac{3(12LC^2 \log(B/\delta))^{1/3}}{B^{1/3}} \end{aligned}$$

where C is an absolute constant.

This result shows that the Algorithm 2 converges at a $O(1/B^{1/3})$ rate, or equivalently that $O(1/\epsilon^3)$ samples suffice to find an ϵ -suboptimal point for a Lipschitz objective F , matching classical rates for this setting [Kleinberg et al.,

2008]. The proof proceeds by observing that since Algorithm 2 never discards x_* , eventually it will explore a fine enough grid of points that there will be an interval whose lower-confidence-bound is within $O(1/B^{1/3})$ of $F(x_*)$, so that picking a random \hat{x} from the interval with largest lower-confidence-bound is guaranteed to be $O(1/B^{1/3})$ suboptimal.

Next, we consider the case of *smooth* rather than Lipschitz F . For smooth F , we can improve upon the Lipschitz analysis. The key idea is that the average value of F over an interval of width W is within $O(W^2)$ of the value of F at the midpoint of the interval. In contrast, if we assume only Lipschitz F , the average value may be $O(W)$ away from any given point in the interval. This improvement translates into a suboptimality of $O(1/B^{2/5})$ rather than $O(1/B^{1/3})$, as described formally in Theorem 2.

Definition 4. A function F is β -smooth if $F(x)$ is differentiable and $F'(x)$ is β -Lipschitz.

Theorem 2. Suppose F is β -smooth and that $F'(x_*) = 0$ (i.e. x_* is not on the boundary). Let \hat{x} be the output of Algorithm 2 with total sample budget B and input failure probability δ . Then, there is a constant C such that for any $\delta < 1/2$ there is an event E that occurs with probability at least $1 - \delta$ such that:

$$\mathbb{E}[F(\hat{x})|E] \geq F(x_*) - \frac{5(\beta C^4 \log(B/\delta)^2)^{1/5}}{B^{2/5}}$$

Theorem 1 and 2 establish baseline convergence rates for Lipschitz or smooth F . They have the desirable property that the algorithm need not use any knowledge of the Lipschitz or smoothness parameters L and β , so that the bounds hold with the tightest possible (unknown) values. On the other hand, the result is somewhat naive: we did not use any of the “elimination” power of the algorithm, and in fact the same result would hold if we simply applied one single round of Algorithm 2 and did not even require unimodality.

The power of our algorithm is that it automatically enjoys a better convergence rate if the function meets certain conditions, without any changes to the input or prior knowledge. The first such condition we explore is the following:

Assumption 1 ((γ, M) -Lipschitz lower bound). There is some γ and M such that for all x, y with $|x - x_*| \geq \gamma$ and $|y - x_*| \geq \gamma$, and $\text{sign}(x - x_*) = \text{sign}(y - x_*)$, we have $|F(x) - F(y)| \geq M|x - y|$ (i.e. if F is differentiable, then $|F'(x)| \geq M$ whenever $|x - x_*| \geq \gamma$).

Intuitively, these functions have a “difficult” interval of size 2γ around the optimum x_* . The lower bound on the slope of F outside this region allows our algorithm to rapidly eliminate subintervals and converge to a small interval near the optimum. We can improve upon Theorem 1, and establish that the asymptotic error is $\tilde{O}((\gamma/B)^{1/3} + 1/\sqrt{B})$. Notice

that the algorithm at no point requires knowledge of the parameter γ .

Theorem 3. *Suppose F is L -Lipschitz and satisfies Assumption 1. Let \hat{x} be the output of Algorithm 2 with total sample budget $B > 78$ and input failure probability δ . Then, there is a constant C such that for any $\delta < 1/2$, there is an event E that occurs with probability at least $1 - \delta$ such that:*

$$\begin{aligned} \mathbb{E}[F(\hat{x})|E] &\geq F(x_*) - 3(\text{LGAP}_T C^2 \log(B/\delta))^{1/3} \\ &\geq F(x_*) - \tilde{O}\left(\frac{L\gamma^{1/3}}{B^{1/3}} + \frac{1}{\sqrt{B}M^{1/3}}\right) \end{aligned}$$

where C is an absolute constant.

Finally, we consider the class of functions satisfying Assumption 2, which is a strict superset of function class studied by Combes et al. [2020] (denoted $\mathcal{U}_{[0,1]}$ in their paper). We show that our Algorithm 2 recovers the same $\tilde{O}(1/\sqrt{B})$ convergence rate in this setting, again without requiring any knowledge of problem parameters.

Assumption 2. *There exists A_1, A_2 and z such that for all x , $A_2|x - x_*|^z \geq F(x_*) - F(x) \geq A_1|x - x_*|^z$. Such functions need be neither Lipschitz nor smooth.*

Theorem 4. *Suppose F satisfies Assumption 2. Then, there is a constant C such that for all $\delta < 1/2$, for any B satisfying $B \geq \frac{24A_2^{1/z}}{C^{1/z} \log(B/\delta)^{1/2z}}$, with probability at least $1 - \delta$ Algorithm 2 with $\tau = 0$ guarantees:*

$$\begin{aligned} F(\hat{x}) &\geq F(x_*) - F(x_*) \\ &\quad - \max\left(12^z \frac{C^{\frac{2z}{2z+1}} A_2^{\frac{2z+2}{2z+1}} \log(B/\delta)^{\frac{z}{2z+1}}}{A_1 2^{\frac{zB}{24z+12}}}, \right. \\ &\quad \left. \frac{12^z 2^{\frac{z}{2z+1}} \sqrt{72} C A_2^{1+\frac{1}{2z}} \sqrt{\log(B/\delta)}}{\sqrt{2^{\frac{2z}{2z+1}} - 1} A_1^{1+\frac{1}{2z}} \sqrt{B}}\right) \\ &\geq F(x_*) - \tilde{O}(1/\sqrt{B}) \end{aligned}$$

It should be noted that the dependencies of Theorem 4 on the parameter z are slightly worse than in Combes et al. [2020]. However, this is mitigated by the fact that our function class is actually somewhat larger (Combes et al. [2020] consider a stronger version of Assumption 2 that controls local behavior of F even far from x_*). Further, due to the greater range of exploration of our algorithm, we are able to handle more general Lipschitz or smooth losses as described previously.

We would also like to note that our algorithm would be better than an algorithm which provides a third of the budget to CAB1 [Kleinberg, 2004], a third to [Combes et al., 2020] and picks the best between the two returned answers by testing on the remaining trials. In order to see this, let us look at the following informal function class.

Assumption 3 (Informal). *Let F be the class of unimodal functions that satisfy Assumption 2 only on a sub-interval $R \in [0, 1]$ containing the optimum, and outside of this sub-interval we make no further assumptions other than unimodality.*

Under Assumption 3, our algorithm will eliminate everything outside of R in $\text{poly}(1/|R|)$ time, using essentially a discretization style analysis. However, once this elimination occurs, the algorithm will automatically switch to the faster $O(1/\sqrt{B})$ rate. Note that for any $|R| > 0$, CAB1 will not achieve a $O(1/\sqrt{B})$ asymptotic rate (due to its discretization it cannot hope for better than $O(1/B^{1/3})$), while for $|R| < 1/4$, the algorithm in [Combes et al., 2020] may not converge. Thus, this is a class of functions for which we outperform the minimum of the two baselines. Moreover, this class of functions that are “well-behaved near the optimum” is reasonably natural: indeed our Figure 1 seems to exhibit such behavior.

4 MULTIDIMENSIONAL PROBLEMS VIA COORDINATE ASCENT

In this section, we describe a simple approach to generalize our unimodal optimization algorithm to higher dimensions. While in the previous section the threshold parameter τ was essentially a nuisance factor and could be safely set to zero, here we will need the threshold to overcome a technical difficulty: we would like to ensure that any time a point x is eliminated by Algorithm 2, it is possible to identify a x^+ for which $F(x_*) - F(x^+)$ is significantly smaller than $F(x_*) - F(x)$. We conjecture that τ is not actually required for this task, but this is left as an open question.

We consider functions $F : [0, 1]^d \rightarrow [0, 1]$ such that for every point $x \in [0, 1]^d$, the d 1-D functions given by restricting F to each coordinate axis through x are all unimodal. This assumption suggests a natural iterative strategy: we will initialize our algorithm at some $x_1 \in [0, 1]^d$. Then, we run a copy of our 1-D algorithm on each coordinate independently, choosing samples along each coordinate in a round-robin fashion. If one of the copies is able to successfully eliminate the original point w_1 , then we let the output of this copy be a next iterate w_2 and repeat the process.

The analysis of this method proceeds by showing that (1) $F(x_{t+1}) \geq F(x_t) + \tau$ for all t , and that (2) the number of samples required to identify x_{t+1} is not too large, so long as an appropriate x_{t+1} exists. Combined, this means that after at most roughly $1/\tau$ iterations of this procedure, we will converge to a τ -approximate local optimum (Definition 2)

We provide an analysis of this algorithm for the case of Lipschitz F in Theorem 5. It follows from an alternative analysis of Algorithm 2, which tells us that Algorithm 2 will quickly eliminate any non-local minimum point. For details, please see Lemma 5 in the appendix. Informally, this result tells us that with a budget of B samples, we will

Algorithm 3: Unimodal Coordinate Ascent

```

1:  $t = 1$ 
2: Choose an arbitrary  $w^t = (w_1^t, \dots, w_d^t) \in [0, 1]^d$ 
3: while  $N \leq B - d$  do
4:   Given a starting point  $w^t \in [0, 1]^d$ , initialize  $d$ 
   copies of Algorithm 2 with  $\delta = 6/dt^2\pi^2$ , budget set
   to  $\infty$  and threshold  $\tau$  where the  $i$ th copy considers
   the function  $F_i : [0, 1] \rightarrow [0, 1]$ 
    $F_i(x) = F(w_1^t, \dots, w_{i-1}^t, x, w_{i+1}^t, \dots, w_d^t)$ .
   Let  $\Delta_{ti}^t$  be the active interval associated with the  $i$ th
   copy.
   Let  $I_{\text{best},i}^t$  be the current value of  $I_{\text{best}}$  maintained by
   the  $i$ th copy.
5:   while  $w_i^t \in \Delta_{ti}^t$  for all  $i$  do
6:     if  $N \leq B - d$  then
7:       Returns  $w^t$ .
8:     end if
9:     Let each copy sample one additional point.
10:     $N = N + d$ .
11:  end while
12:  Suppose the  $j$ th copy eliminated  $w_j^t$  ( $w_j^t \notin \Delta_j^t$ ).
13:  Let  $w_j^{t+1}$  be a randomly selected point in  $I_{\text{best},j}^t$ .
14:  Set  $w^{t+1} = (w_1^t, \dots, w_{j-1}^t, w_j^{t+1}, w_{j+1}^t, \dots)$ .
15:   $t \leftarrow t + 1$ .
16: end while
17: Return  $w^t$ .

```

find an $\tilde{O}(d^{1/4}/B^{1/4})$ -approximate local minimum for a d -dimensional Lipschitz unimodal function.

Theorem 5. *There is a constant C such that, given a budget of B and failure probability $\delta < 1/2$, if we set $\tau = \max\left(\frac{2d \log(B/\delta)}{B}, \frac{37 \cdot (dC^2 L \log(B/\delta)^2)^{1/4}}{B^{1/4}}\right)$, Then with probability at least $1 - 2\delta$, Algorithm 3, returns a point \hat{w} that is a 3τ -approximate local minimum. That is, for each coordinate i , for any w_i^* that differs from \hat{w} only in the i th coordinate, we have $F(w_i^*) \leq F(\hat{w}) + 3\tau$*

5 EXPERIMENTS

We tested our algorithm on machine learning hyperparameter optimization (HPO) problems. We used the continuous variant of HPO-B benchmark by [Arango et al., 2021, Sec 5.4]. The blackbox functions being optimized are the (approximated) validation accuracy of models for different hyper-parameter settings. The models vary across several tasks like training SVMs, GLMNs etc on various datasets. The approximation is done through XGBoost surrogate functions. We used the "HPO-B-v3 Meta-test split" as detailed in Arango et al. [2021]. It consists of 86 tasks in 16 distinct search spaces. We followed the exact test protocols as defined in their open source package.

Benchmark algorithms. The state of the art for HPO often involves using Gaussian Process (GP) regressors to derive

surrogate functions and optimize them. We used the GP-UCB implementation included in the HPO-B open source package. It uses the GP implementation in BoTorch [Balandat et al., 2019] with UCB coefficient of 0.1 (this method was shown to have strong performance on the benchmarks [Arango et al., 2021]). We also use random search over the domain as another baseline. In order to be fair to all the algorithms, we initialize each of them at a random point for any of the benchmark tasks. All algorithms are run five times for every task.


Algorithm 4: Modified Unimodal Coordinate Ascent

```


1: Set  $N, t = 10$ . Choose 10 random points  $[0, 1]^d$  and set
   arbitrary  $w^t$  as the one with the best observed value.
2: while  $N \leq B - d$  do
3:   Given a starting point  $w^t \in [0, 1]^d$ , initialize  $d$ 
   copies of Algorithm 6 with  $\delta = 6/dt^2\pi^2$ , budget set
   to  $\infty$  where the  $i$ th copy considers the function
    $F_i : [0, 1] \rightarrow [0, 1]$ 
    $F_i(x) = F(w_1^t, \dots, w_{i-1}^t, x, w_{i+1}^t, \dots, w_d^t)$ .
   Let  $\Delta_{ti}^t$  be the active interval associated with the  $i$ th
   copy.
   Let  $I_{\text{best},i}^t$  be the current value of  $I_{\text{best}}$  maintained by
   the  $i$ th copy.
4:   while  $w_i^t \in \Delta_{ti}^t$  for all  $i$  do
5:     if  $N \leq B - d$  then
6:       Return  $w^t$ .
7:     end if
8:     Let  $i$  be the dimension sampled from distribution
     proportional to  $\{\exp(s_1), \dots, \exp(s_d)\}$ . Here  $s_i$ 
     is the observed standard deviation of the values
     for dimension  $i$  so far.
9:     Run one more epoch of Algorithm 6 for
     dimension  $i$ .
10:     $N = N + \text{<\#new points sampled>}$ .
11:  end while
12:  Suppose the  $j$ -th copy is one such copy that
  eliminated  $w_j^t$  ( $w_j^t \notin \Delta_j^t$ ) and  $\Delta_{t,j}$  is smallest
  among all eliminating dimensions.
13:  Let  $w_j^{t+1}$  be the best point selected in  $I_{\text{best},j}^t$ , in
  terms of observed value.
14:  Set  $w^{t+1} = (w_1^t, \dots, w_{j-1}^t, w_j^{t+1}, w_{j+1}^t, \dots)$  and
  begin a new round (reuse points if available).
15: end while
16: Return  $w^t$ .

```


Our implementation. We make some practical modifications to Algorithm 3 for the real world benchmarks, both to improve runtime and also to acknowledge that the true black-box functions might not exactly adhere to our assumptions. First, instead of comparing every possible pair of disjoint intervals (I_1, I_2) , we compare only pairs where $|I_1| = |I_2|$ and both are powers of 2. This significantly reduces the number of pairs to compare, and an inspection of our proof techniques shows that it will only harm constants




(a) Regret across all tasks



(b) Cumulative runtime



(c) Regret vs Dimension



(d) Regret on lower dimensional tasks

Figure 2: In (a), we plot the normalized regret as a function of number of trials (observations) averaged over 86 benchmark tasks. Each task is repeated 5 times and we plot the corresponding standard errors. It can be seen that our algorithm (Unimodal Ascent) outperforms the baselines over the majority of the x-axis. In (b), we compare the run-times of our algorithm and BoTorch GP-UCB as a function of number of trials. In (c), we plot the average regret of our algorithm and BoTorch-GP as a function of problem dimension. It can be seen that we outperform the GP on all dimensions except 8 and 18. In (d), we plot the normalized regret as a function of number of trials for problems with dimensions ≤ 6 .

in the bounds. Further, we never sample the same point twice, but instead *re-use* the previous sample every time it is requested in Algorithm 1. This again only harms constants in the bounds. Next, in all tasks, we start our algorithm with 10 random points in the domain, and then choose the point with the best observed value as the starting point for the unimodal coordinate ascent. In practice this leads to a good initialization of our algorithm. Note that we count those ten points as trials expended by our algorithm so that it is a fair comparison with the benchmarks. In Algorithm 3, if at any point there are multiple coordinates that can eliminate the current point (line 12), we choose the coordinate along which the maximum area has been eliminated. Finally, instead of letting each coordinate sample a point in line 9, we choose the coordinate that can sample the next point from a distribution that up-weights coordinates that have shown more variation in past trials. We also set the threshold parameter τ to zero. In the interest of space, we provide the pseudo-code of this modified algorithm as Algorithm 4 (more details in Appendix D).

Regret on the full benchmark. In Figure 2a, we plot the *normalized regret* as a function of number of trials/ observations. We adopt the definition of normalized regret from [Arango et al., 2021] i.e given the points $\{x_1, x_2, \dots, x_t\}$ chosen by the algorithm till time t , the normalized regret is

$$r(t) = \frac{F(x_*) - \max_{s \in [t]} F(x_s)}{F(x_*) - F(x_{min})}$$

where x_{min} is the point with the lowest objective value. It can be observed that our simple unimodal ascent algorithm outperforms a state of the art GP averaged over all the 86 real world benchmark HPO tasks. This is quite remarkable since the benchmark tasks might not strictly adhere to our assumptions and the GP has been carefully tuned by Arango et al. [2021] for good performance on the HPO-B tasks. We also provide a comparison w.r.t a version of the Zooming Algorithm [Kleinberg et al., 2008] in Appendix D.1.

Runtime. Figure 2b shows the cumulative runtime of our algorithm against the BoTorch GP as a function of number of observations. The logarithmic scale in the y-axis shows that our algorithm is orders of magnitude faster than the GP, owing to its simplicity. All our experiments were performed on Intel(R) Xeon(R) CPU E5-2620 v4 @ 2.10GHz machines with NVIDIA Tesla P100 GPU.

Indeed, it is well-known that generating the N -th suggestion from the GP takes $O(A(DN^2 + N^3) + B(DN + N^2))$ time, where D is search space dimension, A is the number of kernel hyper-parameter optimization (also known as *ARD*) iterations, and B is the number of acquisition function optimization iterations [Garnett, 2022].

In comparison, generating the N -th suggestion from Algorithm 4 takes $O(N^2)$ time, up to logarithmic factors. This

can be a significant advantage when tuning a lightweight model or optimizing with a large budget on the number of evaluations.

Dimension of search-space. Now we dive deeper into the behavior of the algorithms in search-spaces of varying dimensions. The HPO-B benchmarks contain search spaces of dimensions ranging from 2 to 18. We expect our algorithm to work significantly better than the GP for lower dimensional problems in the presence of approximately unimodal structure. This phenomenon can be observed in Figure 2c where we plot the average normalized regret achieved at the end of 100 trials, broken down by search-space dimension. Our gains are much more significant for lower dimensions. The BoTorch GP is only better than us for problems with dimension 18 and comparable to us on problems with dimension 8.

In Figure 2d, we plot regret as a function of number of trials, for tasks with dimension less than or equal to 6 (more than 38% of the search-spaces). This shows that we are significantly better than the GP for these tasks. In conclusion our algorithm is simpler and faster than the state of the art BoTorch GP, while being more performant on average on the HPO-B benchmark.

6 CONCLUSION

We describe a new algorithm for black-box optimization of unimodal functions. Our algorithm is based upon the intuitive idea that one can estimate the objective value at any given point by sampling several nearby points, rather than sampling the same point many times. This allows us to uniformly cover the input space and gain more information about the shape of the function while still employing a simple elimination test to remove suboptimal regions of the input quickly. We demonstrate theoretically that this method matches prior work, and empirically we are able to outperform more advanced methods based upon Gaussian processes. We conjecture that this capability arises from the unimodality assumption providing an informative prior about the overall shape of the objective (even if it is not strictly true in all practical settings).

Our work suggests at least three natural avenues for future investigation. Perhaps the most pressing is improving the method of extension to high dimensional problems. We adopt a relatively simple coordinate ascent strategy, while one might hope for a more refined algorithm that views all dimensions at once. Nevertheless, even this naive approach performs surprisingly well in our experimental study. Next, although our theoretical development matches prior work in the 1D setting, there are classes of functions missing from our analysis, notably the *convex* F , which are of course also unimodal. Thus, it would be valuable to improve the analysis or algorithm to achieve the optimal $O(1/\sqrt{T})$ convergence rates for convex functions.

References

- Alekh Agarwal, Dean P Foster, Daniel J Hsu, Sham M Kakade, and Alexander Rakhlin. Stochastic convex optimization with bandit feedback. *Advances in Neural Information Processing Systems*, 24, 2011.
- Rajeev Agrawal. The continuum-armed bandit problem. *SIAM journal on control and optimization*, 33(6):1926–1951, 1995.
- Sebastian Pineda Arango, Hadi S Jomaa, Martin Wistuba, and Josif Grabocka. Hpo-b: A large-scale reproducible benchmark for black-box hpo based on openml. *arXiv preprint arXiv:2106.06257*, 2021.
- Peter Auer, Ronald Ortner, and Csaba Szepesvári. Improved rates for the stochastic continuum-armed bandit problem. In *International Conference on Computational Learning Theory*, pages 454–468. Springer, 2007.
- Maximilian Balandat, Brian Karrer, Daniel R Jiang, Samuel Daulton, Benjamin Letham, Andrew Gordon Wilson, and Eytan Bakshy. Botorch: Programmable bayesian optimization in pytorch. 2019.
- Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- Herman Chernoff. Sequential design of experiments. *The Annals of Mathematical Statistics*, 30(3):755–770, 1959.
- Richard Combes and Alexandre Proutiere. Unimodal bandits: Regret lower bounds and optimal algorithms. In *International Conference on Machine Learning*, pages 521–529. PMLR, 2014.
- Richard Combes, Alexandre Proutière, and Alexandre Faulet. Unimodal bandits with continuous arms: Order-optimal regret without smoothness. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 4(1):1–28, 2020.
- John C Duchi, Michael I Jordan, Martin J Wainwright, and Andre Wibisono. Optimal rates for zero-order convex optimization: The power of two function evaluations. *IEEE Transactions on Information Theory*, 61(5):2788–2806, 2015.
- Matthias Feurer and Frank Hutter. Hyperparameter optimization. In *Automated machine learning*, pages 3–33. Springer, Cham, 2019.
- Roman Garnett. *Bayesian Optimization*. Cambridge University Press, 2022. in preparation.
- Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Karro, and David Sculley. Google vizier: A service for black-box optimization. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1487–1495, 2017.
- Elad Hazan, Adam Klivans, and Yang Yuan. Hyperparameter optimization: a spectral approach. In *International Conference on Learning Representations*, 2018.
- Kevin G Jamieson, Robert Nowak, and Ben Recht. Query complexity of derivative-free optimization. *Advances in Neural Information Processing Systems*, 25, 2012.
- Kwang-Sung Jun, Aniruddha Bhargava, Robert Nowak, and Rebecca Willett. Scalable generalized linear bandits: Online computation and hashing. *Advances in Neural Information Processing Systems*, 30, 2017.
- Jack Kiefer. Sequential minimax search for a maximum. *Proceedings of the American mathematical society*, 4(3): 502–506, 1953.
- Robert Kleinberg. Nearly tight bounds for the continuum-armed bandit problem. *Advances in Neural Information Processing Systems*, 17, 2004.
- Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. Multi-armed bandits in metric spaces. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 681–690, 2008.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.
- Luis Miguel Rios and Nikolaos V Sahinidis. Derivative-free optimization: a review of algorithms and comparison of software implementations. *Journal of Global Optimization*, 56(3):1247–1293, 2013.
- Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.
- Ohad Shamir. On the complexity of bandit and derivative-free stochastic convex optimization. In *Conference on Learning Theory*, pages 3–24. PMLR, 2013.
- Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems*, 25, 2012.
- Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.
- Kei Terayama, Masato Sumita, Ryo Tamura, and Koji Tsuda. Black-box optimization for automated discovery. *Accounts of Chemical Research*, 54(6):1334–1346, 2021.

Jia Yuan Yu and Shie Mannor. Unimodal bandits. In *International Conference on Machine Learning*. PMLR, 2011.

A PROOFS FOR SECTION 3

A.1 PROOF OF THEOREM 1

In this section we provide the missing proofs for the analysis of the 1-D unimodal optimization algorithm. Results that appear in the main text are restated in this section for reference.

We begin with a technical observation that allows us to show Theorem 1:

Lemma 2. *Using the notation of Algorithm 1, there exists a universal constant h such that with probability at least $1 - \delta$, for all $i, j \in \{1, \dots, N\}$,*

$$\begin{aligned} \frac{1}{j-i+1} \sum_{k=i}^j F(x^{(k)}) &\leq \frac{1}{j-i+1} \sum_{k=i}^j F(x^{(k)}) + h\sqrt{\log(N/\delta)}/\sqrt{j-i+1} \\ \frac{1}{j-i+1} \sum_{k=i}^j F(x^{(k)}) &\geq \frac{1}{j-i+1} \sum_{k=i}^j F(x^{(k)}) - h\sqrt{\log(N/\delta)}/\sqrt{j-i+1} \end{aligned}$$

Proof. Notice that the sequence $f(x^{(k)}, z^{(k)}) - F(x^{(k)})$ forms a bounded martingale difference. The Lemma then follows directly from Azuma-Hoeffding inequality followed by union bound over the $\frac{N(N-1)}{2}$ possible choices for i, j . \square

With this, we can prove:

Lemma 1. *There is a universal constant h that can be provided to Algorithm 1 such that with probability at least $1 - \delta$, if $x_\star \in \Delta$ and $[S_l, S_r]$ is the output of Algorithm 1 with any GAP and τ , then $x_\star \in [S_l, S_r]$, and $\frac{1}{j-i+1} \sum_{k=i}^j F(x^{(k)}) \in [L(i, j), U(i, j)]$ for all i, j .*

Proof. By Lemma 2, we have that $\frac{1}{j-i+1} \sum_{k=i}^j F(x^{(k)}) \in [L(i, j), U(i, j)]$ for all i, j with probability $1 - \delta$. The rest of the proof conditions on this high probability event.

Now, we will show that $x_\star \geq S_l$. The argument that $x_\star \leq S_r$ is completely symmetric. Thus, this will establish $x_\star \in [S_l, S_r]$.

First, if $S_l = x^{(1)} = u$, then $x_\star \geq S_l$ by assumption since $x_\star \in \Delta$. Otherwise, there must exist $i \leq j \leq k \leq l$ such that $S_l = x^{(i)}$ and $U(i, j) < L(k, l) - \tau$. Now, since

$$\begin{aligned} \frac{1}{j-i+1} \sum_{k=i}^j F(x^{(k)}) &\leq U(i, j) \leq L(k, l) \\ \frac{1}{l-k+1} \sum_{v=k}^l F(x^{(v)}) &\geq L(k, l) \end{aligned}$$

we have that there must exist some $b \in \{x^{(i)}, \dots, x^{(j)}\}$ and $a \in \{x^{(k)}, \dots, x^{(l)}\}$ such that $a \geq b$ and $F(a) \geq F(b)$. Thus, since F is unimodal we must have $x_\star \geq b \geq x^{(i)} = S_l$ as desired. \square

Next, we prove Lemma 3:

Lemma 3. *For all t , Algorithm 2 guarantees:*

1. *Let $x^{(1,t)}, \dots, x^{(N,t)}$ be the input points sampled by the t th call to Algorithm 1 and let $L^t(i, j)$ and $U^t(i, j)$ be the $L(i, j)$ and $U(i, j)$ values computed using the samples. Then with probability at least $1 - \delta$ we have $x_\star \in \Delta_t$ and for all t and $\frac{1}{j-i+1} \sum_{k=i}^j F(x^{(k,t)}) \in [L^t(i, j), U^t(i, j)]$ for all i, j, t .*
2. *There is a constant h such that the confidence widths s_{ij} computed in each round of Algorithm 2 satisfy $s_{ij} \leq C\sqrt{\log(B/\delta)}/\sqrt{j-i+1}$ for all ij for all rounds for some constant C for any $\delta < 1/2$.*
3. *$|\Delta_t|/\epsilon_t \in \mathbb{N}$ for all t (with probability 1).*
4. *At least $B/3$ budget is consumed (with probability 1).*
5. *The value of GAP_t after consuming N samples is at most $4/N$ (with probability 1). In particular, the final value of GAP_t is at most $12/B$.*

Proof. 1. The first statement follows from Theorem 1 combined with union bound and the fact that $\sum_{i=1}^{\infty} 1/i^2 = \pi^2/6$.

2. For the second statement, in the t th round, the confidence width is

$$\begin{aligned}
s_{ij} &= \frac{h\sqrt{\log(N_t\pi^2/3\delta)}}{\sqrt{j-i+1}} \\
&\leq \frac{h\sqrt{\log(B\pi^2/3\delta)}}{\sqrt{j-i+1}} \\
&\leq \frac{h\sqrt{\log(\pi^2/3) + \log(B/\delta)}}{\sqrt{j-i+1}} \\
&\leq \frac{h\sqrt{\log(2\pi^2/6)/\log(2)}\sqrt{\log(2) + \log(B/\delta)}}{\sqrt{j-i+1}} \\
&\leq \frac{h\sqrt{2\log(2\pi^2/6)/\log(2)}\sqrt{\log(B/\delta)}}{\sqrt{j-i+1}}
\end{aligned}$$

which shows the claim.

3. Observe that since $\Delta_{t+1} = [x^{(i)}, x^{(j)}]$ for some $x^{(i)}$ and $x^{(j)}$ produced by the t th call to Algorithm 1, we have that $|\Delta_{t+1}| = k\epsilon_t$ for some integer k . Now, since $\epsilon_t = 2^{-t}$ for all t , we clearly have $|\Delta_{t+1}|/\epsilon_{t+1} = 2k \in \mathbb{N}$.
4. Now, to show that Algorithm 2 consumes at least $B/3$ budget, we claim that the t th call to Algorithm 2 uses at most twice as many samples as the $t-1$ st call: $N_{t+1} \leq 2N_t$. This claim suffices to show the desired result: the algorithm runs until $N = \sum_{i=1}^{t+1} N_i \geq B$ and consumes $\sum_{i=1}^t N_i$ budget. If $\sum_{i=1}^t N_i < B/3$, then $N_t < B/3$. However, since $N_{t+1} \leq 2N_t$, this implies $N_{t+1} < 2B/3$ so that $\sum_{i=1}^{t+1} N_i < B$, which is a contradiction.

So, it remains to establish the claim. Notice that $N_t = 1 + \frac{|\Delta_t|}{\text{GAP}_t} = 1 + |\Delta_t|2^t$ for all t . Further, $|\Delta_{t+1}| \leq |\Delta_t|$ for all t . Therefore:

$$\begin{aligned}
N_{t+1} &= 1 + |\Delta_{t+1}|2^{t+1} \\
&\leq 1 + |\Delta_t|2^{t+1} \\
&= 2(1 + |\Delta_t|2^t) - 1 \\
&\leq 2N_t
\end{aligned}$$

as desired.

5. Finally, we show that when N samples are consumed, the smallest GAP_t value is at most $4/N$. Since $\text{GAP}_t = 2^{-t}$, this means that we need to show $2^t \geq B/4$. To see this, observe that $N_t \leq 1 + |\Delta_t|2^t \leq 1 + 2^t$ for all t . Therefore $N \leq \sum_{i=1}^t N_i \leq t + 2^{t+1} \leq 2^{t+2}$ so that indeed $B/4 \leq 2^t$.

□

A.2 PROOF OF THEOREM 1

Theorem 1. Suppose F is L -lipschitz. Let \hat{x} be the output of Algorithm 2 with total sample budget B and input failure probability δ . Then, for any $\delta < 1/2$, there is a constant C and an event E that occurs with probability at least $1 - \delta$ such that:

$$\begin{aligned}
\mathbb{E}[F(\hat{x})|E] &\geq F(x_*) - 3(LC^2\text{GAP}_T \log(B/\delta))^{1/3} \\
&\geq F(x_*) - \frac{3(12LC^2 \log(B/\delta))^{1/3}}{B^{1/3}}
\end{aligned}$$

where C is an absolute constant.

Proof. By Lemma 3, we have that $x_* \in \Delta_t$ for all t with probability at least $1 - \delta$. Then, let E be the event that $x_* \in \Delta_t$ for all t so that $P(E) \geq 1 - \delta$ as required. Let C be the same absolute constant as in Lemma 3.

Let T be the number of calls to Algorithm 1 made in Algorithm 2. Then, Lemma 3 also implies that $\text{GAP}_T \leq 12/B$. Let $x^{(1)}, \dots, x^{(N_T)}$ be the points sampled by Algorithm 1 in the T th and final round of Algorithm 2. Suppose $i \leq j$ are such

that $x^{(i)} \leq x^* \leq x^{(j)}$. Define $l = j - i + 1$. Note that such an interval exists for all $l \leq N_T$. Then we have the following identity:

$$\begin{aligned}
L(i, j) &\geq \frac{1}{j - i + 1} \sum_{k=i}^j f(x^{(k)}, z^{(j)}) - \frac{C\sqrt{\log(B/\delta)}}{\sqrt{l}} \\
&\geq \frac{1}{j - i + 1} \sum_{k=i}^j F(x^{(k)}) - \frac{2C\sqrt{\log(B/\delta)}}{\sqrt{l}} \\
&\geq F(x_*) - L|x^{(i)} - x^{(j)}| - \frac{2C\sqrt{\log(B/\delta)}}{\sqrt{l}} \\
&\geq F(x_*) - LGAP_T(l - 1) - \frac{2C\sqrt{\log(B/\delta)}}{\sqrt{l}}
\end{aligned}$$

Now, suppose that there is no such i, j satisfying $l = \lceil C^{2/3} \log(B/\delta)^{1/3} / (LGAP_t)^{2/3} \rceil$. Then we must have $N_T < \lceil C^{2/3} \log(B/\delta)^{1/3} / (LGAP_t)^{2/3} \rceil$, which implies $N_T - 1 < C^{2/3} \log(B/\delta)^{1/3} / (LGAP_t)^{2/3}$. Therefore, every point \hat{x} in the interval Δ satisfies:

$$\begin{aligned}
F(\hat{x}) &\geq F(x_*) - LGAP_T(N_T - 1) \\
&\geq F(x_*) - (LGAP_T C^2 \log(B/\delta))^{1/3}
\end{aligned}$$

Then, since Lemma 3 tells us $GAP_T \leq \frac{12}{B}$, this yields $F(\hat{x}) \leq F(x_*) - \frac{(12LC^2 \log(B/\delta))^{1/3}}{B^{1/3}}$ which would establish the desired result.

On the other hand, suppose that such an i, j exists. Then for such an i, j we have:

$$\begin{aligned}
L(i, j) &\geq F(x_*) - LGAP_T(l - 1) - \frac{2C\sqrt{\log(B/\delta)}}{\sqrt{l}} \\
&\geq F(x_*) - 3(LGAP_T C^2 \log(B/\delta))^{1/3}
\end{aligned}$$

Finally, since we return a random selection from the interval $[x^{(a)}, x^{(b)}]$ with largest value of $L(a, b)$, we have:

$$\begin{aligned}
\mathbb{E}[F(\hat{x})] &= \frac{1}{b - a + 1} \sum_{k=a}^b F(x^{(k)}) \\
&\geq L(a, b) \\
&\geq L(i, j) \\
&\geq F(x_*) - 3(LGAP_T C^2 \log(B/\delta))^{1/3}
\end{aligned}$$

Now, apply Lemma 3 to see $GAP_T \geq 12/B$ to establish the final result. □

A.3 PROOF OF THEOREM 2

Theorem 2. Suppose F is β -smooth and that $F'(x_*) = 0$ (i.e. x_* is not on the boundary). Let \hat{x} be the output of Algorithm 2 with total sample budget B and input failure probability δ . Then, there is a constant C such that for any $\delta < 1/2$ there is an event E that occurs with probability at least $1 - \delta$ such that:

$$\mathbb{E}[F(\hat{x})|E] \geq F(x_*) - \frac{5(\beta C^4 \log(B/\delta)^2)^{1/5}}{B^{2/5}}$$

Proof. This proof follows nearly identical reasoning to that of Theorem 1. The only significant change is that instead of bounding the average value of F over an interval containing x_* by the length of the interval, we leverage smoothness to bound it by the length of the interval squared, which will allow for a tighter tradeoff at the end of the proof.

Following the proof of Theorem 1, by Lemma 3, we have that $x_\star \in \Delta_t$ for all t with probability at least $1 - \delta$. Then, we again let E be the event that $x_\star \in \Delta_t$ for all t so that $P(E) \geq 1 - \delta$ as required.

Let T be the number of calls to Algorithm 1 made in Algorithm 2. Then, Lemma 3 also implies that $\text{GAP}_T \leq 12/B$. Let $x^{(1)}, \dots, x^{(N_T)}$ be the points sampled by Algorithm 1 in the T th and final round of Algorithm 2. Suppose $i \leq j$ are such that $x^{(i)} \leq x^\star \leq x^{(j)}$. Define $l = j - i + 1$. Note that such an interval exists for all $l \in [2, N_T]$.

Next, we make use of a standard identity for smooth losses (e.g. see Bubeck et al. [2015]). For all $x \in [0, 1]$, since $F'(x_\star) = 0$, we have:

$$F(x) \geq F(x_\star) - \frac{\beta}{2}|x - x_\star|^2$$

From this, we see that:

$$\begin{aligned} L(i, j) &\geq \frac{1}{j - i + 1} \sum_{k=i}^j f(x^{(k)}, z^{(j)}) - \frac{C\sqrt{\log(B/\delta)}}{\sqrt{l}} \\ &\geq \frac{1}{j - i + 1} \sum_{k=i}^j F(x^{(k)}) - \frac{2C\sqrt{\log(B/\delta)}}{\sqrt{l}} \\ &\geq F(x_\star) - \frac{\beta}{2}|x^{(i)} - x^{(j)}|^2 - \frac{2C\sqrt{\log(B/\delta)}}{\sqrt{l}} \\ &\geq F(x_\star) - \frac{\beta \text{GAP}_T^2 (l - 1)^2}{2} - \frac{2C\sqrt{\log(B/\delta)}}{\sqrt{l}} \end{aligned}$$

Now, suppose that there is no such i, j satisfying $l = \lceil C^{2/5} \log(B/\delta)^{1/5} / (\beta \text{GAP}_T^2)^{2/5} \rceil$. Then we must have $N_T < \lceil C^{2/5} \log(B/\delta)^{1/5} / (\beta \text{GAP}_T^2)^{2/5} \rceil$ so that $N_T - 1 < C^{2/5} \log(B/\delta)^{1/5} / (\beta \text{GAP}_T^2)^{2/5}$. Therefore, *every* point \hat{x} in the interval Δ satisfies:

$$\begin{aligned} F(\hat{x}) &\geq F(x_\star) - \frac{\beta \text{GAP}_T^2 (N_T - 1)^2}{2} \\ &\geq F(x_\star) - \frac{(\beta C^4 \log(B/\delta)^2 \text{GAP}_T^2)^{1/5}}{2} \end{aligned}$$

Next, since Lemma 3 tells us $\text{GAP}_T \leq \frac{12}{B}$, this yields $F(\hat{x}) \leq F(x_\star) - \frac{(144\beta C^4 \log(B/\delta)^2)^{1/5}}{2B^{2/5}} \geq F(x_\star) - \frac{2(\beta C^4 \log(B/\delta)^2)^{1/5}}{B^{2/5}}$ which would establish the desired result.

On the other hand, suppose that such an i, j exists. Then for such an i, j we have:

$$\begin{aligned} L(i, j) &\geq F(x_\star) - \frac{\beta \text{GAP}_T^2 (l - 1)^2}{2} - \frac{2C\sqrt{\log(B/\delta)}}{\sqrt{l}} \\ &\geq F(x_\star) - \frac{3(144\beta C^4 \log(B/\delta)^2)^{1/5}}{2B^{2/5}} \\ &\quad F(x_\star) - \frac{5(\beta C^4 \log(B/\delta)^2)^{1/5}}{B^{2/5}} \end{aligned}$$

Finally, since we return a random selection from the interval $[x^{(a)}, x^{(b)}]$ with largest value of $L(a, b)$, we have:

$$\begin{aligned} \mathbb{E}[F(\hat{x})] &= \frac{1}{b - a + 1} \sum_{k=a}^b F(x^{(k)}) \\ &\geq L(a, b) \\ &\geq L(i, j) \\ &\geq F(x_\star) - \frac{5(\beta C^4 \log(B/\delta)^2)^{1/5}}{B^{2/5}} \end{aligned}$$

which again establishes the desired final result. \square

A.4 PROOF OF LEMMA 6

Lemma 4 shows that indeed Algorithm 2 quickly eliminates the easy region.

Lemma 4. *Suppose F satisfies Assumption 1. Let $\Delta_t = [l_t, r_t]$ be the t th interval produced by Algorithm 2. Then with probability at least $1 - \delta$, for all $t > 1$ we have:*

$$\begin{aligned} l_t &\geq x_\star - 6 \cdot 2^{-t} - 3 \left[\max \left(\gamma, \frac{2^{5/3} C^{2/3} \log(B/\delta)^{1/3} 2^{-t/3}}{M^{2/3}} \right) + \frac{\tau}{M} \right] \\ r_t &\leq x_\star + 6 \cdot 2^{-t} + 3 \left[\max \left(\gamma, \frac{2^{5/3} C^{2/3} \log(B/\delta)^{1/3} 2^{-t/3}}{M^{2/3}} \right) + \frac{\tau}{M} \right] \end{aligned}$$

Proof. As before, we consider the event E that $x_\star \in \Delta_t$ for all t , and in each call to Algorithm 1, for all i, j , $\sum_{k=i}^j F(x^{(k)}) \in [L(i, j), U(i, j)]$ for all i, j . By Lemma 3, $P(E) \geq 1 - \delta$. The entire argument is conditioned on this event E . We will prove the statement for r_t . The argument for l_t is completely symmetric. For $t = 1$ the statement is trivially true since $6 \cdot 2^{-1} \geq 1$.

Fix some value of t . Suppose for purposes of induction that $r_k \leq x_\star + 6 \cdot 2^{-k} + 3 \left[\max \left(\gamma, \frac{2^{5/3} C^{2/3} \log(B/\delta)^{1/3} 2^{-k/3}}{M^{2/3}} \right) + \frac{\tau}{M} \right]$ for all $k \leq t$. We will show $r_{t+1} \leq x_\star + 6 \cdot 2^{-(t+1)} + 3 \left[\max \left(\gamma, \frac{2^{5/3} C^{2/3} \log(B/\delta)^{1/3} 2^{-(t+1)/3}}{M^{2/3}} \right) + \frac{\tau}{M} \right]$. First, suppose that already we have $r_t \leq x_\star + 6 \cdot 2^{-(t+1)} + 3 \left[\max \left(\gamma, \frac{2^{5/3} C^{2/3} \log(B/\delta)^{1/3} 2^{-(t+1)/3}}{M^{2/3}} \right) + \frac{\tau}{M} \right]$. Then there is nothing left to prove since $r_{t+1} \leq r_t$. Thus, we may suppose $r_t > x_\star + 6 \cdot 2^{-(t+1)} + 3 \left[\max \left(\gamma, \frac{2^{5/3} C^{2/3} \log(B/\delta)^{1/3} 2^{-(t+1)/3}}{M^{2/3}} \right) + \frac{\tau}{M} \right] = x_\star + 3 \cdot 2^{-t} + 3 \left[\max \left(\gamma, \frac{2^{4/3} C^{2/3} \log(B/\delta)^{1/3} 2^{-t/3}}{M^{2/3}} \right) + \frac{\tau}{M} \right]$.

Next, define $k = \left\lceil \max \left(\frac{\gamma}{\text{GAP}_t}, \frac{2^{4/3} C^{2/3} \log(B/\delta)^{1/3}}{M^{2/3} \text{GAP}_t^{2/3}} \right) + \frac{\tau}{M \text{GAP}_t} \right\rceil$. Notice that

$$\begin{aligned} k \text{GAP}_t &\geq \max \left(\gamma, \frac{2^{4/3} C^{2/3} \log(B/\delta)^{1/3} \text{GAP}_t^{1/3}}{M^{2/3}} \right) + \frac{\tau}{M} \\ 3k \text{GAP}_t &\leq 3 \max \left(\gamma, \frac{2^{4/3} C^{2/3} \log(B/\delta)^{1/3} \text{GAP}_t^{1/3}}{M^{2/3}} \right) + \frac{3\tau}{M} + 3\text{GAP}_t \end{aligned}$$

Therefore, since $\text{GAP}_t = 2^{-t}$ and $r_t > x_\star + 3 \cdot 2^{-t} + 3 \left[\max \left(\gamma, \frac{2^{4/3} C^{2/3} \log(B/\delta)^{1/3} 2^{-t/3}}{M^{2/3}} \right) + \frac{\tau}{M} \right]$, we have $r_t \geq x_\star + 3k \text{GAP}_t$.

Let $x^{(1)}, \dots, x^{(N_t)}$ be the points queried during the t th call to Algorithm 1. Then, since $r_t > x_\star + 3k \text{GAP}_t$, there exists i such that $i + 3k \leq N_t$ and $x^{(i)} \leq x_\star \leq x^{(i+1)}$ and $x^{(i+k)} \leq x_\star + k \text{GAP}_t$.

Then, letting $U(i, j)$ and $L(i, j)$ be the bounds defined in the t th call to Algorithm 1, we have:

$$\begin{aligned} L(i+1, i+k) &\geq F(x_\star + k \text{GAP}_t) - 2 \frac{C \sqrt{\log(B/\delta)}}{\sqrt{k}} \\ &\geq F(x_\star + k \text{GAP}_t) - 2^{1/3} C^{2/3} M^{1/3} \log(B/\delta)^{1/3} \text{GAP}_t^{1/3} \end{aligned}$$

Next, observe that $x_\star + k \text{GAP}_t \geq x_\star + \gamma + \frac{\tau}{M}$. Therefore, by Assumption 1, we have:

$$\begin{aligned} F(x_\star + 2k \text{GAP}_t) &\leq F(x_\star + k \text{GAP}_t) - k \text{GAP}_t M \\ &\leq F(x_\star + k \text{GAP}_t) - M \max \left(\gamma, \frac{2^{4/3} C^{2/3} \log(B/\delta)^{1/3} \text{GAP}_t^{1/3}}{M^{2/3}} \right) - \tau \\ &\leq F(x_\star + k \text{GAP}_t) - 2^{4/3} C^{2/3} M^{1/3} \log(B/\delta)^{1/3} \text{GAP}_t^{1/3} - \tau \end{aligned}$$

Finally, we have:

$$\begin{aligned}
U(i+1+2k, i+3k) &\leq F(x_\star + 2k\text{GAP}_t) + 2\frac{C\sqrt{\log(B/\delta)}}{\sqrt{k}} \\
&\leq F(x_\star + k\text{GAP}_t) - 2^{4/3}C^{2/3}M^{1/3}\log(B/\delta)^{1/3}\text{GAP}_t^{1/3} - \tau + 2\frac{C\sqrt{\log(B/\delta)}}{\sqrt{k}} \\
&\leq F(x_\star + k\text{GAP}_t) - 2^{4/3}C^{2/3}M^{1/3}\log(B/\delta)^{1/3}\text{GAP}_t^{1/3} + 2^{1/3}C^{2/3}M^{1/3}\log(B/\delta)^{1/3}\text{GAP}_t^{1/3} - \tau \\
&\leq F(x_\star + k\text{GAP}_t) - 2^{1/3}C^{2/3}M^{1/3}\log(B/\delta)^{1/3}\text{GAP}_t^{1/3} - \tau
\end{aligned}$$

Thus, we have $U(i+1+2k, i+3k) < L(i+1, i+k) - \tau$ so that

$$\begin{aligned}
r_{t+1} &\leq x_\star + 3k\text{GAP}_t \\
&\leq x_\star + 3\text{GAP}_t + 3\left[\max\left(\gamma, \frac{2^{4/3}C^{2/3}\log(B/\delta)^{1/3}\text{GAP}_t^{1/3}}{M^{2/3}}\right) + \frac{\tau}{M}\right] \\
&\leq x_\star + 6\text{GAP}_{t+1} + 3\left[\max\left(\gamma, \frac{2^{5/3}C^{2/3}\log(B/\delta)^{1/3}\text{GAP}_{t+1}^{1/3}}{M^{2/3}}\right) + \frac{\tau}{M}\right]
\end{aligned}$$

□

A.5 PROOF OF THEOREM 6

Here we prove Theorem 6. The full version of this result is stated with all constants included below.

Theorem 6. Suppose F is L -Lipschitz and satisfies Assumption 1. Let \hat{x} be the output of Algorithm 2 with total sample budget $B > 78$ and input failure probability δ . Then, for any $\delta < 1/2$, there is an event E that occurs with probability at least $1 - \delta$ such that:

$$\begin{aligned}
\mathbb{E}[F(\hat{x})|E] &\geq F(x_\star) - 3(L\text{GAP}_T C^2 \log(B/\delta))^{1/3} \\
&\geq F(x_\star) - 3\max\left(\frac{(LC^2 \log(B/\delta))^{1/3}}{2^{B/351}}, \frac{(54LC^2 \log(B/\delta)(\gamma + \frac{\tau}{M}))^{1/3}}{B^{1/3}}, \frac{C^{1/3}\sqrt{\log(B/\delta)}27^{1/3}2^{5/3}}{\sqrt{B}M^{1/3}\sqrt{2^{2/3}-1}}\right)
\end{aligned}$$

where C is an absolute constant.

Proof. As previously, we define E to be the event that the statements in Lemma 3 hold and let C be the absolute constant implied by the same Lemma.

Now, we improve upon Theorem 1 by using Lemma 4 to show that GAP_T can become smaller than $O(1/B)$, decreasing to roughly $O(\gamma/B)$.

To see this, notice that at each round we have $N_t = 1 + |\Delta_t|2^t$ and by Lemma 4, we have:

$$\begin{aligned}
|\Delta_t| &\leq 12 \cdot 2^{-t} + 6\left[\max\left(\gamma, \frac{2^{5/3}C^{2/3}\log(B/\delta)^{1/3}2^{-t/3}}{M^{2/3}}\right) + \frac{\tau}{M}\right] \\
&\leq 12 \cdot 2^{-t} + 6\gamma + 6\frac{\tau}{M} + 6\frac{2^{5/3}C^{2/3}\log(B/\delta)^{1/3}2^{-t/3}}{M^{2/3}}
\end{aligned}$$

Further, by Lemma 3, $\sum_{i=1}^t N_t \geq B/3$. Therefore if Algorithm 2 runs for T rounds,

$$\begin{aligned}
B/3 &\leq 13T + 6\left(\gamma + \frac{\tau}{M}\right)2^{T+1} + \sum_{t=1}^T 6\frac{2^{5/3}C^{2/3}\log(B/\delta)^{1/3}2^{2t/3}}{M^{2/3}} \\
&\leq 13T + 6\left(\gamma + \frac{\tau}{M}\right)2^{T+1} + 6\frac{2^{7/3}C^{2/3}\log(B/\delta)^{1/3}2^{2T/3}}{M^{2/3}(2^{2/3}-1)}
\end{aligned}$$

Now, from this we see that any value of T that satisfies $B/9 \geq 13T$, $B/9 \geq 6(\gamma + \frac{\tau}{M})2^{T+1}$ and $B/9 \geq 6 \frac{2^{7/3} C^{2/3} \log(B/\delta)^{1/3} 2^{2T/3}}{M^{2/3}(2^{2/3}-1)}$ must be a lower-bound for the true value. Therefore,

$$\frac{1}{\text{GAP}_T} = 2^T \geq \min \left(2^{B/117}, \frac{B}{54(\gamma + \frac{\tau}{M})}, \frac{B^{3/2} M (2^{2/3} - 1)^{3/2}}{27^{3/2} \cdot 2^5 C \sqrt{\log(B/\delta)}} \right)$$

The result now follows from Theorem 1. \square

B PROOF OF THEOREM 4

Theorem 4. Suppose F satisfies Assumption 2. Then, there is a constant C such that for all $\delta < 1/2$, for any B satisfying $B \geq \frac{24A_2^{1/z}}{C^{1/z} \log(B/\delta)^{1/2z}}$, with probability at least $1 - \delta$ Algorithm 2 with $\tau = 0$ guarantees:

$$\begin{aligned} F(\hat{x}) &\geq F(x_\star) - F(x_\star) \\ &\quad - \max \left(12^z \frac{C^{\frac{2z}{2z+1}} A_2^{\frac{2z+2}{2z+1}} \log(B/\delta)^{\frac{z}{2z+1}}}{A_1 2^{\frac{zB}{24z+12}}}, \right. \\ &\quad \left. \frac{12^z 2^{\frac{z}{2z+1}} \sqrt{72} C A_2^{1+\frac{1}{2z}} \sqrt{\log(B/\delta)}}{\sqrt{2^{\frac{2z}{2z+1}} - 1} A_1^{1+\frac{1}{2z}} \sqrt{B}} \right) \\ &\geq F(x_\star) - \tilde{O}(1/\sqrt{B}) \end{aligned}$$

Proof. As before, we consider the event E that $x_\star \in \Delta_t$ for all t , and in each call to Algorithm 1, for all i, j , $\sum_{k=i}^j F(x^{(k)}) \in [L(i, j), U(i, j)]$ for all i, j . By Lemma 3, $P(E) \geq 1 - \delta$. The entire argument is again conditioned on this event E .

In the i th round, we make at most $1 + 2^i$ queries. Therefore, since we must consume at least $B/3$ budget (by Lemma 6), the total number of rounds t must satisfy $t + 2^{t+1} \geq B/3$ so that $2^{t+2} \geq B/3$. In particular, since $B \geq \frac{24A_2^{1/z}}{C^{1/z} \log(B/\delta)^{1/2z}}$, we have

$$t \geq \log_2 \left(\frac{2A_2^{1/z}}{C^{1/z} \log(B/\delta)^{1/2z}} \right)$$

Now, let $\Delta_t = [l_t, r_t]$. Our bound on t implies:

$$\frac{C^{\frac{z}{2z+1}} \log(B/\delta)^{\frac{1}{2z+1}}}{A_2^{\frac{2}{2z+1}} \text{GAP}_t^{\frac{2z}{2z+1}}} \geq 2$$

For all such t , we will show:

$$\begin{aligned} r_t &\leq x_\star + 6 \frac{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z+2}} \log(B/\delta)^{\frac{1}{2z+1}} \text{GAP}_t^{\frac{1}{2z+1}}}{A_1^{\frac{1}{z}}} \\ l_t &\geq x_\star - 6 \frac{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z+2}} \log(B/\delta)^{\frac{1}{2z+1}} \text{GAP}_t^{\frac{1}{2z+1}}}{A_1^{\frac{1}{z}}} \end{aligned}$$

We will write the proof for r_t , and for notational convenience we will establish the bound for r_{t+1} rather than r_t . The argument for l_t is symmetric.

Suppose $r_t \leq x_\star + 6 \frac{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z+2}} \log(B/\delta)^{\frac{1}{2z+1}} \text{GAP}_{t+1}^{\frac{1}{2z+1}}}{A_1^{\frac{1}{z}}}$. Then clearly $r_{t+1} \leq r_t \leq x_\star + 6 \frac{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z+2}} \log(B/\delta)^{\frac{1}{2z+1}} \text{GAP}_{t+1}^{\frac{1}{2z+1}}}{A_1^{\frac{1}{z}}}$

and so we are done. So, let us suppose $r_t > x_\star + 6 \frac{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z+2}} \log(B/\delta)^{\frac{1}{2z+1}} \text{GAP}_{t+1}^{\frac{1}{2z+1}}}{A_1^{\frac{1}{z}}} \geq x_\star +$

$3 \frac{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z^2+z}} \log(B/\delta)^{\frac{1}{2z+1}} \text{GAP}_t^{\frac{1}{2z+1}}}{A_1^{\frac{1}{z}}}$. Now, define the constants:

$$\begin{aligned} y &= \frac{C^{\frac{2}{2z+1}} \log(B/\delta)^{\frac{1}{2z+1}}}{A_2^{\frac{2}{2z+1}} \text{GAP}_t^{\frac{2z}{2z+1}}} \\ k &= \lceil y - 1 \rceil \\ d &= \left\lceil \left(\frac{A_2}{A_1} \right)^{1/z} y + 1 \right\rceil = \left\lceil \frac{A_2^{\frac{1}{2z^2+z}} C^{\frac{2}{2z+1}} \log(B/\delta)^{\frac{1}{2z+1}}}{A_1^{\frac{1}{z}} \text{GAP}_t^{\frac{2z}{2z+1}}} + 1 \right\rceil \end{aligned}$$

Notice that $k \geq 1$ by our assumption on t and that $k \leq d \leq 2 \left(\frac{A_2}{A_1} \right)^{1/z} y$.

Let $x^{(1)}, \dots, x^{(N_t)}$ be the points queried in round t . Then, since $r_t > x_\star + 3 \frac{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z^2+z}} \log(B/\delta)^{\frac{1}{2z+1}} \text{GAP}_t^{\frac{1}{2z+1}}}{A_1^{\frac{1}{z}}} = x_\star + 3 \text{GAP}_t \left(\frac{A_2}{A_1} \right)^{1/z} y$, there exists i such that $i + k + d \leq N_t$ and $x^{(i)} \leq x_\star \leq x^{(i+1)} \leq \dots \leq x^{(i+k)}$. Therefore, we have:

$$\begin{aligned} L(i, j) &\geq F(x_\star) - A_2(k \text{GAP}_t)^z - 2 \frac{C \sqrt{\log(B/\delta)}}{\sqrt{k+1}} \\ &\geq F(x_\star) - A_2(y \text{GAP}_t)^z - 2 \frac{C \sqrt{\log(B/\delta)}}{\sqrt{y}} \\ &= F(x_\star) - 3 A_2^{\frac{1}{2z+1}} C^{\frac{2z}{2z+1}} \log(B/\delta)^{\frac{z}{2z+1}} \text{GAP}_T^{\frac{z}{2z+1}} \end{aligned}$$

Similarly, we have

$$\begin{aligned} U(i+d, i+k+d) &\leq F(x_\star) - A_1(d \text{GAP}_t)^z + 2 \frac{C \sqrt{\log(B/\delta)}}{\sqrt{k+1}} \\ &\geq F(x_\star) - A_1((A_2/A_1)^{1/z} y \text{GAP}_t)^z - 2 \frac{C \sqrt{\log(B/\delta)}}{\sqrt{y}} \\ &= F(x_\star) - 3 A_2^{\frac{1}{2z+1}} C^{\frac{2z}{2z+1}} \log(B/\delta)^{\frac{z}{2z+1}} \text{GAP}_T^{\frac{z}{2z+1}} \end{aligned}$$

Thus, we have that $U(i+d, i+k+d) \leq L(i, i+k)$ so that in the next round we have

$$\begin{aligned} r_{t+1} &\leq x^{(i+k+d)} \\ &\leq x_\star + \text{GAP}_t(k+d) \\ &\leq x_\star + 3 \text{GAP}_t \left(\frac{A_2}{A_1} \right)^{1/z} y \\ &\leq x_\star + 3 \frac{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z^2+z}} \log(B/\delta)^{\frac{1}{2z+1}} \text{GAP}_t^{\frac{1}{2z+1}}}{A_1^{\frac{1}{z}}} \\ &\leq x_\star + 6 \frac{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z^2+z}} \log(B/\delta)^{\frac{1}{2z+1}} \text{GAP}_{t+1}^{\frac{1}{2z+1}}}{A_1^{\frac{1}{z}}} \end{aligned}$$

Thus, so long as $t \geq \log_2 \left(\frac{2 A_2^{1/z}}{C^{1/z} \log(B/\delta)^{1/2z}} \right)$, we have that

$$|\Delta_t| \leq 12 \frac{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z^2+z}} \log(B/\delta)^{\frac{1}{2z+1}} \text{GAP}_t^{\frac{1}{2z+1}}}{A_1^{\frac{1}{z}}}$$

Further, since $\text{GAP}_t = 2^{-t}$, we have $t \geq \log_2(B/12)$ so that the condition on T is implied by $B \geq \frac{24A_2^{1/z}}{C^{1/z} \log(B/\delta)^{1/2z}}$.

Now, for any $\hat{x} \in \Delta_T$, we have:

$$\begin{aligned} F(\hat{x}) &\geq F(x_\star) - A_2 |\Delta_T|^z \\ &\geq F(x_\star) - 12^z \frac{C^{\frac{2z}{2z+1}} A_2^{\frac{2z+2}{2z+1}} \log(B/\delta)^{\frac{z}{2z+1}} \text{GAP}_T^{\frac{z}{2z+1}}}{A_1} \end{aligned}$$

So, it remains to compute GAP_T for the final round T . To this end, notice that in the rounds with $t < \log_2\left(\frac{2A_2^{1/z}}{C^{1/z} \log(B/\delta)^{1/2z}}\right)$,

we consume $N_1 + \dots + N_t \leq \frac{4A_2^{1/z}}{C^{1/z} \log(B/\delta)^{1/2z}}$ samples. In subsequent rounds, we have

$$\begin{aligned} N_t &= 1 + |\Delta_t| 2^t \\ &\leq 1 + 12 \frac{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z^2+z}} \log(B/\delta)^{\frac{1}{2z+1}} 2^{\frac{2zt}{2z+1}}}{A_1^{\frac{1}{z}}} \end{aligned}$$

Now, we have

$$\sum_{t=1}^T 2^{\frac{2zt}{2z+1}} \leq \frac{2^{\frac{2z}{2z+1}}}{2^{\frac{2z}{2z+1}} - 1} 2^{\frac{2zT}{2z+1}}$$

so that in T rounds our total number of samples can be bounded:

$$\sum_{t=1}^T N_t \leq \frac{4A_2^{1/z}}{C^{1/z} \log(B/\delta)^{1/2z}} + T + 12 \frac{2^{\frac{2z}{2z+1}}}{2^{\frac{2z}{2z+1}} - 1} \frac{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z^2+z}} \log(B/\delta)^{\frac{1}{2z+1}} 2^{\frac{2zT}{2z+1}}}{A_1^{\frac{1}{z}}}$$

Now, since we have $\sum_{t=1}^T N_t \geq B/3$ by Lemma 3, this means that any T such that

$$\begin{aligned} \frac{4A_2^{1/z}}{C^{1/z} \log(B/\delta)^{1/2z}} + T &\leq B/6 \\ 12 \frac{2^{\frac{2z}{2z+1}}}{2^{\frac{2z}{2z+1}} - 1} \frac{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z^2+z}} \log(B/\delta)^{\frac{1}{2z+1}} 2^{\frac{2zT}{2z+1}}}{A_1^{\frac{1}{z}}} &\leq B/6 \end{aligned}$$

is a lower bound for the true number of rounds. From this, we obtain that for all $B \geq \frac{48A_2^{1/z}}{C^{1/z} \log(B/\delta)^{1/2z}}$,

$$2^{\frac{zT}{2z+1}} \geq \min\left(2^{\frac{zB}{24z+12}}, \sqrt{\frac{2^{\frac{2z}{2z+1}} - 1}{72 \cdot 2^{\frac{2z}{2z+1}}} \frac{BA_1^{\frac{1}{z}}}{C^{\frac{2}{2z+1}} A_2^{\frac{1}{2z^2+z}} \log(B/\delta)^{\frac{1}{2z+1}}}}\right)$$

Thus, putting all together we have:

$$\begin{aligned} F(\hat{x}) &\geq F(x_\star) - 12^z \frac{C^{\frac{2z}{2z+1}} A_2^{\frac{2z+2}{2z+1}} \log(B/\delta)^{\frac{z}{2z+1}} \text{GAP}_T^{\frac{z}{2z+1}}}{A_1} \\ &\geq F(x_\star) - \max\left(12^z \frac{C^{\frac{2z}{2z+1}} A_2^{\frac{2z+2}{2z+1}} \log(B/\delta)^{\frac{z}{2z+1}}}{A_1 2^{\frac{zB}{24z+12}}}, \frac{12^z 2^{\frac{z}{2z+1}} \sqrt{72} C A_2^{1+\frac{1}{2z}} \sqrt{\log(B/\delta)}}{A_1^{1+\frac{1}{2z}} \sqrt{B}}\right) \end{aligned}$$

□

C PROOFS FOR SECTION 4

The first step is to show that the 1-D algorithm will be able to eliminate the point w^t (i.e. have $w_j^t \notin \Delta_j^t$ for some j) in a reasonable number of samples, provided that $F(w_t)$ is not a 3τ -approximate local minimum point (see Definition 2).

To this end, we provide the following additional result for the 1-D Algorithm 2:

Lemma 5. *Suppose $F : [0, 1] \rightarrow [0, 1]$ is L -Lipschitz. Let x be a point such that $F(x) < F(x_\star) - 3\tau$. Then there is an even E that occurs with probability at least $1 - \delta$ such that conditioned on E , after at most $O\left(\frac{Lh^2 \log(N/\delta)}{(F(x_\star) - F(x))^3}\right)$, Algorithm 2 eliminates x . Moreover, regardless of when x is eliminated, if \hat{x} is a randomly selected point in I_{best} ,*

$$\mathbb{E}[F(\hat{x})|E] \geq F(x) + \tau$$

Proof. As usual, we condition on the probability $1 - \delta$ event E that all confidence intervals created by Algorithm 2 are valid, as described by Lemma 3.

Without loss of generality, suppose $x \leq x^\star$ (the proof is completely symmetric if $x \geq x^\star$). Let $w = x^\star - \frac{F(x^\star) - F(x)}{4L}$ and $y = x + \frac{F(x^\star) - F(x)}{4L}$. Observe that by L -lipschitzness and unimodality, $F(x^\star) \geq F(w) \geq F(x^\star) - \frac{1}{4}(F(x^\star) - F(x))$ and $F(x) \leq F(y) \leq F(x) + \frac{1}{4}(F(x^\star) - F(x))$ so that $F(w) \geq F(y) + \frac{1}{2}(F(x^\star) - F(x))$.

Now, by Lemma 3, after taking $N = \lceil \frac{20Lk}{F(x^\star) - F(x)} \rceil$ samples for any $k \in \mathbb{N}$, we must have $\epsilon < \frac{F(x^\star) - F(x)}{5Lk}$. Thus, there is an interval $[x^{(a)}, x^{(b)}]$ with $b - a + 1 = k$ contained in (x, y) and an $[x^{(c)}, x^{(d)}]$ with $d - c + 1 = k$ contained in (w, x^\star) .

Now, the upper bound $U(a, b)$ for the interval $[a, b]$ satisfies (where C is the universal constant from Lemma 3):

$$\begin{aligned} U(a, b) &\leq \frac{1}{k} \sum_{i=a}^b f(x^{(i)}, z^{(i)}) + C \frac{\sqrt{\log(N/\delta)}}{\sqrt{k}} \\ &\leq \frac{1}{k} \sum_{i=a}^b F(x^{(i)}) + 2C \frac{\sqrt{\log(N/\delta)}}{\sqrt{k}} \\ &\leq F(x^{(b)}) + C \frac{\sqrt{\log(N/\delta)}}{\sqrt{k}} \\ &\leq F(y) + 2C \frac{\sqrt{\log(N/\delta)}}{\sqrt{k}} \\ &< F(x) + \frac{1}{4}(F(x^\star) - F(x)) + 2C \frac{\sqrt{\log(N/\delta)}}{\sqrt{k}} \end{aligned}$$

Similarly, the lower bound $L(c, d)$ for the interval $[c, d]$ satisfies:

$$\begin{aligned} L(c, d) &= \frac{1}{k} \sum_{i=c}^d f(x^{(i)}, z^{(i)}) - C \frac{\sqrt{\log(N/\delta)}}{\sqrt{k}} \\ &\geq \frac{1}{k} \sum_{i=c}^d F(x^{(i)}) - 2C \frac{\sqrt{\log(N/\delta)}}{\sqrt{k}} \\ &\geq F(x^{(c)}) - 2C \frac{\sqrt{\log(N/\delta)}}{\sqrt{k}} \\ &\geq F(w) - 2C \frac{\sqrt{\log(N/\delta)}}{\sqrt{k}} \\ &> F(x^\star) - \frac{1}{4}(F(x^\star) - F(x)) - 2C \frac{\sqrt{\log(N/\delta)}}{\sqrt{k}} \end{aligned}$$

Now, set

$$k \geq \frac{576C^2 \log(N/\delta)}{(F(x^\star) - F(x))^2}$$

to obtain

$$\begin{aligned} L(c, d) &> F(x^*) - \frac{1}{3}(F(x^*) - F(x)) \\ U(a, b) &< F(x) + \frac{1}{3}(F(x^*) - F(x)) \end{aligned}$$

Now, since $\tau \leq \frac{F(x^*) - F(x)}{3}$, we will eliminate the region to the left of a , which includes x . Putting all together, we have used $N \leq \lceil \frac{20Lk}{F(x^*) - F(x)} \rceil = \lceil \frac{11520 \cdot C^2 L \log(N/\delta)}{(F(x^*) - F(x))^3} \rceil$ samples.

Now, we prove the second part of the claim. On the round that x is eliminated, there must exist $a \leq b \leq c \leq d$ with $x < x^{(a)}$ and $U(a, b) < L(c, d) - \tau$.

Also, we must have that $F(x^{(b)}) \geq F(x)$, for otherwise this implies $x^* \leq x^{(b)}$ by unimodality so that $F(x) \geq F(x^{(c)}) \geq F(x^{(d)})$ again by unimodality and so it cannot hold that $F(x^{(c)}) = \max_{i \in \{c, d\}} F(x^{(i)}) \geq L(c, d) \geq U(a, b) \geq \min_{i \in \{a, b\}} F(x^{(i)}) = F(x^{(b)})$. Therefore, we have:

$$F(x) \leq \frac{1}{b - a + 1} \sum_{i=a}^b F(x^{(i)}) \leq U(a, b) \leq L(c, d) - \tau$$

Now, let $I_{\text{best}} = [e, f]$. We have $L(e, f) \geq L(c, d)$ by definition. Therefore:

$$\begin{aligned} \frac{1}{f - e + 1} \sum_{i=e}^f F(x^{(i)}) &\geq L(e, f) \\ &\geq L(c, d) \\ &\geq U(a, b) + \tau \\ &\geq F(x) + \tau \end{aligned}$$

□

Theorem 5. *There is a constant C such that, given a budget of B and failure probability $\delta < 1/2$, if we set $\tau = \max\left(\frac{2d \log(B/\delta)}{B}, \frac{37 \cdot (dC^2 L \log(B/\delta)^2)^{1/4}}{B^{1/4}}\right)$, Then with probability at least $1 - 2\delta$, Algorithm 3, returns a point \hat{w} that is a 3τ -approximate local minimum. That is, for each coordinate i , for any w_i^* that differs from \hat{w} only in the i th coordinate, we have $F(w_i^*) \leq F(\hat{w}) + 3\tau$*

Proof. For each of the d copies of Algorithm 2 instantiated in the t th iteration of Algorithm 3, by Lemma 3, with probability at least $1 - 6/t^2 d\pi$, all confidence bounds created internal by Algorithm 2 and the calls it makes to Algorithm 1 are valid. Thus, by union bound all such bounds are valid for all copies of Algorithm 2 with probability at least $1 - \delta$. Let E be this event, so that $P(E) \geq 1 - \delta$.

Now, suppose that in some iteration t , w^t is *not* a 3τ -stationary point. For each $i \in \{1, \dots, d\}$, let $w_{*,i}^t$ be the point that maximizes F while differing from w^t only in the i th coordinate. Then, since w^t is not a 3τ stationary point, there is some i such that $F(w_{*,i}^t) - 3\tau \geq F(w^t)$.

Thus, by Lemma 5, after at most $\lceil \frac{11520 \cdot C^2 L \log(B/\delta)}{(F(w_{*,i}^t) - F(w^t))^3} \rceil \leq \lceil \frac{311040 \cdot C^2 L \log(N/\delta)}{\tau^3} \rceil$ samples, the i th copy of Algorithm 2 will eliminate w^t , and select a point w^{t+1} such that $\mathbb{E}[F(w^{t+1})] \geq F(w^t) + \tau$. Since there are d total copies, this consumes $O\left(\frac{d \log(B/\delta)}{\tau^3}\right)$ total samples.

Now, at an intuitive level, since $\mathbb{E}[F(w^t)]$ can increase by τ at most $1/\tau$ times, this process repeats at most $1/\tau$ times, after which we must have found a 3τ -critical point.

To make this formal, let T be the index (if it exists) at which $F(w^t)$ is a 3τ stationary point. Define $X_t = F(w^t)$ for $t \leq T$ and $X_t = F(w_T)$ for $t > T$. Then, we can apply Lemma 6 to see that $P\left[T \geq \frac{3 \log(\delta)}{\tau}\right] < \delta$. Therefore, with

probability at least $1 - 2\delta$, after

$$\begin{aligned} \frac{3d \log(1/\delta)}{\tau} \left\lceil \frac{311040 \cdot C^2 L \log(N/\delta)}{\tau^3} \right\rceil &\leq \frac{3d \log(1/\delta)}{\tau} + \frac{933120 \cdot dC^2 L \log(N/\delta)^2}{\tau^4} \\ &\leq \frac{B}{2} + \frac{B}{2} = B \end{aligned}$$

samples, we find a 3τ stationary point, as desired. \square

C.1 A TECHNICAL STOPPING TIME LEMMA

Lemma 6. *Let $0 = X_0, X_1, X_2, \dots$ be a sub-martingale in $[0, 1]$. Let T be the stopping time $T = \{\min t \mid X_t \geq 1 - 3\tau\}$. Suppose that for $t \geq T$, $X_{t+1} = X_t$ and for $t < T$, $\mathbb{E}[X_{t+1}] \geq X_t + \tau$. Then for any $\delta > 0$, with probability at least $1 - \delta$, $T \leq \frac{3 \log(1/\delta)}{\tau}$.*

Proof. We consider an alternative non-negative sub-martingale Y_1, Y_2, \dots defined by $Y_t = X_t$ for $t \leq T$ and $Y_{t+1} = Y_t + \tau$ for $t \geq T$. Let $d_t = Y_t - Y_{t-1}$ be the associated martingale difference sequence. Clearly we have $\mathbb{E}[d_t] \geq \tau$ for all t . Then, we have:

$$\begin{aligned} 1 &\geq Y_T \geq \sum_{t=1}^{\infty} d_t \mathbb{1}[t \leq T] \\ &= \sum_{t=1}^{\infty} d_t (1 - \mathbb{1}[t > T]) \\ &= \sum_{t=1}^{\infty} d_t - \sum_{i=1}^{t-1} d_t \mathbb{1}[i = T] \\ &= \sum_{t=1}^{\infty} d_t - \tau \sum_{i=1}^{t-1} \mathbb{1}[i = T] \\ &= \mathbb{E} \left[\sum_{t=1}^{\infty} d_t - \tau \sum_{i=1}^{t-1} \mathbb{1}[i = T] \right] \\ &= \sum_{t=1}^{\infty} \mathbb{E}[d_t] - \tau P[T < t] \\ &\geq \sum_{t=1}^{\infty} \tau P[T \geq t] \\ &= \tau \mathbb{E}[T] \end{aligned}$$

Therefore, we have $\mathbb{E}[T] \leq \frac{1}{\tau}$. Thus, by Markov inequality, $P[T > 3/\tau] \leq 1/3 \leq 1/e$.

Now, let L_i be the event that $T > \frac{3i}{\tau}$. Let $T_i = T - \frac{3(i-1)}{\tau}$ for $i \geq 1$. Observe that the above argument also shows that $\mathbb{E}[T_i | L_{i-1}] \leq \frac{1}{\tau}$, so that $P[T_i > 3/\tau | L_i] \leq \frac{1}{e}$. Further, notice that

$$\begin{aligned} P[L_i] &= P[T_i > 3/\tau | L_{i-1}] P[L_{i-1}] \\ &\leq \frac{1}{e^i} \end{aligned}$$

The result now follows by seeing that $P\left[T > \frac{3 \log(1/\delta)}{\tau}\right] = L_{\log(1/\delta)}$ \square

D MORE EXPERIMENTAL DETAILS

Algorithm 4 contains some small but useful modifications of Algorithm 3. We need to first introduce a modified version of Algorithm 1 in Algorithm 5 that has τ set to zero and only compares intervals of size in a power of 2 among each other. This reduces the running time of each round to $O(N)$, where N is the number of points sampled. It can also be shown that this would only affect lower order terms in our final bounds. Algorithm 6 is a version of Algorithm 2 that uses Algorithm 5 as the sub-routine.

Apart from using the modified Algorithm 6 in the inner loop per dimension, we make some other practical changes. For the first ten slots we sample a random point. The the best out of these 10 points (in terms of observed value) is set as the starting point w^t . Next, out of many possible dimensions that can eliminate the current point in line 12, we select the one that has the smallest surviving interval. Lastly, in the loop at line 4 (when none of the dimensions have managed to eliminate the current point), we choose a dimension from a distribution that gives higher weights to dimensions with more standard deviation of observed values so far. The sampled dimension is the one that gets to sample the next epoch of points for Algorithm 6.

D.1 COMPARISON WITH ZOOMING ALGORITHM

In order to show the advantage of using unimodal structure in real hyper-parameter tuning tasks we implement a version of zooming algorithm [Kleinberg et al., 2008] under our framework and run it on the HPO-B benchmark. In particular, in our algorithm we do not eliminate regions from the left and right (as we would like to not use unimodal structure in the zooming baseline). Instead we eliminate the current best point along a dimension if any interval containing the current best point has an UCB less than highest LCB of any interval. The rest of the algorithm is identical to ours with all the improvements mentioned in Section 5. Note that this algorithm can adapt to unknown Lipschitz constant unlike the original algorithm in [Kleinberg et al., 2008] and therefore is arguably a stronger baseline. The results are shown in Fig. 4. It can be seen that Zooming Ascent is worse than GP-UCB while our algorithm that uses unimodal structure is better.

E EMPIRICAL EVIDENCE FOR UNIMODALITY

In order to find evidence for approximate unimodality across coordinate dimensions we run a grid search of five hyper-parameters on the Cifar 10 dataset. The hyper-parameters are tuned for AlexNet architecture [Krizhevsky, 2009] where the number of convolutional filters can be changed for the first two layers. The other three hyper-parameters we were allowed to vary were batch size, learning rate and dropout level of the last linear layer. We had a grid of 5 points per dimension and each point was sampled 3 times to estimate the noise in experiments with the same hyper-parameter. Thus the experiment involved training $5^5 \times 3$ models and getting the test accuracy for them.

Figure 3 shows the resulting plots. Each row corresponds to varying one parameter (coordinate) while all other points are held fixed. each column corresponds to a randomly chosen point in the search-space from which one dimension is varied at a time, while keeping all other fixed. The x-axis denotes the hyperparam that is varied. We can see that the black-box function is approximately unimodal per coordinate dimension. We believe that our algorithm works well on benchmarks because such approximate unimodal structure might be present in many real world hyper-parameter tuning tasks.

Algorithm 5: One Elimination Round of Unimodal Optimization - Power of 2 Intervals

Require: Confidence parameter δ , interval $\Delta = [l, r]$, confidence scaling constant $h > 0$ to be set by Lemma 2.

- 1: Define the points $l = x^{(1)} < x^{(2)} < \dots < x^{(N)} = r$ for $N = 1 + \frac{l-r}{\text{GAP}}$ equally spaced in the interval Δ .
 - 2: Compute $f(x^{(k)}, z^{(k)})$ for i.i.d. $z^{(1)}, \dots, z^{(N)}$.
 - 3: Given any $i, j \in \{1, \dots, N\}$, define $m_{ij} = \sum_{k=i}^j f(x^{(k)}, z^{(k)}) / (j - i + 1)$, $s_{ij} = h \sqrt{\log(2N/\delta)} / \sqrt{j - i + 1}$.
 - 4: Define upper confidence bound $U(i, j) = m_{ij} + s_{ij}$ for each $i, j \in \{1, \dots, N\}$
 - 5: Define lower confidence bound $L(i, j) = m_{ij} - s_{ij}$ for each $i, j \in \{1, \dots, N\}$
 - 6: **for** s in $\{2^2, 2^3, \dots, 2^{\log N}\}$ **do**
 - 7: Let $S_l^{(s)} = \max\{x^{(i)} \mid \exists i \leq j \leq k \leq l \text{ s.t. } U(i, j) < L(k, l) \text{ and } j - i + 1 = 2^s, l - k + 1 = 2^s\} \cup \{x^{(1)}\}$.
 - 8: Let $S_r^{(s)} = \min\{x^{(l)} \mid \exists i \leq j \leq k \leq l \text{ s.t. } U(k, l) < L(i, j) \text{ and } j - i + 1 = 2^s, l - k + 1 = 2^s\} \cup \{x^{(N)}\}$.
 - 9: **end for**
 - 10: Let I_{best} be the best LCB for intervals considered in the above for loop.
 - 11: Return New interval $[\max_s S_l^{(s)}, \min_s S_r^{(s)}] \subset \Delta$ and the best interval I_{best} .
-

Algorithm 6: Modified 1D Unimodal Optimization

Require: Confidence parameter δ , budget B .

- 1: Initialize $t = 1$, $\Delta_1 = [0, 1]$, $\text{GAP}_1 = 0.5$, $N = N_1 = 3$.
 - 2: **while** $N \leq B$ **do**
 - 3: Call Algorithm 5 with input $6\delta/\pi^2 t^2$, Δ_t , GAP_t to obtain outputs $[S_l, S_r]$, I_{best} .
 - 4: Let $\Delta_{t+1} = [S_l, S_r]$.
 - 5: Set $\text{GAP}_{t+1} = \text{GAP}_t/2$.
 - 6: Set $N_{t+1} = 1 + \frac{S_r - S_l}{\text{GAP}}$ //Budget to be consumed by next iteration.
 - 7: Set $N = N + N_t$. //Total budget consumed at end of next iteration.
 - 8: Set $t = t + 1$.
 - 9: **end while**
 - 10: Return a random element of I_{best} .
-

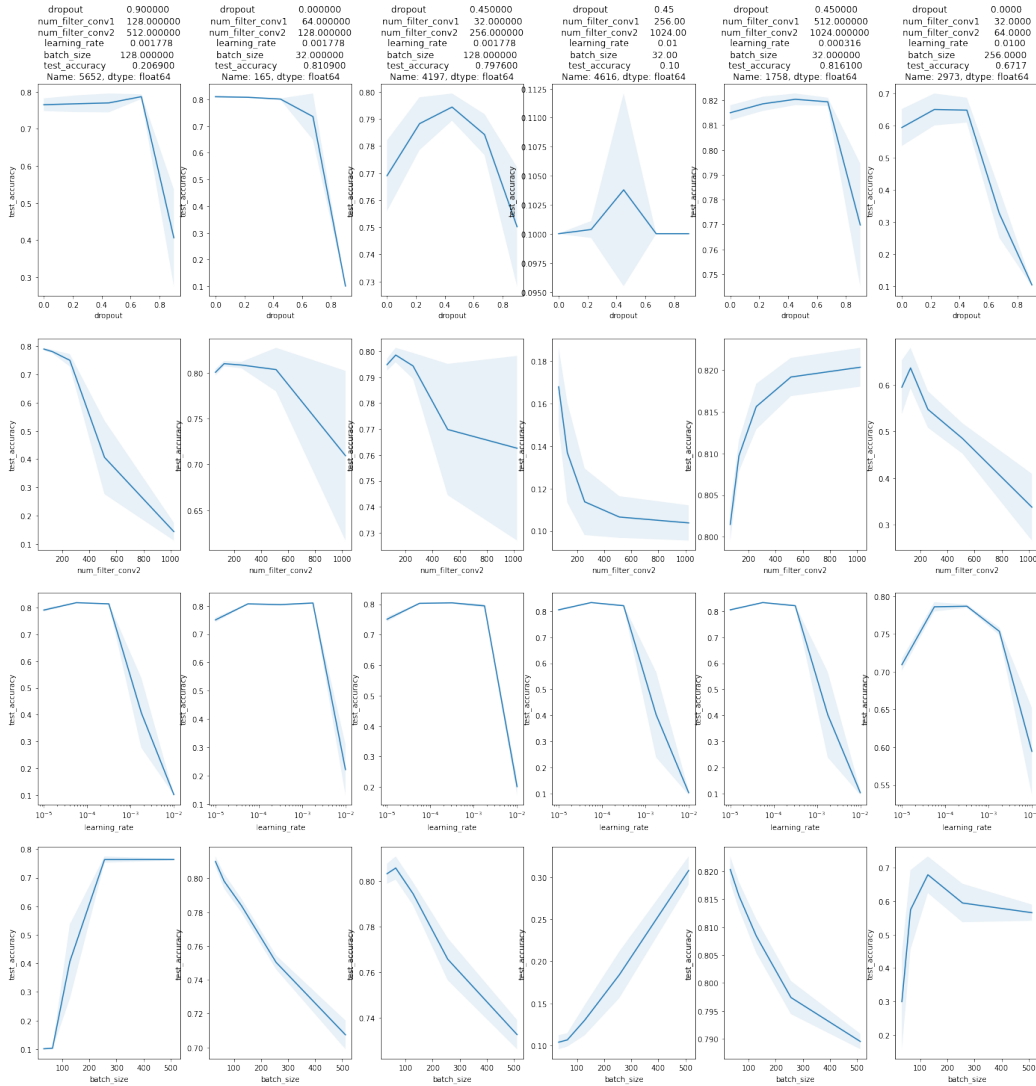


Figure 3: We illustrate approximate unimodality per coordinate dimensions for the task of tuning a CNN on Cifar 10. Each row corresponds to varying one parameter (coordinate) while all other points are held fixed. each column corresponds to a randomly chosen point in the search-space from which one dimension is varied at a time, while keeping all other fixed. The x-axis denotes the hyper-param that is varied. We plot the standard error over three runs per point, as demonstrated by the shaded area.

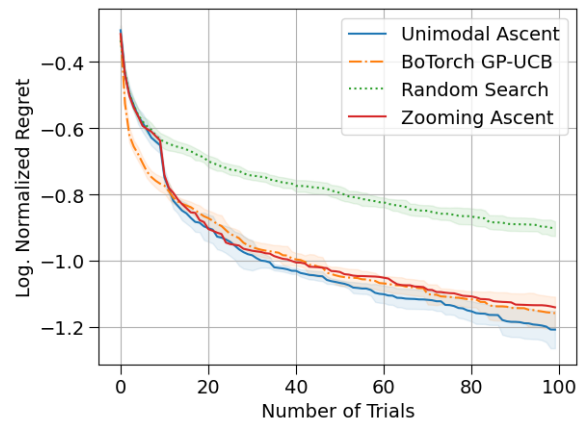


Figure 4: Comparison with respect to a version of the zooming algorithm [Kleinberg et al., 2008]