
Bayesian Optimization with Exponential Convergence

Kenji Kawaguchi
MIT
Cambridge, MA, 02139
kawaguch@mit.edu

Leslie Pack Kaelbling
MIT
Cambridge, MA, 02139
lpk@csail.mit.edu

Tomás Lozano-Pérez
MIT
Cambridge, MA, 02139
tlp@csail.mit.edu

Abstract

This paper presents a Bayesian optimization method with exponential convergence *without* the need of auxiliary optimization and *without* the δ -cover sampling. Most Bayesian optimization methods require auxiliary optimization: an additional non-convex global optimization problem, which can be time-consuming and hard to implement in practice. Also, the existing Bayesian optimization method with exponential convergence [1] requires access to the δ -cover sampling, which was considered to be impractical [1, 2]. Our approach eliminates both requirements and achieves an exponential convergence rate.

1 Introduction

We consider a general global optimization problem: maximize $f(x)$ subject to $x \in \Omega \subset \mathbb{R}^D$ where $f: \Omega \rightarrow \mathbb{R}$ is a non-convex black-box deterministic function. Such a problem arises in many real-world applications, such as parameter tuning in machine learning [3], engineering design problems [4], and model parameter fitting in biology [5]. For this problem, one performance measure of an algorithm is the *simple regret*, r_n , which is given by $r_n = \sup_{x \in \Omega} f(x) - f(x^+)$ where x^+ is the best input vector found by the algorithm. For brevity, we use the term “regret” to mean simple regret.

The general global optimization problem is known to be intractable if we make no further assumptions [6]. The simplest additional assumption to restore tractability is to assume the existence of a bound on the slope of f . A well-known variant of this assumption is Lipschitz continuity with a known Lipschitz constant, and many algorithms have been proposed in this setting [7, 8, 9]. These algorithms successfully guaranteed certain bounds on the regret. However appealing from a theoretical point of view, a practical concern was soon raised regarding the assumption that a tight Lipschitz constant is known. Some researchers relaxed this somewhat strong assumption by proposing procedures to estimate a Lipschitz constant during the optimization process [10, 11, 12].

Bayesian optimization is an efficient way to relax this assumption of complete knowledge of the Lipschitz constant, and has become a well-recognized method for solving global optimization problems with non-convex black-box functions. In the machine learning community, Bayesian optimization—especially by means of a Gaussian process (GP)—is an active research area [13, 14, 15]. With the requirement of the access to the δ -cover sampling procedure (it samples the function uniformly such that the density of samples doubles in the feasible regions at each iteration), de Freitas et al. [1] recently proposed a theoretical procedure that maintains an exponential convergence rate (exponential regret). However, as pointed out by Wang et al. [2], one remaining problem is to derive a GP-based optimization method with an exponential convergence rate *without* the δ -cover sampling procedure, which is computationally too demanding in many cases.

In this paper, we propose a novel GP-based global optimization algorithm, which maintains an exponential convergence rate and converges rapidly *without* the δ -cover sampling procedure.

2 Gaussian Process Optimization

In Gaussian process optimization, we estimate the distribution over function f and use this information to decide which point of f should be evaluated next. In a parametric approach, we consider a parameterized function $f(x; \theta)$, with θ being distributed according to some prior. In contrast, the non-parametric GP approach directly puts the GP prior over f as $f(\cdot) \sim GP(m(\cdot), \kappa(\cdot, \cdot))$ where $m(\cdot)$ is the mean function and $\kappa(\cdot, \cdot)$ is the covariance function or the kernel. That is, $m(x) = \mathbb{E}[f(x)]$ and $\kappa(x, x') = \mathbb{E}[(f(x) - m(x))(f(x') - m(x'))^T]$. For a finite set of points, the GP model is simply a joint Gaussian: $\mathbf{f}(\mathbf{x}_{1:N}) \sim \mathcal{N}(\mathbf{m}(\mathbf{x}_{1:N}), \mathbf{K})$, where $\mathbf{K}_{i,j} = \kappa(x_i, x_j)$ and N is the number of data points. To predict the value of f at a new data point, we first consider the joint distribution over f of the old data points and the new data point:

$$\begin{pmatrix} \mathbf{f}(\mathbf{x}_{1:N}) \\ f(x_{N+1}) \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mathbf{m}(\mathbf{x}_{1:N}) \\ m(x_{N+1}) \end{pmatrix}, \begin{bmatrix} \mathbf{K} & \mathbf{k} \\ \mathbf{k}^T & \kappa(x_{N+1}, x_{N+1}) \end{bmatrix} \right)$$

where $\mathbf{k} = \kappa(\mathbf{x}_{1:N}, \mathbf{x}_{N+1}) \in \mathbb{R}^{N \times 1}$. Then, after factorizing the joint distribution using the Schur complement for the joint Gaussian, we obtain the conditional distribution, conditioned on observed entities $\mathcal{D}_N := \{\mathbf{x}_{1:N}, \mathbf{f}(\mathbf{x}_{1:N})\}$ and x_{N+1} , as:

$$f(\mathbf{x}_{N+1}) | \mathcal{D}_N, x_{N+1} \sim \mathcal{N}(\mu(x_{N+1} | \mathcal{D}_N), \sigma^2(x_{N+1} | \mathcal{D}_N))$$

where $\mu(x_{N+1} | \mathcal{D}_N) = m(x_{N+1}) + \mathbf{k}^T \mathbf{K}^{-1}(\mathbf{f}(\mathbf{x}_{1:N}) - \mathbf{m}(\mathbf{x}_{1:N}))$ and $\sigma^2(x_{N+1} | \mathcal{D}_N) = \kappa(\mathbf{x}_{N+1}, \mathbf{x}_{N+1}) - \mathbf{k}^T \mathbf{K}^{-1} \mathbf{k}$. One advantage of GP is that this closed-form solution simplifies both its analysis and implementation.

To use a GP, we must specify the mean function and the covariance function. The mean function is usually set to be zero. With this zero mean function, the conditional mean $\mu(x_{N+1} | \mathcal{D}_N)$ can still be flexibly specified by the covariance function, as shown in the above equation for μ . For the covariance function, there are several common choices, including the Matern kernel and the Gaussian kernel. For example, the Gaussian kernel is defined as $\kappa(x, x') = \exp\left(-\frac{1}{2}(x - x')^T \Sigma^{-1}(x - x')\right)$ where Σ^{-1} is the kernel parameter matrix. The kernel parameters or hyperparameters can be estimated by empirical Bayesian methods [16]; see [17] for more information about GP.

The flexibility and simplicity of the GP prior make it a common choice for continuous objective functions in the Bayesian optimization literature. Bayesian optimization with GP selects the next query point that optimizes the acquisition function generated by GP. Commonly used acquisition functions include the upper confidence bound (UCB) and expected improvement (EI). For brevity, we consider Bayesian optimization with UCB, which works as follows. At each iteration, the UCB function \mathcal{U} is maintained as $\mathcal{U}(x | \mathcal{D}_N) = \mu(x | \mathcal{D}_N) + \zeta \sigma(x | \mathcal{D}_N)$ where $\zeta \in \mathbb{R}$ is a parameter of the algorithm. To find the next query x_{n+1} for the objective function f , GP-UCB solves an additional non-convex optimization problem with \mathcal{U} as $x_{n+1} = \arg \max_x \mathcal{U}(x | \mathcal{D}_N)$. This is often carried out by other global optimization methods such as DIRECT and CMA-ES. The justification for introducing a new optimization problem lies in the assumption that the cost of evaluating the objective function f dominates that of solving additional optimization problem.

For deterministic function, de Freitas et al. [1] recently presented a theoretical procedure that maintains exponential convergence rate. However, their own paper and the follow-up research [1, 2] point out that this result relies on an impractical sampling procedure, the δ -cover sampling. To overcome this issue, Wang et al. [2] combined GP-UCB with a hierarchical partitioning optimization method, the SOO algorithm [18], providing a regret bound with polynomial dependence on the number of function evaluations. They concluded that creating a GP-based algorithm with an *exponential* convergence rate *without* the impractical sampling procedure remained an open problem.

3 Infinite-Metric GP Optimization

3.1 Overview

The GP-UCB algorithm can be seen as a member of the class of bound-based search methods, which includes Lipschitz optimization, A* search, and PAC-MDP algorithms with optimism in the

face of uncertainty. Bound-based search methods have a common property: the tightness of the bound determines its effectiveness. The tighter the bound is, the better the performance becomes. However, it is often difficult to obtain a tight bound while maintaining correctness. For example, in A* search, admissible heuristics maintain the correctness of the bound, but the estimated bound with admissibility is often too loose in practice, resulting in a long period of global search.

The GP-UCB algorithm has the same problem. The bound in GP-UCB is represented by UCB, which has the following property: $f(x) \leq \mathcal{U}(x|\mathcal{D})$ with some probability. We formalize this property in the analysis of our algorithm. The problem is essentially due to the difficulty of obtaining a tight bound $\mathcal{U}(x|\mathcal{D})$ such that $f(x) \leq \mathcal{U}(x|\mathcal{D})$ and $f(x) \approx \mathcal{U}(x|\mathcal{D})$ (with some probability). Our solution strategy is to first admit that the bound encoded in GP prior may not be tight enough to be useful by itself. Instead of relying on a single bound given by the GP, we leverage the existence of an *unknown* bound encoded in the continuity at a global optimizer.

Assumption 1. (Unknown Bound) There exists a global optimizer x^* and an *unknown* semi-metric ℓ such that for all $x \in \Omega$, $f(x^*) \leq f(x) + \ell(x, x^*)$ and $\ell(x, x^*) < \infty$.

In other words, we do not expect the *known* upper bound due to GP to be tight, but instead expect that there exists some *unknown* bound that might be tighter. Notice that in the case where the bound by GP is as tight as the unknown bound by semi-metric ℓ in Assumption 1, our method still maintains an exponential convergence rate and an advantage over GP-UCB (no need for auxiliary optimization). Our method is expected to become relatively much better when the *known* bound due to GP is less tight compared to the unknown bound by ℓ .

As the semi-metric ℓ is unknown, there are infinitely many possible candidates that we can think of for ℓ . Accordingly, we simultaneously conduct global and local searches based on all the candidates of the bounds. The bound estimated by GP is used to reduce the number of candidates. Since the bound estimated by GP is known, we can ignore the candidates of the bounds that are looser than the bound estimated by GP. The source code of the proposed algorithm is publicly available at <http://lis.csail.mit.edu/code/imgpo.html>.

3.2 Description of Algorithm

Figure 1 illustrates how the algorithm works with a simple 1-dimensional objective function. We employ hierarchical partitioning to maintain hyperintervals, as illustrated by the line segments in the figure. We consider a hyperrectangle as our hyperinterval, with its center being the evaluation point of f (blue points in each line segment in Figure 1). For each iteration t , the algorithm performs the following procedure *for each interval size*:

- (i) Select the interval with the maximum center value among the intervals of the same size.
- (ii) Keep the interval selected by (i) if it has a center value greater than that of any *larger* interval.
- (iii) Keep the interval accepted by (ii) if it contains a UCB greater than the center value of any *smaller* interval.
- (iv) If an interval is accepted by (iii), divide it along with the longest coordinate into three new intervals.
- (v) For each new interval, if the UCB of the evaluation point is less than the best function value found so far, skip the evaluation and use the UCB value as the center value until the interval is accepted in step (ii) on some future iteration; otherwise, evaluate the center value.
- (vi) Repeat steps (i)–(v) until every size of intervals are considered

Then, at the end of each iteration, the algorithm updates the GP hyperparameters. Here, the purpose of steps (i)–(iii) is to select an interval that might contain the global optimizer. Steps (i) and (ii) select the possible intervals based on the unknown bound by ℓ , while Step (iii) does so based on the bound by GP.

We now explain the procedure using the example in Figure 1. Let n be the number of divisions of intervals and let N be the number of function evaluations. t is the number of iterations. Initially, there is only one interval (the center of the input region $\Omega \subset \mathbb{R}$) and thus this interval is divided, resulting in the first diagram of Figure 1. At the beginning of iteration $t = 2$, step (i) selects the third

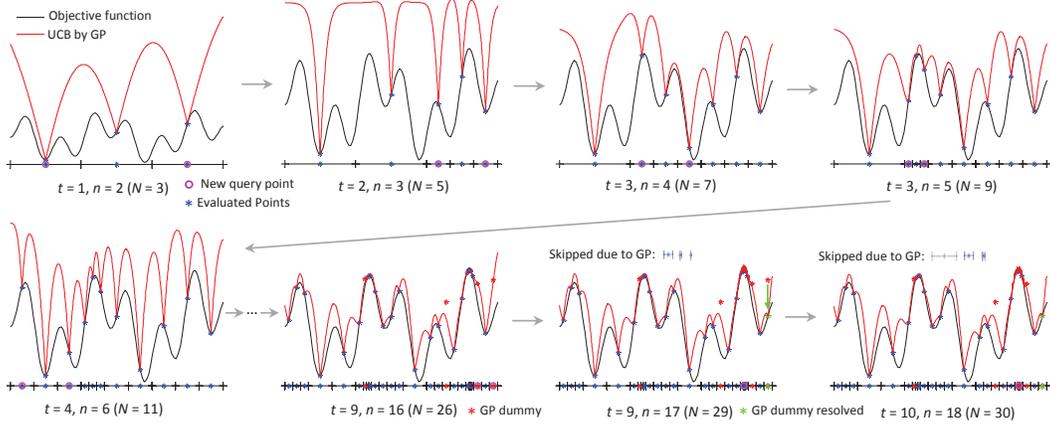


Figure 1: An illustration of IMGPO: t is the number of iteration, n is the number of divisions (or splits), N is the number of function evaluations.

interval from the left side in the first diagram ($t = 1, n = 2$), as its center value is the maximum. Because there are no intervals of different size at this point, steps (ii) and (iii) are skipped. Step (iv) divides the third interval, and then the GP hyperparameters are updated, resulting in the second diagram ($t = 2, n = 3$). At the beginning of iteration $t = 3$, it starts conducting steps (i)–(v) for the largest intervals. Step (i) selects the second interval from the left side and step (ii) is skipped. Step (iii) accepts the second interval, because the UCB within this interval is no less than the center value of the smaller intervals, resulting in the third diagram ($t = 3, n = 4$). Iteration $t = 3$ continues by conducting steps (i)–(v) for the smaller intervals. Step (i) selects the second interval from the left side, step (ii) accepts it, and step (iii) is skipped, resulting in the fourth diagram ($t = 3, n = 4$). The effect of the step (v) can be seen in the diagrams for iteration $t = 9$. At $n = 16$, the far right interval is divided, but no function evaluation occurs. Instead, UCB values given by GP are placed in the new intervals indicated by the red asterisks. One of the temporary dummy values is resolved at $n = 17$ when the interval is queried for division, as shown by the green asterisk. The effect of step (iii) for the rejection case is illustrated in the last diagram for iteration $t = 10$. At $n = 18$, t is increased to 10 from 9, meaning that the largest intervals are first considered for division. However, the three largest intervals are all rejected in step (iii), resulting in the division of a very small interval near the global optimum at $n = 18$.

3.3 Technical Detail of Algorithm

We define h to be the depth of the hierarchical partitioning tree, and $c_{h,i}$ to be the center point of the i^{th} hyperrectangle at depth h . N_{gp} is the number of the GP evaluations. Define $\text{depth}(\mathcal{T})$ to be the largest integer h such that the set \mathcal{T}_h is not empty. To compute UCB \mathcal{U} , we use $\varsigma_M = \sqrt{2 \log(\pi^2 M^2 / 12\eta)}$ where M is the number of the calls made so far for \mathcal{U} (i.e., each time we use \mathcal{U} , we increment M by one). This particular form of ς_M is to maintain the property of $f(x) \leq \mathcal{U}(x|\mathcal{D})$ during an execution of our algorithm with probability at least $1 - \eta$. Here, η is the parameter of IMGPO. Ξ_{max} is another parameter, but it is only used to limit the possibly long computation of step (iii) (in the worst case, step (iii) computes UCBs $3^{\Xi_{max}}$ times although it would rarely happen).

The pseudocode is shown in Algorithm 1. Lines 8 to 23 correspond to steps (i)–(iii). These lines compute the index i_h^* of the candidate of the rectangle that may contain a global optimizer for each depth h . For each depth h , non-null index i_h^* at Line 24 indicates the remaining candidate of a rectangle that we want to divide. Lines 24 to 33 correspond to steps (iv)–(v) where the remaining candidates of the rectangles for all h are divided. To provide a simple executable division scheme (line 29), we assume Ω to be a hyperrectangle (see the last paragraph of section 4 for a general case).

Lines 8 to 17 correspond to steps (i)–(ii). Specifically, line 10 implements step (i) where a single candidate is selected for each depth, and lines 11 to 12 conduct step (ii) where some candidates are screened out. Lines 13 to 17 resolve the temporary dummy values computed by GP. Lines 18 to 23 correspond to step (iii) where the candidates are further screened out. At line 21, $\mathcal{T}'_{h+\xi}(c_{h,i_h^*})$

Algorithm 1 Infinite-Metric GP Optimization (IMGPO)

Input: an objective function f , the search domain Ω , the GP kernel κ , $\Xi_{max} \in \mathbb{N}^+$ and $\eta \in (0, 1)$

- 1: Initialize the set $\mathcal{T}_h = \{\emptyset\} \forall h \geq 0$
- 2: Set $c_{0,0}$ to be the center point of Ω and $\mathcal{T}_0 \leftarrow \{c_{0,0}\}$
- 3: Evaluate f at $c_{0,0}$: $g(c_{0,0}) \leftarrow f(c_{0,0})$
- 4: $f^+ \leftarrow g(c_{0,0}), \mathcal{D} \leftarrow \{(c_{0,0}, g(c_{0,0}))\}$
- 5: $n, N \leftarrow 1, N_{gp} \leftarrow 0, \Xi \leftarrow 1$
- 6: **for** $t = 1, 2, 3, \dots$ **do**
- 7: $v_{max} \leftarrow -\infty$
- 8: **for** $h = 0$ to $depth(\mathcal{T})$ **do** # for-loop for steps (i)-(ii)
- 9: **while** true **do**
- 10: $i_h^* \leftarrow \arg \max_{i: c_{h,i} \in \mathcal{T}_h} g(c_{h,i})$
- 11: **if** $g(c_{h,i_h^*}) < v_{max}$ **then**
- 12: $i_h^* \leftarrow \emptyset$, **break**
- 13: **else if** $g(c_{h,i_h^*})$ is not labeled as *GP-based* **then**
- 14: $v_{max} \leftarrow g(c_{h,i_h^*})$, **break**
- 15: **else**
- 16: $g(c_{h,i_h^*}) \leftarrow f(c_{h,i_h^*})$ and remove the *GP-based* label from $g(c_{h,i_h^*})$
- 17: $N \leftarrow N + 1, N_{gp} \leftarrow N_{gp} - 1$
- 18: $\mathcal{D} \leftarrow \{\mathcal{D}, (c_{h,i_h^*}, g(c_{h,i_h^*}))\}$
- 19: **for** $h = 0$ to $depth(\mathcal{T})$ **do** # for-loop for step (iii)
- 20: **if** $i_h^* \neq \emptyset$ **then**
- 21: $\xi \leftarrow$ the smallest positive integer s.t. $i_{h+\xi}^* \neq \emptyset$ and $\xi \leq \min(\Xi, \Xi_{max})$ if exists, and 0 otherwise
- 22: $z(h, i_h^*) = \max_{k: c_{h+\xi,k} \in \mathcal{T}'_{h+\xi}(c_{h,i_h^*})} \mathcal{U}(c_{h+\xi,k} | \mathcal{D})$
- 23: **if** $\xi \neq 0$ and $z(h, i_h^*) < g(c_{h+\xi, i_{h+\xi}^*})$ **then**
- 24: $i_h^* \leftarrow \emptyset$, **break**
- 25: $v_{max} \leftarrow -\infty$
- 26: **for** $h = 0$ to $depth(\mathcal{T})$ **do** # for-loop for steps (iv)-(v)
- 27: **if** $i_h^* \neq \emptyset$ and $g(c_{h,i_h^*}) \geq v_{max}$ **then**
- 28: $n \leftarrow n + 1$.
- 29: Divide the hyperrectangle centered at c_{h,i_h^*} along with the longest coordinate into three new hyperrectangles with the following centers:
 $\mathcal{S} = \{c_{h+1,i(left)}, c_{h+1,i(center)}, c_{h+1,i(right)}\}$
- 30: $\mathcal{T}_{h+1} \leftarrow \{\mathcal{T}_{h+1}, \mathcal{S}\}$
- 31: $\mathcal{T}_h \leftarrow \mathcal{T}_h \setminus c_{h,i_h^*}, g(c_{h+1,i(center)}) \leftarrow g(c_{h,i_h^*})$
- 32: **for** $i_{new} = \{i(left), i(right)\}$ **do**
- 33: **if** $\mathcal{U}(c_{h+1,i_{new}} | \mathcal{D}) \geq f^+$ **then**
- 34: $g(c_{h+1,i_{new}}) \leftarrow f(c_{h+1,i_{new}})$
- 35: $\mathcal{D} \leftarrow \{\mathcal{D}, (c_{h+1,i_{new}}, g(c_{h+1,i_{new}}))\}$
- 36: $N \leftarrow N + 1, f^+ \leftarrow \max(f^+, g(c_{h+1,i_{new}})), v_{max} = \max(v_{max}, g(c_{h+1,i_{new}}))$
- 37: **else**
 $g(c_{h+1,i_{new}}) \leftarrow \mathcal{U}(c_{h+1,i_{new}} | \mathcal{D})$ and label $g(c_{h+1,i_{new}})$ as *GP-based*.
 $N_{gp} \leftarrow N_{gp} + 1$
- 38: Update Ξ : if f^+ was updated, $\Xi \leftarrow \Xi + 2^2$, and otherwise, $\Xi \leftarrow \max(\Xi - 2^{-1}, 1)$
- 39: Update GP hyperparameters by an empirical Bayesian method

indicates the set of *all* center points of a fully expanded tree until depth $h + \xi$ *within* the region covered by the hyperrectangle centered at c_{h,i_h^*} . In other words, $\mathcal{T}'_{h+\xi}(c_{h,i_h^*})$ contains the nodes of the fully expanded tree rooted at c_{h,i_h^*} with depth ξ and can be computed by dividing the current rectangle at c_{h,i_h^*} and recursively divide all the resulting new rectangles until depth ξ (i.e., depth ξ from c_{h,i_h^*} , which is depth $h + \xi$ in the whole tree).

3.4 Relationship to Previous Algorithms

The most closely related algorithm is the BaMSOO algorithm [2], which combines SOO with GP-UCB. However, it only achieves a polynomial regret bound while IMGPO achieves a exponential regret bound. IMGPO can achieve exponential regret because it utilizes the information encoded in the GP prior/posterior to reduce the degree of the unknownness of the semi-metric ℓ .

The idea of considering a set of infinitely many bounds was first proposed by Jones et al. [19]. Their DIRECT algorithm has been successfully applied to real-world problems [4, 5], but it only maintains the consistency property (i.e., convergence in the limit) from a theoretical viewpoint. DIRECT takes an input parameter ϵ to balance the global and local search efforts. This idea was generalized to the case of an unknown semi-metric and strengthened with a theoretical support (finite regret bound) by Munos [18] in the SOO algorithm. By limiting the depth of the search tree with a parameter h_{max} , the SOO algorithm achieves a finite regret bound that depends on *the near-optimality dimension*.

4 Analysis

In this section, we prove an exponential convergence rate of IMGPO and theoretically discuss the reason why the novel idea underling IMGPO is beneficial. The proofs are provided in the supplementary material. To examine the effect of considering infinitely many possible candidates of the bounds, we introduce the following term.

Definition 1. (Infinite-metric exploration loss). The infinite-metric exploration loss ρ_t is the number of intervals to be divided during iteration t .

The infinite-metric exploration loss ρ_τ can be computed as $\rho_t = \sum_{h=1}^{depth(\mathcal{T})} \mathbb{1}(i_h^* \neq \emptyset)$ at line 25. It is the cost (in terms of the number of function evaluations) incurred by not committing to any particular upper bound. If we were to rely on a specific bound, ρ_τ would be minimized to 1. For example, the DOO algorithm [18] has $\rho_t = 1 \forall t \geq 1$. *Even if we know a particular upper bound, relying on this knowledge and thus minimizing ρ_τ is not a good option unless the known bound is tight enough compared to the unknown bound leveraged in our algorithm.* This will be clarified in our analysis. Let $\bar{\rho}_t$ be the maximum of the averages of $\rho_{1:t'}$ for $t' = 1, 2, \dots, t$ (i.e., $\bar{\rho}_t \equiv \max(\{\frac{1}{t'} \sum_{\tau=1}^{t'} \rho_\tau; t' = 1, 2, \dots, t\})$.

Assumption 2. For some pair of a global optimizer x^* and an *unknown* semi-metric ℓ that satisfies Assumption 1, both of the following, (i) shape on ℓ and (ii) lower bound constant, conditions hold:

- (i) there exist $L > 0, \alpha > 0$ and $p \geq 1$ in \mathbb{R} such that for all $x, x' \in \Omega, \ell(x', x) \leq L \|x' - x\|_p^\alpha$.
- (ii) there exists $\theta \in (0, 1)$ such that for all $x \in \Omega, f(x^*) \geq f(x) + \theta \ell(x, x^*)$.

In Theorem 1, we show that the exponential convergence rate $O(\lambda^{N+N_{gp}})$ with $\lambda < 1$ is achieved. We define $\Xi_n \leq \Xi_{max}$ to be the largest ξ used so far with n total node expansions. For simplicity, we assume that Ω is a square, which we satisfied in our experiments by scaling original Ω .

Theorem 1. Assume Assumptions 1 and 2. Let $\beta = \sup_{x, x' \in \Omega} \frac{1}{2} \|x - x'\|_\infty$. Let $\lambda = 3^{-\frac{\alpha}{2CD\bar{\rho}_t}} < 1$. Then, with probability at least $1 - \eta$, the regret of IMGPO is bounded as

$$r_N \leq L(3\beta D^{1/p})^\alpha \exp\left(-\alpha \left[\frac{N + N_{gp}}{2CD\bar{\rho}_t} - \Xi_n - 2\right] \ln 3\right) = O(\lambda^{N+N_{gp}}).$$

Importantly, our bound holds for the best values of the unknown L, α and p even though these values are not given. The closest result in previous work is that of BaMSOO [2], which obtained $\tilde{O}(n^{-\frac{2\alpha}{D(4-\alpha)}})$ with probability $1 - \eta$ for $\alpha = \{1, 2\}$. As can be seen, we have improved the regret bound. Additionally, in our analysis, we can see how L, p , and α affect the bound, allowing us to view the inherent difficulty of an objective function in a theoretical perspective. Here, C is a constant in N and is used in previous work [18, 2]. For example, if we conduct 2^D or $3^D - 1$ function evaluations per node-expansion and if $p = \infty$, we have that $C = 1$.

We note that λ can get close to one as input dimension D increases, which suggests that there is a remaining challenge in scalability for higher dimensionality. One strategy for addressing this problem would be to leverage additional assumptions such as those in [14, 20].

Remark 1. (The effect of the tightness of UCB by GP) If UCB computed by GP is “useful” such that $N/\bar{\rho}_t = \Omega(N)$, then our regret bound becomes $O\left(\exp\left(-\frac{N+N_{gp}}{2CD}\alpha \ln 3\right)\right)$. If the bound due to UCB by GP is too loose (and thus useless), $\bar{\rho}_t$ can increase up to $O(N/t)$ (due to $\bar{\rho}_t \leq \sum_{i=1}^t i/t \leq$

$O(N/t)$), resulting in the regret bound of $O\left(\exp\left(-\frac{t(1+N_{gp}/N)}{2CD}\alpha\ln 3\right)\right)$, which can be bounded by $O\left(\exp\left(-\frac{N+N_{gp}}{2CD}\max\left(\frac{1}{\sqrt{N}}, \frac{t}{N}\right)\alpha\ln 3\right)\right)^1$. This is still better than the known results.

Remark 2. (The effect of GP) Without the use of GP, our regret bound would be as follows: $r_N \leq L(3\beta D^{1/p})^\alpha \exp(-\alpha[\frac{N}{2CD}\frac{1}{\bar{\rho}_t}-2]\ln 3)$, where $\bar{\rho}_t \leq \tilde{\rho}_t$ is the infinite-metric exploration loss without GP. Therefore, the use of GP reduces the regret bound by increasing N_{gp} and decreasing $\bar{\rho}_t$, but may potentially increase the bound by increasing $\Xi_n \leq \Xi$.

Remark 3. (The effect of infinite-metric optimization) To understand the effect of considering all the possible upper bounds, we consider the case without GP. If we consider all the possible bounds, we have the regret bound $L(3\beta D^{1/p})^\alpha \exp(-\alpha[\frac{N}{2CD}\frac{1}{\bar{\rho}_t}-2]\ln 3)$ for the best unknown L , α and p . For standard optimization with a estimated bound, we have $L'(3\beta D^{1/p'})^{\alpha'} \exp(-\alpha'[\frac{N}{2C'D}-2]\ln 3)$ for an estimated L' , α' , and p' . By algebraic manipulation, considering all the possible bounds has a better regret when $\tilde{\rho}_t^{-1} \geq \frac{2CD}{N\ln 3^\alpha}((\frac{N}{2C'D}-2)\ln 3^{\alpha'} + 2\ln 3^\alpha - \ln \frac{L'(3\beta D^{1/p'})^{\alpha'}}{L(3\beta D^{1/p})^\alpha})$. For an intuitive insight, we can simplify the above by assuming $\alpha' = \alpha$ and $C' = C$ as $\tilde{\rho}_t^{-1} \geq 1 - \frac{C_2 D}{N} \ln \frac{L' D^{\alpha/p'}}{L D^{\alpha/p}}$. Because L and p are the ones that achieve the lowest bound, the logarithm on the right-hand side is always non-negative. Hence, $\tilde{\rho}_t = 1$ always satisfies the condition. When L' and p' are not tight enough, the logarithmic term increases in magnitude, allowing $\tilde{\rho}_t$ to increase. For example, if the second term on the right-hand side has a magnitude of greater than 0.5, then $\tilde{\rho}_t = 2$ satisfies the inequality. Therefore, even if we know the upper bound of the function, we can see that it may be better not to rely on this, but rather take the infinite many possibilities into account.

One may improve the algorithm with different division procedures than one presented in Algorithm 1 as discussed in the supplementary material.

¹This can be done by limiting the depth of search tree as $depth(T) = O(\sqrt{N})$. Our proof works with this additional mechanism, but results in the regret bound with N being replaced by \sqrt{N} . Thus, if we assume to have at least “not useless” UCBs such that $N/\bar{\rho}_t = \Omega(\sqrt{N})$, this additional mechanism can be disadvantageous. Accordingly, we do not adopt it in our experiments.

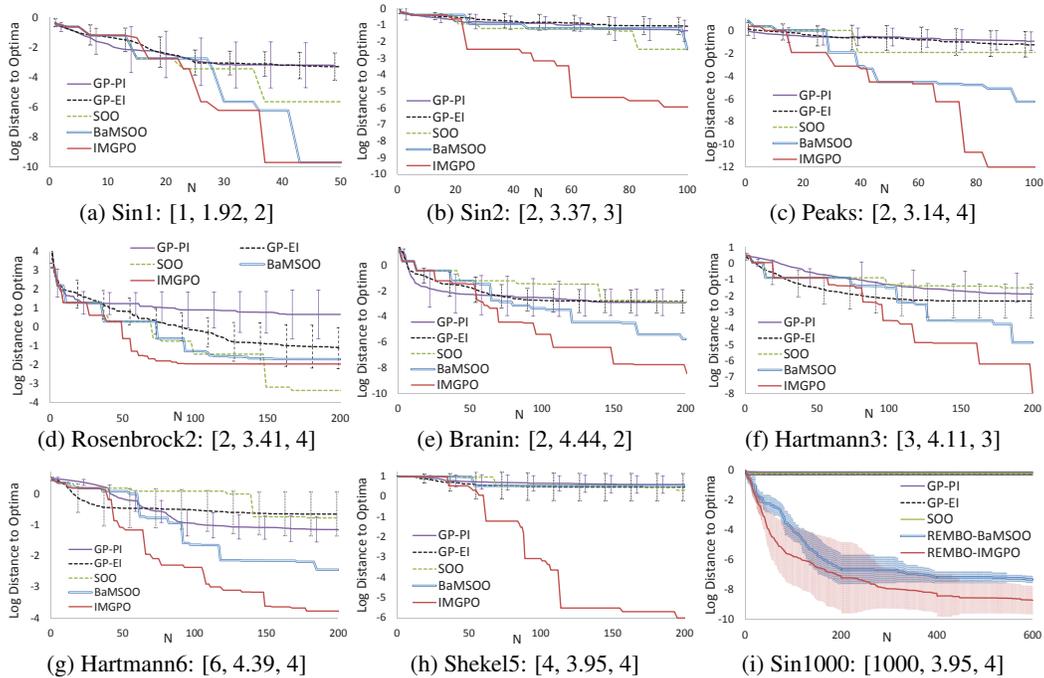


Figure 2: Performance Comparison: in the order, the digits inside of the parentheses [] indicate the dimensionality of each function, and the variables $\bar{\rho}_t$ and Ξ_n at the end of computation for IMGPO.

Table 1: Average CPU time (in seconds) for the experiment with each test function

Algorithm	Sin1	Sin2	Peaks	Rosenbrock2	Branin	Hartmann3	Hartmann6	Shekel5
GP-PI	29.66	115.90	47.90	921.82	1124.21	573.67	657.36	611.01
GP-EI	12.74	115.79	44.94	893.04	1153.49	562.08	604.93	558.58
SOO	0.19	0.19	0.24	0.744	0.33	0.30	0.25	0.29
BaMSOO	43.80	4.61	7.83	12.09	14.86	14.14	26.68	371.36
IMGPO	1.61	3.15	4.70	11.11	5.73	6.80	13.47	15.92

5 Experiments

In this section, we compare the IMGPO algorithm with the SOO, BaMSOO, GP-PI and GP-EI algorithms [18, 2, 3]. In previous work, BaMSOO and GP-UCB were tested with a pair of a handpicked good kernel and hyperparameters for each function [2]. In our experiments, we assume that the knowledge of good kernel and hyperparameters is unavailable, which is usually the case in practice. Thus, for IMGPO, BaMSOO, GP-PI and GP-EI, we simply used one of the most popular kernels, the isotropic Matern kernel with $\nu = 5/2$. This is given by $\kappa(x, x') = g(\sqrt{5}\|x - x'\|^2/l)$, where $g(z) = \sigma^2(1 + z + z^2/3) \exp(-z)$. Then, we blindly initialized the hyperparameters to $\sigma = 1$ and $l = 0.25$ for all the experiments; these values were updated with an empirical Bayesian method after each iteration. To compute the UCB by GP, we used $\eta = 0.05$ for IMGPO and BaMSOO. For IMGPO, Ξ_{max} was fixed to be 2^2 (the effect of selecting different values is discussed later). For BaMSOO and SOO, the parameter h_{max} was set to \sqrt{n} , according to Corollary 4.3 in [18]. For GP-PI and GP-EI, we used the SOO algorithm and a local optimization method using gradients to solve the auxiliary optimization. For SOO, BaMSOO and IMGPO, we used the corresponding deterministic division procedure (given Ω , the initial point is fixed and no randomness exists). For GP-PI and GP-EI, we randomly initialized the first evaluation point and report the mean and one standard deviation for 50 runs.

The experimental results for eight different objective functions are shown in Figure 2. The vertical axis is $\log_{10}(f(x^*) - f(x^+))$, where $f(x^*)$ is the global optima and $f(x^+)$ is the best value found by the algorithm. Hence, the lower the plotted value on the vertical axis, the better the algorithm’s performance. The last five functions are standard benchmarks for global optimization [21]. The first two were used in [18] to test SOO, and can be written as $f_{sin1}(x) = (\sin(13x) \sin + 1)/2$ for Sin1 and $f_{sin2}(x) = f_{sin1}(x_1)f_{sin1}(x_2)$ for Sin2. The form of the third function is given in Equation (16) and Figure 2 in [22]. The last function is Sin2 embedded in 1000 dimension in the same manner described in Section 4.1 in [14], which is used here to illustrate a possibility of using IMGPO as a main subroutine to scale up to higher dimensions with additional assumptions. For this function, we used REMBO [14] with IMGPO and BaMSOO as its Bayesian optimization subroutine. All of these functions are multimodal, except for Rosenbrock2, with dimensionality from 1 to 1000.

As we can see from Figure 2, IMGPO outperformed the other algorithms in general. SOO produced the competitive results for Rosenbrock2 because our GP prior was misleading (i.e., it did not model the objective function well and thus the property $f(x) \leq \mathcal{U}(x|\mathcal{D})$ did not hold many times). As can be seen in Table 1, IMGPO is much faster than traditional GP optimization methods although it is slower than SOO. For Sin 1, Sin2, Branin and Hartmann3, increasing Ξ_{max} does not affect IMGPO because Ξ_n did not reach $\Xi_{max} = 2^2$ (Figure 2). For the rest of the test functions, we would be able to improve the performance of IMGPO by increasing Ξ_{max} at the cost of extra CPU time.

6 Conclusion

We have presented the first GP-based optimization method with an exponential convergence rate $O(\lambda^{N+N_{gp}})$ ($\lambda < 1$) without the need of auxiliary optimization and the δ -cover sampling. Perhaps more importantly in the viewpoint of a broader global optimization community, we have provided a practically oriented analysis framework, enabling us to see why *not* relying on a particular bound is advantageous, and how a non-tight bound can still be useful (in Remarks 1, 2 and 3). Following the advent of the DIRECT algorithm, the literature diverged along two paths, one with a particular bound and one without. GP-UCB can be categorized into the former. Our approach illustrates the benefits of combining these two paths.

As stated in Section 3.1, our solution idea was to use a bound-based method but rely less on the estimated bound by considering all the possible bounds. It would be interesting to see if a similar principle can be applicable to other types of bound-based methods such as planning algorithms (e.g., A* search and the UCT or FSSS algorithm [23]) and learning algorithms (e.g., PAC-MDP algorithms [24]).

Acknowledgments

The authors would like to thank Dr. Remi Munos for his thoughtful comments and suggestions. We gratefully acknowledge support from NSF grant 1420927, from ONR grant N00014-14-1-0486, and from ARO grant W911NF1410433. Kenji Kawaguchi was supported in part by the Funai Overseas Scholarship. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of our sponsors.

References

- [1] N. De Freitas, A. J. Smola, and M. Zoghi. Exponential regret bounds for Gaussian process bandits with deterministic observations. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*, 2012.
- [2] Z. Wang, B. Shakibi, L. Jin, and N. de Freitas. Bayesian Multi-Scale Optimistic Optimization. In *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics (AISTAT)*, pages 1005–1014, 2014.
- [3] J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, pages 2951–2959, 2012.
- [4] R. G. Carter, J. M. Gablonsky, A. Patrick, C. T. Kelley, and O. J. Eslinger. Algorithms for noisy problems in gas transmission pipeline optimization. *Optimization and engineering*, 2(2):139–157, 2001.
- [5] J. W. Zwolak, J. J. Tyson, and L. T. Watson. Globally optimised parameters for a model of mitotic control in frog egg extracts. *IEEE Proceedings-Systems Biology*, 152(2):81–92, 2005.
- [6] L. C. W. Dixon. *Global optima without convexity*. Numerical Optimisation Centre, Hatfield Polytechnic, 1977.
- [7] B. O. Shubert. A sequential method seeking the global maximum of a function. *SIAM Journal on Numerical Analysis*, 9(3):379–388, 1972.
- [8] D. Q. Mayne and E. Polak. Outer approximation algorithm for nondifferentiable optimization problems. *Journal of Optimization Theory and Applications*, 42(1):19–30, 1984.
- [9] R. H. Mladineo. An algorithm for finding the global maximum of a multimodal, multivariate function. *Mathematical Programming*, 34(2):188–200, 1986.
- [10] R. G. Strongin. Convergence of an algorithm for finding a global extremum. *Engineering Cybernetics*, 11(4):549–555, 1973.
- [11] D. E. Kvasov, C. Pizzuti, and Y. D. Sergeyev. Local tuning and partition strategies for diagonal GO methods. *Numerische Mathematik*, 94(1):93–106, 2003.
- [12] S. Bubeck, G. Stoltz, and J. Y. Yu. Lipschitz bandits without the Lipschitz constant. In *Algorithmic Learning Theory*, pages 144–158. Springer, 2011.
- [13] J. Gardner, M. Kusner, K. Weinberger, and J. Cunningham. Bayesian Optimization with Inequality Constraints. In *Proceedings of The 31st International Conference on Machine Learning (ICML)*, pages 937–945, 2014.
- [14] Z. Wang, M. Zoghi, F. Hutter, D. Matheson, and N. De Freitas. Bayesian optimization in high dimensions via random embeddings. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 1778–1784. AAAI Press, 2013.

- [15] N. Srinivas, A. Krause, M. Seeger, and S. M. Kakade. Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, pages 1015–1022, 2010.
- [16] K. P. Murphy. *Machine learning: a probabilistic perspective*. MIT press, page 521, 2012.
- [17] C. E. Rasmussen and C. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [18] R. Munos. Optimistic optimization of deterministic functions without the knowledge of its smoothness. In *Proceedings of Advances in neural information processing systems (NIPS)*, 2011.
- [19] D. R. Jones, C. D. Perttunen, and B. E. Stuckman. Lipschitzian optimization without the Lipschitz constant. *Journal of Optimization Theory and Applications*, 79(1):157–181, 1993.
- [20] K. Kandasamy, J. Schneider, and B. Póczos. High dimensional Bayesian optimisation and bandits via additive models. *arXiv preprint arXiv:1503.01673*, 2015.
- [21] S. Surjanovic and D. Bingham. Virtual library of simulation experiments: Test functions and datasets. Retrieved November 30, 2014, from <http://www.sfu.ca/~ssurjano>, 2014.
- [22] D. B. McDonald, W. J. Grantham, W. L. Tabor, and M. J. Murphy. Global and local optimization using radial basis function response surface models. *Applied Mathematical Modelling*, 31(10):2095–2110, 2007.
- [23] T. J. Walsh, S. Goschin, and M. L. Littman. Integrating Sample-Based Planning and Model-Based Reinforcement Learning. In *Proceedings of the 24th AAAI conference on Artificial Intelligence (AAAI)*, 2010.
- [24] A. L. Strehl, L. Li, and M. L. Littman. Reinforcement learning in finite MDPs: PAC analysis. *The Journal of Machine Learning Research (JMLR)*, 10:2413–2444, 2009.

Supplementary Material

In this supplementary material, we provide the proofs of the theoretical results. Along the way, we also prove regret bounds for a general class of algorithms, the result of which may be used to design a new algorithm.

We first provide a known property of the upper confidence bound of GP.

Lemma 1. (Bound Estimated by GP) According to the belief encoded in the GP prior/posterior², for any x , $f(x) \leq \mathcal{U}(x|\mathcal{D})$ holds during the execution of Algorithm 1 with probability at least $1 - \eta$.

Proof. It follows the proof of lemma 5.1 of [15]. From the property of the standard gaussian distribution, $\Pr(f(x) > \mathcal{U}(x|\mathcal{D})) < \frac{1}{2}e^{-\varsigma_M^2/2}$. Taking union bound on the entire execution of Algorithm 1, $\Pr(f(x) > \mathcal{U}(x|\mathcal{D}) \forall M \geq 1) < \frac{1}{2} \sum_{M=1}^{\infty} e^{-\varsigma_M^2/2}$. Substituting $\varsigma_M = \sqrt{2 \log(\pi^2 M^2 / 12\eta)}$, we obtain the statement. \square

Our algorithm has a concrete division procedure in line 27 of Algorithm 1. However, one may improve the algorithm with different division procedures. Accordingly, we first derive abstract version of regret bound for the IMGPO (Algorithm 1) under a family of division procedures that satisfy Assumptions 3 and 4. After that, we provide a proof for the main results in the paper.

A With Family of Division Procedure

In this section, we modify the result obtained by [18]. Let $x_{h,i}$ to be any point in the region covered by the i^{th} hyperinterval at depth h , and $x_{h,i}^*$ be the global optimizer that may exist in the i^{th} hyperinterval at depth h . The previous work provided the regret bound of the SOO algorithm with a family of division procedure that satisfies the following two assumptions.

Assumption 3. (Decreasing diameter) There exists a diameter function $\delta(h) > 0$ such that, for any hyperinterval $\omega_{h,i} \subset \Omega$ and its center $c_{h,i} \in \omega_{h,i}$ and any $x_{h,i} \in \omega_{h,i}$, we have $\delta(h) \geq \sup_{x_{h,i}} \ell(x_{h,i}, c_{h,i})$ and $\delta(h-1) \geq \delta(h)$ for all $h \geq 1$.

Assumption 4. (Well-shaped cell) There exists $\nu > 0$ such that any hyperinterval $\omega_{h,i}$ contains at least an ℓ -ball of radius $\nu\delta(h)$ centered in $\omega_{h,i}$.

Thus, in this section, hyperinterval is not restricted to hyperrectangle. We now revisit the definitions of several terms and variables used in [18]. Let the ϵ -optimal space X_ϵ be defined as $X_\epsilon := \{x \in \Omega : f(x) + \epsilon \geq f(x^*)\}$. That is, the ϵ -optimal space is the set of input vectors whose function value is at least ϵ -close to the global optima. To bound the number of hyperintervals relevant to this ϵ -optimal space, we define a near-optimality dimension as follows.

Definition 3. (Near-optimality dimension) The near-optimality dimension is the smallest $d > 0$ such that, there exists $C > 0$, for all $\epsilon > 0$, the maximum number of disjoint ℓ -balls of radius $\nu\epsilon$ with center in the ϵ -optimal space X_ϵ is less than $C\epsilon^{-d}$.

Finally, we define the set of δ -optimal hyperintervals $I_{\delta(h)}$ as $I_{\delta(h)} := \{c \in \Omega : f(c) + \delta(h) \geq f(x^*), c \text{ is the center point of the interval, } \omega_{h,i}, \text{ for some } (h, i)\}$. The δ -optimal hyperinterval $I_{\delta(h)}$ is used to relate the hyperintervals to the ϵ -optimal space. Indeed, the δ -optimal hyperinterval $I_{\delta(h)}$ is almost identical to the $\delta(h)$ -optimal space $X_{\delta(h)}$, except that $I_{\delta(h)}$ is focused on the center points whereas $X_{\delta(h)}$ considers the whole input vector space. In the following, we use $|I_{\delta(h)}|$ to denote the number of $I_{\delta(h)}$ and derive its upper bound.

Lemma 2. (Lemma 3.1 in [18]) Let d be the near-optimality dimension and C denote the corresponding constant in Definition 1. Then, the number of δ -optimal hyperintervals is bounded by $|I_{\delta(h)}| \leq C\delta(h)^{-d}$.

²Thus, the probability in this analysis should be seen as that of *the subjective view*. If we assume that f is indeed a sample from the GP, we have the same result with *the objective view* of probability.

We are now ready to present the main result in this section. In the following, we use the term *optimal hyperinterval* to indicate a hyperinterval that contains a global optimizer x^* . We say a hyperinterval is *dominated* by other intervals when it is rejected or not selected in step (i)-(iii). In Lemma 3, we bound the maximum size of the optimal hyperinterval. From Assumption 1, this can be translated to the regret bound, as we shall see in Theorem 2.

Lemma 3. Let $\Xi_n \leq \min(\Xi, \Xi_{max})$ be the largest ξ used so far with n total node expansions. Let h_n^* be the depth of the deepest expanded node that contains a global optimizer x^* after n total node expansions (i.e., $h_n^* \leq n$ determines the size of the *optimal hyperinterval*). Then, with probability at least $1 - \eta$, h_n^* is bounded below by some h' that satisfies

$$n \geq \sum_{\tau=1}^{h'+\Xi} \rho_\tau.$$

Proof. Let T_h denote the time at which the optimal hyperinterval is further divided. We prove the statement by showing that the time difference $T_{h+1} - T_h$ is bounded by the number of δ -optimal hyperintervals. To do so, we first note that there are three types of hyperinterval that can dominate an optimal hyperinterval $c_{h+1,*}$ during the time $[T_h, T_{h+1} - 1]$, all of which belong to δ -optimal hyperintervals I_δ . The first type has the same size (i.e., same depth h), $c_{h+1,i}$. In this case,

$$f(c_{h+1,i}) \geq f(c_{h+1,*}) \geq f(x_{h+1,*}^*) - \delta(h+1),$$

where the first inequality is due to line 10 (step (i)) and the second follows Assumptions 1 and 2. Thus, it must be $c_{h+1,i} \in I_{h+1}$. The second case is where the optimal hyperinterval may be dominated by a hyperinterval of larger size (depth $l < h+1$), $c_{l,i}$. In this case, similarly,

$$f(c_{l,i}) \geq f(c_{h+1,*}) \geq f(x_{h+1,*}^*) - \delta(l),$$

where the first inequality is due to lines 11 to 12 (step (ii)) and thus $c_{l,i} \in I_l$. In the final scenario, the optimal hyperinterval is dominated by a hyperinterval of smaller size (depth $h+1+\xi$), $c_{h+1+\xi,i}$. In this case,

$$f(c_{h+1+\xi,i}) \geq z(h+1,*) \geq f(x_{h+1,*}^*) - \delta(h+1+\xi)$$

with probability at least $1 - \eta$ where $z(\cdot, \cdot)$ is defined in line 21 of Algorithm 1. The first inequality is due to lines 19 to 23 (step (iii)) and the second inequality follows Lemma 1 and Assumptions 1 and 3. Hence, we can see that $c_{h+1+\xi,i} \in I_{h+1+\xi}$.

For all of the above arguments, the temporarily assigned \mathcal{U} under GP has no effect. This is because the algorithm still covers the above three types of δ -optimal hyperintervals I_δ , as $\mathcal{U} \geq f$ with probability at least $1 - \eta$ (Lemma 1). However, these are only expanded based on f because of the temporary nature of \mathcal{U} . Putting these results together,

$$T_{h+1} - T_h \leq \sum_{\tau=1}^{h+1+\Xi_n} |I_{\delta(\tau)}| \rho_\tau.$$

Since if one of the I_δ is divided during $[T_h, T_{h+1} - 1]$, it cannot be divided again during another time period,

$$\sum_{h=0}^{h_n^*} T_{h+1} - T_h \leq \sum_{\tau=1}^{h_n^*+1+\Xi_n} |I_\tau| \rho_\tau,$$

where on the right-hand side, we could combine the summation $\sum_{h=0}^{h_n^*}$ and $\sum_{\tau=1}^{h+1+\Xi_n} |I_{\delta(\tau)}|$ into the one, because each h in the summation refers to the same δ -optimal interval $I_{\delta(l)}$ with $l \leq h_n^*+1+\Xi_n$, and should not be double-counted. As $\sum_{h=0}^{h_n^*} T_{h+1} - T_h = T_{h_n^*+1} - T_0$, $T_0 = 1$ and $|I_{\delta(0)}| = 1$,

$$T_{h_n^*+1} \leq 1 + \sum_{\tau=1}^{h_n^*+1+\Xi_n} |I_\tau| \rho_\tau \leq \sum_{\tau=1}^{h_n^*+1+\Xi_n} |I_\tau| \rho_\tau.$$

As $T_{h_n^*+1} > n$ by definition, for any h' such that $\sum_{\tau=1}^{h'+\Xi_n} |I_\tau| \rho_\tau \leq n < \sum_{\tau=1}^{h_n^*+1+\Xi_n} |I_\tau| \rho_\tau$, we have $h_n^* > h'$. \square

With Lemmas 2 and 3, we are ready to present a finite regret bound with the family of division procedures.

Theorem 2. Assume Assumptions 1, 3, and 4. Let $h(n)$ be the smallest integer h such that

$$n \leq \sum_{\tau=1}^{h+\Xi_n} \delta(\tau)^{-d} \rho_\tau.$$

Then, with probability at least $1 - \eta$, the regret of the IMGPO with any general division procedure is bounded as

$$r_n \leq \delta(h(n) - 1).$$

Proof. Let $c(n)$ and $c_{h_n^*,*}$ be the center point expanded at the n th expansion and the optimal hyperinterval containing a global optimizer x^* , respectively. Then, from Assumptions 1, 3, and 4, $f(c(n)) \geq f(c_{h_n^*,*}) \geq f^* - \delta(h_n^*)$, where f^* is the global optima. Hence, the regret bound is $r_h \leq \delta(h_n^*)$. To find a lower bound for the quantity h_n^* , we first relate $h(n)$ to Lemma 3 by

$$n > \sum_{\tau=1}^{h(n)+\Xi_n-1} \delta(\tau)^{-d} \rho_\tau \geq \sum_{\tau=1}^{h(n)+\Xi_n-1} |I_\tau| \rho_\tau,$$

where the first inequality comes from the definition of $h(n)$, and the second follows from Lemma 2. Then, from Lemma 3, we have $h_n^* \geq h(n) - 1$. Therefore, $r_n \leq \delta(h_n^*) \leq \delta(h(n) - 1)$. \square

Assumption 5. (Decreasing diameter revisit) The decreasing diameter defined in Assumption 3 can be written as $\delta(h) = c_1 \gamma^{h/D}$ for some $c_1 > 0$ and $\gamma < 1$ with a division procedure that requires c_2 function evaluations per node expansion.

Corollary 1. Assume Assumptions 1, 3, 4, and 5. Then, if $d = 0$, with probability at least $1 - \eta$,

$$r_N \leq O\left(\exp\left(-\frac{N + N_{gp}}{c_2 C D \bar{\rho}_t}\right)\right).$$

If $d > 0$, with probability at least $1 - \eta$,

$$r_N \leq O\left(\left(\frac{1}{N + N_{gp}}\right)^{1/d} \left(-\frac{c_2 C \bar{\rho}_t}{1 - \gamma^{d/D}}\right)^{1/d} \gamma^{-\frac{1}{D}}\right).$$

Proof. For the case $d = 0$, we have $n \leq \sum_{\tau=1}^{h(n)+\Xi_n} \delta(\tau)^{-d} \rho_\tau \leq \sum_{\tau=1}^{h(n)+\Xi_n+1} \bar{\rho}_t$, where the first inequality follows from the definition of $h(n)$, and the second comes from the definition of $\bar{\rho}_t$ and the assumption $d = 0$. The second inequality holds for $\bar{\rho}_t$ that only considers ρ_τ with $\tau \leq t$. This is computable, because $\tau \leq t$ by construction. Indeed, the condition of Lemma 3 implies $t \geq \sum_{l=0}^{h'+\Xi_n} |I_l|$. Therefore, the two inequalities hold, and we can deduce that $h(n) \geq \frac{n}{C \bar{\rho}_t} - \Xi_n - 1$ by algebraic manipulation. By Assumption 5, $n = (N + N_{gp})/c_2$. With this, substituting the lower bound of $h(n)$ into the statement of Theorem 2 with Assumption 5,

$$r_N \leq c_1 \exp\left(-\left[\frac{N + N_{gp}}{c_2 D} \frac{1}{C \bar{\rho}_t} - \Xi_n - 2\right] \ln \frac{1}{\gamma}\right).$$

Similarly, for the case $d > 0$,

$$n \leq \sum_{\tau=1}^{h(n)+\Xi_n} \delta(\tau)^{-d} \rho_\tau \leq \sum_{\tau=1}^{h(n)+\Xi_n+1} \frac{c^{-d} C \gamma^{-(h(n)+\Xi_n+1)d/D-1}}{\gamma^{-d/D-1}} \bar{\rho}_t,$$

and hence $c \gamma^{\frac{h(n)+\Xi_n}{D}} \leq \left(\frac{n(1-\gamma^{d/D})}{C \bar{\rho}_t}\right)^{-1/d}$ by algebraic manipulation. Substituting this into the result of Theorem 2, we arrive at the desired result. \square

B With a Concrete Division Procedure

In this section, we prove the main result in the paper. In Theorem 1, we show that the exponential convergence rate bound $O(\lambda^{N+N_{gp}})$ with $\lambda < 1$ is achieved *without* Assumptions 3, 4 and 5 and *without* the assumption that $d = 0$.

Theorem 1. Assume Assumptions 1 and 2. Let $\beta = \sup_{x, x' \in \Omega} \frac{1}{2} \|x - x'\|_\infty$. Let $\lambda = 3^{-\frac{\alpha}{2C\bar{\rho}_t D}} < 1$. Then, *without* Assumptions 3, 4 and 5 and *without* the assumption on d , with probability at least $1 - \eta$, the regret of IMGPO with the division procedure in Algorithm 1 is bounded as

$$r_N \leq L(3\beta D^{1/p})^\alpha \exp\left(-\alpha \left[\frac{N + N_{gp}}{2C\bar{\rho}_t D} - \Xi_n - 2\right] \ln 3\right) = O(\lambda^{N+N_{gp}}).$$

Proof. To prove the statement, we show that Assumptions 3, 4, and 5 can all be satisfied while maintaining $d = 0$.

From Assumption 2 (i), and based on the division procedure that the algorithm uses,

$$\sup_{x \in \omega_{h,i}} \ell(x, c_{h,i}) \leq \sup_{x \in \omega_{h,i}} L \|x - c_{h,i}\|_p^\alpha \leq L \left(3^{-\lfloor h/D \rfloor} \beta D^{1/p}\right)^\alpha.$$

This upper bound corresponds to the diagonal length of each hyperrectangle with respect to p -norm, where $3^{-\lfloor h/D \rfloor} \beta$ corresponds to the length of the longest side. We fix the form of δ as $\delta(h) = L3^\alpha D^{\alpha/p} 3^{-h\alpha/D} \beta^\alpha \geq L(3^{-\lfloor h/D \rfloor} \beta D^{1/p})^\alpha$, which satisfies Assumption 3.

This form of $\delta(h)$ also satisfies Assumption 5 with $\gamma = 3^{-\alpha}$ and $c_1 = L3^\alpha D^{\alpha/p} \beta^\alpha$.

Every hyperrectangle contains at least one ℓ -ball with a radius corresponding to the length of the shortest side of the hyperrectangle. Thus, we have at least one ℓ -ball of radius $\nu\delta(h) = L3^{-\alpha\lceil h/D \rceil} \geq L3^{-\alpha} 3^{-\alpha h/D}$ for every hyperrectangle with $\nu \geq 3^{-2\alpha} D^{-\alpha/p}$. This satisfies Assumption 4.

Finally, we show that $d = 0$. The set of δ -optimal hyperintervals $I_{\delta(h)}$ is contained by the $\delta(h)$ -optimal space $X_{\delta(h)}$ as

$$\begin{aligned} I_{\delta(h)} &= \{c \in \Omega : f(x^*) - f(c) \leq \delta(h), c \text{ is the center point of the interval, } \omega_{h,i}, \text{ for some } (h, i)\} \\ &\subseteq \{x \in \Omega : f(x^*) - f(x) \leq \delta(h)\} = X_{\delta(h)} \end{aligned}$$

Let θ be a value that satisfies Assumption 2 (ii) (which is nonzero). Consider an ℓ -ball of radius $\frac{\delta(h)}{\theta}$ at x^* , which is a set $\{x \in \Omega \mid \theta\ell(x, x^*) \leq \delta(h)\}$. Since $\theta\ell(x, x^*) \leq f(x^*) - f(x)$ by Assumption 2 (ii), the $\delta(h)$ -optimal space $X_{\delta(h)}$ is covered by an ℓ -ball of radius $\frac{\delta(h)}{\theta}$. Therefore, $I_{\delta(h)} \subseteq X_{\delta(h)} \subseteq$ (an ℓ -ball of radius $\frac{\delta(h)}{\theta}$ at x^*). By Assumption 2 (i), the volume V of an ℓ -ball of radius $\nu\delta(h)$ is proportional to $(\nu\delta(h))^D$ as $V_D^p(\nu\delta(h)) = (2\nu\delta(h)\Gamma(1+1/p))^D / \Gamma(1+D/p)$. Thus, the number of disjoint ℓ -balls of radius $\nu\delta(h)$ that fit in $X_{\delta(h)}$ is at most $\lceil (\frac{\delta(h)}{\theta\nu\delta(h)})^D \rceil = \lceil (\theta\nu)^{-D} \rceil$. Therefore, the number of ℓ -balls does not depend on $\delta(h)$ in this case, which means $d = 0$.

Now that we have satisfied Assumptions 3, 4, and 5 with $d = 0$, $\gamma = 3^{-\alpha}$, and $c_1 = L3^\alpha D^{\alpha/p} \beta^\alpha$, we follow the proof of Corollary 1 and deduce the desired statement. \square