

A Lipschitz Exploration-Exploitation Scheme for Bayesian Optimization

Ali Jalali, Javad Azimi, Xiaoli Fern, and Ruofei Zhang

Turn Inc., Microsoft Inc., Oregon State University and Microsoft Inc.
ajalali@turn.com, xfern@eecs.oregonstate.edu,
{jaazimi,bzhang}@microsoft.com

Abstract. The problem of optimizing unknown costly-to-evaluate functions has been studied extensively in the context of Bayesian optimization. Algorithms in this field aim to find the optimizer of the function by requesting only a few function evaluations at carefully selected locations. An ideal algorithm should maintain a perfect balance between exploration (probing unexplored areas) and exploitation (focusing on promising areas) within the given evaluation budget. In this paper, we assume the unknown function is Lipschitz continuous. Leveraging the Lipschitz property, we propose an algorithm with a distinct exploration phase followed by an exploitation phase. The exploration phase aims to select samples that shrink the search space as much as possible, while the exploitation phase focuses on the reduced search space and selects samples closest to the optimizer. We empirically show that the proposed algorithm significantly outperforms the baseline algorithms.

Keywords: Bayesian Optimization, Exploration, Exploitation, Lipschitz Continuity.

1 Introduction

In many applications, we would like to optimize an unknown function $f(\cdot)$ that is costly to evaluate over a compact input space. Classic optimization methods, such as gradient descent, cannot be applied to this type of problems since they need to evaluate the function frequently. In contrast, Bayesian Optimization (BO) [1,2] algorithms try to solve this problem with a small number of function evaluations. Bayesian optimization algorithms, generally, have two key components: 1) A posterior model to predict the output value of the function at any arbitrary input point, and 2) A selection criterion to determine which point to be evaluated next.

The first step of a BO algorithm is to learn a posterior probabilistic model over unobserved points of the function. Gaussian processes (GP) [3] have been used in the literature of Bayesian optimization as the probabilistic posterior model. GP models the function output for any unobserved point in the input space as a normal random variable, whose mean and variance depend on the location of the point in relation to a set of given observed samples. Based on

the learned posterior model, a selection criterion is then used to choose the next sample to be evaluated. A number of selection criteria have been proposed in the literature of Bayesian optimization. They typically work by selecting an example that optimizes some objective function designed to balance between exploring unobserved area and exploiting areas that are promising based on existing observations. Maximum probability of improvement [4,5] and maximum expected improvement (EI) [6] are two successful examples.

In this paper, we focus on the design of the selection criterion for Bayesian optimization. In particular, we study BO in a sequential setting [1,7], where the samples are chosen sequentially and a selection is made only after the function evaluations of the previous samples are revealed. We make a mild assumption that the unknown function is Lipschitz-continuous. Leveraging the Lipschitz property, we design a selection algorithm that operates in two distinct phases: the exploration phase and the exploitation phase. In general, in the context of Bayesian optimization [1] and bandit problems [8], the exploration phase selects sample from unexplored area while the exploitation focuses on promising area. In this paper, we introduce a new interpretation of exploration and exploitation.

The exploration phase of the proposed algorithm, at each step, selects a sample that eliminates the largest possible portion of the input space while guaranteeing, with high probability, that the eliminated part does not include the maximizer of the function. Hence, the exploration stage of the algorithm tries to shrink the search space of the function as much as possible. In contrast, the exploitation phase of our algorithm selects the point which is believed to be the closest sample to the optimal point with high probability.

Experimental results over 8 real and synthetic benchmarks indicate that the proposed approach is able to outperform the Expected Improvement (EI) criterion, one of the current state-of-the-art BO selection methods. In particular, we show that our algorithm is better than EI both in terms of the mean and variance of the performance. We also investigate whether combining our exploration stage with EI can boost the performance of EI. However, the results were negative. Sometimes it helps and sometimes it hurts and on average we observe little to no improvement to EI. This is possibly because our exploration method actively aims to eliminate regions from the input space and the EI criterion does not take that into consideration when selecting samples.

The remainder of the paper is organized as follows. In Section 2, we motivate the use of exploration-exploitation Bayesian optimization by analyzing the behavior of EI. Section 3 introduces our algorithm and provides insights into both theoretical and practical aspects of the algorithm. Experimental evaluation of our algorithm is shown in Section 4. Finally, the paper is concluded in Section 5.

2 Motivating Observation

In this section, we motivate our approach by revealing a key observation about the well known Expected Improvement (EI) algorithm [6]. The original EI is defined as

$$EI(x) = \mathbb{E} \left[(f(x) - y_{\max}) \mathbb{I}_{\{f(x) - y_{\max} > 0\}} \right], \quad (1)$$

where $\mathbb{I}_{\{\cdot\}}$ is the indicator function. Hence, it measures the expected improvement of the choice of x over the current maximum function evaluations y_{\max} over observed samples. Using Gaussian Process (GP) [3] as the posterior model of the unknown function, the EI objective can be represented by

$$EI(x|\mathcal{O}) = (\mu_{x|\mathcal{O}} - y_{\max})\Phi\left(\frac{\mu_{x|\mathcal{O}} - y_{\max}}{\sigma_{x|\mathcal{O}}}\right) + \sigma_{x|\mathcal{O}}\phi\left(\frac{\mu_{x|\mathcal{O}} - y_{\max}}{\sigma_{x|\mathcal{O}}}\right), \quad (2)$$

where, $\mu_{x|\mathcal{O}}$ and $\sigma_{x|\mathcal{O}}$ are the mean and standard deviation associated with the point x by GP, and, $\Phi(\cdot)$ and $\phi(\cdot)$ are standard Gaussian CDF and PDF, respectively. Here, $\mathcal{O} = \{(x_i, f(x_i))\}_{i=1}^n$ is the set of n observed samples $x_{\mathcal{O}}$ with their function evaluations $f(x_{\mathcal{O}})$ and define $y_{\max} = \max_{x_i \in x_{\mathcal{O}}} f(x_i)$. Further, the means and variances are defined as follows:

$$\begin{aligned} \mu_{x|\mathcal{O}} &= k(x, x_{\mathcal{O}}) k(x_{\mathcal{O}}, x_{\mathcal{O}})^{-1} f(x_{\mathcal{O}}) \\ \sigma_{x|\mathcal{O}}^2 &= k(x, x) - k(x, x_{\mathcal{O}}) k(x_{\mathcal{O}}, x_{\mathcal{O}})^{-1} k(x_{\mathcal{O}}, x), \end{aligned}$$

where $k(\cdot, \cdot)$ is some kernel function. In this paper, we consider Gaussian kernel $k(x_1, x_2) = \exp(-\frac{1}{\ell} \|x_1 - x_2\|_2^2)$.

EI has been widely used and studied; however, there has been always a concern about balancing the exploration and exploitation of EI. The main reason for this concern is that even though the asymptotic convergence of EI is guaranteed under certain conditions [9], EI tries to exploit the information and potentially can request a lot of samples if it hits a local optimum region, while we have a limited number of experiments. There has been some attempts in the literature to address this concern with varying degrees of success, which we briefly discuss here.

- (a) Considering the original definition of EI, researchers have proposed to replace y_{\max} with a smaller value to make EI more exploitative and with a larger value to make it more explorative. In particular, [10] suggested $y_{\max} + \xi$ and [11] suggested $(1 + \xi)y_{\max}$ to replace y_{\max} . However, this approach has not seen much empirical success. [10] showed that starting with large values of ξ (to be explorative in the beginning) and cooling it down (to make it more and more exploitative) makes little or no difference in the performance of EI.
- (b) On a separate line of work, [12] proposed to consider a surrogate function

$$EI_{\xi}(x) = \mathbb{E}[(f(x) - y_{\max})^{\xi} \mathbb{I}_{\{f(x) - y_{\max} > 0\}}].$$

For $\xi = 1$, this objective tries to improve over y_{\max} (exploiting mode) and if we decrease ξ it starts to explore uncertain areas (exploration mode). This method is very sensitive to small changes in ξ and except for very specific setup like the one used in [13], there is no systematic way to choose ξ . This makes it nearly impossible to use this method.

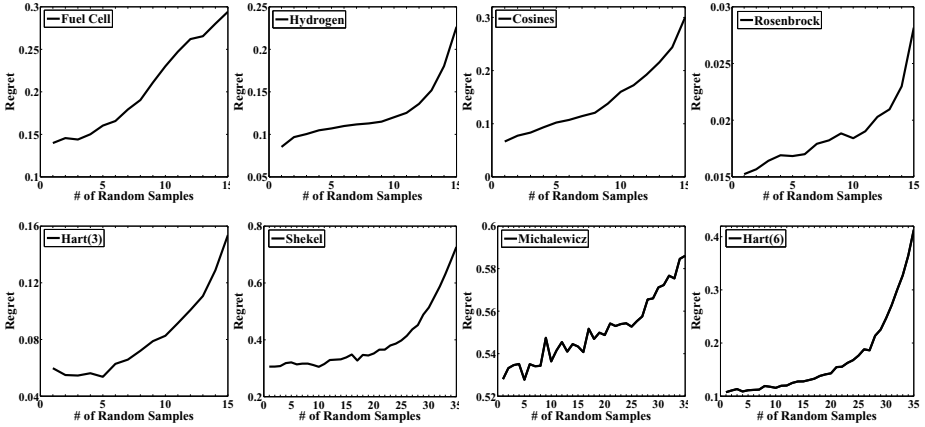


Fig. 1. Plot of regret versus the number of random exploration for EI algorithm. For a fixed budget n_b , we run a number of experiments as follows: first we consider the case where there are 1 random samples followed by $n_b - 1$ EI samples, next we consider the case where there are 2 random samples followed by $n_b - 2$ EI samples and so on. For 2D and 3D functions, we let $n_b = 15$ and for high-dimensional functions, we let $n_b = 35$. This result shows that the best EI performance is when we do not do random exploration.

(c) The third proposal is to have a “random” exploration phase preceding EI. In this approach, we take a number of random samples before switching to EI. We analyzed this method in Fig. 1. For a fixed budget n_b , we run n_b experiments as follows: first we consider the case where there is 1 random sample followed by $n_b - 1$ samples selected by the EI criterion, next we consider the case where there are 2 random samples followed by $n_b - 2$ EI samples and so on. The purpose of this investigation is to understand whether exploring with random samples prior to selecting with EI can improve the performance of EI, and if so how much exploring is necessary. We run this experiments on a number of different functions introduced in Section 4. These experiments reveal that “random” exploration never helps EI, since the regret monotonically increases as we increase the number of random samples from 1 to n_b . One possible explanation for this behavior is that the values of the function are highly correlated and hence, uniform sampling does not efficiently represent the skewness of the data points.

Based on the existing literature as well as our empirical investigation of EI discussed above, we would like to know whether or not it is possible to design an algorithm that operates in two naturally defined phases of exploration and exploitation and achieves consistently better performance than EI. We devote the next section to answer this question and introduce our proposed algorithm.

Algorithm 1. Next Best exploRative Sample (NBRs)

Input: Maximum M , Lipschitz Constant L and Set of observed samples $\{(x_1, f(x_1)), \dots, (x_t, f(x_t))\}$ **Output:** Next best explorative sample x

$$\mathbb{D}_t = \mathbb{D} - \bigcup_{i=1}^t \mathbb{S}(x_i, r_{x_i})$$

$$x \leftarrow \operatorname{argmax}_{x \in \mathbb{D}_t} \operatorname{Vol} \left(\mathbb{D}_t \cap \mathbb{S} \left(x, \frac{|M - \mu_x| - 1.5\sigma_x}{L} \right) \right)$$

3 Finite Horizon Bayesian Optimization

Not being able to balance the exploration-exploitation, EI might have poor performance especially when the query budget is small. In this section, we propose a two-phase exploration/exploitation algorithm that outperforms EI with its smart exploration and exploitation.

3.1 Exploration

Generally, a good exploration algorithm should be able to shrink the search space, so that we are left with a small region to focus on during the exploit stage. Let $\mathbb{D} = \bigotimes [a_i, b_i] \in \mathbb{R}^d$ be the Cartesian product of intervals $[a_i, b_i]$ for some $a_i < b_i$ and $i \in \{1, 2, \dots, d\}$. Suppose the unknown function $f : \mathbb{D} \mapsto [m, M]$ (with $f(x^*) = M$) is a Lipschitz function over \mathbb{D} with constant L , that is for all $x_1, x_2 \in \mathbb{D}$, we have

$$|f(x_1) - f(x_2)| \leq L \|x_1 - x_2\|_2.$$

Notice that if the function is not Lipschitz, then there is no hope that we can find the global optimum of $f(\cdot)$ even with infinitely countable evaluations. Thus, the Lipschitz continuity assumption is not a strong assumption. Moreover, functions with larger L are harder to optimize since they change more abruptly over the space.

For any point $x \in \mathbb{D}$, let $r_x = \frac{M - f(x)}{L}$ be the associated radius to the point x . By Lipschitz continuity assumption, we know that $x^* \notin \mathbb{S}(x, r_x)$, where, $\mathbb{S}(x, r_x)$ is the set of all points inside the sphere (or circle) with radius r_x centered at x (and single point x if $r_x \leq 0$); otherwise, the Lipschitz assumption is violated. This means if we have a sample at point x , then we do not need any more samples inside $\mathbb{S}(x, r_x)$.

The expected value of r_x satisfies $\mathbb{E}[r_x] = \frac{|M - \mu_x|}{L}$. Since $f(x)$ is a normal random variable $\mathcal{N}(\mu_x, \sigma_x^2)$, using Hoeffding inequality for all $\epsilon > 0$, we have

$$\mathbb{P} \left[r_x < \frac{|M - \mu_x|}{L} - \epsilon \right] \leq \exp \left(-\frac{2\epsilon^2 L^2}{\sigma_x^2} \right).$$

Replacing ϵ with $1.5\frac{\sigma_x}{L}$, the above inequality entails that with high probability (99%), $r_x \geq \frac{|M-\mu_x|-1.5\sigma_x}{L}$. Hence, a “good” algorithm for exploration should try to find x that maximizes the lower bound on r_x . This choice of x will remove a large volume of points from the search space. Note, however, if x is close to the boundaries of \mathbb{D} , then it might be the case that most of the volume of the sphere lies outside \mathbb{D} . Also, the sphere associated with x might have significant overlap with spheres of other points that are already selected. To fix this issue, we pick the point whose sphere has the largest intersection with unexplored search space in terms of its volume. The pseudo code of this method is described in Algorithm 1, which we refer to as the Next Best explorative Sample (NBRs) algorithm. NBRs achieves the optimal exploration in the sense that it maximizes the *expected explored* volume.

The value of $|M - \mu_x| - 1.5\sigma_x$ might be negative, especially for large values of σ_x . This artifact happens at points x that are “far” from previously observed samples. To prevent/minimize this, we need to make sure that the observed samples affect the mean and variance of all points in the space. For example, if we use the Gaussian kernel $k(x_1, x_2) = \exp(-\frac{1}{\ell_r}\|x_1 - x_2\|_2^2)$ for exploration, then we need to choose ℓ_r large enough to make sure each observed sample affects all the points in the space, e.g., $\ell_r \geq \sum_{i=1}^d (b_i - a_i)^2$. If we pick small ℓ_r , then the exploration algorithm starts exploring around the previous samples and extend the explored area gradually to reach to the other side of the search space. This strategy is not optimal if we have limited samples for exploration.

To implement NBRs, we need to maximize the volume

$$g(x) = \mathbf{Vol} \left(\mathbb{D}_t \cap \mathbb{S} \left(x, \frac{|M - \mu_x| - 1.5\sigma_x}{L} \right) \right)$$

where \mathbb{D}_t represents the current unexplored input space. To evaluate $g(x)$, we take a large number of points N inside the sphere $\mathbb{S}(x, \frac{|M - \mu_x| - 1.5\sigma_x}{L})$ uniformly at random. Then, for each point, we check if it crosses the borders $[a_i, b_i]$ or falls into the spheres of previously observed samples. If not, we count that point as a newly explored point. Finally, if there are n newly explored points, then we set $g(x) \approx \frac{n}{N} \left(\frac{|M - \mu_x| - 1.5\sigma_x}{L} \right)^d$.

To optimize $g(x)$, one can use deterministic and derivative free optimizers like DIRECT [14]. The problem is that DIRECT only optimizes Lipschitz continuous functions; however, $g(x)$ is not necessarily Lipschitz continuous. In our implementation, we take a large number of points inside \mathbb{D}_t and evaluate $g(\cdot)$ at those points and pick the maximum. This method might be slower than DIRECT, but avoids inaccurate results of DIRECT especially when \mathbb{D}_t describes a small region.

3.2 Exploitation

In the exploitation phase of the algorithm, we would like to use the information gained in the exploration phase to find the optimal point of $f(\cdot)$. Suppose we

Algorithm 2. Next Best exploitive Sample (NBIS)

Input: Maximum M , Lipschitz Constant L and Set of observed samples $\{(x_1, f(x_1)), \dots, (x_q, f(x_q))\}$ **Output:** Next best exploitive sample x

$$\mathbb{D}_q = \mathbb{D} - \bigcup_{i=1}^q \mathbb{S}(x_i, r_{x_i})$$

$$x \leftarrow \underset{x \in \mathbb{D}_q}{\operatorname{argmin}} \operatorname{Vol} \left(\mathbb{S} \left(x, \frac{|M - \mu_x|_{\mathcal{O}} + 1.5\sigma_x|_{\mathcal{O}}}{L} \right) \right)$$

have explored the search space with t samples and we want to find $x^* \in \mathbb{D}_t$. In order to exploit, we would like to find points x whose sphere is small. The reason is that if $r_x = \frac{M-f(x)}{L} \leq \gamma$ is small enough, then by *local* strong convexity of $f(\cdot)$ around x^* , for some constant κ we have

$$\frac{\kappa}{2} \|x - x^*\|_2^2 \leq M - f(x) \leq L\gamma.$$

Following the argument in Section 3.1, we estimate r_x by its mean $\mathbb{E}[r_x] = \frac{|M - \mu_x|}{L}$. By Hoeffding inequality, for all $\epsilon > 0$, we have

$$\mathbb{P} \left[r_x > \frac{|M - \mu_x|}{L} + \epsilon \right] \leq \exp \left(-\frac{2\epsilon^2 L^2}{\sigma_x^2} \right).$$

Similarly, replacing ϵ with $1.5\frac{\sigma_x}{L}$, the above inequality entails that with high probability (99%), $r_x \leq \frac{|M - \mu_x| + 1.5\sigma_x}{L}$. Hence, a “good” algorithm for exploitation should try to find the point x that minimizes the upper bound on r_x . This choice of x introduces the expected closest point to x^* . We present the pseudo code of this method in Algorithm 2.

The optimization in Algorithm 2 is nothing but minimizing

$$h(x) = \frac{|M - \mu_x|_{\mathcal{O}} + 1.5\sigma_x|_{\mathcal{O}}}{L}.$$

To optimize $h(x)$, again we take a large number of points in \mathbb{D}_q (the current unexplored space) uniformly at random and evaluate $h(\cdot)$ on those and pick the minimum.

3.3 Exploration-Exploitation Trade-Off

The main algorithm consists of an initial exploration phase followed by exploitation. Notice that we are using GP as an estimate of the unknown function and our method, like EI, highly relies on the quality of this estimation. On a high level, if the function is very complex, i.e., has large Lipschitz constant L , then

Table 1. Benchmark Functions

| | | | |
|--------------|--|---|---|
| Cosines(2) | $1 - (u^2 + v^2 - 0.3 \cos(3\pi u) - 0.3 \cos(3\pi v))$ $u = 1.6x - 0.5, v = 1.6y - 0.5$ | Rosenbrock(2) | $10 - 100(y - x^2)^2 - (1 - x)^2$ |
| Hartman(3,6) | $\sum_{i=1}^4 \Omega_i \exp\left(-\sum_{j=1}^d A_{ij}(x_j - P_{ij})^2\right)$ $\Omega_{1 \times 4}, A_{4 \times d}, P_{1 \times d}$ are constants | Michalewicz(5) | $-\sum_{i=1}^5 \sin(x_i) \sin\left(\frac{i x_i^2}{\pi}\right)^{20}$ |
| Shekel(4) | $\frac{1}{\sum_{i=1}^{10} \frac{1}{\omega_i + \sum_{j=1}^4 (x_j - B_{ji})^2}}$ | $\omega_{1 \times 10}, B_{4 \times 10}$ are constants | |

we need more exploration to fit better with GP. Small values of L correspond to flatter functions that are easier to optimize. Thus, in general, we expect the number of exploration steps to scale up with L . As a rule of thumb, functions we normally deal with satisfy $2 < L < 20$, for which we spend 20% of our budget in exploration and the rest in exploitation.

We use different kernel widths for the exploration and exploitation phases. In the case of exploration for complex functions, if we have enough budget (and hence, enough explorative samples), the kernel width can be set to a small value to fit a better local GP model. However, if we do not have enough budget, we need to take the kernel width to be large. In the case of exploitation, we pick the kernel width under which EI achieves its best performance.

Note that the choice of M and L plays a crucial role in this algorithm. If we pick L larger than the true Lipschitz function, then the radius of our spheres shrink and hence we might need more budget to achieve a certain performance. Choosing L smaller than the true Lipschitz is dangerous since it makes the spheres large and increases the chance of including the optimal point in a sphere and hence removing it. Thus, it is better to choose L slightly larger than our estimate of the true Lipschitz to be on the safe side.

The method is less sensitive to the choice of M , since the derivative of the radius with respect to M is proportional to $\frac{1}{L}$. Thus, as long as we do not over estimate M significantly, the $\frac{1}{L}$ factor prevents the spheres to become very large (and include/remove the optimal point). Small values of M , make the spheres smaller and hence, if we underestimate M , we would need more budget to achieve certain performance. However, if M is significantly (proportional to L) smaller than the true maximum of the function, then the algorithm will look for the point that achieves M and hence will perform poorly.

4 Experimental Results

In this section, we compare our algorithm with EI under different scenarios for different functions. We consider six well-known synthetic benchmark functions:

- (1,2) Cosines [15] and Rosenbrock [16] over $[0, 1]^2$
- (3,4) Hartman(3,6) [17] over $[0, 1]^{3,6}$
- (5) Shekel [17] over $[3, 6]^4$
- (6) Michalewicz [18] over $[0, \pi]^5$

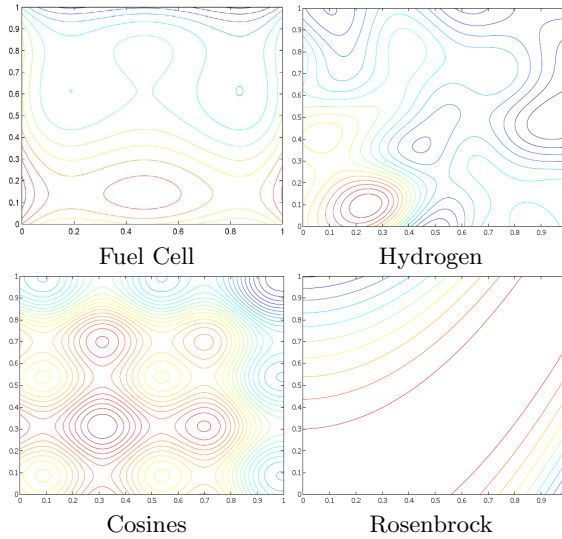


Fig. 2. The contour plots for the four 2–dimension proposed benchmarks

The mathematical expression of these functions are shown in Table 1. Moreover, we use two benchmarks derived from real-world applications:

- (1) Hydrogen [19] over $[0, 1]^2$
- (2) Fuel Cell [20] over $[0, 1]^2$

The contour plots of these two benchmarks along with the Cosines and Rosenbrock benchmarks are shown in Fig 2. The Fuel Cell benchmark is based on optimizing electricity output of microbial fuel cell by modifying some nano structure properties of the anodes. In particular, the inputs that we try to adjust are the average area and average circularity of the nano tube and the output that we try to maximize is the power output of the fuel cell. We fit a regression model on a set of observed samples to simulate the underlying function $f(\cdot)$ for evaluation. The Hydrogen benchmark is based on maximizing the Hydrogen production of a particular bacteria by varying the PH and Nitrogen levels of its growth medium. A GP is fitted to a set of observed samples to simulate the underlying function $f(\cdot)$. We consider a Lipschitz constant $L \approx 3$ for all of the benchmarks, except for Cosines and Michalewicz with $L \approx 6$ and Rosenbrock with $L \approx 45$. For the sake of comparison, we consider the normalized versions of all these functions and hence $M = 1$ in all cases. As mentioned previously, we spend 20% of the budget on exploration and 80% on exploitation.

4.1 Comparison to EI

In the first set of experiments, we would like to compare our algorithm with the best possible performance of EI. For each benchmark, we search over different

Table 2. Comparison of the best results of EI, NBRSEI and NBRSENBIS. This result shows that our algorithm outperforms the other two counterparts significantly in most cases both in terms of the mean and variance of the performance.

| | EI | EI_M | NBRSEI | NBRSENBIS |
|-----------|--------------|--------------|---------------------|---------------------|
| Cosines | .0736 ± .016 | .2938 ± .020 | .1057 ± .029 | .0270 ± .009 |
| Fuel Cell | .1366 ± .006 | .2232 ± .007 | .1357 ± .004 | .0965 ± .004 |
| Hydrogen | .0902 ± .004 | .1689 ± .012 | .1149 ± .004 | .0475 ± .006 |
| Rosen | .0134 ± .001 | .0153 ± .003 | .0163 ± .001 | .0034 ± .000 |
| Hart(3) | .0618 ± .006 | .0837 ± .001 | .0450 ± .003 | .0384 ± .003 |
| Shekel | .3102 ± .017 | .4104 ± .021 | .3011 ± .018 | .3240 ± .030 |
| Michal | .5173 ± .010 | .5210 ± .008 | .5011 ± .010 | .4554 ± .019 |
| Hart(6) | .1212 ± .002 | .2207 ± .006 | .1235 ± .002 | .1020 ± .003 |

values of the kernel width and find the one that optimizes EI’s performance. Fig. 1 is plotted using these optimal kernel widths and shows that the best performance of EI happens when we take only one random sample from a given budget. This performance is then used as the baseline for comparison in Table 2. In addition to EI, we introduced a new version of EI, called EI_M . Instead of taking the expectation of improvement I from 0 to infinity, (equation 2), we calculate the expectation of improvement from 0 to $M - y_{max}$ assuming M is given. This simple change decreases the level of exploration of EI and changes its behavior to be more exploitative than explorative. Using GP as our posterior model, the following lemma represents the EI_M . The proof is in supplementary document.

Lemma 1. Let $u_1 = \frac{y_{max} - \mu_x}{\sigma_x}$ and $u_2 = \frac{M - \mu_x}{\sigma_x}$, then

$$\begin{aligned}
EI_M(x) &= \mathbb{E} \left[(f(x) - y_{max}) \mathbb{I}_{\{0 \leq f(x) - y_{max} \leq M - y_{max}\}} \right] \\
&= \sigma(x) \left(-u_1 \Phi(u_2) + u_1 \Phi(u_1) + \phi(u_1) \right).
\end{aligned} \tag{3}$$

In light of the results of Fig. 1, we are also interested in whether our exploration algorithm can be used to improve the performance of EI. To this end, we replace the proposed exploitation algorithm with EI to examine if our exploration strategy helps EI. We refer to this setting as NBRSEI.

Table 2 summarizes the mean and variance of the performance, measured as the “Regret” = $M - \max f(x_{\mathcal{O}})$, for different benchmarks estimated over 1000 random runs. Interestingly, EI can consistently outperform the EI_M in all benchmarks. This shows that decreasing the exploration rate of EI could degrade the performance.

It is easy to see that in all benchmarks, our algorithm (NBRSENBIS) outperforms EI consistently except for the Shekel benchmark where EI and NBRSEI have slightly better performances. We suspect this is due to the fact that we have not optimized our kernel widths, where as the EI kernel width is optimized.

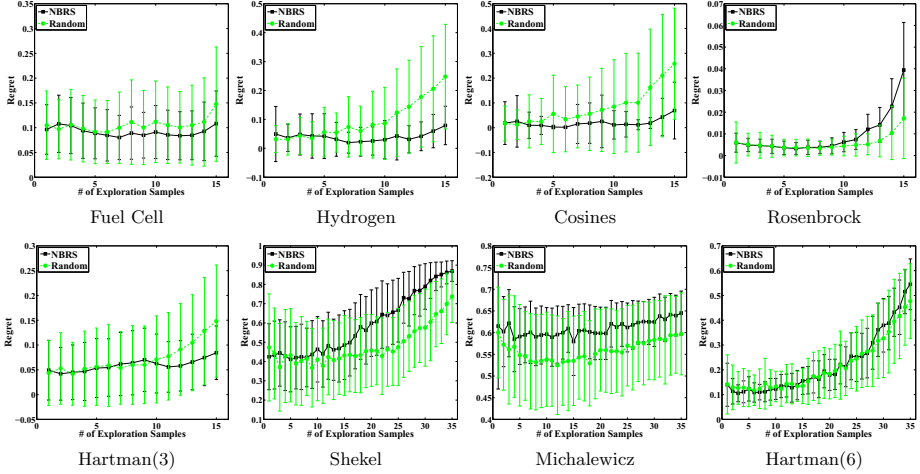


Fig. 3. Plot of regret versus the number of explorations for NBIS algorithm. For a fixed budget n_b , we run a number of experiments as follows: first we consider the case where there are 1 explorative sample (either random or NBRBS) followed by $n_b - 1$ EI samples, next we consider the case where there are 2 explorative samples followed by $n_b - 2$ EI samples and so on. For 2D and 3D functions, we let $n_b = 15$ and for high-dimensional functions, we let $n_b = 35$. This result shows that in most cases, our exploration is a) better than random, and b) necessary, since the regret achieves its minimum somewhere apart from zero. On average, we need to explore 20% of our budget, however, this portion can be optimized if we consider any specific function. The error bar here is the variance of the regret over different runs. This shows that our regret variance is smaller.

We also note that NBRBS+EI does not lead to any consistent improvement over EI. This is possibly due to the fact that EI does not take advantage of the reduced search space produced by NBRBS during selection.

4.2 Exploration Analysis

In the second set of experiments, we would like to compare our exploration algorithm NBRBS with random exploration when using NBIS for exploitation. As discussed previously, both random exploration and NBRBS fail to produce better performance when used with EI. Thus, it is interesting to see whether they can help NBIS in terms of the overall regret, and if so which one is more effective. Figure 3 summarizes this result for all benchmarks. For a fixed budget n_b , we start with 1 explorative sample (either using NBRBS or random) followed by $n_b - 1$ NBIS samples; next, we start with 2 explorative samples followed by $n_b - 2$ NBIS samples and so on. In each case, we average the regret over 1000 runs. The black line corresponds to the NBRBS exploration and the green line corresponds to the random exploration. We will discuss each function in more

details later, but in general, this result shows that our exploration algorithm is a) better than random exploration and b) necessary. To see why it is necessary, notice that the minimum regret on all curves is achieved for a non-zero number of NBRs samples. This means unlike EI, our exploitation algorithm benefits from NBRs.

Looking closer into the results, we see that NBRs always lead to a smaller regret comparing to the random exploration. On the Shekel benchmark, we see that random exploration has better performance if we spend majority of the budget to *explore*. However, for a *reasonable* amount of exploration that leads to the minimum regret (5 to 10 experiments), random exploration and NBRs achieve similar performance.

On our 6-dimensional benchmark Hartman(6), we notice that random exploration and NBRs behave very similarly. This shows that the input space is so large that no matter how clever you explore, you will not likely to improve the performance for the limited budget of 35.

NBRs starts from an initial point and explores the input space step by step. Imagine you are in a dark room with a torch in your hand and you want to explore the room. You start from an initial point and little by little walk through the space until you explore the whole space. This is exactly how NBRs does the exploration. Roughly speaking, NBRs minimizes $\mu_{x|\mathcal{O}} + 1.5\sigma_{x|\mathcal{O}}$ and hence, if a point is far from previous observations, i.e., $\sigma_{x|\mathcal{O}}$ is large, it is unlikely to be chosen. We see this effect in all functions, but most clearly in the Michalewicz benchmark. When the number of explorative samples is smaller than 10, the step-by-step explore procedure cannot explore the whole space and the exploitation can be trapped in local minima. For 10 – 15 explorative samples, NBRs can walk through the entire space fairly well and hence we get a minimum regret. For more than 15 explorative samples, since the space is well explored, we are wasting the samples that could be potentially used to improve our exploitation and hence, the performance becomes worse.

Finally, this investigation suggests that the result in Table 2 can be further improved by taking different number of explorative samples for different functions. To minimize parameter tuning, we chose to explore 20% of our budget. In general, this ratio can be adjusted according to the property of the function (e.g., the Lipschitz constant).

5 Conclusion

In this paper, we consider the problem of maximizing an unknown costly-to-evaluate function when we have a small evaluation budget. Using the Bayesian optimization framework, we proposed a two-phase exploration-exploitation algorithm that finds the maximizer of the function with few function evaluations by leveraging the Lipschitz property of the unknown function. In the exploration phase, our algorithm tries to remove as many points as possible from the search space and hence shrinks the search space. In the exploitation phase, the algorithm tries to find the point that is closest to the optimal. Our empirical results show that our algorithm outperforms EI (even in its best condition).

References

1. Jones, D.R.: A taxonomy of global optimization methods based on response surfaces. *Journal of Global Optimization* 21, 345–383 (2001)
2. Brochu, E., Cora, M., de Freitas, N.: A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. Technical Report TR-2009-23, Department of Computer Science, University of British Columbia (2009)
3. Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. MIT (2006)
4. Elder IV, J.F.: Global rd optimization when probes are expensive: the grope algorithm. In: *IEEE International Conference on Systems, Man and Cybernetics*, pp. 577–582 (1992)
5. Stuckman, B.E.: A global search method for optimizing nonlinear systems. *IEEE Transactions on Systems, Man, and Cybernetic* 18, 965–977 (1988)
6. Locatelli, M.: Bayesian algorithms for one-dimensional global optimization. *Journal of Global Optimization* 10(1), 57–76 (1997)
7. Moore, A., Schneider, J., Boyan, J., Lee, M.S.: Q2: Memory-based active learning for optimizing noisy continuous functions. In: *ICML*, pp. 386–394 (1998)
8. Li, W., Wang, X., Zhang, R., Cui, Y., Mao, J., Jin, R.: Exploitation and exploration in a performance based contextual advertising system. In: *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2010*, pp. 27–36. ACM (2010)
9. Vazquez, E., Bect, J.: Convergence properties of the expected improvement algorithm with fixed mean and covariance functions. *Journal of Statistical Planning and Inference* 140(11), 3088–3095 (2010)
10. Lizotte, D.: *Practical Bayesian Optimization*. PhD thesis, University of Alberta, Edmonton, Alberta, Canada (2008)
11. Azimi, J., Fern, A., Fern, X.: Batch bayesian optimization via simulation matching. In: *NIPS* (2010)
12. Schonlau, M.: *Computer Experiments and Global Optimization*. PhD thesis, University of Waterloo, Waterloo, Ontario, Canada (1997)
13. Sasena, M.J.: *Flexibility and Efficiency Enhancement for Constrained Global Design Optimization with Kriging Approximations*. PhD thesis, University of Michigan, Michigan, MI (2002)
14. Jones, D.R., Perttunen, C.D., Stuckman, B.E.: Lipschitzian optimization without the lipschitz constant. *Journal of Optimization Theory and Applications* 79(1), 157–181 (1993)
15. Anderson, B.S., Moore, A., Cohn, D.: A nonparametric approach to noisy and costly optimization. In: *ICML* (2000)
16. Brunato, M., Battiti, R., Pasupuleti, S.: A memory-based rash optimizer. In: *AAAI 2006 Workshop on Heuristic Search, Memory Based Heuristics and Their Applications* (2006)
17. Dixon, L., Szeg, G.: *The Global Optimization Problem: An Introduction Toward Global Optimization*. North-Holland, Amsterdam (1978)
18. Michalewicz, Z.: *Genetic algorithms + data structures = evolution programs*, 2nd edn. Springer-Verlag New York, Inc., New York (1994)

19. Burrows, E.H., Wong, W.K., Fern, X., Chaplen, F.W., Ely, R.L.: Optimization of ph and nitrogen for enhanced hydrogen production by *synechocystis* sp. pcc 6803 via statistical and machine learning methods. *Biotechnology Progress* 25, 1009–1017 (2009)
20. Azimi, J., Fern, X., Fern, A., Burrows, E., Chaplen, F., Fan, Y., Liu, H., Jaio, J., Schaller, R.: Myopic policies for budgeted optimization with constrained experiments. In: *AAAI* (2010)

Appendix: Proof of Lemma 1

Let $f(x)$ be our function prediction at any point x distributed as a normal random variable with mean μ_x and variance σ_x^2 ; i.e $f(x) \sim \mathcal{N}(\mu(x), \sigma_x^2)$ where μ_x and σ_x^2 obtained from Gaussian process. Suppose y_{max} is the best current observation, the probability of improvement of $I \in [0, M - y_{max}]$ can be calculated as $p(f(x) = y_{max} + I)$:

$$p(f(x) = y_{max} + I) = \frac{1}{\sqrt{2\pi}\sigma_x} \exp\left(-\frac{(y_{max} + I - \mu_x)^2}{2\sigma_x^2}\right). \tag{4}$$

Therefore we define $EI_M(x)$ as is simply the expectation of likelihood over $I \in [0, M]$ at any given point x :

$$\begin{aligned} EI_M(x) &= \int_{I=0}^{I=M-y_{max}} I \left\{ \frac{1}{\sqrt{2\pi}\sigma_x} \exp\left(-\frac{(y_{max} + I - \mu_x)^2}{2\sigma_x^2}\right) \right\} dI \\ &= \frac{1}{\sqrt{2\pi}\sigma_x} \exp\left(-\frac{(y_{max} - \mu_x)^2}{2\sigma_x^2}\right) \int_0^{M-y_{max}} I \exp\left(-\frac{2I(y_{max} - \mu_x) + I^2}{2\sigma_x^2}\right) dI. \end{aligned} \tag{5}$$

Let define

$$\begin{aligned} T &= \exp\left(-\frac{2I(y_{max} - \mu_x) + I^2}{2\sigma_x^2}\right) \\ \frac{\partial T}{\partial I} &= -\frac{1}{\sigma_x^2} (IT + (y_{max} - \mu_x)T), \end{aligned} \tag{6}$$

therefore we can get

$$IT = -(y_{max} - \mu_x)T - \frac{\partial T}{\partial I} \sigma_x^2. \tag{7}$$

Using equations 7,6,5 we can get

$$\begin{aligned} EI_M(x) &= \frac{1}{\sqrt{2\pi}\sigma_x} \exp\left(-\frac{(y_{max} - \mu_x)^2}{2\sigma_x^2}\right) \int_0^{M-y_{max}} IT dI \\ &= \sigma_x \phi\left(\frac{y_{max} - \mu_x}{\sigma_x}\right) \\ &\quad - (y_{max} - \mu_x) \int_0^{M-y_{max}} \frac{1}{\sqrt{2\pi}\sigma_x} \exp\left(-\frac{1}{2} \left(\frac{y_{max} + I - \mu_x}{\sigma_x}\right)^2\right) dI. \end{aligned} \tag{8}$$

Let

$$I^* = \frac{y_{max} + I - \mu_x}{\sigma_x}, \quad \text{then} \quad dI^* = \frac{dI}{\sigma_x}, \quad (9)$$

then the equation 8 can be written as

$$\begin{aligned} EI_M(x) &= \sigma_x \phi\left(\frac{y_{max} - \mu_x}{\sigma_x}\right) \\ &\quad - (y_{max} - \mu_x) \int_{\frac{y_{max} - \mu_x}{\sigma_x}}^{\frac{M - \mu_x}{\sigma_x}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}I^{*2}\right) dI^* \\ &= \sigma_x \phi\left(\frac{y_{max} - \mu_x}{\sigma_x}\right) - \left[(y_{max} - \mu_x) \left(\Phi\left(\frac{M - \mu_x}{\sigma_x}\right) - \Phi\left(\frac{y_{max} - \mu_x}{\sigma_x}\right)\right)\right]. \end{aligned} \quad (10)$$

Let

$$u_1 = \frac{y_{max} - \mu_x}{\sigma_x}, u_2 = \frac{M - \mu_x}{\sigma_x},$$

then we can finally drive the maximum expected improvement at any given point x as

$$MEI(x) = \sigma_x \left(-u_1 \Phi(u_2) + u_1 \Phi(u_1) + \phi(u_1) \right), \quad (11)$$

where $\Phi(\cdot)$ is the normal cumulative distribution function and $\phi(\cdot)$ is the standard normal distribution.