

Regional Expected Improvement for Efficient Trust Region Selection in High-Dimensional Bayesian Optimization

Nobuo Namura, Sho Takemori

Fujitsu Limited, Kawasaki, Japan
 {namura.nobuo, takemori.sho}@fujitsu.com

Abstract

Real-world optimization problems often involve complex objective functions with costly evaluations. While Bayesian optimization (BO) with Gaussian processes is effective for these challenges, it suffers in high-dimensional spaces due to performance degradation from limited function evaluations. To overcome this, simplification techniques like dimensionality reduction have been employed, yet they often rely on assumptions about the problem characteristics, potentially underperforming when these assumptions do not hold. Trust-region-based methods, which avoid such assumptions, focus on local search but risk stagnation in local optima. In this study, we propose a novel acquisition function, regional expected improvement (REI), designed to enhance trust-region-based BO in medium to high-dimensional settings. REI identifies regions likely to contain the global optimum, improving performance without relying on specific problem characteristics. We provide a theoretical proof that REI effectively identifies optimal trust regions and empirically demonstrate that incorporating REI into trust-region-based BO outperforms conventional BO and other high-dimensional BO methods in medium to high-dimensional real-world problems.

1 Introduction

Real-world optimization problems often involve objective functions where gradients are difficult to obtain, and evaluations are time-consuming. Typical examples include the design optimization of industrial products like aircraft (Liem, Mader, and Martins 2015) and automobiles (Moustapha et al. 2015), drug discovery (Shields et al. 2021), materials development (Zhang, Apley, and Chen 2020), and the tuning of hyperparameters in machine learning models (Perrone et al. 2021). For such optimization problems, surrogate-based methods that can optimize with only tens to a few thousand of objective function evaluations are widely used. In particular, Bayesian optimization (BO) (Mockus, Tiesis, and Zilinskas 1978) using Gaussian processes (GPs) (Matheron 1963) is frequently employed because it effectively balances the exploitation of the best-known solutions and the exploration of uncharted regions. This method constructs a GP model to represent the probability distribution of objective function values using a small number of evaluated sample points. It then iteratively generates new sample points by

optimizing an acquisition function, which quantifies the utility of any point in the design space based on the GP model.

However, conventional BO using global GP models suffers from performance degradation as the number of design variables increases, making it primarily suitable for problems up to about 10 dimensions. For medium to high-dimensional black-box optimization problems exceeding 10 dimensions, the exponentially increasing design space poses a challenge due to the limited number of samples available for objective function evaluations (Binois and Wycoff 2022). As a result, problem simplification approaches are often applied. Common techniques include local exploration using GP models within trust regions (Regis 2016), dimensionality reduction assuming a sparse set of effective design variables (Wang et al. 2016), reducing sensitivity in non-essential dimensions (Eriksson and Jankowiak 2021), simplifying GP models to capture only global trends under the assumption of a simple objective function (Hvarfner, Hellsten, and Nardi 2024), and solving subproblems under the assumption of separable design variables (Duvenaud, Nickisch, and Rasmussen 2011).

These methods have enabled improvements in objective function values even for problems with hundreds of dimensions. However, methods that rely on assumptions about problem characteristics may not perform well on problems that deviate from these assumptions. In practice, the characteristics of a problem are often unknown until it is solved, which raises concerns about the generalization performance of these methods. Even for problems with tens of dimensions, they underperform compared to standard BO with global GP models (Li et al. 2017). High-dimensional BO methods that do not rely on such assumptions are limited (Oh, Gavves, and Welling 2018; Gupta et al. 2020). Among these, methods that use trust regions have been gaining attention. However, these methods primarily focus on local search, which raises concerns about stagnation in local optima. For instance, trust region BO (TuRBO) (Eriksson et al. 2019) explores multiple randomly defined trust regions simultaneously to capture the global optimum, but efficiently obtaining the global optimum remains challenging due to the vast design space in high-dimensional problems.

In this study, we propose a novel acquisition function, regional expected improvement (REI), to enhance trust-region-based BO for global optimization in high-

dimensional problems. By introducing REI into trust-region-based BO, we can efficiently solve medium to high-dimensional problems without making assumptions about the characteristics of the problem. The specific contributions of this study are as follows:

- We propose REI, an acquisition function designed to identify regions likely to include the global optimum and to escape from local optima.
- We theoretically prove that REI effectively identifies an optimal trust region.
- We empirically show that incorporating REI into TuRBO outperforms conventional BO and various high-dimensional BO methods, without relying on assumptions about problem characteristics.

The code for REI is available at: <https://github.com/nobuonamura/regional-expected-improvement>. Appendices for this paper can be found in Namura and Takemori (2024).

2 Related Works

In high-dimensional BO, the sample density decreases across the extensive design space, necessitating approaches that reduce the complexity of GP models.

A common strategy assumes that in high-dimensional design spaces, some dimensions contribute minimally to the variation in the objective function. This strategy compresses the design space into an active subspace comprising only the effective dimensions. By constructing a GP model on the compressed design space, model complexity is reduced. The random embedding BO (REMBO) (Wang et al. 2016) and hashing-enhanced subspace BO (HeSBO) (Nayebi, Munteanu, and Poloczek 2019) use random matrices for dimensionality reduction and demonstrate theoretical utility when effective dimensions are limited. Alternatively, the sparse axis-aligned subspace BO (SAASBO) (Eriksson and Jankowiak 2021) simplifies the model by using a sparse prior distribution that reduces the sensitivity of the kernel function to ineffective dimensions by decreasing their inverted length scale. The vanilla BO with D -scaled $p(\ell)$ (DSP) (Hvarfner, Hellsten, and Nardi 2024) assumes that even in high-dimensional problems, the landscape of the objective function is simple. By setting a prior distribution that increases the length scale with the dimension size and simplifying the variance parameters of a GP model, it prevents model complexity from escalating. This approach achieves a simplification similar to SAASBO but with reduced computational cost.

Another notable approach involves using local GP models instead of a global GP model. This approach constructs GP models using sample points distributed within full-dimensional subregions of the design space or restricts sample addition to specific subregions. This local modeling approach suppresses complexity increases and accurately captures the local shape of the objective function. TuRBO employs a trust region centered around the best existing sample point, progressively reducing the size of the trust region when no improvement in the objective function is observed. This method captures detailed local structures incrementally

while advancing the optimization process. When the trust region becomes too small, the method restarts with a new trust region centered around the best randomly generated sample point. An extension of TuRBO, BO with adaptively expanding subspaces (BAxUS) (Papenmeier, Nardi, and Poloczek 2022), combines dimensionality reduction with TuRBO’s search, gradually increasing the dimensions used for model construction instead of restarting the trust region. Additionally, CASMOPOLITAN (Wan et al. 2021) extends TuRBO for discrete optimization and mixed-integer problems. Similarly, Bounce (Papenmeier, Nardi, and Poloczek 2023) extends BAxUS to handle discrete variables and incorporates strategies to reduce the number of samples needed for optimization in lower-dimensional spaces.

3 Preliminaries

We consider obtaining the global optimum of the minimization problem for the objective function f as shown below:

$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x} \in \Omega} f(\mathbf{x}), \quad (1)$$

where $\Omega = [0, 1]^D$ represents a D -dimensional design space. This study focuses on problems with $D > 10$, where the performance of standard GP-based BO is known to deteriorate. For maximization problems of an objective function f' , we consider solving the corresponding minimization problem by setting $f = -f'$. In this section, we discuss the foundational concepts of GP, BO, and TuRBO, which are prerequisites for introducing REI.

3.1 Gaussian Process

For a mean $m : \Omega \rightarrow \mathbb{R}$ and a covariance (kernel) function $k : \Omega \times \Omega \rightarrow \mathbb{R}$, a GP $\hat{f} \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$ is a random function (or stochastic process) characterized by the following properties: (i) for any sequence $\mathbf{x}_1, \dots, \mathbf{x}_i \in \Omega$ with $i \geq 1$, $\hat{f}(\mathbf{x}_1), \dots, \hat{f}(\mathbf{x}_i)$ are jointly Gaussian, (ii) for any $\mathbf{x}, \mathbf{x}' \in \Omega$, the mean $\mathbb{E}[\hat{f}(\mathbf{x})]$ and covariance $\operatorname{Cov}[\hat{f}(\mathbf{x}), \hat{f}(\mathbf{x}')] = \mathbb{E}[\hat{f}(\mathbf{x})\hat{f}(\mathbf{x}') - m(\mathbf{x})m(\mathbf{x}')] = k(\mathbf{x}, \mathbf{x}') - m(\mathbf{x})m(\mathbf{x}')$ are given as $m(\mathbf{x})$ and $k(\mathbf{x}, \mathbf{x}')$, respectively. By providing sample points $\mathcal{D}_t = \{\mathbf{x}_i, f_i\}_{i=1}^t$ consisting of pairs of specific design variables \mathbf{x}_i and objective function values f_i to this model, we can derive the functions representing the posterior mean $\hat{f}(\mathbf{x}|\mathcal{D}_t)$ and variance $\hat{\sigma}^2(\mathbf{x}|\mathcal{D}_t)$ as a surrogate of the objective function. In this study, assuming the mean function is zero, it is necessary to standardize the objective function values of the sample points before generating the GP model. Additionally, we assume that the observed values of the function include noise following $\mathcal{N}(0, \sigma_n^2)$. The noise variance σ_n^2 and the kernel hyperparameters are determined by maximum a posteriori (MAP) estimation.

3.2 Bayesian Optimization

In BO, the process iteratively approximates the objective function using the GP model based on existing sample points and adds new sample points by optimizing the acquisition function. The acquisition function is designed to balance exploitation, which involves adding new sample points in regions likely to contain the optimal solution to improve the

objective function, and exploration, which involves adding new sample points in regions with low sample density to improve the approximation accuracy of the GP model.

Prominent acquisition functions include expected improvement (EI) (Mockus, Tiesis, and Zilinskas 1978) and Thompson sampling (TS) (Thompson 1933). The EI at a location \mathbf{x} on the GP model, assuming a normal distribution, represents the expected value of the objective function improvement over the best-known objective function value $f_{\text{ref}} \in \mathcal{D}_t$, and is calculated as follows.

$$\begin{aligned} \text{EI}(\mathbf{x}) &= \int_{-\infty}^{f_{\text{ref}}} (f_{\text{ref}} - f)\phi(f)df \\ &= (f_{\text{ref}} - \hat{f}(\mathbf{x}|\mathcal{D}_t))\Phi(f_{\text{ref}}) + \hat{\sigma}(\mathbf{x}|\mathcal{D}_t)\phi(f_{\text{ref}}) \end{aligned}$$

where ϕ and Φ denote the probability density function and cumulative distribution function, respectively, of the normal distribution following $\mathcal{N}(\hat{f}(\mathbf{x}|\mathcal{D}_t), \hat{\sigma}^2(\mathbf{x}|\mathcal{D}_t))$. EI is often maximized using gradient methods, but the gradient can vanish in regions with small improvement. Therefore, LogEI (Ament et al. 2023), which uses the logarithm of the improvement, has been proposed. Additionally, a Monte Carlo acquisition function qEI (Wang et al. 2020; Balandat et al. 2020), which calculates EI via Monte Carlo sampling, is utilized to add multiple sample points simultaneously.

In contrast, TS involves randomly sampling functions $\text{TS} \sim \mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}'|\mathcal{D}_t))$ from the posterior distribution of the GP and then adding new samples by minimizing the sampled function in the case of minimization problems. Although the latter part is greedy, the functions sampled from the posterior distribution are more likely to have low objective function values in regions where the GP uncertainty $\hat{\sigma}^2(\mathbf{x})$ is high, thus balancing exploration and exploitation.

3.3 Trust Region Bayesian Optimization

TuRBO is a representative high-dimensional Bayesian optimization method that uses trust regions. It is denoted as TuRBO- m where m represents the number of trust regions used in parallel. For simplicity, Algorithm 1 illustrates the procedure for TuRBO-1, where there is only one trust region and a single sample point is added in each iteration. Note that while the original TuRBO-1 uses TS as the acquisition function, this study also implements LogEI.

In TuRBO, n initial sample points are generated uniformly at random across the entire design space, and a hypercube with side length $l = l_0$ centered around the sample point \mathbf{c} with the minimum objective function value is defined as the trust region. A GP model is then constructed using the initial sample points, and the next sample point is selected by minimizing the TS acquisition function or maximizing the LogEI acquisition function within the trust region.

To update the trust region, the objective function value of the added sample point is compared with that of the existing sample points for the current trust region. If the minimum value of the objective function does not improve for τ_{fail} consecutive iterations, l is halved, reducing the search area. Conversely, if the minimum value of the objective function improves for τ_{succ} consecutive iterations, l is doubled, expanding the search area. If the minimum objective function

Algorithm 1: TuRBO-1

Input: Objective function f , Design space Ω , Initial budget n , Total budget N

Parameter: Initial trust region length l_0 , Minimum trust region length l_{min}

Output: Samples $\mathcal{D}_N = \{\mathbf{x}_i, f_i\}_{i=1}^N$

```

1:  $\mathcal{D}_0 = \emptyset, l = l_0, t = 0$ 
2: while  $t < N$  do
3:   if  $t = 0$  or  $l < l_{\text{min}}$  then
4:     Generate initial samples  $\mathcal{D}^{TR} = \{\mathbf{x}_i, f_i\}_{i=1}^n$ 
5:     Initialize trust region  $\mathcal{TR}$  with length  $l = l_0$  centered around  $\mathbf{c} = \text{argmin}_{\mathbf{x} \in \mathcal{D}^{TR}} f(\mathbf{x})$ 
6:     Augment  $\mathcal{D}_{t+n} \leftarrow \mathcal{D}_t \cup \mathcal{D}^{TR}$ 
7:      $t \leftarrow t + n$ 
8:   else
9:     Generate a GP model  $\hat{f}(\mathbf{x}|\mathcal{D}^{TR}), \hat{\sigma}^2(\mathbf{x}|\mathcal{D}^{TR})$ 
10:    Select a next sample using an acquisition function
     $\mathbf{x}_{t+1} = \text{argmin}_{\mathbf{x} \in \mathcal{TR}} \text{TS}(\mathbf{x})$ 
    or
     $\mathbf{x}_{t+1} = \text{argmax}_{\mathbf{x} \in \mathcal{TR}} \text{LogEI}(\mathbf{x})$ 
11:    Evaluate objective function  $f_{t+1} = f(\mathbf{x}_{t+1})$ 
12:    Update  $\mathbf{c}$  and  $l$  of  $\mathcal{TR}$ 
13:    Augment  $\mathcal{D}^{TR} \leftarrow \mathcal{D}^{TR} \cup \{(\mathbf{x}_{t+1}, f_{t+1})\}$ 
14:    Augment  $\mathcal{D}_{t+1} \leftarrow \mathcal{D}_t \cup \mathcal{D}^{TR}$ 
15:     $t \leftarrow t + 1$ 
16:   end if
17: end while
18: return  $\mathcal{D}_t$ 

```

value is updated by the added sample point, the center of the trust region is also updated to the location of that sample point. The actual side length of the trust region is adjusted for each dimension according to the length scale of the GP model.

If l becomes smaller than a predetermined minimum length l_{min} , the current trust region's search is terminated. The optimization is then restarted in a new trust region by generating initial sample points randomly, without using any information from previous sample points.

4 Proposed Method

While TuRBO demonstrates excellent performance for high-dimensional optimization problems without requiring special assumptions, it predominantly focuses on local search and has limited strategies for obtaining the global optimum. In this section, we introduce REI, an approach specifically designed to efficiently discover trust regions that are likely to contain the global optimum. We describe the method for selecting the trust region when applying REI to TuRBO. It is important to note that REI is a general approach that can be applied to any optimization method utilizing trust regions, not just TuRBO.

4.1 Region-Averaged Acquisition Functions

Since TuRBO evaluates solutions hundreds of times within a single trust region once it is selected, the trust region should

be determined based on macro features within the region, rather than the local landscape details of the GP. Therefore, instead of using conventional acquisition functions such as EI or upper confidence bound (UCB), which strongly depend on the local structure of the GP, a specialized acquisition function that aggregates the GP structure within the region is needed.

This acquisition function requires converting the GP distribution within the trust region into a scalar value. We propose simply averaging commonly used acquisition functions within the trust region to achieve this scalarization. Given the center of the trust region as x_d and its length as l_d for the d -th dimension, the trust region around the center \mathbf{x} is defined as follows:

$$\mathcal{TR}_{\mathbf{x}} = \prod_{d=1}^D [\max(x_d - l_d/2, 0), \min(x_d + l_d/2, 1)]. \quad (2)$$

The region-averaged acquisition function is then defined as the mean value of the acquisition function within this trust region. While averaging the objective function is common in robust optimization (Ribaud et al. 2020), this study instead uses the average of the acquisition function, which captures the uncertainty of the GP model rather than the objective function itself.

4.2 Regional Expected Improvement

A practical region-averaged acquisition function, REI, is defined as the mean value of EI within the trust region. However, the integral of EI within the trust region is analytically challenging to obtain, so this study employs Monte Carlo sampling as a substitute:

$$\text{REI}(\mathbf{x}) = \frac{1}{V_{\mathcal{TR}_{\mathbf{x}}}} \int_{\mathcal{TR}_{\mathbf{x}}} \text{EI}(\hat{\mathbf{x}}) d\hat{\mathbf{x}} \approx \frac{1}{N_x} \sum_{j=1}^{N_x} \text{EI}(\hat{\mathbf{x}}_j), \quad (3)$$

where $\hat{\mathbf{x}}_j \in \mathcal{TR}_{\mathbf{x}}$ is a Monte Carlo sample in the trust region. We primarily use TuRBO-1; however, to prepare for selecting multiple trust regions in TuRBO- m , we employ the Monte Carlo acquisition function variant qREI by replacing EI with qEI:

$$\text{qREI}(\mathbf{x}) = \frac{1}{N_x N_f} \sum_{j=1}^{N_x} \sum_{k=1}^{N_f} \max(f_{\text{ref}} - \text{TS}_k(\hat{\mathbf{x}}_j), 0), \quad (4)$$

where TS_k is the k -th randomly sampled function from the posterior distribution of the GP.

By using qREI, regions with greater potential for yielding good solutions can be identified. The left side of Fig. 1 shows the results of calculating qREI and EI after uniformly generating initial samples for a one-dimensional minimization problem. EI exhibits a sharp peak on the left end and a broader peak on the right end. When using EI to select the trust region, aiming to locate points with the largest objective function improvement, the left end, where EI is maximized, is selected. In contrast, qREI averages EI within a defined range centered on each point (the orange-shaded region centered at the dashed line) to search for spaces with

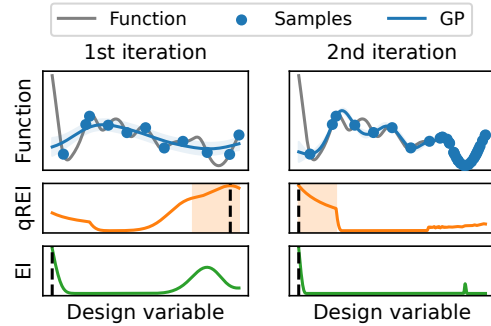


Figure 1: Comparison of qREI and EI distributions

Algorithm 2: Trust region selection

Input: Objective function f , Design space Ω , Initial budget n , Samples \mathcal{D}_t

Parameter: Initial trust region length l_0

Output: Samples \mathcal{D}^{TR}

- 1: **if** $\mathcal{D}_t = \emptyset$ **then**
 - 2: Generate initial samples at random
 $\mathcal{D}_t = \mathcal{D}^{TR} = \{\mathbf{x}_i, f_i\}_{i=1}^n$
 - 3: $t = n$
 - 4: **else**
 - 5: $\mathcal{D}^{TR} = \emptyset$
 - 6: **end if**
 - 7: Generate a GP model $\hat{f}(\mathbf{x}|\mathcal{D}_t), \hat{\sigma}^2(\mathbf{x}|\mathcal{D}_t)$
 - 8: Select a trust region $\mathcal{TR}_{\mathbf{c}}$ with length $l = l_0$ using qREI
 $\mathbf{c} = \text{argmax}_{\mathbf{x} \in \Omega} \text{qREI}(\mathbf{x})$
 - 9: Generate samples in the selected trust region at random
 $\{\mathbf{x}_i\}_{i=t+1}^{t+n-1} \in \mathcal{TR}_{\mathbf{c}}$
 - 10: Evaluate the objective function and augment a dataset
 $\mathcal{D}^{TR} \leftarrow \mathcal{D}^{TR} \cup \{\mathbf{c}, f(\mathbf{c})\} \cup \{\mathbf{x}_i, f(\mathbf{x}_i)\}_{i=t+1}^{t+n-1}$
 - 11: **return** \mathcal{D}^{TR}
-

robust improvement. Consequently, qREI achieves its maximum at the right end, where EI has a broader peak. Once the local search in the trust region near the right end is completed, a second region selection occurs (right side of Fig. 1), and local search is conducted in the left-end region.

In actual problems, EI is likely to have many local maxima, and qREI selects regions not only based on the width of each EI peak but also on the number of local optima within the region. Region selection using qREI initially prioritizes areas with broad valleys, ensuring good solutions even within local optima. As the search progresses, qREI shifts to selecting regions with deep valleys that have the potential to update known local optima.

4.3 Trust Region Selection

The trust region selection using qREI is performed during the initialization and restart phases of TuRBO. Specifically, this is achieved by replacing the initial sample generation in line 4 of Algorithm 1 with the procedure outlined in Algorithm 2. During these phases, a global GP model is gener-

ated using all existing sample points. The next trust regions are determined by fixing the trust region length to the initial value $l = l_0$ and varying the centers $\mathbf{x} \in \Omega$ to maximize qREI. Once the trust region is selected, TuRBO proceeds as described in Algorithm 1 using initial samples that include one point at the center of the selected trust region and $(n-1)$ points randomly generated within the trust region.

5 Theoretical Analysis

In the last section, we have defined region-averaged acquisition functions (specifically the qREI acquisition function) for identifying a trust region likely to include the global optimum. In this section, we prove that these region-averaged acquisition functions enable effective exploration of an optimal trust region (i.e., region likely including the global optimum). We provide proofs in Appendix A.

Notation and Assumptions. Let $k : \mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}$ be a positive definite kernel of the form $k(\mathbf{x}, \mathbf{x}') = F(\mathbf{x} - \mathbf{x}')$ for any $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^D$, where F is a continuous function on \mathbb{R}^D with $F \in L^1(\mathbb{R}^D)$ and $L^1(\mathbb{R}^D)$ is the space of L^1 integrable functions. This represents a wide class of kernels including Matérn, Gaussian, and rational quadratic kernels (Wendland 2004, Chapter 6). We denote the associated reproducing kernel Hilbert space (RKHS) on Ω by $\mathcal{H}_F(\Omega)$ and the RKHS norm of $f \in \mathcal{H}_F(\Omega)$ by $\|f\|_{\mathcal{H}_F(\Omega)}$.

Region-Averaging Operator. To discuss an optimization problem regarding the region-averaged acquisition functions, we define an operator \mathcal{S} that transforms the objective function to a region-averaged objective function as follows. For $f \in \mathcal{H}_F(\Omega)$, it is known that there exists a natural extension $\tilde{f} : \mathbb{R}^D \rightarrow \mathbb{R}$ of f (i.e., $\tilde{f}(\mathbf{x}) = f(\mathbf{x})$ for any $\mathbf{x} \in \Omega$) with $\tilde{f} \in \mathcal{H}_F(\mathbb{R}^D)$ (we refer to Theorem 10.46 in Wendland (2004). See also Proposition A.2). For $f \in \mathcal{H}_F(\Omega)$, we define a region-averaged objective function $\mathcal{S}f : \Omega \rightarrow \mathbb{R}$ as

$$\mathcal{S}f(\mathbf{x}) := \mathbb{E}_{\mathbf{u} \sim \text{Unif}(\mathcal{T}\mathcal{R}_0)} [\tilde{f}(\mathbf{x} + \mathbf{u})], \quad \mathbf{x} \in \Omega.$$

Here $\mathcal{T}\mathcal{R}_0$ is a trust region centered at $0 \in \mathbb{R}^D$ defined as $(\prod_{i=1}^D [-l_i/2, l_i/2])$ and $\text{Unif}(\mathcal{T}\mathcal{R}_0)$ denotes the uniform distribution on it. In Fig. 2, we illustrate a simple example of the objective function f and $\mathcal{S}f$ in the case when $D = 1$. If $a : \mathbb{R}^D \rightarrow \mathbb{R}$ is a lower confidence bound (LCB) of the (extended) objective function \tilde{f} , then by definition, the region-averaged acquisition function $\mathcal{S}a$ is a LCB of $\mathcal{S}f$. More generally, for any acquisition function a , we regard the region-averaged acquisition function $\mathcal{S}a$ as an acquisition function for $\mathcal{S}f$.

Problem Complexity of BO. For a Bayesian optimization algorithm \mathcal{A} and an objective function $f \in \mathcal{H}_F(\Omega)$, we define (cumulative) regret $R_N(f, \mathcal{A})$ by $\sum_{t=1}^N (f(\mathbf{x}_t) - f(\mathbf{x}^*))$, where $\mathbf{x}_t \in \Omega$ is a point selected by the algorithm \mathcal{A} . Following (Gupta et al. 2022; Chowdhury and Gopalan 2017), we define

$$U_N(f) = \beta_N(f) \sqrt{N\gamma_N}, \quad (5)$$

where $\beta_N(f) := \sigma_n \sqrt{2(\gamma_{N-1} + 1 + \log(1/\delta))} + \|f\|_{\mathcal{H}_F(\Omega)}$, $\delta \in (0, 1)$, and γ_N is the maximum information




Figure 2: A simple example of the objective function f and its region-averaged function $\mathcal{S}f$.

gain (Chowdhury and Gopalan 2017). Then, it is known that popular algorithms such as GP-EI and GP-UCB (or IGP-UCB) have the following regret upper bound (Gupta et al. 2022; Chowdhury and Gopalan 2017; Srinivas et al. 2010): That is, the following inequality holds with a probability at least $1 - \delta$:

$$R_N(f, \mathcal{A}) = O(U_N(f)),$$

for $\mathcal{A} = \text{GP-EI}$ and IGP-UCB , where $\delta \in (0, 1)$ is a parameter of the algorithms. Therefore, we can regard $U_N(f)$ as the problem complexity of the Bayesian optimization problem for the objective function f .

Main Theorem. The example illustrated in Fig. 2 indicates that the complexity of the objective function f is reduced if we take the operator \mathcal{S} . The following theorem formalizes this intuition for a wide class of kernels.

Theorem 5.1. For any objective function $f \in \mathcal{H}_F(\Omega)$, let $U_N(f)$ be the regret upper bound defined as (5). Then, for any $N \geq 1$, $U_N(f)$ is reduced if we take the operator \mathcal{S} , i.e., we have $\mathcal{S}f \in \mathcal{H}_F(\Omega)$ and the following inequality holds:

$$U_N(\mathcal{S}f) \leq U_N(f).$$

Theorem 5.1 indicates that with region-averaged acquisition functions such as qREI, one can explore a globally optimal trust region more effectively compared to the original optimization problem regarding f and shows a potential superiority over standard methods using global GP models.

Wan et al. (2021) also extended TuRBO and provided cumulative regret analysis (Theorems 3.3, 3.4, Wan et al. (2021)). However, their regret analysis followed the standard proof techniques developed by (Srinivas et al. 2010), and they have not provided a theoretical comparison to standard methods using global GP models such as GP-UCB (Srinivas et al. 2010).

6 Experiment

6.1 Implementation Details

We evaluate the effectiveness of qREI using TuRBO-1 implemented in Botorch (Balandat et al. 2020) with LogEI as the acquisition function for local search (TuRBO-1-LogEI). Results for TuRBO- m and TuRBO-1 with TS as the acquisition function are provided in Appendices E and F. For simplicity, qREI is always used for trust region selection with $N_x = 128$ and $N_f = 256$ used for Monte Carlo sampling. We use two versions of TuRBO with qREI: one where qREI is implemented as in Algorithm 2, used for both initialization and restarts, and another where standard TuRBO initialization is used, with qREI applied only during restarts. The




Figure 3: Histories and violin plots of objective function values for the medium-dimensional problems.

former is denoted as TuRBO-1-LogEI-qREI, while the latter is referred to as TuRBO-1-LogEI-qREI(restart).

The parameters for TuRBO follow those of Botorch implementation with $\tau_{\text{succ}} = 10$. Other key parameters, including $l_0 = 0.8$, $l_{\text{min}} = 0.5^7$, and $\tau_{\text{fail}} = D$, are consistent with those in the original paper. Details of the other parameter settings are provided in Appendix B.

6.2 Baselines

The baselines for comparison include TuRBO-1-LogEI (without qREI), along with DSP, Bounce, SAASBO, GP-LogEI, and CMA-ES (Hansen and Ostermeier 2001). DSP and Bounce were implemented using the authors’ original repositories. SAASBO and GP-LogEI adopted the Botorch implementation, with SAASBO using LogEI instead of EI as the acquisition function, and GP-LogEI using a global GP and LogEI for standard BO. For CMA-ES, we used the `pycma` (Hansen, Akimoto, and Baudis 2019) implementation through `pymoo` (Blank and Deb 2020).

6.3 Benchmarks

To address high-dimensional optimization problems where assumptions about problem characteristics do not hold, we use benchmark problems based on real-world applications. In evaluating the effectiveness of qREI, the number of restarts is considered crucial. To confirm the impact of qREI within a limited number of function evaluations, we mainly focus on design optimization problems for vehicles where noise influence is minimal. Specifically, we use human-powered aircraft (HPA) (Namura 2024), rover trajectory planning (Wang et al. 2018), and MOPTA08 problems. The HPA test suite includes nine single-objective problems. We used HPA101-0 ($D = 17$), 102-1 ($D = 32$), and 103-1 ($D = 32$) as medium dimensional problems and 101-2 ($D = 108$) as a high-dimensional problem. The rover trajectory planning problem ($D = 60$) was used with a fixed noise term to ensure reproducibility, in order to clearly assess the impact of qREI. The MOPTA08 problem ($D = 124$) was

used in its unconstrained form, consistent with the previous study (Eriksson and Jankowiak 2021). These two problems were used as high-dimensional problems.

6.4 Setup

For all problems and methods except CMA-ES, the number of initial samples is $n = 30$, and the batch size is 1. The population sizes of CMA-ES are determined by its own algorithm, implemented with the default settings. Due to computational resource constraints, the maximum number of function evaluations varies by method: $N = 2000$ for the proposed methods, TuRBO, Bounce, and CMA-ES; $N = 1000$ for DSP and GP-LogEI; and $N = 500$ for SAASBO. The number of independent runs with different random seeds is 11 for all problems. To compare the optimization results, we use the mean histories of the best objective function values with respect to the number of evaluations and violin plots of the final results. These visualizations allow us to evaluate the best method based on the number of evaluations and assess the variability in the performance of each method.

6.5 Results

Medium-Dimensional Problems. First, we evaluate the effectiveness of qREI in mid-dimensional problems, where TuRBO excels and trust regions are restarted many times. As illustrated in Fig. 3, the use of qREI enhances the performance of TuRBO-1-LogEI across all problems. Both variants of TuRBO-1-LogEI-qREI exhibit the best performance in HPA101-0 and HPA102-1. In these two problems, many methods underperform relative to GP-EI. However, TuRBO-1-LogEI-qREI, which significantly surpasses both TuRBO-1-LogEI and GP-EI (see Appendix C), emerges as a practical approach for medium-dimensional problems. In HPA103-1, unlike DSP, SAASBO, and GP-EI, which show significantly reduced performance, TuRBO-1-LogEI-qREI demonstrates excellent performance on par with Bounce and CMA-ES. It consistently exhibits high performance across different problems. In Appendix D, we experimentally demonstrate

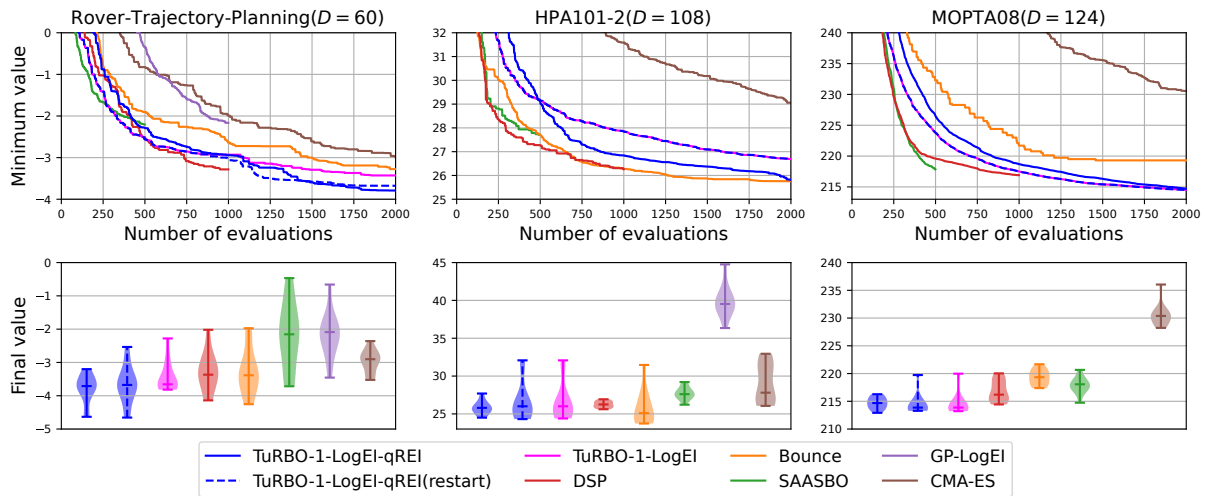


Figure 4: Histories and violin plots of objective function values for the high-dimensional problems.

that region-averaged acquisition functions, including qREI, yield better solutions than those obtained through trust region selection based on conventional acquisition functions, such as EI or UCB.

High-Dimensional Problems. Subsequently, we assess the performance of TuRBO-1-LogEI-qREI in high-dimensional problems where restarts are infrequent. In the rover trajectory planning problem, as shown in Fig. 4, one or two restarts barely occurred, and with the performance improvements brought by qREI, both variants of TuRBO-1-LogEI-qREI achieved the best results after more than 1000 evaluations. For the HPA101-2 and MOPTA08 problems, almost no restarts occurred, and even when a restart did happen, the number of evaluations reached 2000 in the early stages of local search in the new trust region. Therefore, the effect of qREI was limited to the initial trust region selection, resulting in nearly identical performance between TuRBO-1-LogEI-qREI(restart) and TuRBO-1-LogEI.

In the HPA101-2 problem, however, the introduction of qREI increased the probability of selecting a reliable initial trust region, and in the later stages of the optimization process, TuRBO-1-LogEI-qREI outperformed TuRBO-1-LogEI. It should be noted that by increasing the batch size and the number of evaluations, the performance improvement effect of qREI during restarts has been confirmed, as detailed in Appendix G. The MOPTA problem exhibits notably smaller variability across trials than other problems and converges to solutions of similar quality regardless of the chosen trust region, thus resulting in a very limited effect of qREI. Additionally, because computation was terminated when TuRBO-1-LogEI-qREI caught up with TuRBO-1-LogEI, the effect of qREI is evident only in the reduction of the maximum values in the violin plot for TuRBO-1-LogEI-qREI compared to TuRBO-1-LogEI.

The approach using qREI is considered effective even for high-dimensional problems. However, in high-dimensional problems, the number of evaluations required before restart-

ing increases, leading to fewer restarts, which is a limitation of this method. To address this, similar to the dimensionality management in Bounce, it may be effective to limit the number of function evaluations in each trust region and to shrink the trust region in a planned manner.

DSP and SAASBO showed good performance in the early stages of exploration in these high-dimensional problems although they had modest performance in medium-dimensional problems. While Bounce demonstrated excellent performance on HPA101-2, it did not perform as well on other problems. Since many of the methods in prior studies make assumptions about the problem characteristics, they may not always achieve sufficient performance depending on the nature of the real-world problem. Therefore, TuRBO-1-LogEI-qREI, which consistently shows stable and superior performance regardless of individual problem characteristics, can be considered a powerful method for medium to high-dimensional black-box optimization where the problem characteristics are unknown in advance.

7 Conclusions

In this paper, we introduced the use of regional expected improvement (REI) and its variant qREI within the TuRBO framework to enhance high-dimensional Bayesian optimization. We theoretically demonstrated the effectiveness of the region-averaged acquisition function, including qREI. Our experiments showed that when the trust region was restarted multiple times, the proposed method outperformed existing methods due to the optimal trust region selection by qREI. They also demonstrated robustness across different problem settings, in contrast to other methods that are more dependent on specific problem characteristics. These findings highlight qREI coupled with trust-region-based Bayesian optimization as a powerful and versatile tool for medium to high-dimensional black-box optimization.

References

- Ament, S.; Daulton, S.; Eriksson, D.; Balandat, M.; and Bakshy, E. 2023. Unexpected Improvements to Expected Improvement for Bayesian Optimization. In *Proceedings of the 36th Advances in Neural Information Processing Systems*.
- Balandat, M.; Karrer, B.; Jiang, D.; Daulton, S.; Letham, B.; Wilson, A. G.; and Bakshy, E. 2020. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. *Proceedings of the 33th Advances in Neural Information Processing Systems*, 21524–21538.
- Binois, M.; and Wycoff, N. 2022. A Survey on High-dimensional Gaussian Process Modeling with Application to Bayesian Optimization. *ACM Transactions on Evolutionary Learning and Optimization*, 2(2): 1–26.
- Blank, J.; and Deb, K. 2020. pymoo: Multi-Objective Optimization in Python. *IEEE Access*, 8: 89497–89509.
- Chowdhury, S. R.; and Gopalan, A. 2017. On Kernelized Multi-Armed Bandits. In *Proceedings of the 34th International Conference on Machine Learning*, 844–853. PMLR.
- Duvenaud, D. K.; Nickisch, H.; and Rasmussen, C. 2011. Additive Gaussian Processes. *Proceedings of the 24th Advances in Neural Information Processing System*.
- Eriksson, D.; and Jankowiak, M. 2021. High-Dimensional Bayesian Optimization with Sparse Axis-Aligned Subspaces. In *Proceedings of the 37th Conference on Uncertainty in Artificial Intelligence*, 493–503. PMLR.
- Eriksson, D.; Pearce, M.; Gardner, J.; Turner, R. D.; and Poloczek, M. 2019. Scalable Global Optimization via Local Bayesian Optimization. *Proceedings of the 32nd Advances in Neural Information Processing Systems*.
- Gupta, S.; Rana, S.; Venkatesh, S.; et al. 2020. Trading Convergence Rate with Computational Budget in High Dimensional Bayesian Optimization. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, 2425–2432.
- Gupta, S.; Rana, S.; Venkatesh, S.; et al. 2022. Regret Bounds for Expected Improvement Algorithms in Gaussian Process Bandit Optimization. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, 8715–8737. PMLR.
- Hansen, N.; Akimoto, Y.; and Baudis, P. 2019. CMA-ES/pycma on Github. Zenodo, DOI:10.5281/zenodo.2559634.
- Hansen, N.; and Ostermeier, A. 2001. Completely Derandomized Self-Adaptation in Evolution Strategies. *Evolutionary Computation*, 9(2): 159–195.
- Hvarfner, C.; Hellsten, E. O.; and Nardi, L. 2024. Vanilla Bayesian Optimization Performs Great in High Dimensions. In *Proceedings of the 41st International Conference on Machine Learning*, 20793–20817. PMLR.
- Li, C.; Gupta, S.; Rana, S.; Nguyen, V.; Venkatesh, S.; and Shilton, A. 2017. High Dimensional Bayesian Optimization Using Dropout. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 2096–2102.
- Liem, R. P.; Mader, C. A.; and Martins, J. R. 2015. Surrogate Models and Mixtures of Experts in Aerodynamic Performance Prediction for Aircraft Mission Analysis. *Aerospace Science and Technology*, 43: 126–151.
- Matheron, G. 1963. Principles of Geostatistics. *Economic Geology*, 58(8): 1246–1266.
- Mockus, J.; Tiesis, V.; and Zilinskas, A. 1978. The Application of Bayesian Methods for Seeking the Extremum. *Towards Global Optimization*, 2(117-129): 2.
- Moustapha, M.; Sudret, B.; Bourinet, J.-M.; and Guillaume, B. 2015. Adaptive Kriging Reliability-Based Design Optimization of an Automotive Body Structure Under Crashworthiness Constraints. In *Proceedings of the 12th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP12)*. Vancouver, Canada: University of British Columbia.
- Namura, N. 2024. Single and Multi-Objective Optimization Benchmark Problems Focusing on Human-Powered Aircraft Design. arXiv:2312.08953.
- Namura, N.; and Takemori, S. 2024. Regional Expected Improvement for Efficient Trust Region Selection in High-Dimensional Bayesian Optimization. arXiv:2412.11456.
- Nayebi, A.; Munteanu, A.; and Poloczek, M. 2019. A Framework for Bayesian Optimization in Embedded Subspaces. In *Proceedings of the 36th International Conference on Machine Learning*, 4752–4761. PMLR.
- Oh, C.; Gavves, E.; and Welling, M. 2018. BOCK: Bayesian Optimization with Cylindrical Kernels. In *Proceedings of the 35th International Conference on Machine Learning*, 3868–3877. PMLR.
- Papenmeier, L.; Nardi, L.; and Poloczek, M. 2022. Increasing the Scope as You Learn: Adaptive Bayesian Optimization in Nested Subspaces. *Proceedings of the 35th Advances in Neural Information Processing Systems*, 11586–11601.
- Papenmeier, L.; Nardi, L.; and Poloczek, M. 2023. Bounce: Reliable High-Dimensional Bayesian Optimization for Combinatorial and Mixed Spaces. *Proceedings of the 36th Advances in Neural Information Processing Systems*, 1764–1793.
- Perrone, V.; Donini, M.; Zafar, M. B.; Schmucker, R.; Kenthapadi, K.; and Archambeau, C. 2021. Fair Bayesian Optimization. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 854–863.
- Regis, R. G. 2016. Trust Regions in Kriging-Based Optimization with Expected Improvement. *Engineering Optimization*, 48(6): 1037–1059.
- Ribaud, M.; Blanchet-Scalliet, C.; Helbert, C.; and Gillot, F. 2020. Robust Optimization: A Kriging-Based Multi-Objective Optimization Approach. *Reliability Engineering & System Safety*, 200: 106913.
- Shields, B. J.; Stevens, J.; Li, J.; Parasram, M.; Damani, F.; Alvarado, J. I. M.; Janey, J. M.; Adams, R. P.; and Doyle, A. G. 2021. Bayesian Reaction Optimization as a Tool for Chemical Synthesis. *Nature*, 590(7844): 89–96.
- Srinivas, N.; Krause, A.; Kakade, S.; and Seeger, M. W. 2010. Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design. In *Proceedings of the 27th International Conference on Machine Learning*, 1015–1022.

- Thompson, W. R. 1933. On the Likelihood That One Unknown Probability Exceeds Another in View of the Evidence of Two Samples. *Biometrika*, 25(3-4): 285–294.
- Wan, X.; Nguyen, V.; Ha, H.; Ru, B.; Lu, C.; and Osborne, M. A. 2021. Think Global and Act Local: Bayesian Optimisation over High-Dimensional Categorical and Mixed Search Spaces. In *Proceedings of the 38th International Conference on Machine Learning*, 10663–10674. PMLR.
- Wang, J.; Clark, S. C.; Liu, E.; and Frazier, P. I. 2020. Parallel Bayesian Global Optimization of Expensive Functions. *Operations Research*, 68(6): 1850–1865.
- Wang, Z.; Gehring, C.; Kohli, P.; and Jegelka, S. 2018. Batched Large-Scale Bayesian Optimization in High-Dimensional Spaces. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*, 745–754. PMLR.
- Wang, Z.; Hutter, F.; Zoghi, M.; Matheson, D.; and De Feitas, N. 2016. Bayesian Optimization in a Billion Dimensions via Random Embeddings. *Journal of Artificial Intelligence Research*, 55: 361–387.
- Wendland, H. 2004. *Scattered data approximation*, volume 17. Cambridge university press.
- Zhang, Y.; Apley, D. W.; and Chen, W. 2020. Bayesian Optimization for Materials Design with Mixed Quantitative and Qualitative Variables. *Scientific Reports*, 10(1): 4924.