
Discovering Many Diverse Solutions with Bayesian Optimization

Natalie Maus

University of Pennsylvania

Kaiwen Wu

University of Pennsylvania

David Eriksson

Meta

Jacob Gardner

University of Pennsylvania

Abstract

Bayesian optimization (BO) is a popular approach for sample-efficient optimization of black-box objective functions. While BO has been successfully applied to a wide range of scientific applications, traditional approaches to single-objective BO only seek to find a single best solution. This can be a significant limitation in situations where solutions may later turn out to be intractable. For example, a designed molecule may turn out to violate constraints that can only be reasonably evaluated after the optimization process has concluded. To address this issue, we propose Rank-Ordered Bayesian Optimization with Trust-regions (ROBOT) which aims to find a portfolio of high-performing solutions that are diverse according to a user-specified diversity measure. We evaluate ROBOT on several real-world applications and show that it can discover large sets of high-performing diverse solutions while requiring few additional function evaluations compared to finding a single best solution.

1 INTRODUCTION

Bayesian optimization (BO) Jones et al. (1998); Shahriari et al. (2015) is a general framework for optimizing black-box functions $\arg \min_{\mathbf{x}^*} f(\mathbf{x}^*)$ in a sample-efficient fashion. BO has been successfully applied to hyperparameter tuning Snoek et al. (2012b); Turner et al. (2021), A/B testing Letham et al. (2019), chemical engineering Hernández-Lobato et al. (2017), drug discovery Negoescu et al. (2011), and more. For example, f may measure the antibiotic activity of a molecule \mathbf{x} , and we might therefore apply BO to design a molecule with high antibiotic activity.

In many of these settings, however, the fact that BO traditionally seeks a single best optimizer \mathbf{x}^* may be a significant

limitation. This “all-or-nothing” attribute of BO is particularly undesirable for problems where the returned \mathbf{x}^* may indeed optimize some useful objective, but is later found to be unsuitable for unforeseen reasons. For example, a molecule \mathbf{x}^* may have strong *in vitro* antibiotic activity, but may later be found unsafe or ineffective for use in humans through clinical testing. Worse, incorporating constraints like human safety directly into the optimization procedure of unknown molecules seems intractably expensive at best and unethical at worst.

In these settings where the risk of wasting the evaluation budget in search of \mathbf{x}^* is high, practitioners benefit from being given—in addition to the best single optimizer we can find—a series of alternative solutions to the problem: a set of “back up plans.” Formally, we might seek a set $S^* = \{\mathbf{x}_1^*, \mathbf{x}_2^*, \dots, \mathbf{x}_M^*\}$ of solutions that are all of high objective value, but we may require that these solutions are sufficiently *diverse* to ensure that this set of solutions is less likely to later fail for unrelated reasons. The practitioner may therefore provide us with a *symmetric diversity measure* $\delta(\mathbf{x}, \mathbf{x}')$ that must exceed some threshold τ for all pairs of solutions in the set S^* . For example, a biochemist may require that molecules in S^* have sufficiently low fingerprint similarity Brown et al. (2019). Solving this problem efficiently equips practitioners with large sets of *potential* solutions to their true problem, which can mitigate the risk that a single optimizer fails to be useful. This diverse solution search problem is challenging to cast under existing BO frameworks. The diversity constraints $\delta(\mathbf{x}_i^*, \mathbf{x}_j^*) \geq \tau$ are challenging to view as traditional black-box constraints, as the constraint functions for the j th point in S^* depend on the locations of the other points in S^* .

In this paper, we propose ROBOT, a method to solve the above problem and find a diverse set of high scoring solutions S^* so that $\forall \mathbf{x}_i^*, \mathbf{x}_j^* \in S^*, \delta(\mathbf{x}_i^*, \mathbf{x}_j^*) \geq \tau$. Across a variety of problem settings ranging from reinforcement learning to molecule design, ROBOT is able to discover large sets of diverse solutions S^* with little loss in efficiency over finding a single best solution.

Contributions

1. We introduce the problem setting of finding a set of high-performing solutions under a *user-defined diver-*

sity measure. While prior work outside the BO literature has considered a similar setting where diversity is taken to be distance in input space (e.g. Vassiliades et al. (2016); Mouret and Clune (2015)), this is the first work we are aware of to consider arbitrary, user-defined diversity measures like fingerprint similarity for molecules.

2. We propose a local Bayesian optimization solution, ROBOT, to this problem that extends to large sample sizes, high dimensional inputs, and structured inputs.
3. We provide empirical results across challenging, high-dimensional optimization tasks to show that our algorithm can consistently produce large populations of diverse, high-performing solutions with *virtually no loss in efficiency* compared to finding a *single* solution.
4. We additionally demonstrate results on structured drug discovery tasks using the widely used fingerprint similarity function as a diversity measure, demonstrating the value of our approach to practitioners in the physical sciences.
5. We prove global consistency of ROBOT in subsection A.5.
6. We release an open-source implementation of ROBOT using BoTorch (Balandat et al., 2020).

2 BACKGROUND AND RELATED WORK

Bayesian optimization. Bayesian optimization (BO) Moćkus (1975); Snoek et al. (2012a) is an approach to sample-efficient black-box optimization that utilizes a probabilistic *surrogate model*—commonly a Gaussian process (GP) Rasmussen (2003)—and an *acquisition function* that leverages the surrogate to find the most promising candidates to evaluate next. BO is a sequential optimization algorithm that proceeds in iterations. In each iteration, a surrogate model is trained on data collected from evaluating the black-box objective function. The acquisition function, defined given the surrogate model’s predictive posterior, is then maximized to select one or more candidates to evaluate next, trading off exploration and exploitation.

Parametric Gaussian process regressors (PPGPR). Because we consider tasks with large function evaluation budgets, we use an approximate GP surrogate model. Approximate GP models use inducing point methods in combination with variational inference to allow approximate GP inference on large data sets Hensman et al. (2013); Titsias (2009). In this work, we use the PPGPR approximate GP model proposed by Jankowiak et al. (2020), which we found provides substantial improvements in Bayesian optimization performance on the molecule optimization tasks we consider.

Constrained Bayesian optimization. While it may be tempting to attempt to formulate our problem as a constrained BO problem Gardner et al. (2014); Hernández-Lobato et al. (2016); Eriksson and Poloczek (2021), this is challenging as the constraints depend on the set S^* and are therefore non-stationary. In particular, the i th point in S^* , \mathbf{x}_i^* , must satisfy $i - 1$ constraints that depend on $\mathbf{x}_1^*, \dots, \mathbf{x}_{i-1}^*$ which are unknown in advance. One potential solution is to acquire the points in S^* sequentially. Specifically, unconstrained optimization can be used to obtain \mathbf{x}_1^* , the second point can then be acquired subject to the single constraint that $d(\mathbf{x}_1^*, \mathbf{x}_2^*) \geq \tau$, and so on. As a baseline, we adapt the work of Eriksson and Poloczek (2021) to utilize this strategy, which we refer to as `Sequential SCBO` in section 4.

Multi-objective Bayesian optimization (MOBO). There has been much work in developing new methods for MOBO in recent years Hernández-Lobato et al. (2015); Belakaria et al. (2019); Turchetta et al. (2019); Daulton et al. (2021). However, these methods cannot be readily applied since our problem setting is quite different. Diversity in our setting is not a second objective since we do not try to maximize the diversity, but instead we require that the diversity between pairs of solutions exceed some threshold τ . For example, if two molecules are sufficiently diverse, they can be expected to have relatively unrelated chemical properties and further increasing their diversity may not provide much value to the practitioner. If one desires to simultaneously maximize the diversity between solutions, this becomes a very different multi-objective problem. Using existing MOBO methods for even this different problem setting is itself non-trivial because any diversity measure cannot be measured for a single point in isolation and the Pareto frontier therefore does not exist in the usual sense. This prevents the direct application of methods such as that of Konakovic Lukovic et al. (2020) which seek to maximize diversity along the Pareto frontier.

Generative modeling for molecules. Many generative modeling approaches have been proposed to generate populations novel molecules. This includes variational autoencoder models such as the Junction Tree Variational Auto Encoder Jin et al. (2018) and the SELFIES-VAE Maus et al. (2022). Populations of molecules generated by sampling from these models can be evaluated for diversity and validity using methods such as Xie et al. (2021). However, while these models can successfully generate diverse populations of molecules, they are not designed to generate molecules with any particular user specified characteristics. Thus, when we desire a diverse population of molecules which also each have some set of desirable traits, it becomes necessary to use black box optimization tools on top of these generative models.

Bayesian optimization for molecular design. BO has been utilized extensively in recent years for molecular design problems, both over fixed pre-defined lists of existing molecules (Williams et al., 2015; Hernández-Lobato et al., 2017; Graff et al., 2021) and by utilizing the continuous latent spaces of variational autoencoders (Gómez-Bombarelli et al., 2018; Eissman et al., 2018; Tripp et al., 2020; Grosnit et al., 2021; Siivola et al., 2021; Jin et al., 2018). In latent space optimization, an encoder $\Phi(\mathbf{x})$ is used to map molecules \mathbf{x} to real-valued latent vectors \mathbf{z} . BO is then applied in the continuous latent space, and candidate latent vectors are decoded using the decoder $\Gamma(\mathbf{z})$ to generate candidate molecules. Maus et al. (2022) introduce an extension of TuRBO Eriksson et al. (2019) to the latent space optimization setting where the surrogate model and VAE are trained jointly using variational inference – in our molecular design results in section 4 we will make this same adaptation for our method.

Population generation algorithms. Although some general frameworks have been proposed to extend Bayesian optimization techniques to problems outside of optimization Neiswanger et al. (2021), to the best of our knowledge, this work is the first to consider the setting of finding multiple solutions under user specified diversity constraints using Bayesian optimization. We therefore compare to population generation algorithms designed to generate populations of solutions, some of which are designed for “quality diversity” Hansen (2016); Mouret and Clune (2015); Kent and Branke (2020); Vassiliades et al. (2016); Mouret and Maguire (2020); Gaier et al. (2018); Wessing and Preuss (2017). Most relevant is the CVT-MAP-Elites algorithm Vassiliades et al. (2016) which extends much of this work to high-dimensional optimization tasks by avoiding constructing exponential discretizations as in e.g. Mouret and Clune (2015); Kent and Branke (2020); Gaier et al. (2018). However, these algorithms measure diversity via distance in the search space, and are not straightforward to adapt to user-specified notions of diversity. In section 4 we show these approaches can fail to find diverse solutions for many optimization tasks when using semantically meaningful notions of diversity.

Trust Region Bayesian Optimization (TuRBO) Traditional approaches to BO are generally limited to low-dimensional problems with at most twenty tunable parameters Frazier (2018). Many methods tailored for high-dimensional BO are only suitable for small evaluation budgets and generally make strong assumptions on the underlying black-box objective function Kandasamy et al. (2015); Wang et al. (2016); Letham et al. (2020); Mutny and Krause (2018).

Eriksson et al. (2019) proposed TuRBO- M which maintains M local optimization runs, each with its own dataset \mathcal{D}_i and surrogate model. Each local optimizer i proposes can-

didates from within a hyper-rectangular *trust region* \mathcal{T}_i and a batch of candidates is selected from across all local optimizers in each iteration. Because acquisition is performed globally across all trust regions, local optimizers with the most promising evaluations of the objective receive a larger fraction of the evaluation budget. Each trust region \mathcal{T}_i is a rectangular subset of the input space \mathcal{X} centered at the best point found by the i th local optimizer—the *incumbent*— \mathbf{x}_i^+ and has a side-length $\ell_i \in [\ell_{min}, \ell_{max}]$. If a local optimizer improves upon its own incumbent ρ_{succ} times in a row, ℓ_i is increased to $\min(2\ell_i, \ell_{max})$. Similarly, when a local optimizer fails to make progress ρ_{fail} times in a row, the length ℓ_i is reduced to $\ell/2$. If $\ell_i < \ell_{min}$, that local optimizer is restarted. While TuRBO- M is not directly applicable to our problem setting, we will also use multiple trust regions with the same length adjustment dynamics. Additionally, while TuRBO- M keeps a separate data history for each trust region, other trust region methods such as MORBO Daulton et al. (2021) allow data sharing such that trust regions can be recentered on candidates from the data history of other trust regions.

3 METHODS

We consider the task of finding a diverse set of M solutions for some high-dimensional objective function $f(\cdot)$. For a given input \mathbf{x} , we can evaluate $f(\mathbf{x})$ to obtain a (possibly noisy) objective value y . To measure diversity, we use a symmetric, user-supplied function $\delta(\mathbf{x}_1, \mathbf{x}_2)$ defined on pairs of points in the search space \mathcal{X} . Formally, we seek a sequence $S^* := \{\mathbf{x}_1^*, \dots, \mathbf{x}_M^*\}$ so that:

$$\begin{aligned} \mathbf{x}_1^* &= \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \\ \mathbf{x}_i^* &= \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \text{ s.t. } \delta(\mathbf{x}_i^*, \mathbf{x}_j^*) \geq \tau \text{ for } j = 1, \dots, i-1 \end{aligned} \quad (1)$$

Under this formalization of the problem setting, the optima in S^* are defined *hierarchically*. In particular, we explicitly still wish to recover the best possible optimizer \mathbf{x}_1^* of the original objective. This choice of formalization is deliberate: our goal in this paper is to develop a method that still optimizes the given objective function $f(\cdot)$ as a practitioner may expect, but also produce alternative high quality solutions that are meaningfully different to the optimum as a by-product with as few additional evaluations as possible.

3.1 Rank-Ordered Bayesian Optimization with Trust Regions (ROBOT)

In this section, we propose ROBOT - an algorithm which extends Bayesian optimization to the problem setting above. We demonstrate the global consistency of our approach in subsection A.5. In order to find a set of M solutions, ROBOT maintains M simultaneous local optimization runs using M individual trust regions. As in prior work, trust

regions are defined as rectangular regions of the input space, e.g. $\mathcal{T}_i \subseteq \mathcal{X}$, with side lengths defined using standard Euclidean distance. We note that it would be interesting to explore the setting where trust region side lengths are instead defined with respect to the diversity measure $\delta(\cdot, \cdot)$; however, the problem of even sampling candidate \mathbf{x} locations from within such a trust region becomes challenging in the general setting.

Each local run i aims to find a single diverse solution \mathbf{x}_i^* , which together form the desired set S^* . In our problem definition, the solution \mathbf{x}_i^* is only constrained with respect to prior solutions, e.g. for $j < i$. Mirroring this, we assign a hierarchical rank-ordering to the M trust regions, $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_M$, so that the local optimization run \mathcal{T}_i responsible for finding \mathbf{x}_i^* is only constrained by local optimizers \mathcal{T}_j with $j < i$.

Acquiring candidates. Because the diversity constraints are non-stationary, they must be handled in an online fashion as we explore the input space. A natural way to accomplish this is to enforce diversity with respect to all candidates chosen by the optimization procedure. Mirroring the optimization problem in Equation 1, in each iteration of optimization we select candidates $\hat{\mathbf{x}}_i$ from each trust region \mathcal{T}_i to improve over its own incumbent (\mathbf{x}_i^+, y_i^+) using a similarly hierarchically-constrained acquisition function:

$$\begin{aligned} \hat{\mathbf{x}}_1 &= \arg \max_{\mathbf{x} \in \mathcal{T}_1} \alpha(\mathbf{x}; y_1^+) \\ \hat{\mathbf{x}}_i &= \arg \max_{\mathbf{x} \in \mathcal{T}_i} \alpha(\mathbf{x}; y_i^+) \text{ s.t. } \delta(\mathbf{x}, \hat{\mathbf{x}}_j) \geq \tau \quad \forall j < i \end{aligned} \quad (2)$$

Here, α may be a standard maximization acquisition function such as Expected Improvement (EI), and y_i^+ refers to the best objective value observed so far by the i th optimizer. By asymmetrically constraining candidates, we select diverse sets of candidates. Furthermore, because high ranking trust regions \mathcal{T}_i are less constrained they are unimpeded by lower-rank trust regions \mathcal{T}_j where $i < j$. As a consequence of this, the highest ranking trust region, \mathcal{T}_1 , is never impeded by any other trust region. For an illustration of this, see Figure 1. Since δ is an arbitrary user-defined black-box function, the above optimization problem is challenging. However, when the acquisition function maximization is approximated via a discretization of the input space \mathcal{X} —a relatively common approach—the above optimization remains straightforward. Nevertheless, a reliance on discretization motivates the usage of a modified Thompson sampling procedure that we describe below.

Thompson Sampling Another approach to acquisition is to use hierarchically constrained Thompson Sampling (TS). To accomplish this, we select a candidate $\hat{\mathbf{x}}_i$ for each trust region \mathcal{T}_i one at a time, in rank-order. To select a candidate $\hat{\mathbf{x}}_i$ from \mathcal{T}_i , we sample r points $C_i = \{\mathbf{c}_{i1}, \mathbf{c}_{i2}, \dots, \mathbf{c}_{ir}\}$ from \mathcal{T}_i . We then sample a realization $\hat{f}(\mathbf{c}_{ij})$ for each of these.

Denote by P_i the set of all candidate points which have already been selected from each of the higher-ranking trust regions $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_{i-1}$. We select a batch of candidates from among those points in C_i that are feasible with respect to all points in P_i :

$$\hat{F} = \{\mathbf{c} \mid \mathbf{c} \in C_i, \forall \mathbf{p} \in P_i \delta(\mathbf{c}, \mathbf{p}) \geq \tau\}$$

If \hat{F} is empty, then no point in the discretization C_i of the interior of \mathcal{T}_i was feasible, and we select no candidate from \mathcal{T}_i in this round. Because our experiments are run mostly on high dimensional settings where the trust regions are separated by relatively large distances, we found this to be an extremely rare occurrence empirically, happening only a handful of times across all experiments.

Trust region modifications. In each iteration, ROBOT recenters the trust regions such that the current set S_t^+ after iteration t approximating S^* is equivalent to the set of all centers $\{\mathbf{x}_1^+, \dots, \mathbf{x}_M^+\}$. When trust regions select new centers, all diversity constraints change, and since higher-ranking trust regions are unconstrained by subordinate ones, they may re-center on candidates that cause subordinate trust region incumbents to violate these new constraints. This would render some points in S_t^+ suddenly infeasible. To remedy this, in each step t , we greedily reconstruct the feasible set S_t^+ using the full set of data points D_t evaluated so far by all optimizers. In particular, we set $S_t^+ = \{\mathbf{x}'_1, \dots, \mathbf{x}'_M\}$, where:

$$\begin{aligned} \mathbf{x}'_1 &= \arg \max_{(\mathbf{x}, y) \in D_t} y \\ \mathbf{x}'_i &= \arg \max_{(\mathbf{x}, y) \in D_t} y \text{ s.t. } \forall j < i \delta(\mathbf{x}, \mathbf{x}'_j) \geq \tau \end{aligned} \quad (3)$$

We then re-center trust region \mathcal{T}_i on \mathbf{x}'_i . As a result, \mathcal{T}_1 is always centered on the best-scoring point found by any trust region, \mathcal{T}_2 is centered on the best remaining point which is sufficiently diverse from the new center of \mathcal{T}_1 , and so on. In addition to recentering, we note that optimizer i is only attempting to improve on its own current incumbent objective value (subject to its own diversity constraints), and not trying to improve over the globally best value observed so far (e.g. y_1^+). Therefore, trust region successes and failures as described in section 2 are defined with respect to each trust region’s own incumbent.

Global surrogate model The recentering procedure described above can recenter a trust region \mathcal{T}_i on any data point in the entire optimization history D_t , not merely its own local optimization history. This makes the use of local GP surrogate models for each trust region ill suited to the task, as \mathcal{T}_i may move to locations that were not acquired by its own local surrogate. ROBOT therefore instead maintains a single global surrogate model across all M trust regions. The benefit of this is isolated in Figure 5.




Figure 1: Diagram of three rank-ordered subordinate trust regions, Green \succ Blue \succ Purple. **(Left)** Each trust region generates b candidates. **(Middle)** Starting with \mathcal{T}_1 , we discard candidates in subordinate trust regions that violate diversity constraints with candidates in \mathcal{T}_1 . **(Right)** We repeat this procedure with \mathcal{T}_2 , removing infeasible candidates from \mathcal{T}_3 .

4 EXPERIMENTS

We apply ROBOT to five high-dimensional BO tasks for which finding a diverse set of solutions is desirable. We additionally optimize diverse S&P 500 investment portfolios in the appendix. Three of these tasks are continuous problems that enable direct application of ROBOT as described in section 3. The last three are structured molecule optimization tasks.

Implementation details and hyperparameters. We implement ROBOT leveraging BoTorch Balandat et al. (2020) and GPyTorch Gardner et al. (2018), with code available at <https://github.com/nataliemaus/robot>. All trust region hyperparameters are set to the TuRBO defaults as used in Eriksson et al. (2019). Particular choices of new task-specific parameters, M , τ , and $\delta(\cdot)$ are motivated in the corresponding section for each task. Since we consider large numbers of function evaluations for several tasks, we use an approximate GP surrogate model. In particular, we use a PPGPR for all tasks Jankowiak et al. (2020). The number of random points used to initialize optimization is kept consistent across all methods compared for each task and is included in the x-axis of all plots. We use Thompson sampling for all experiments. See subsection A.2 for additional implementation details.

Extending ROBOT to the structured BO setting. To extend ROBOT to the structured setting for the molecule optimization tasks, we use the pre-trained SELFIES VAE introduced by Maus et al. (2022) to map from the structured molecule space to a continuous search space where Bayesian optimization can be directly applied. Additionally, as in LOL-BO (Maus et al., 2022), we periodically train the surrogate model jointly with the VAE end-to-end in order to encourage the continuous latent space to be organized in a way that is more amenable to optimization.

We refer to this extension of ROBOT as LOL-ROBOT.

Plots. For each task, we plot the objective value of the current feasible set of M solutions obtained after a certain number of function evaluations. All plots with $M = 1$ show the objective value of the single best solution found, and are included to highlight the loss in efficiency incurred by all methods when seeking larger sets of solutions. For baselines such as Standard BO, TuRBO, and TuRBO- M which are designed to find a single solution rather than a diverse set, we plot the mean objective value of the best M feasible solutions found in the history of the run.

All plots are averaged over multiple runs and show standard errors. The expensive Guacamol experiments used 10 repetitions, while all others used 20. On many plots we include runs of ROBOT- k —i.e., our method seeking k solutions—for $k > M$. While one would in practice always run ROBOT- M when seeking M solutions, these results demonstrate the loss of efficiency of discovering M solutions when seeking more. For plots where we show figures with different M , we do not plot ROBOT- k for $k < M$.

Baselines. In all plots, we compare ROBOT against TuRBO, TuRBO- M , and an alternative diverse optimization algorithm involving sequential runs of constrained TuRBO (see Sequential SCBO description below). Since these algorithms are variants of TuRBO, each can be adapted to the latent space BO setting using an analogous version of LOL-BO.

We denote these latent space adaptations using notation: LOL-BO, LOL-BO- M , and Sequential LOL-SCBO. Applying end-to-end training to each baseline allows them to be directly compared to LOL-ROBOT for the three molecular tasks.

Note that when $M = 1$, TuRBO- M , and Sequential SCBO are the same algorithm so we only plot TuRBO.

Although baselines such as Standard BO, TuRBO, and TuRBO- M are designed to find a single solution rather than M diverse solutions, we compare to them in plots with

$M > 1$ by plotting the mean of the best M diverse solutions found by the method along the way.

Sequential SCBO Baseline As discussed in section 2, we can cast our problems as a constrained optimization problem if the solutions in S^* are generated sequentially—e.g., the constraints for \mathbf{x}_2^* are well defined given a fixed \mathbf{x}_1^* . To directly compare against this alternative, we run SCBO M times in a row, where the i th run of SCBO has diversity constraints against all solutions found from runs $j < i$. We additionally make several modifications to improve the efficiency of this algorithm. We start each sequential run from the best point observed on any previous run which meets the new set of constraints. Additionally, we maintain the same surrogate model across sequential runs rather than discard data.

4.1 Continuous BO Tasks

In this section, we consider two optimization tasks for which finding a diverse set of solutions is useful—optimizing the trajectory of a mars rover, and optimizing the policy used by a lunar landing device.

4.1.1 Rover

The rover trajectory optimization task consists of finding a 60 dimensional policy that allows a rover to move along some trajectory while avoiding a set of obstacles Wang et al. (2018). This optimization problem is useful as it allows us to directly visualize the diverse paths found by optimization.

Diversity function δ and threshold τ A meaningful diversity measure is one that requires the resulting trajectories to take distinct routes around the obstacles. To measure the distance between two trajectories, we use the one-way-distance (δ_{OWD}) metric from Su et al. (2020). The obstacles used in the rover environment are squares of side length 0.05 and all four sides of each obstacle are 0.1 from the side of some other obstacle (see Figure 2). Since our goal is to find diverse trajectories which take different routes around the obstacles, we therefore set $\tau = 0.15$.

Results. Results on this task for $M = 3$ trajectories are displayed in Figure 2. The leftmost figure depicts convergence speed of the top trajectory optimized only. ROBOT- M incurs no decrease in efficiency for finding the best trajectory for $M = 3$, despite also finding (middle left) three diverse trajectories of equivalent reward to the first. Although Standard BO fails to find a good single solution (leftmost figure), it outperforms TuRBO when finding $M > 1$ solutions (middle left figure). Standard BO is less myopic, and therefore finds a larger number of diverse trajectories with positive reward. CVT-MAP-Elites performs worse than all other baselines when we take the average of the top three diverse solutions found (middle left figure). Likely,

this is due to the usage of input space distance as a diversity measure, which does not necessarily correlate with the more semantically meaningful diversity measure δ . Trajectories found by multiple runs of TuRBO and a single run of ROBOT are depicted in the middle right and rightmost plots, clearly demonstrating diversity.

4.1.2 Learning Robust Lunar Lander Policies

The lunar lander optimization task seeks a control policy that enables an autonomous lander to land without crashing on a randomly generated set of terrain environments. The objective function is defined as the average reward of the policy obtained on a set of environments. We optimize the same controller as in (Eriksson et al., 2019). Although TuRBO finds policies that land on the training environments, we find that these policies sometimes crash when tested on unseen environments.

Diversity function δ and threshold τ . Since there is no obvious semantically meaningful measure of diversity between two policies for this task, we define δ to be the euclidean distance between two policies. We choose $\tau = 0.6$ since, for larger values of τ , the random set of 1024 policies used to initialize optimization often did not contain a sufficient number of feasible policies to start from.

Results. To demonstrate the value of this notion of diversity in this setting, we use ROBOT to find a diverse set of 20 policies S^* and then construct a single robust policy which simply takes the majority vote action of the diverse policies \mathbf{x}_i^* at every step. For comparison, we generate twenty policies by running TuRBO sequentially twenty times (requiring $20\times$ as many evaluations). In Figure 3 (right), we plot a histogram of rewards obtained by each of these strategies on 200 unseen environments. Without diversity constraints, the policies obtained by TuRBO occasionally achieve catastrophically low rewards. However, the ensembled 20 diverse policies never fail to land across this larger set of environments.

To demonstrate optimization efficiency, we plot function evaluations versus mean objective value found in Figure 3. We show results for optimizing a set of $M = 1$ and 20 feasible policies. Despite distributing its evaluation budget to find twenty diverse solutions, ROBOT incurs only a $3\times$ slowdown. Although CVT-MAP-Elites converges faster, ROBOT eventually obtains a higher mean objective value. This task is particularly well suited for CVT-MAP-Elites as it uses input space diversity measures. This task is one of the least successful for ROBOT, as discovering 20 policies requires roughly $6\times$ as many evaluations, where as most tasks in this section require significantly fewer additional evaluations. Nevertheless, this is still better than linear slowdown.

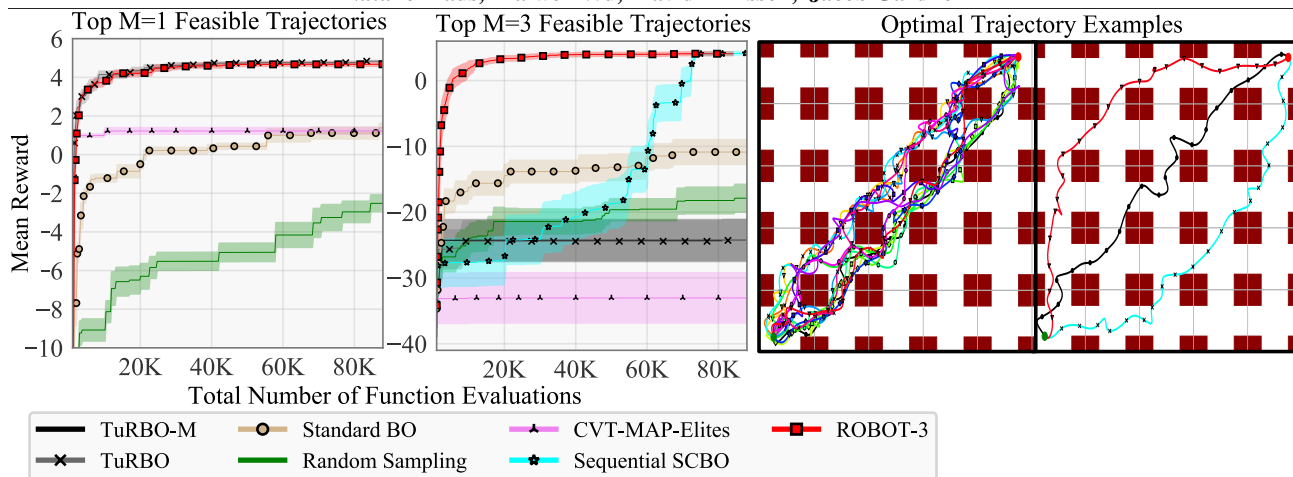


Figure 2: **(Left, Middle Left)** Rover optimization results where feasible trajectories have a minimum one-way-distance (OWD) of 0.15. Note that in the leftmost figure, TuRBO and ROBOT preform similarly so the curves overlap and are difficult to differentiate. Similarly for TuRBO and TuRBO-M in the middle left figure. **(Middle Right)** 15 optimized trajectories found by 15 individual runs of TuRBO (without diversity constraints). **(Right)** Three diverse trajectories found by ROBOT.




Figure 3: **(Left, Middle Left)** Lunar lander optimization, feasible policies have a minimum Euclidean distance of 0.6 in parameter space. Note that in the leftmost figure, CVT-MAP-ELITES and ROBOT-5 preform similarly so the curves overlap and are difficult to differentiate. **(Middle Right, Right)** Rewards obtained by majority vote ensembling 20 policies for 200 new environments. **(Middle Right)** Set of policies obtained by 20 independent runs of TuRBO. **(Right)** Set of policies obtained by a single run of ROBOT-20.

4.2 Molecular BO Tasks

The Guacamol benchmark suite Brown et al. (2019) contains scoring oracles for a variety of molecule design tasks, with scores ranging between 0 and 1. Of these tasks, we select two, Sitagliptin MPO and Valsartan SMARTS, for which high-scoring molecules found by LOL-BO tended to have high fingerprint similarity, making a search for diverse solutions particularly desirable. The task definitions are discussed in Brown et al. (2019). Because it is arguably the most commonly studied molecule optimization problem, we include Penalized Log P results in the appendix.

Additional baseline. In the molecule design setting, additional methods outside of BO have been proposed that specifically produce populations of solutions. While these populations are not specifically constrained to be diverse in any way, the Guacamol scoring procedure often scores the entire population generated. Therefore, we additionally compare to Graph GA Jensen (2019), one of the top performing methods on the Guacamol benchmark leaderboard.

Diversity function δ and threshold τ Fingerprint similarity (FPS) measures how similar two molecules are Brown et al. (2019). We therefore define the diversity function δ between two molecules to be the negative of their fingerprint




Figure 4: Comparing to baselines for finding solution sets of $M = 5, 20, 50, 100$ solutions to GuacaMol molecule design tasks. Tight constraint ($\tau = -0.4$) used for $M = 5$ and relaxed constraint ($\tau = -0.53$) used for $M = 20, 50, 100$. Note that in the top row figures, Random Sampling, CVT-MAP-ELITES, and Standard LS-BO all fail to make any progress so the curves overlap..

similarity. We evaluate finding solutions in two settings: one where we seek a small set of highly diverse solutions, and one where we seek a very large set of moderately diverse solutions. Since random pairs of molecules in the ZINC database Irwin et al. (2020) have FPS of ~ 0.4 on average, we use $\tau = -0.4$ for the highly diverse setting. For the moderately diverse setting, we relax the constraint to $\tau = -0.53$ and sets of up to 100 molecules.

Results In Figure 4 we depict optimization results for finding $M = 5, 20, 50, 100$ diverse solutions on the Valsartan SMARTS and Sitagliptin MPO optimization tasks. ROBOT consistently outperforms all other baselines. Notably, across all regimes, Graph-GA produces significantly fewer high scoring yet diverse molecules despite returning an entire population of solutions by default.

4.3 Ablation Studies.

In this section, we perform ablation studies evaluating the effect of K on the convergence of ROBOT, as well as the various components of ROBOT.

4.3.1 How Does K Affect the Convergence Rate of ROBOT- K ?

As ROBOT- K is asked to find more diverse solutions (i.e., as K increases), it is reasonable to expect that the convergence in the top $M < K$ solutions becomes slower. This is because a fixed evaluation budget must be split K ways for ROBOT- K . Thus, when seeking M solutions, the best strategy is to run ROBOT- M . However, ROBOT includes a number of features that try to mitigate the loss of efficiency as K increases by sharing information across the K solutions.

To study the overall impact of these features, in this section we run ROBOT to find M solutions, but run ROBOT- K with various $K \geq M$ (e.g., we ask ROBOT to find more than M solutions). We then evaluate the performance of these ROBOT- K runs on *only the top M solutions*. If the performance loss as K increases is small, we expect that ROBOT- K should be comparable in convergence to ROBOT- M , even when $K > M$.

Results of this study are in Figure 5 **Left, Middle Left**, where we ablate ROBOT with various M on Sitagliptin MPO. Even for $K = 50$, ROBOT incurs

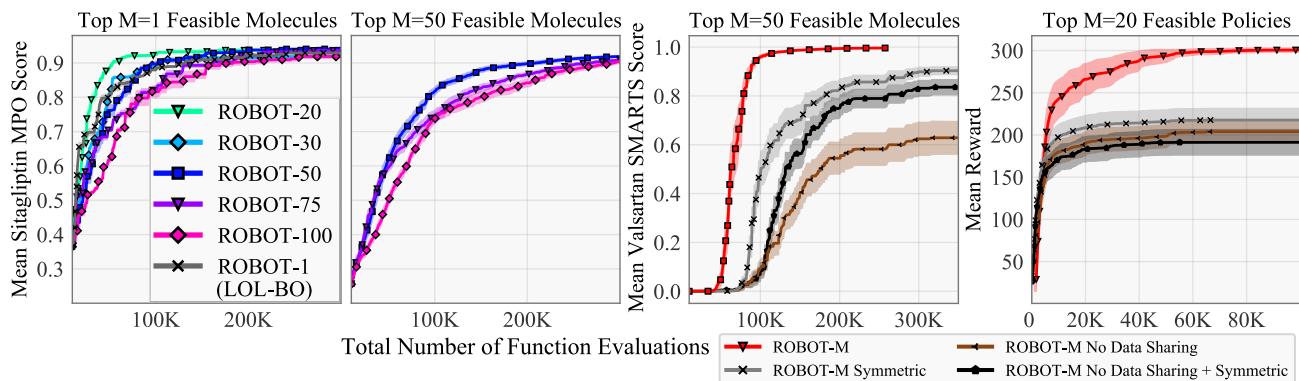


Figure 5: Ablations. **(Left, Middle Left)** Comparing convergence speed of ROBOT- K on Sitagliptin MPO with different K . ROBOT loses minimal efficiency up to $K = 50$, even when searching for $K > M$ molecules. Here we use the relaxed constraint $\tau = -0.53$. See appendix for similar ablation with the tighter constraint $\tau = -0.4$, and repeated ablation on Valsartan SMARTS. **(Middle Right, Right)** Comparing ROBOT with various components removed for Valsartan SMARTS and Lunar Lander tasks respectively.

negligible slowdown compared to finding a single solution ($M = 1$), despite the evaluation budget being split between 50 solutions. In Figure 11 we provide results on Valsartan SMARTS, and in Figure 12, we provide an ablation with $\tau = -0.4$, where we observe that ROBOT-5 converges at the same speed in terms of finding the single best 1 value, but is able to find 5 molecules with nearly the same score with negligible additional evaluations.

4.3.2 How Do the Different Components of ROBOT Affect Performance?

In Figure 5 **(Middle Right, Right)**, we evaluate the components of ROBOT.

To evaluate the gain in optimization performance achieved by using a trust region hierarchy to asymmetrically throw out candidates in acquisition, we compare to the simpler approach of symmetrically discarding all pairs of infeasible candidates. In particular, this involves discarding all pairs of candidates that do not satisfy the diversity constraints from across all trust regions regardless of the trust region hierarchy. We refer to this version of ROBOT as ROBOT Symmetric. Results in Figure 5 **(Middle Right, Right)** show that ROBOT consistently outperforms ROBOT Symmetric. This indicates that asymmetrically discarding candidates according to the trust region hierarchy is essential for good performance.

To evaluate the gain in optimization performance achieved by allowing collaboration (data sharing) between trust regions, we compare to a version of ROBOT without data sharing. In particular, for this version of ROBOT we do not allow trust regions to recenter on data points found other trust regions. Additionally, we use a separate surrogate model for each trust region rather than a single global surrogate model. In this case, each surrogate model is only

updated on the data found by its corresponding trust region. We refer to this version of ROBOT as ROBOT No Data Sharing. Results in Figure 5 **(Middle Right, Right)** show that ROBOT consistently outperforms ROBOT No Data Sharing. This indicates that using a global surrogate model and sharing data across the trust regions is essential for good performance.

Additionally, we compare to a version of our method where we both do not allow data sharing and symmetrically discard candidates. We refer to this version of our method as ROBOT No Data Sharing + Symmetric. Results in Figure 5 **(Middle Right, Right)** show that ROBOT consistently outperforms ROBOT No Data Sharing + Symmetric.

5 DISCUSSION

In real world settings, the objective function fed to the optimization routine often tells only part of the story. Practitioners often have preferences for solutions beyond simple objective value. By discovering solutions under *arbitrary user specified* diversity measures rather than input space measures alone, we believe this work may help make Bayesian optimization applicable in scenarios where an “all-or-nothing” approach may not be viable. In these scenarios, rather than being a final source of truth, the optimizer can be instead deployed as a tool to generate *suggestions* to a practitioner in a human-in-the-loop fashion, who may ultimately rely on domain expertise to choose the best solution.

References

M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In

- Advances in Neural Information Processing Systems* 33, 2020.
- S. Belakaria, A. Deshwal, and J. R. Doppa. Max-value entropy search for multi-objective bayesian optimization. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/82edc5c9e21035674d481640448049f3-Paper.pdf>.
- N. Brown, M. Fiscato, M. H. Segler, and A. C. Vaucher. Guacamol: Benchmarking models for de novo molecular design. *Journal of Chemical Information and Modeling*, 59(3):1096–1108, Mar 2019.
- S. Daulton, D. Eriksson, M. Balandat, and E. Bakshy. Multi-objective bayesian optimization over high-dimensional search spaces. *arXiv preprint arXiv:2109.10964*, 2021.
- S. Eissman, D. Levy, R. Shu, S. Bartzsch, and S. Ermon. Bayesian optimization and attribute adjustment. In *Proc. 34th Conference on Uncertainty in Artificial Intelligence*, 2018.
- D. Eriksson and M. Poloczek. Scalable constrained Bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 730–738. PMLR, 2021.
- D. Eriksson, M. Pearce, J. Gardner, R. D. Turner, and M. Poloczek. Scalable global optimization via local Bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- P. I. Frazier. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.
- A. Gaier, A. Asteroth, and J.-B. Mouret. Data-efficient design exploration through surrogate-assisted illumination. *Evolutionary computation*, 26(3):381–410, 2018.
- J. R. Gardner, M. J. Kusner, Z. E. Xu, K. Q. Weinberger, and J. P. Cunningham. Bayesian optimization with inequality constraints. In *ICML*, pages 937–945, 2014.
- J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson. Gpytorch: Blackbox matrix-matrix Gaussian process inference with gpu acceleration. *arXiv preprint arXiv:1809.11165*, 2018.
- R. Gómez-Bombarelli, J. N. Wei, D. Duvenaud, J. M. Hernández-Lobato, B. Sánchez-Lengeling, D. Sheberla, J. Aguilera-Iparraguirre, T. D. Hirzel, R. P. Adams, and A. Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2):268–276, 2018.
- D. E. Graff, E. I. Shakhnovich, and C. W. Coley. Accelerating high-throughput virtual screening through molecular pool-based active learning. *Chemical science*, 12(22):7866–7881, Apr. 2021.
- A. Grosnit, R. Tutunov, A. M. Maraval, R. Griffiths, A. I. Cowen-Rivers, L. Yang, L. Zhu, W. Lyu, Z. Chen, J. Wang, J. Peters, and H. Bou-Ammar. High-dimensional Bayesian optimisation with variational autoencoders and deep metric learning. *CoRR*, abs/2106.03609, 2021.
- N. Hansen. The cma evolution strategy: A tutorial, 2016. URL <https://arxiv.org/abs/1604.00772>.
- J. Hensman, N. Fusi, and N. D. Lawrence. Gaussian processes for big data. *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, 2013.
- J. M. Hernández-Lobato, M. A. Gelbart, R. P. Adams, M. W. Hoffman, and Z. Ghahramani. A general framework for constrained Bayesian optimization using information-based search. *The Journal of Machine Learning Research*, 17(1):5549–5601, 2016. Code available at: <https://github.com/HIPS/Spearmint/tree/PESC>. Last accessed on 02/03/2020.
- J. M. Hernández-Lobato, J. Requeima, E. O. Pyzer-Knapp, and A. Aspuru-Guzik. Parallel and distributed Thompson sampling for large-scale accelerated exploration of chemical space. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 1470–1479. PMLR, 2017.
- D. Hernández-Lobato, J. M. Hernández-Lobato, A. Shah, and R. P. Adams. Predictive entropy search for multi-objective bayesian optimization, 2015. URL <https://arxiv.org/abs/1511.05467>.
- J. J. Irwin, K. G. Tang, J. Young, C. Dandarchuluun, B. R. Wong, M. Khurelbaatar, Y. S. Moroz, J. Mayfield, and R. A. Sayle. Zinc20—a free ultralarge-scale chemical database for ligand discovery. *Journal of chemical information and modeling*, 60(12):6065–6073, 2020.
- M. Jankowiak, G. Pleiss, and J. R. Gardner. Parametric gaussian process regressors. In *Proceedings of the 37th International Conference on Machine Learning, ICML'20*. JMLR.org, 2020.
- J. Jensen. A graph-based genetic algorithm and generative model/Monte Carlo tree search for the exploration of chemical space. *chem sci* 10 (12): 3567–3572, 2019.
- W. Jin, R. Barzilay, and T. S. Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *International Conference on Machine Learning*. PMLR, 2018.
- D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.
- K. Kandasamy, J. Schneider, and B. Póczos. High dimensional Bayesian optimisation and bandits via additive

- models. In *International conference on machine learning*, pages 295–304. PMLR, 2015.
- P. Kent and J. Branke. Bop-elites, a bayesian optimization algorithm for quality-diversity search, 2020. URL <https://arxiv.org/abs/2005.04320>.
- M. Konakovic Lukovic, Y. Tian, and W. Matusik. Diversity-guided multi-objective bayesian optimization with batch evaluations. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 17708–17720. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/cd3109c63bf4323e6b987a5923becb96-Paper.pdf>.
- B. Letham, B. Karrer, G. Ottoni, and E. Bakshy. Constrained Bayesian optimization with noisy experiments. *Bayesian Analysis*, 14(2):495–519, 2019.
- B. Letham, R. Calandra, A. Rai, and E. Bakshy. Re-examining linear embeddings for high-dimensional Bayesian optimization. *Advances in neural information processing systems*, 33:1546–1558, 2020.
- N. Maus, H. T. Jones, J. S. Moore, M. J. Kusner, J. Bradshaw, and J. R. Gardner. Local latent space bayesian optimization over structured inputs. *arXiv preprint arXiv:2201.11872*, 2022.
- J. Moćkus. On bayesian methods for seeking the extremum. In *Optimization Techniques IFIP Technical Conference: Novosibirsk, July 1–7, 1974*, pages 400–404. Springer, 1975.
- J.-B. Mouret and J. Clune. Illuminating search spaces by mapping elites, 2015.
- J.-B. Mouret and G. Maguire. Quality diversity for multi-task optimization. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*. ACM, jun 2020.
- M. Mutny and A. Krause. Efficient high dimensional Bayesian optimization with additivity and quadrature Fourier features. In *Advances in Neural Information Processing Systems 31*, pages 9019–9030, 2018.
- D. M. Negoescu, P. I. Frazier, and W. B. Powell. The knowledge-gradient algorithm for sequencing experiments in drug discovery. *INFORMS Journal on Computing*, 23(3):346–363, 2011.
- W. Neiswanger, K. A. Wang, and S. Ermon. Bayesian algorithm execution: Estimating computable properties of black-box functions using mutual information. In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8005–8015. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/neiswanger21a.html>.
- C. E. Rasmussen. Gaussian processes in machine learning. In *Summer School on Machine Learning*, pages 63–71. Springer, 2003.
- R. G. Regis and C. A. Shoemaker. A stochastic radial basis function method for the global optimization of expensive functions. *INFORMS Journal on Computing*, 19(4):497–509, 2007.
- B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.
- E. Siivola, A. Paleyes, J. González, and A. Vehtari. Good practices for Bayesian optimization of high dimensional structured spaces. *Applied AI Letters*, 2(2):e24, 2021.
- J. Snoek, H. Larochelle, and R. P. Adams. Practical bayesian optimization of machine learning algorithms, 2012a. URL <https://arxiv.org/abs/1206.2944>.
- J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. *Advances in neural information processing systems*, 25, 2012b.
- H. Su, S. Liu, B. Zheng, X. Zhou, and K. Zheng. A survey of trajectory distance measures and performance evaluation. *The VLDB Journal*, 29:13–32, 2020.
- M. Titsias. Variational learning of inducing variables in sparse gaussian processes. In *Artificial intelligence and statistics*, pages 567–574. PMLR, 2009.
- A. Tripp, E. A. Daxberger, and J. M. Hernández-Lobato. Sample-efficient optimization in the latent space of deep generative models via weighted retraining. In *Advances in Neural Information Processing Systems 33*, 2020.
- M. Turchetta, A. Krause, and S. Trimpe. Robust model-free reinforcement learning with multi-objective bayesian optimization, 2019. URL <https://arxiv.org/abs/1910.13399>.
- R. Turner, D. Eriksson, M. McCourt, J. Kiili, E. Laaksonen, Z. Xu, and I. Guyon. Bayesian optimization is superior to random search for machine learning hyperparameter tuning: Analysis of the black-box optimization challenge 2020. In *NeurIPS 2020 Competition and Demonstration Track*, pages 3–26, 2021.
- V. Vassiliades, K. Chatzilygeroudis, and J.-B. Mouret. Using centroidal voronoi tessellations to scale up the multi-dimensional archive of phenotypic elites algorithm, 2016.
- Z. Wang, F. Hutter, M. Zoghi, D. Matheson, and N. De Freitas. Bayesian optimization in a billion dimensions via random embeddings. *J. Artif. Int. Res.*, 55(1):361–387, Jan. 2016.
- Z. Wang, C. Gehring, P. Kohli, and S. Jegelka. Batched large-scale Bayesian optimization in high-dimensional spaces. In *International Conference on Artificial Intelligence and Statistics*, volume 84, pages 745–754, 2018.

- S. Wessing and M. Preuss. The true destination of ego is multi-local optimization, 2017. URL <https://arxiv.org/abs/1704.05724>.
- K. Williams, E. Bilsland, A. Sparkes, W. Aubrey, M. Young, L. N. Soldatova, K. De Grave, J. Ramon, M. de Clare, W. Sirawaraporn, S. G. Oliver, and R. D. King. Cheaper faster drug development validated by the repositioning of drugs against neglected tropical diseases. *Journal of the Royal Society, Interface / the Royal Society*, 12(104): 20141289, Mar. 2015.
- Y. Xie, Z. Xu, J. Ma, and Q. Mei. How much space has been explored? measuring the chemical space covered by databases and machine-generated molecules, 2021. URL <https://arxiv.org/abs/2112.12542>.

Discovering Many Diverse Solutions with Bayesian Optimization Supplementary Materials

A APPENDIX

A.1 ROBOT Initialization Details

ROBOT is initialized with a small set of quasi-random data. The M trust regions are initialized one at a time (in rank-order). Each trust region is centered on the highest-scoring point in the initial data which is sufficiently diverse from the centers of all higher-ranking trust regions. See Table 1 for the number of initialization data points used for each task. All baseline methods use the same number of initialization data points for each task.

Table 1: Number of random initialization data points used to initialize all methods for each task in section 4

TASK	NUMBER OF POINTS
Rover	1024
Lunar Lander	1024
Stock Portfolio Diversification	1024
Molecule Tasks	10,000

A.2 Additional Implementation Details

In this section, we provide additional implementation details for ROBOT and baseline methods.

A.2.1 Surrogate Model

As discussed in section 4, we use a PPGPR Jankowiak et al. (2020) surrogate model. To maintain fair comparison, we use the same surrogate model with the same configuration for ROBOT and all baseline Bayesian optimization methods. We use a PPGPR with a constant mean and standard RBF kernel. Due to the high dimensionality of our chosen tasks, use a deep kernel (several fully connected layers between the search space and the GP kernel). In particular, we use two fully connected layers with 32 nodes each. We update the parameters of the PPGPR during optimization by training it on collected data using the Adam optimizer with a learning rate of 0.001. The PPGPR is initially trained on a small set of initialization data for 20 epochs. For the number of initialization data points used for each task, see Table 1. On each step of optimization, the model is updated on the newly collected data for 2 epochs. This is kept consistent across all Bayesian optimization methods.

A.2.2 Details for Non-BO Methods

For the CVT-Map-Elits method, we use all default hyper-parameters from Vassiliades et al. (2016) with the "Number of Niches" parameter set equal to M (the number of diverse solutions we want to find). For the Graph GA method we use the implementation provided by Brown et al. (2019) which returns a population of molecules by default. We use the default hyper-parameters from Jensen (2019). For random sampling, we select query points in the search space uniformly at random, and check to make sure that we never select the same point twice.

A.3 Additional Experiments

In this section, we provide some additional experimental results, including analysis for ROBOT on the two additional optimization tasks mentioned in section 4.

A.3.1 Stock Portfolio Diversification

In portfolio optimization, the goal is to find a 500-dimensional weight vector \mathbf{x} giving the optimal fraction of a principal that should be invested in each company in the S&P 500 in order to maximize return while minimizing volatility. We quantify the return/volatility trade-off by directly maximizing the Sharpe ratio of the portfolio, which we compute using the past three years of data from the S&P 500. For our toy setting, we make the simplifying assumption that risk-free return is always zero, such that the Sharpe ratio is defined by the equation: $\frac{ROI}{\sigma\sqrt{252}}$. Here, the ROI is the total return on the investment, σ is the total volatility (the standard deviation of day-to-day changes in return) and $\sqrt{252}$ is a standard normalizing constant which accounts for the number of trading days in a year.

Diversity function δ and threshold τ We define the diversity function δ between two portfolios to be the maximum integer k such that the two portfolios have no stocks in common in their top k invested stocks. We require feasible portfolios to have $\tau = 10$ different companies in the top 10 – that is, the top 10 stocks must be disjoint.




Figure 6: S&P 500 optimization task. Feasible portfolios have a minimum diversity of $\tau = 10$.

Results. Results on this task for $M = 1$ and $M = 3$ trajectories are displayed in Figure 6. ROBOT-3 converges at nearly the same speed in terms of finding the single best 1 portfolio, and is simultaneously able to find 3 portfolios with nearly the same Sharpe ratio with negligible additional evaluations.

A.3.2 Penalized Log P Optimization

In this section, we provide results for a third molecular optimization task - Penalized Log P. As in subsection 4.2, we use negative fingerprint similarity for our diversity function δ and consider the two settings with tight and relaxed constraints ($\tau = -0.4$ and $\tau = -0.53$ respectively). Since Maus et al. (2022) showed that the Penalized Log P oracle can be exploited to produce high-scoring molecules which are wholly unrealistic, we constrain the decoder of the VAE to producing SELFIES strings of 400 tokens or fewer. With this added constraint, LOL-BO achieves Penalized Log P scores of ~ 100 rather than ~ 500 Figure 7. Since the Penalized Log P scores reported by Jensen (2019) are not competitive with ~ 100 (Jensen (2019) reports maximum Penalized Log P scores of ~ 12), we do not compare to Graph-GA for this task.

Results with smaller M and tighter constraints. In Figure 7 we depict optimization results for finding $M = 1$ and $M = 5$ diverse solutions on the Penalized Log P optimization task. In the $M = 1$ case (left panel in Figure 7), ROBOT-3, ROBOT-5, and ROBOT-10 surprisingly appear to find the best molecule a bit faster than LOL-BO despite simultaneously searching for other diverse solutions. A possible explanation for this is that LOL-BO uses only one trust region and is therefore more myopic than ROBOT- M . It appears that when M is small, the additional exploration done by ROBOT- M can actually allow for a single best solution to be found more quickly. Furthermore, despite searching for 5 diverse molecules, ROBOT-5 is simultaneously able to find 5 molecules with nearly the same score with negligible additional evaluations (right panel in Figure 7).




Figure 7: Penalized Log P molecular optimization task. Feasible molecules have a maximum fingerprint similarity of 0.4.

Results with larger M and relaxed constraints. In Figure 8 we depict optimization results for finding larger sets of $M = 1, 20,$ and 50 diverse molecules. Even up to asking ROBOT to find 50 solutions, we incur very little slowdown compared to finding a single solution. Furthermore, ROBOT is able to find a full set of 20 and 50 high-scoring molecules with a small number of additional evaluations.




Figure 8: Penalized Log P molecular optimization task. Feasible molecules have a maximum fingerprint similarity of 0.53.

A.3.3 Fingerprint Similarity Thresholds

In Figure 9, we provide visualization of the distribution of fingerprint similarities between molecules in the Zinc20 data-set. This distribution informed our choices $\tau = -0.4$ and $\tau = -0.53$ for molecular optimization tasks.

A.3.4 Additional Ablations

In Figure 10 we plot the best objective value found by individual rank-ordered trust regions during optimization runs of ROBOT- M on Valsartan SMARTS and Lunar Lander tasks. We observe that higher-ranking trust regions are able to converge more quickly since they are unimpeded by lower-ranking trust regions.

In Figure 5 (Left, Middle Left) we ablate ROBOT with various M on Sitagliptin MPO with $\tau = -0.53$. In this section, we also provide the same ablation on Valsartan SMARTS with $\tau = -0.53$ (see Figure 11), and on both Sitagliptin MPO and Valsartan SMARTS with $\tau = -0.4$ (see Figure 12).

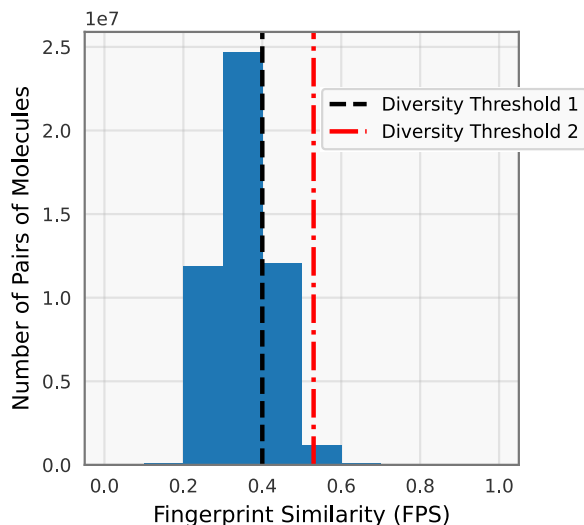


Figure 9: Fingerprint similarities between all pairs of 10,000 randomly selected molecules from the Zinc20 data-set. Our chosen diversity constraint thresholds of $-\tau = 0.4$ (threshold 1) $-\tau = 0.53$ (threshold 2) are shown with vertical lines.

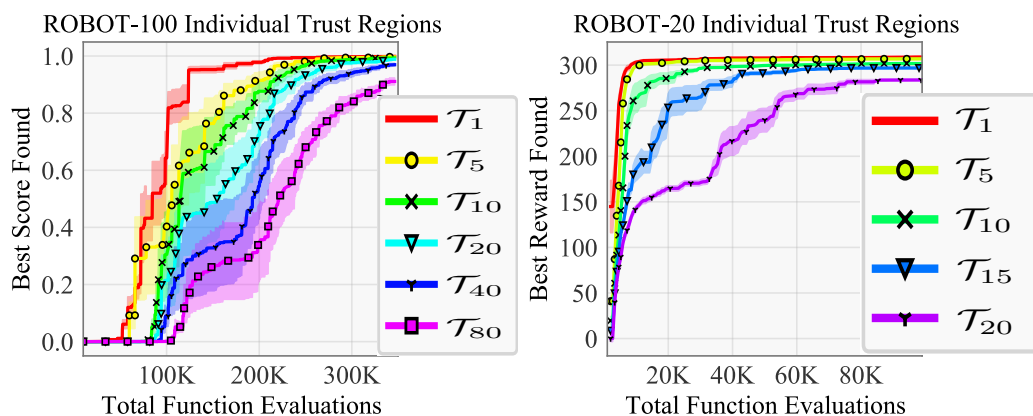


Figure 10: Objective found by individual trust regions during optimization with ROBOT. Higher-rank trust regions converge faster than lower-ranking trust regions. **(Left)** Optimization of Valsartan SMARTS with ROBOT-100, $\tau = -0.53$. **(Right)** Optimization of Lunar Lander with ROBOT-20, $\tau = 0.6$.

A.4 Run-Time Considerations

Our method scales to millions of evaluations by leveraging minibatch training made possible by the variational GP approximation used by Jankowiak et al. (2020). The total running times for the experiments in this paper is around half a day. The one exception to this is the stock portfolio optimization task (subsubsection A.3.1) - since this task requires more than 10x as many function evaluations to converge, ROBOT takes roughly one week to run on this task.




Figure 11: Valsartan SMARTS molecular optimization task. Feasible molecules have a maximum fingerprint similarity of 0.53. Ablating ROBOT with various M .

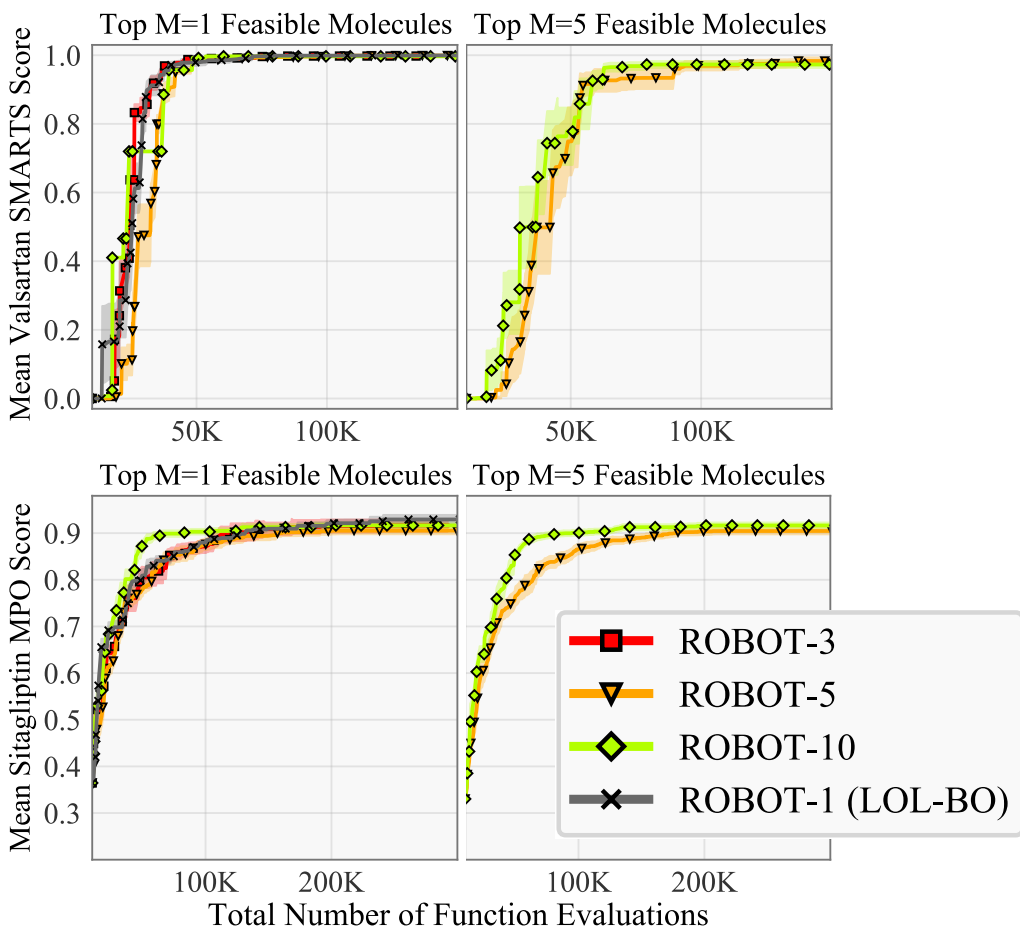


Figure 12: Valsartan SMARTS and Sitagliptin MPO molecular optimization tasks. Feasible molecules have a maximum fingerprint similarity of 0.4. Ablating ROBOT with various M .

A.5 Global Consistency of ROBOT

In this section we prove that ROBOT converges to the optimal set $S^* = \{\mathbf{x}_1^*, \dots, \mathbf{x}_M^*\}$ as the number of samples tends to infinity. Note that assumptions 4 and 5 correspond to trivial modifications of ROBOT itself, rather than assumptions

about the underlying objective. Furthermore, assumption three is a fairly common setting in Bayesian optimization. The primary non-trivial assumption needed in the Theorem below is assumption 2. This assumption requires a certain degree of smoothness from the objective function and the diversity constraint near points in S^* . In particular, we assume points in an epsilon ball centered at \mathbf{x}_1^* , $B(\mathbf{x}_1^*, \epsilon)$, have higher objective value than points outside the ball. Furthermore, points in $B(\mathbf{x}_2^*, \epsilon)$ have higher objective value than other points in the input domain \mathcal{X} that are *feasible* with respect to points in $B(\mathbf{x}_1^*, \epsilon)$. Note that assumption 2 is also implied under assumption 1 if both $f(\cdot)$ and $\delta(\cdot)$ are continuous and $\delta(\mathbf{x}_1^*, \mathbf{x}_2^*) > \tau$.

We start by showing the result for $M = 2$:

Theorem 1 ($M=2$). *Suppose that ROBOT with default parameters is used under the following assumptions:*

1. $S^* = \{\mathbf{x}_1^*, \mathbf{x}_2^*\}$ exists and is non-empty. In particular, for each i , \mathbf{x}_i^* is the unique optimizer to the optimization problem described in section 3 by Equation 1.
2. There exists an $\epsilon > 0$ such that:
 - (a) For any $\mathbf{w}_1 \in B(\mathbf{x}_1^*, \epsilon) \cap \mathcal{X}$ we have $f(\mathbf{w}_1) > f(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X} \setminus B(\mathbf{x}_1^*, \epsilon)$.
 - (b) For any $\mathbf{w}_1 \in B(\mathbf{x}_1^*, \epsilon) \cap \mathcal{X}$, $\mathbf{w}_2 \in B(\mathbf{x}_2^*, \epsilon) \cap \mathcal{X}$ we have $\delta(\mathbf{w}_1, \mathbf{w}_2) \geq \tau$.
 - (c) For any $\mathbf{w}_1 \in B(\mathbf{x}_1^*, \epsilon) \cap \mathcal{X}$, $\mathbf{w}_2 \in B(\mathbf{x}_2^*, \epsilon) \cap \mathcal{X}$ we have $f(\mathbf{w}_2) > f(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X} \setminus B(\mathbf{x}_2^*, \epsilon)$ s.t. $\delta(\mathbf{x}, \mathbf{w}_1) \geq \tau$.
3. The objective f is bounded and the input domain \mathcal{X} is a compact hypercube.
4. ROBOT generates new initial points when a trust region restarts. These initial points are chosen such that for any $\delta > 0$ and $x \in \mathcal{X}$ there exists $\nu(x, \delta) > 0$ such that the probability that at least one point ends up in a ball centered at x with radius δ is at least $\nu(x, \delta)$.
5. ROBOT considers any sampled point an improvement only if it improves the current best solution by at least some constant $\gamma > 0$.

Then, ROBOT converges to the unique global minimizing set S^* .

Proof. We start by observing that each trust region will collect only a finite number of samples before restarting due to conditions (3) and (5) as well as the fact that we similarly to SCBO use a finite success and failure tolerance. This means that each trust region in ROBOT will restart infinitely often with a fresh trust region and hence there is an infinite subsequence of initial points that satisfy (4). Now, denote the infinite sequence of initial points collected by ROBOT by $\{\hat{\mathbf{x}}(t)\}_{t \geq 1}$ where t is the iteration number. We will construct a new sequence $\{\hat{\mathbf{x}}_1(t)\}_{t \geq 1}$ as follows:

$$\hat{\mathbf{x}}_1(t) := \begin{cases} \hat{\mathbf{x}}(1) & \text{if } t = 1 \\ \hat{\mathbf{x}}(t) & \text{if } f(\hat{\mathbf{x}}(t)) > f(\hat{\mathbf{x}}_1(t-1)) \\ \hat{\mathbf{x}}_1(t-1) & \text{otherwise} \end{cases}$$

It now follows directly from the arguments made by Regis and Shoemaker (2007) in Theorem 1 and assumption (2) that $\hat{\mathbf{x}}_1(t) \rightarrow \mathbf{x}_1^*$ almost surely. Next, assumptions (2) and (4) allows us to find a point $\hat{\mathbf{z}}_1 \in \{\hat{\mathbf{x}}_1(t)\}_{t \geq 1}$ that is arbitrarily close to \mathbf{x}_1^* that is both better than any point outside $B(\mathbf{x}_1^*, \epsilon)$ and also satisfies the diversity constraint w.r.t. any point $\hat{\mathbf{z}}_2 \in B(\mathbf{x}_2^*, \epsilon) \cap \mathcal{X}$. We can then construct the following subsequence

$$\hat{\mathbf{x}}_2(t) := \begin{cases} \hat{\mathbf{x}}(1) & \text{if } t = 1 \\ \hat{\mathbf{x}}(t) & \text{if } \delta(\hat{\mathbf{x}}(t), \hat{\mathbf{z}}_1) \geq \tau \text{ and } \delta(\hat{\mathbf{x}}_2(t-1), \hat{\mathbf{z}}_1) < \tau, \\ \hat{\mathbf{x}}(t) & \text{if } f(\hat{\mathbf{x}}(t)) > f(\hat{\mathbf{x}}_2(t-1)) \text{ and } \delta(\hat{\mathbf{x}}(t), \hat{\mathbf{z}}_1) \geq \tau \text{ and } \delta(\hat{\mathbf{x}}_2(t-1), \hat{\mathbf{z}}_1) \geq \tau, \\ \hat{\mathbf{x}}_2(t-1) & \text{otherwise} \end{cases}$$

Following the same argument as before we have that $\hat{\mathbf{x}}_2(t) \rightarrow \mathbf{x}_2^*$ almost surely. □

We can extend these ideas to cover any finite M by extending these assumptions.

Corollary 1.1. *Assume 3-5 from Theorem 1 are satisfied. In addition, we also assume that the following is true:*

1. $S^* = \{\mathbf{x}_1^*, \dots, \mathbf{x}_M^*\}$ and is non-empty, where $M > 2$ is finite.

2. There exists an $\epsilon > 0$ such that:

(a) For any $\mathbf{w}_1 \in B(\mathbf{x}_1^*, \epsilon) \cap \mathcal{X}$ we have $f(\mathbf{w}_1) > f(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X} \setminus B(\mathbf{x}_1^*, \epsilon)$.

(b) The following holds for $j = 2, \dots, M$: For any $\mathbf{w}_1 \in B(\mathbf{x}_1^*, \epsilon) \cap \mathcal{X}, \dots, \mathbf{w}_j \in B(\mathbf{x}_j^*, \epsilon) \cap \mathcal{X}$ we have $\min_{i \in \{1, \dots, j-1\}} \delta(\mathbf{w}_i, \mathbf{w}_j) \geq \tau$.

(c) The following holds for $j = 2, \dots, M$: For any $\mathbf{w}_1 \in B(\mathbf{x}_1^*, \epsilon) \cap \mathcal{X}, \dots, \mathbf{w}_j \in B(\mathbf{x}_j^*, \epsilon) \cap \mathcal{X}$ we have $f(\mathbf{w}_j) > f(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X} \setminus B(\mathbf{x}_j^*, \epsilon)$ s.t. $\min_{i \in \{1, \dots, j-1\}} \delta(\mathbf{x}, \mathbf{w}_i) \geq \tau$.

Then ROBOT converges to the unique set S^* .

A.6 Limitations and Assumptions of ROBOT

Similar to any Bayesian optimization method, ROBOT assumes that the probabilistic surrogate model can obtain a good fit to the black box function. Additionally, ROBOT assumes that the diversity function δ is cheap to compute relative to the black box function.

A.6.1 Selection of τ

ROBOT assumes that the diversity constraints are mild enough and M is small enough that finding M diverse points in the search space is possible. We suggest choosing a τ value which represents the most relaxed constraint possible that still enforces enough diversity to be meaningful to the practitioner. We assume that the practitioner can evaluate the diversity function to determine what values indicate sufficient diversity for their particular application.

A.7 Extending TuRBO to ROBOT

In this section, we describe the modifications necessary to modify an implementation of TuRBO into an implementation of ROBOT. Note that a full BoTorch implementation of ROBOT is also available, as linked to in the experimental results section of the main text. We give a pseudocode algorithm for ROBOT (1). Lines of the algorithm in blue show the parts of ROBOT that differ from TuRBO- M . We provide additional clarification on the differences highlighted between ROBOT and TuRBO- M here. Both methods use M hyper-rectangular trust regions and dynamically update the length and center of each trust region using the dynamics of Eriksson et al. (2019). However, ROBOT performs acquisition to find a set of M diverse solutions according to a diversity function δ and threshold τ , with each trust region responsible for one solution, while TuRBO- M finds a single solution from across the trust regions. Accordingly, TuRBO- M allocates budget related to the strength of each trust region’s incumbent, and some trust regions may see very few allocations if they start in poor regions of the search space.

In contrast, ROBOT selects the same number of candidates from each trust region on each step of optimization, maintains a rank-ordering of the M trust regions and re-centers the M trust regions in rank-order on best diverse set of points from the entire shared data history $S_t^+ = \{\mathbf{x}_1^{(t)}, \dots, \mathbf{x}_M^{(t)}\}$. TuRBO- M re-centers each trust region on the best candidate queried by that individual trust region. Finally, ROBOT discards candidates from lower-ranking trust regions if they are not sufficiently diverse from the candidates from higher-ranking trust regions while TuRBO- M does not discard candidates.

Algorithm 1 ROBOT Algorithm

Inputs: f, D_0, M, δ, τ

- 1: **for** $i = 1, \dots, M$ **do**
- 2: Initialize trust region \mathcal{T}_i
- 3: **end for**
- 4: **for** Every time step t **do**
- 5: Update surrogate model on D_t
- 6: $S_t^+ = \{\mathbf{x}_1^{(t)}, \dots, \mathbf{x}_M^{(t)}\}$, where:

$$\mathbf{x}_1^{(t)} = \arg \max_{(\mathbf{x}, y) \in D_t} y$$

$$\mathbf{x}_i^{(t)} = \arg \max_{(\mathbf{x}, y) \in D_t} y \text{ s.t. } \forall j < i \delta(\mathbf{x}, \mathbf{x}_j^{(t)}) \geq \tau \quad (4)$$

- 7: Set center of trust region \mathcal{T}_i to $\mathbf{x}_i^{(t)}$
 - 8: **for** $i = 1, \dots, M$ **do**
 - 9: Select candidate x_i from \mathcal{T}_i using acquisition function
 - 10: **end for**
 - 11: **for** $i = 1, \dots, M$ **do**
 - 12: **if** $\delta(x_i, x_j) < \tau$ for any $j < i$ **then**
 - 13: Discard x_i
 - 14: **else**
 - 15: $y_i = f(x_i)$
 - 16: Update length of \mathcal{T}_i based on y_i
 - 17: **end if**
 - 18: **end for**
 - 19: $D_{t+1} \leftarrow D_t \cup (\mathbf{x}, \mathbf{y})$
 - 20: **end for**
-