Batched Energy-Entropy acquisition for Bayesian Optimization

Felix Teufel¹² Carsten Stahlhut¹ Jesper Ferkinghoff-Borg¹

¹Machine Intelligence, Novo Nordisk A/S ²Department of Biology, University of Copenhagen {fegt,ctqs,jfgb}@novonordisk.com

Abstract

Bayesian optimization (BO) is an attractive machine learning framework for performing sample-efficient global optimization of black-box functions. The optimization process is guided by an acquisition function that selects points to acquire in each round of BO. In batched BO, when multiple points are acquired in parallel, commonly used acquisition functions are often high-dimensional and intractable, leading to the use of sampling-based alternatives. We propose a statistical physics inspired acquisition function for BO with Gaussian processes that can natively handle batches. Batched Energy-Entropy acquisition for BO (BEEBO) enables tight control of the explore-exploit trade-off of the optimization process and generalizes to heteroskedastic black-box problems. We demonstrate the applicability of BEEBO on a range of problems, showing competitive performance to existing methods.

1 Introduction

Bayesian Optimization (BO) has since its inception [1, 2] made a profound contribution to the realm of global optimization of black-box functions through the usage of Bayesian statistics. For global optimization problems pursuing $x_* = \operatorname{argmax}_{x \in \mathcal{X}} f_{\text{true}}(x)$, BO has surfaced as a premier strategy for efficiently handling especially complex and costly unknown functions, $f_{\text{true}}(x)$. While BO is traditionally formulated in a single-point scenario, where individual points are queried and results are observed sequentially, there are situations where batched acquisition is needed. Such situations arise when $f_{true}(x)$ is expensive to evaluate in either time or cost, but can be effectively evaluated in parallel by dispatching multiple experiments, reducing the overall optimization time. This is often the case in e.g. drug discovery, materials design or hyperparameter tuning for deep models [3, 4, 5, 6, 7].

The realization that BO could be employed for the training of deep neural networks, as suggested by [3], sparked renewed research interest, with advancements encompassing a variety of areas, including the generalization to accommo-

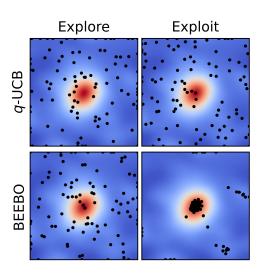


Figure 1: q-UCB does not allow for controlling its explore-exploit trade-off with large batches. A GP surrogate (background) was initialized with 100 random points of the Ackley function. q-UCB was run with $\kappa=0.1$ and $\kappa=100$, BEEBO with T'=0.05 and T'=0.05 Batch size Q=100.

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

date noisy inputs [8, 9], heteroskedastic noise [10, 11], multi-task problems [12], multi-fidelity [13], high-dimensional input spaces [14], and parallel methods with batch queries [15, 16]. Generally, these desired properties are addressed by customizing one of the two key components in BO, either the *surrogate model* or the *acquisition function*. The surrogate model f approximates the black-box function f_{true} using the available data. In BO, the surrogate is formulated from a Bayesian perspective, allowing us to quantify the model's uncertainty when evaluating new points. Typically, the model of choice is a Gaussian Process (GP) [17]. The acquisition function is responsible for guiding the selection of new input point(s) to evaluate at each optimization step, utilizing the surrogate model to identify promising regions in the input domain and exploring the unknown function further.

Any acquisition process needs to trade off exploration (reducing uncertainty to learn a better surrogate model) against exploitation (selecting points with a high expected $f_{\rm true}(x)$ based on the current surrogate). In this work, we are particularly interested in acquisition processes that make this trade-off controllable using a hyperparameter. Controllability can be a desirable property if e.g. domain knowledge relating to the difficulty of the optimization process and the quality of the surrogate model is available, or if the strategy needs to be adjusted depending on future experimental budgets. Similarly, it can be desirable to acquire multiple x with high $f_{\rm true}(x)$ in a batch (as opposed to just finding the optimum x_* , with the remaining x being considered explorative). This is useful when optima identified in BO can be subject to constraints that are unknown at optimization time, but may render x_* intractable [18]. Such constraints arise when the $f_{\rm true}$ explored in BO is a necessary simplification of the actual objective. Practical examples include e.g. the synthesizability of a material at larger scale, when BO experiments are performed at lab scale; or the $in\ vivo$ activity of a molecule with BO experiments performed $in\ vitro$.

A wide range of batch mode acquisition functions has been proposed, with approaches often leveraging random sampling strategies or Monte Carlo (MC) integration, which can adversely affect controllability for large batches (Figure 1). In contrast, we here introduce BEEBO (Batched Energy-Entropy acquisition for BO), a statistical physics inspired acquisition function for BO with GP surrogate models that natively generalizes to batched acquisition. BEEBO enables

- Parallel gradient-based optimization of the inputs, without requiring sampling or Monte Carlo integrals.
- Tight control of the explore-exploit trade-off in batch mode using a single temperature hyperparameter.
- Risk-averse BO under heteroskedastic noise.

We demonstrate the application of BEEBO on a wide range of test problems, and investigate its behaviour under heteroskedastic noise.

2 Related works

Batch variants of traditional strategies Parallel acquisition in BO has seen a variety of approaches, often starting from established single-point acquisition functions like probability of improvement (PI), expected improvement (EI), knowledge gradient (KG) or upper confidence bound (UCB) [2, 19, 20, 21, 22]. Reformulating these to batch mode with Q query points, we obtain q-PI, q-EI, and q-UCB [23, 24]. While the single-point specifications provide an analytical form and enable gradient-based optimization, batch expressions are more challenging and require different optimization strategies, typically involving greedy algorithms [25] or deriving an integral expression over multiple points.

For instance, in the popular EI acquisition function, a single point is selected by maximizing the expression $a_{\rm EI}(x) = \mathbb{E}[\max{(0,f(x)-f_t^*)}] = \int \max{(0,f(x)-f_t^*)} P(f|x) df$. Here f_t^* represents the best observed evaluation of $f_{\rm true}$ so far. With a surrogate model in the form of a GP, the acquisition function depends only on the predictive mean and variance functions, $\mu(x)$ and C(x). Effectively, we need to evaluate the cumulative normal distribution, which quickly becomes intractable for large batch sizes and approximating the gradient of the q-EI acquisition function typically requires MC estimation [26, 27]. However, proper MC integration can be laborious and is sensitive to both the dimension of the problem and the choice of batch size Q. Specifically, MC methods face the *curse of dimensionality* problem when applied to high-dimensional integrals, as they require an exponentially increasing number of sample points to maintain accuracy, making them

computationally impractical for such tasks [28, 29]. Of particular interest is Wilson et al. [24], in which they adopt the reparameterization trick [30, 31] on acquisition functions integrals, enabling gradient based approaches to the optimization of PI, EI, and UCB. This demonstrates particular usefulness in modest to higher dimensions.

While EI trades off exploration and exploitation, users do not have a direct control over the balance. To alleviate this, Sobester et al. [32] proposed a weighted EI formulation. An alternative strategy with an explicit explore-exploit trade-off is offered by the UCB acquisition function, $a_{\text{UCB}}(x) = \mu(x) + \sqrt{\kappa} \cdot \sqrt{C(x)}$, which directly expresses exploration and exploitation as two terms, traded off by the parameter κ . As we are particularly interested in enabling this direct user control, we focus our primary comparison on q-UCB in the main text, while a more extensive comparison with alternative methods can be found in Appendix B, both theoretically and experimentally.

Greedy strategies As mentioned, a popular approach for leveraging single-point acquisition functions is devising batch filling strategies that score candidate points sequentially. Kriging Believer (KB) [33] uses EI to select points and iteratively updates the GP by fantasizing an observation with the posterior mean. Likewise, GP-BUCB [34] uses fantasized observations to update $\sqrt{C(x)}$ at each step. Local penalization (LP) [35] introduces a penalization function that repulses selection away from already selected points. Contal et al. [36] propose selecting a single point using UCB and dedicating the remainder of the batch budget for exploration in a restricted region around the believed optimum. GLASSES [37] treats batch selection as a multi-step lookahead problem to overcome the myopia of only considering the immediate effect of selecting a point.

Entropy based strategies From an information theory perspective, BO can be interpreted as seeking to reduce uncertainty over the location of optima of the unknown function. This has given rise to entropy-based acquisition functions such as entropy search (ES) [38], predictive entropy search (PES) [39] and max-value entropy search (MES) [13, 40, 41]. MES is distinct in that it seeks to quantify the mutual information between the unknown $f_{\text{true}}(x_*)$ and the observations y|D, rather than the location of x_* . General-purpose Information-Based Bayesian Optimization (GIBBON) [42] provides an extension of MES that enables application to batched acquisition as well as other challenges such as multi-fidelity BO. GIBBON proposes a lower bound formulation for the intractable batch MES criterion, which is then optimized using greedy selection. Despite being formulated to handle a large degree of parallelism, Moss et al. [42] reported that GIBBON fails in practice for large batches with Q > 50. Potentially, this behaviour is a consequence of the accuracy of the lower bound approximation. A heuristic scaling of the batch diversity was proposed to improve performance with large batches. GIBBON may also be interpreted as a determinantal point process (DPP) [4, 43]. In Appendix B we provide a detailed discussion of the relationship of the BEEBO acquisition function to GIBBON and DPPs. Note that while we will also make use of the term entropy in BEEBO, the quantity is distinct from the ones leveraged by the aforementioned approaches in the sense that it does not relate to an unknown optimum.

Thompson sampling Given the challenges of generalizing acquisition functions to batch mode, Thompson sampling (TS), which was originally adopted from bandit problems [44, 45, 46, 47, 48], is a popular alternative strategy for guiding batched BO. While being an attractive approach in general, it has been demonstrated that default TS can become too exploitative, motivating the use of alternatives such as Bayesian Quadrature [49], or advanced strategies on top of TS that ensure diversity [18]. Eriksson et al. [50] demonstrate that overexploration also can be problematic in higher dimensions, and alleviate this using local trust regions in TuRBO. Maintaining such regions with high precision discretization can be memory-expensive, as indicated by [51], who suggest using MCMC-BO with adaptive local optimization to address this by transitioning a set of candidate points towards more promising positions.

3 The BEEBO acquisition function

Assume $f_{\text{true}}: X \to \mathbb{R}$ is some real output associated with the input and a set of data be given $D = \{(x_i, y_i)\}_{i=1}^N$ where $y_i \in \mathbb{R}$ represent some noisy observations of $f_{\text{true}}(x_i)$, say

$$y_i = f_{\text{true}}(x_i) + \epsilon_i \tag{1}$$

with ϵ_i denoting the measurement noise. Let $\mathbf{x}=(x_1,\cdots,x_Q)\in X^Q$ represent a collection of test points we wish to assign an acquisition value to. In keeping with the BO framework, we assume a given posterior probability distribution over the surrogate function \mathbf{f} evaluated at \mathbf{x} ,

$$\mathbf{f}(\mathbf{x}) \sim P(\mathbf{f} \mid D, \mathbf{x})$$
 (2)

The lack of knowledge we have of the surrogate function at x is quantified by the differential *entropy* H:

$$H(\mathbf{f} \mid D, \mathbf{x}) = -\int P(\mathbf{f} \mid D, \mathbf{x}) \ln(P(\mathbf{f} \mid D, \mathbf{x})) d\mathbf{f}$$
(3)

This entropy can be contrasted with the expected entropy of the surrogate function if Q observations $\mathbf{y}=(y_1,\cdots,y_Q)$ were acquired at \mathbf{x} , i.e. if the training data D would be augmented with $D'(\mathbf{y})=\{(x_q,y_q)\}_{q=1}^Q$ to form the joint data set $D_{\mathrm{aug}}(\mathbf{y})=(D,D'(\mathbf{y}))$. We refer to this entropy as H_{aug} :

$$H_{\text{aug}}(\mathbf{f} \mid D, \mathbf{x}) = \int P(\mathbf{y} \mid D, \mathbf{x}) H(\mathbf{f} \mid D_{\text{aug}}(\mathbf{y})) d\mathbf{y}, \tag{4}$$

where $P(\mathbf{y} | D, \mathbf{x})$ represent the posterior predictive distribution at \mathbf{x} . The expected *information gain*, $I(\mathbf{x})$, from acquiring observations at \mathbf{x} is given by the expected reduction of entropy from this process:

$$I(\mathbf{x}) = H(\mathbf{f} \mid D, \mathbf{x}) - H_{\text{aug}}(\mathbf{f} \mid D, \mathbf{x})$$
(5)

We propose to represent the explore component of the acquisition function, $a_{\rm BEEBO}$, by $I(\mathbf{x})$. The information gain $I(\mathbf{x})$ is distinct from the quantities exploited by entropy search approaches, as it quantifies global uncertainty reduction, rather than estimating the information over an unknown x_* . The information gain is directly applicable to multivariate functions and to heteroskedastic settings where $\sigma^2 = \sigma^2(\mathbf{x})$. Since large measurement uncertainties imply smaller information gain, $a_{\rm BEBO}$ exhibits risk-averse behaviour [11] by automatically prioritizing regions of small uncertainties from where more precise information of $f_{\rm true}$ can be obtained, everything else being equal.

The exploit component of BEEBO relies on taking expectation values of a scalar function of the random variable $\mathbf{f}(\mathbf{x})$, $\tilde{E}: \mathbb{R}^Q \to \mathbb{R}$, that summarizes the optimality properties of a given batch \mathbf{x} . Natural choices would be the mean or the maximum of $\mathbf{f}(\mathbf{x})$. Of particular interest is expressing the optimality as a softmax-weighted sum over $\mathbf{f}(\mathbf{x})$, as this allows us to smoothly interpolate between the two regimes:

$$E(\mathbf{x}) = -\mathbb{E}[\tilde{E}(\mathbf{x})] \cdot Q = -\mathbb{E}\left[\sum_{q=1}^{Q} \operatorname{softmax}(\beta \mathbf{f})_{q} f_{q}\right] \cdot Q, \tag{6}$$

where β is the softmax inverse temperature. At $\beta=0$, we recover the mean. We scale the expectation with Q so that both I and E scale linearly with increasing batch size. While the mean provides a closed form expression for its expectation, this is not the case for the general softmax-weighted sum of a multivariate normal. Using Taylor expansion, we introduce an approximation of the expectation of the softmax-weighted sum that is fully differentiable and can be computed in closed form. A detailed derivation is provided in Appendix A. At $\beta=0$, all Q points contribute equally to $E(\mathbf{x})$, whereas at $\beta>0$, points that do not compete for optimality are dynamically *released*. This effect can be quantified as the *effective number of points* via the entropy of the softmax weights. In the following, we will refer to the (exact) $\beta=0$ limit as meanBEEBO, and the (approximated) general case as maxBEEBO.

The BEEBO acquisition function then takes the form

$$a_{\text{BEERO}}(\mathbf{x}) = -E(\mathbf{x}) + T \cdot I(\mathbf{x}),\tag{7}$$

where T sets the balance between exploitation (small T) and exploration (large T). As both E and I scale with the batch size Q, a given choice of T would set the explore-exploit balance in an approximately Q-independent manner. This acquisition function bears a strong similarity to the definition of (negative) *free energies* in statistical physics, where E and I correspond to respectively the thermodynamic energy and entropy of the system and T corresponds to the temperature.

BEEBO with Gaussian processes

Gaussian processes offer a particular convenient framework for BO, due to the availability of closedform expressions for the inference step [17]. Specifically

$$P(\mathbf{f} \mid D, \mathbf{x}) = \mathcal{N}(\mathbf{f} \mid \mu(\mathbf{x}), C(\mathbf{x}))$$

$$\mu(\mathbf{x}) = K(\mathbf{x}, \mathbf{x}_D) \cdot M_D^{-1} \cdot \mathbf{y}_D$$

$$C(\mathbf{x}) = K(\mathbf{x}, \mathbf{x}) - K(\mathbf{x}, \mathbf{x}_D) \cdot M_D^{-1} \cdot K(\mathbf{x}_D, \mathbf{x})$$

$$M_D = K(\mathbf{x}_D, \mathbf{x}_D) + \sigma^2(\mathbf{x}_D)$$
(8)

where $\mathcal{N}(\cdot | \mu, C)$ is the multivariate Gaussian distribution with mean μ , and covariance C, \mathbf{x}_D and \mathbf{y}_D are the x and y values of the acquired data, $\sigma(\mathbf{x}_D) = \operatorname{diag}\left(\sigma_1^2, \cdots, \sigma_N^2\right)$ is a diagonal matrix with the measurement uncertainties in the diagonal and $K(\cdot,\cdot)$ are matrices derived from the GP-kernel, $k(\cdot,\cdot)$, i.e. $K(\mathbf{x},\mathbf{x}')_{ij}=k(x_i,x_i')$. It is worth noting that $C(\mathbf{x})$ only depends on the input location of the test points \mathbf{x} and the data points \mathbf{x}_D with their corresponding measurement uncertainties, $\sigma^2(\mathbf{x}_D)$, but not on the actual observations, y_D . Consequently, the entropy of the posterior distribution

$$H(\mathbf{f} \mid D, \mathbf{x}) = \frac{Q}{2} \ln(2\pi e) + \frac{1}{2} \ln \det(C(\mathbf{x}))$$
(9)

is independent of y_D as well, with $\ln \det$ denoting the log determinant. Similarly, the expected entropy of f if observations at x were acquired, simply reads

$$H_{\text{aug}}(\mathbf{f} \mid D, \mathbf{x}) = \frac{Q}{2} \ln(2\pi e) + \frac{1}{2} \ln \det(C_{\text{aug}}(\mathbf{x})), \tag{10}$$

where

$$C_{\text{aug}}(\mathbf{x}) = K(\mathbf{x}, \mathbf{x}) - K(\mathbf{x}, \mathbf{x}_{\text{aug}}) \cdot M_{\text{aug}}^{-1} \cdot K(\mathbf{x}_{\text{aug}}, \mathbf{x})$$

$$M_{\text{aug}} = K(\mathbf{x}_{\text{aug}}, \mathbf{x}_{\text{aug}}) + \sigma^{2}(\mathbf{x}_{\text{aug}})$$
(11)

and $\mathbf{x}_{\text{aug}} = \mathbf{x}_{D_{\text{aug}}}$. The BEEBO acquisition function is then given by

$$a_{\text{BEERO}}(\mathbf{x}) = \mathbb{E}[\tilde{E}(\mathbf{x})] \cdot Q + T \cdot I(\mathbf{x})$$
 (12)

where the expectation is either the mean, $\frac{1}{Q}\sum_{q=1}^{Q}\mu_q$, or the closed form approximation of the softmax-weighted sum described in Appendix A, and

$$I(\mathbf{x}) = \frac{1}{2} \ln \det(C(\mathbf{x})) - \frac{1}{2} \ln \det(C_{\text{aug}}(\mathbf{x})). \tag{13}$$

All operations needed to compute the acquisition value $a_{\text{BEEBO}}(\mathbf{x})$ are analytical. Using automatic differentiation, the batch of points x can therefore be optimized with gradient-based methods, as laid out for meanBEEBO in Algorithm 1, with learning rate γ . In the pseudocode, \mathcal{GP} denotes a trained GP model object that holds the training data and the kernel function. Using the kernel's learned amplitude A, we can relate BEEBO's T parameter to the κ of UCB. This allows us to configure BEEBO using a scaled temperature T' that ensures both methods have equal gradients at isosurfaces, enabling the user to follow existing guidance and intuition from UCB to control the trade-off. A derivation is provided in Appendix B.1.

Algorithm 1: meanBEEBO optimization

Input: model \mathcal{GP} , initial batch points \mathbf{x} , temperature T

repeat

Calculate $\mu(\mathbf{x}), C(\mathbf{x})$ from Equation 8 using \mathcal{GP}

$$E \leftarrow -\sum_{q=1}^{\infty} \mu_q$$

3D \leftarrow fantasize(CD v

$$\begin{split} E &\leftarrow -\sum_{q=1}^{Q} \mu_q \\ \mathcal{GP}_{\text{aug}} &\leftarrow \text{fantasize}(\mathcal{GP}, \mathbf{x}) \\ \text{Calculate } C_{\text{aug}}(\mathbf{x}) \text{ from Equation 11 using } \mathcal{GP}_{\text{aug}} \end{split}$$

 $\begin{matrix} I \leftarrow \frac{1}{2} \ln \det \left(C(\mathbf{x}) \right) - \frac{1}{2} \ln \det \left(C_{\text{aug}}(\mathbf{x}) \right) \\ a \leftarrow -E + T * I \end{matrix}$

 $\mathbf{x} \leftarrow \mathbf{x} + \gamma \nabla a$

until converged

Output: optimized batch points x

4 Experiments

Test problems We benchmark acquisition function performance on a range of maximization test problems with varying dimensions (Table 1) available in BoTorch [52]. Test problems that are evaluated on multiple dimensions support specifying the respective arbitrary d. As a high-dimensional problem with low inherent dimensionality, we embed the six-dimensional Hartmann function in d=100 [50, 53, 54]. We additionally test on two robot control problems (robot arm pushing and rover trajectory planning) in Appendix D.3 [55, 56].

Table 1: Overview of the test problems used in the experiments.

Function	Dimension
Ackley	2, 10, 20, 50, 100
Shekel	4
Hartmann	6
Cosine	8
Rastrigin	2, 10, 20, 50, 100
Rosenbrock	2, 10, 20, 50, 100
Styblinski-Tang	2, 10, 20, 50, 100
Powell	10, 20, 50, 100
Embedded Hartmann 6	100

On each test problem, we perform 10 rounds of

BO using q-UCB or BEEBO with a given explore-exploit parameter for direct comparison. We use the scaled temperature T' (B.1) to ensure that both methods operate at the same trade-off. In round 0, we seed the surrogate GP with Q random points that were drawn so that each point has a minimum distance of 0.5 to the test problem's true optimum. We perform ten replicate runs for each problem and method, with replicate seeds controlled so that all methods start from the same Q random points in a replicate. As we evaluate performance in a fixed-round, fixed-Q optimization scenario, we set the explore hyperparameter to 0 in the last round (for maxBEEBO, we also set the softmax β to 0). We use Q = 100 for all experiments, which is commonly understood to be a large batch size [50]. Additional results on small batch sizes (5, 10) are provided in Appendix D.2. All experiments use BoTorch's default utilities for acquisition function optimization and GPyTorch [57] GP training (C.1).

Heteroskedastic noise We investigate performance when optimizing under heteroskedastic noise on the 2D Branin function with three global optima. To construct a heteroskedastic problem, we specify noise so that the noise level is maximal at optima 2 and 3, decaying exponentially with distance to any of the two noised optima (C.4). No noise maximum is added at optimum 1. Therefore, while all three optima share the same $f_{\text{true}}(x)$ (Figure A1), only optimum 1 is favorable in terms of heteroskedastic risk. We perform BO for ten rounds with $\beta=0.1$ and Q=10 using a heteroskedastic GP that learns surrogate models for both $f_{\text{true}}(x)$ and $\sigma^2(x)$. We report results over five replicate runs.

Metrics We report the mean best observed objective value after 10 rounds over the five replicates. As test problems have highly varying scales, we normalize the results on each test problem using min-max normalization. Typically, the minimum of a maximization problem is not known explicitly. We therefore set the minimum for normalization to the highest value observed among the random seed points. The maximum is given by the $f_{\text{true}}(x^*)$ of the problem. The metric thus directly quantifies how much progress has been made to the true optimum from the random starting configuration on a 0-1 scale.

As we are not only interested in identifying a single x with good $f_{\text{true}}(x)$, we additionally quantify the overall quality of the final (exploitative) batch. We compute the batch instantaneous regret $R = \sum_{q < Q} f_{\text{true}}(x^*) - f_{\text{true}}(x_q)$ of the last, exploitative batch. To bring results on all test problems to a similar scale, we divide it by the batch instantaneous regret of a batch of Q random points on each problem. We refer to this metric as the relative batch instantaneous regret, $R_{\text{rel}} = R_{t=10}/R_{\text{random}}$.

For BO under heteroskedastic noise, we wish to quantify the preference of a given method for different optima. As optima share the same $f_{\text{true}}(x^*)$, metrics operating on $f_{\text{true}}(x)$ are inherently unsuitable, and preference needs to be evaluated on x directly. For each acquired point x_i , we compute the distances $\|x_i - x_j^*\|_2$ to the J individual optima. We report the mean distance to each optimum over all points in a batch.

Table 2: Highest observed value after 10 rounds of BO with Q=100. The best value at each κ is indicated in blue. BEEBO is configured with $T'=1/2\sqrt{\kappa}$. Full BO curves are provided in D.6, confidence intervals and statistical tests in Tables A2, A3 and A4

Problem	d		$\sqrt{\kappa} = 0.1$			$\sqrt{\kappa} = 1.0$,	$\sqrt{\kappa} = 10.0$	
		meanBEEBO	maxBEEBO	q-UCB	meanBEEBO	maxBEEBO	q-UCB	meanBEEBO	maxBEEBO	q-UCB
Ackley	2	0.993	0.982	0.973	0.985	0.980	0.967	0.975	0.988	0.988
Levy	2	1.000	1.000	1.000	0.999	1.000	1.000	0.999	0.998	0.998
Rastrigin	2	0.981	0.993	0.951	0.989	0.983	0.983	0.983	0.993	0.933
Rosenbrock	2	0.976	0.982	0.949	0.956	0.979	0.943	0.955	0.938	0.962
Styblinski-Tang	2	0.961	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.999
Shekel	4	0.540	0.300	0.244	0.915	0.378	0.330	0.698	0.411	0.264
Hartmann	6	1.000	0.894	0.918	1.000	0.976	0.950	0.986	0.974	0.889
Cosine	8	1.000	0.999	0.934	0.999	0.972	0.924	0.619	0.895	0.621
Ackley	10	0.915	0.819	0.800	0.908	0.736	0.772	0.822	0.546	0.513
Levy	10	0.989	0.966	0.904	0.966	0.953	0.904	0.966	0.914	0.560
Powell	10	0.987	0.951	0.920	0.970	0.949	0.916	0.861	0.909	0.283
Rastrigin	10	0.463	0.558	0.420	0.536	0.573	0.522	0.595	0.590	0.311
Rosenbrock	10	0.994	0.991	0.966	0.991	0.986	0.971	0.904	0.975	0.645
Styblinski-Tang	10	0.837	0.822	0.309	0.835	0.638	0.492	0.289	0.229	0.049
Ackley	20	0.827	0.818	0.741	0.851	0.781	0.753	0.777	0.404	0.474
Levy	20	0.949	0.945	0.926	0.943	0.904	0.900	0.889	0.907	0.819
Powell	20	0.955	0.939	0.948	0.965	0.913	0.913	0.872	0.915	0.845
Rastrigin	20	0.399	0.484	0.423	0.473	0.472	0.480	0.508	0.522	0.401
Rosenbrock	20	0.993	0.992	0.973	0.995	0.983	0.982	0.907	0.933	0.924
Styblinski-Tang	20	0.737	0.667	0.203	0.689	0.394	0.561	0.330	0.274	0.034
Ackley	50	0.235	0.623	0.638	0.342	0.594	0.759	0.823	0.465	0.730
Levy	50	0.940	0.971	0.948	0.965	0.958	0.951	0.943	0.879	0.941
Powell	50	0.954	0.975	0.970	0.982	0.969	0.961	0.950	0.938	0.980
Rastrigin	50	0.322	0.472	0.431	0.476	0.470	0.397	0.432	0.439	0.481
Rosenbrock	50	0.971	0.976	0.968	0.984	0.981	0.986	0.983	0.962	0.981
Styblinski-Tang	50	0.584	0.509	0.312	0.675	0.342	0.694	0.393	0.325	0.356
Ackley	100	0.277	0.417	0.708	0.190	0.540	0.645	0.863	0.682	0.844
Emb. Hartmann 6	100	0.951	0.957	0.896	0.987	0.936	0.913	0.928	0.907	0.916
Levy	100	0.837	0.966	0.950	0.961	0.950	0.934	0.944	0.940	0.964
Powell	100	0.810	0.985	0.982	0.952	0.980	0.981	0.983	0.979	0.984
Rastrigin	100	0.497	0.441	0.446	0.401	0.455	0.442	0.459	0.455	0.443
Rosenbrock	100	0.822	0.953	0.972	0.971	0.976	0.969	0.980	0.970	0.978
Styblinski-Tang	100	0.537	0.423	0.308	0.474	0.353	0.532	0.401	0.296	0.278
Mean		0.795	0.811	0.758	0.828	0.790	0.801	0.788	0.744	0.678
Median		0.940	0.951	0.920	0.961	0.949	0.913	0.889	0.907	0.819

5 Results

5.1 BO on test problems

We benchmark BEEBO against q-UCB at three rates of κ . Overall, we find that the simpler mean-BEEBO variant outperforms maxBEEBO in terms of mean performance on all but the lowest rate of κ (Table 2). As we consider the configuration with the lowest rate to be exploit-dominated, this can be understood as a consequence of maxBEEBO effectively releasing non-contributing points for further exploration, with the low explore rate seeming sufficient to induce the necessary diversity.

While results on individual test problems vary, meanBEEBO shows improved performance over q-UCB especially in the medium dimension range up to 50. For d=100, we find mixed performance, with meanBEEBO gradually becoming more competitive with increasing κ . This is due to the fact that with increasing dimensionality, more exploration is beneficial for learning a good surrogate model before an actual BO process becomes effective. As we find that q-UCB inherently performs more random-like sampling at large Q, irrespective of κ , it benefits in such situations.

On average over all 33 performed experiments, mean BEEBO improves upon q-UCB for large batches at any of the three rates (Table A3). We additionally benchmarked BEEBO against other popular BO strategies without explore-exploit hyperparameters. Interestingly, we found that the Kriging Believer (KB) iterative heuristic [33] can perform very competitively for large batches when using LogEI [58] as the acquisition function (Table A1), especially on the two robot control problems (Appendix D.3), but can be slower to optimize than BEEBO (Table A9).

When evaluating $R_{\rm rel}$ in the ultimate round of BO, we find that BEEBO allows us to effectively acquire a batch with high $f_{\rm true}(x)$, highlighting the controllability of the acquisition function (Table 3). The $R_{\rm rel}$ of q-UCB is only slightly better than the R of a random batch in many cases, even though the explore component was explicitly set to 0. We note that this is not due to the surrogate function being unsuitable - the results in Table 2 indicate that in most cases the location of $f_{\rm true}(x^*)$ is approximately

Table 3: Relative batch instantaneous regret R_{rel} in round 10 ($\kappa = 0$) with Q = 100. The best value at each κ is indicated in blue. BEEBO is configured with $T' = 1/2\sqrt{\kappa}$. Lower means better.

Problem	d		$\sqrt{\kappa} = 0.1$			$\sqrt{\kappa} = 1.0$,	$\sqrt{\kappa} = 10.0$	
		meanBEEBO	maxBEEBO	q-UCB	meanBEEBO	maxBEEBO	q-UCB	meanBEEBO	maxBEEBO	q-UCB
Ackley	2	0.292	0.259	1.006	0.268	0.245	0.999	0.257	0.165	1.002
Levy	2	0.134	0.114	1.236	0.092	0.102	1.046	0.102	0.111	1.114
Rastrigin	2	0.455	0.578	1.010	0.425	0.454	0.999	0.407	0.500	1.020
Rosenbrock	2	0.001	0.004	0.992	0.001	0.004	1.094	0.002	0.002	1.014
Styblinski-Tang	2	0.168	0.172	1.024	0.169	0.170	1.027	0.170	0.170	1.051
Shekel	4	0.810	0.776	0.993	0.688	0.730	0.995	0.688	0.695	0.988
Hartmann	6	0.060	0.229	0.968	0.078	0.086	0.971	0.100	0.098	0.862
Cosine	8	0.045	0.001	0.953	0.001	0.016	0.975	0.222	0.061	0.922
Ackley	10	0.478	0.338	0.931	0.314	0.345	0.943	0.253	0.452	0.950
Levy	10	0.041	0.030	1.188	0.023	0.048	1.011	0.261	0.103	1.111
Powell	10	0.016	0.027	1.037	0.009	0.067	1.101	0.067	0.151	1.215
Rastrigin	10	0.629	0.563	0.920	0.523	0.541	0.907	0.567	0.402	0.905
Rosenbrock	10	0.002	0.013	0.906	0.004	0.015	0.770	0.074	0.052	0.918
Styblinski-Tang	10	0.196	0.220	1.174	0.223	0.337	1.126	0.559	0.496	1.219
Ackley	20	0.629	0.219	0.945	0.282	0.292	0.950	0.226	0.586	0.917
Levy	20	0.128	0.241	0.839	0.063	0.113	0.914	0.140	0.182	1.056
Powell	20	0.093	0.081	0.809	0.010	0.074	0.689	0.028	0.110	0.870
Rastrigin	20	0.686	0.600	0.864	0.610	0.635	0.838	0.541	0.555	0.852
Rosenbrock	20	0.047	0.105	0.591	0.004	0.048	0.578	0.036	0.051	0.903
Styblinski-Tang	20	0.426	0.398	1.113	0.378	0.504	1.107	0.691	0.578	1.177
Ackley	50	0.895	0.464	0.949	0.738	0.606	0.947	0.177	0.530	0.874
Levy	50	0.055	0.029	0.611	0.033	0.085	0.681	0.051	0.268	0.892
Powell	50	0.018	0.021	0.542	0.014	0.078	0.499	0.018	0.064	0.785
Rastrigin	50	0.793	0.573	0.813	0.653	0.592	0.810	0.795	0.585	0.768
Rosenbrock	50	0.016	0.021	0.539	0.049	0.031	0.520	0.010	0.048	0.594
Styblinski-Tang	50	0.463	0.478	1.012	0.676	0.574	1.196	0.681	0.727	0.981
Ackley	100	0.718	0.636	0.948	0.900	0.466	0.935	0.137	0.321	0.863
Emb. Hartmann 6	100	0.068	0.144	0.573	0.035	0.086	0.863	0.175	0.172	0.692
Levy	100	0.119	0.031	0.615	0.044	0.164	0.716	0.042	0.056	0.586
Powell	100	0.094	0.011	0.465	0.027	0.041	0.493	0.013	0.018	0.524
Rastrigin	100	0.506	0.501	0.759	0.604	0.557	0.832	0.540	0.544	0.780
Rosenbrock	100	0.114	0.044	0.518	0.027	0.048	0.589	0.014	0.031	0.507
Styblinski-Tang	100	0.389	0.522	0.924	0.503	0.582	1.203	0.562	0.742	0.930
Mean		0.291	0.256	0.872	0.257	0.265	0.889	0.261	0.292	0.904
Median		0.134	0.219	0.931	0.092	0.164	0.943	0.175	0.172	0.917

known by round 10. Rather, we assume that this a consequence of the challenges of MC-based optimization of the acquisition function at large Q.

5.2 BO under heteroskedastic noise

We compare performance of meanBEEBO and q-UCB on the 3-optimum Branin function. Under heteroskedastic noise, we find that BEEBO preferentially optimizes towards the low-noise optimum 1 at the expense of the noisy optima 2 and 3 (Figure 2) and is therefore risk-averse. In the homoskedastic case, BEEBO does not exhibit this preference and optimizes for multiple optima. As expected, q-UCB, which only uses the model posterior variance $\sqrt{C(x)}$ instead of quantifying the actual information gain, does not display any preference for low-noise optima, showing similar behaviour under heteroskedastic and homoskedastic noise and remaining risk-neutral.

6 Discussion

We introduce BEEBO, an acquisition function for BO with GPs that can be optimized analytically and that scales natively to batched acquisition. By exploiting the independence of the information gain $I(\mathbf{x})$ on measurements \mathbf{y} when using GP surrogates, BEEBO models the interdependence of unknown points \mathbf{x} in a batch and can optimize their positions jointly using gradient descent.

BEEBO enables full control of its explore-exploit trade-off using a hyperparameter T that directly balances two terms, akin to UCB. Unlike in the reparametrization-based q-methods, BEEBO's T has predictable behaviour also at increasing batch sizes.

The numerical complexity of BEEBO is dominated by the need to compute the inverse of $M_{\rm aug}$ in Equation 11, which in a plain implementation scales as $\mathcal{O}((N+Q)^3)$. However, this can be reduced to $\mathcal{O}(N^2Q)$; specifically, the Cholesky decomposition of $M_{\rm aug}$ can be expressed as Q rank-1 updates of the pre-computed Cholesky decomposition of M_D , where each update will have the complexity of $\mathcal{O}(N^2)$. The calculation of the energy, E, and the information gain, I, scales as $\mathcal{O}(N \cdot Q)$ and $\mathcal{O}(Q^3)$, respectively, and are thus sub-dominant to the update needed for $M_{\rm aug}^{-1}$. For large N this

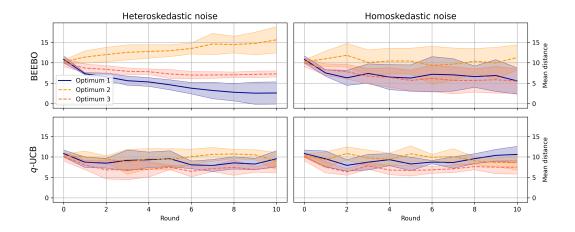


Figure 2: Mean distances of acquired points to the different optima of the Branin function. Under heteroskedastic noise, BEEBO is risk-averse and preferentially optimizes towards the low-noise optimum 1. Under homoskedastic noise, there is no preference. *q*-UCB does not adapt its behaviour to noise, remaining risk-neutral. The means and standard deviations over five replicates are shown.

approach may nevertheless become prohibitively slow. To overcome this limitation, methods for scalable GPs and fast predictive covariances such as LOVE [59] can be considered. The LOVE method allows a further reduction of the complexity of the Cholelsky update of $M_{\rm aug}$ to $\mathcal{O}(N \cdot r \cdot Q)$ [60], where r is the rank of the LOVE approximation for M_D , typically $r \ll N$.

As opposed to q-UCB, BEEBO can take heteroskedastic noise into account when computing the information gain, and preferentially acquires more informative low-noise points. We note that when the noise function $\sigma^2(x)$ is unknown, and needs to be explored at the same time with $f_{\text{true}}(x)$, it is critical that the initial random points sufficiently capture the noise landscape well enough for the information gain component to be useful, as the uncertainty of the surrogate on $\sigma^2(x)$ is not used. This would require a fully Bayesian approach that integrates over the distribution of $\sigma^2(x)$. The problem does not arise if the heteroskedastic noise of an experiment is known beforehand by e.g. instrument calibration. While not the focus of this work, we note that using the information gain could also be beneficial in sequential single-sample BO on heteroskedastic problems.

7 Outlook

In our experiments, we have focused on maintaining consistent explore-exploit ratios throughout the optimization rounds to ensure an equitable experimental comparison with q-UCB and demonstrate the effect of the hyperparameter choice. However, a more dynamic approach involving variable ratios could be more effective in real-world applications with a predetermined number of rounds [61]. Adopting a fully Bayesian perspective, one could consider the temperature hyperparameter T as a random variable. This opens up an intriguing avenue for BEEBO, where T could be drawn from a prior distribution that e.g. varies across optimization rounds, depending on the specific application. By tailoring this distribution, one could encourage a high level of exploration in the initial rounds, gradually transitioning towards a more exploitation-focused approach towards the end. In the presented experiments, we have implemented this as a strict constraint, maintaining a fixed T until the final round, at which point we shift to full exploitation, i.e., T=0.

While not explored in this work, we note that the BEEBO expression could naturally be extended to multi-objective optimization problems by capitalizing on GPs that handle vector-valued functions, such as multi-task GPs [12, 62]. Through e.g. the usage of the *intrinsic model of coregionalization*, we obtain a covariance function k, and thereby a covariance matrix $C(\mathbf{x})$, over all input-task pairs. As the multi-task covariance matrix is jointly Gaussian, the expression of the information gain remains unchanged and can be computed like in the single-task case. The energy $E(\mathbf{x})$ becomes vector-valued, providing an energy term for each of the tasks. This would allow for the introduction of task-specific

weights in the acquisition function. As the extension only affects the surrogate model, the scaling remains cubic in the number of input-task observations.

Beyond GPs, BEEBO could be generalized to work with any probabilistic model. However, GPs are unique in that H_{aug} is available in closed form and can be used to compute $I(\mathbf{x})$ analytically, without solving the integral over \mathbf{y} in Equation 4. Other models may require approximations and sampling-based approaches for computing the information gain.

Availability

A BoTorch implementation of BEEBO is available at

https://github.com/novonordisk-research/BEE-BO.

Acknowledgments and Disclosure of Funding

We thank Christoffer Riis, Jan C. Refsgaard and Kilian W. Conde-Frieboes for helpful discussions. FT was funded in part by the Novo Nordisk Foundation through the Center for Basic Machine Learning Research in Life Science (NNF20OC0062606).

References

- [1] H. J. Kushner. A New Method of Locating the Maximum Point of an Arbitrary Multipeak Curve in the Presence of Noise. *Journal of Basic Engineering*, 86(1):97–106, 03 1964.
- [2] Jonas Močkus. On bayesian methods for seeking the extremum. In *Optimization Techniques IFIP Technical Conference: Novosibirsk, July 1–7, 1974*, pages 400–404. Springer, 1975.
- [3] Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical Bayesian Optimization of Machine Learning Algorithms. *Advances in neural information processing systems*, 25, 2012.
- [4] Jesse Dodge, Kevin Jamieson, and Noah A Smith. Open loop hyperparameter optimization and determinantal point processes. *arXiv preprint arXiv:1706.01566*, 2017.
- [5] Ryan-Rhys Griffiths and José Miguel Hernández-Lobato. Constrained bayesian optimization for automatic chemical design using variational autoencoders. *Chemical science*, 11(2):577–586, 2020.
- [6] Ji Won Park, Samuel Stanton, Saeed Saremi, Andrew Watkins, Henri Dwyer, Vladimir Gligorijevic, Richard Bonneau, Stephen Ra, and Kyunghyun Cho. PropertyDAG: Multi-objective Bayesian optimization of partially ordered, mixed-variable properties for biological sequence design. *arXiv*, 2022.
- [7] Samuel Stanton, Wesley Maddox, Nate Gruver, Phillip Maffettone, Emily Delaney, Peyton Greenside, and Andrew Gordon Wilson. Accelerating Bayesian optimization for biological sequence design with denoising autoencoders. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 20459–20478. PMLR, 17–23 Jul 2022.
- [8] Peter I Frazier. A Tutorial on Bayesian Optimization. arXiv, 2018.
- [9] Samuel Daulton, Sait Cakmak, Maximilian Balandat, Michael A Osborne, Enlu Zhou, and Eytan Bakshy. Robust Multi-Objective Bayesian Optimization Under Input Noise. *arXiv*, 2022.
- [10] Benjamin Letham, Brian Karrer, Guilherme Ottoni, and Eytan Bakshy. Constrained bayesian optimization with noisy experiments, 2018.
- [11] Anastasia Makarova, Ilnura Usmanova, Ilija Bogunovic, and Andreas Krause. Risk-averse heteroscedastic bayesian optimization. *Advances in Neural Information Processing Systems*, 34:17235–17245, 2021.

- [12] Kevin Swersky, Jasper Snoek, and Ryan P Adams. Multi-task bayesian optimization. *Advances in neural information processing systems*, 26, 2013.
- [13] Shion Takeno, Hitoshi Fukuoka, Yuhki Tsukada, Toshiyuki Koyama, Motoki Shiga, Ichiro Takeuchi, and Masayuki Karasuyama. Multi-fidelity bayesian optimization with max-value entropy search and its parallelization. In *International Conference on Machine Learning*, pages 9334–9345. PMLR, 2020.
- [14] Riccardo Moriconi, Marc P Deisenroth, and K S Sesh Kumar. High-dimensional Bayesian optimization using low-dimensional feature spaces. *arXiv*, 2019.
- [15] Javad Azimi, Alan Fern, and Xiaoli Fern. Batch bayesian optimization via simulation matching. *Advances in Neural Information Processing Systems*, 23, 2010.
- [16] Tarun Kathuria, Amit Deshpande, and Pushmeet Kohli. Batched gaussian process bandit optimization via determinantal point processes. Advances in neural information processing systems, 29, 2016.
- [17] Christopher Williams and Carl Rasmussen. Gaussian processes for regression. *Advances in neural information processing systems*, 8, 1995.
- [18] Natalie Maus, Kaiwen Wu, David Eriksson, and Jacob Gardner. Discovering many diverse solutions with bayesian optimization. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent, editors, *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 1779–1798. PMLR, 25–27 Apr 2023.
- [19] Donald R Jones, Matthias Schonlau, and William J Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13:455–492, 1998.
- [20] Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- [21] Peter Frazier, Warren Powell, and Savas Dayanik. The knowledge-gradient policy for correlated normal beliefs. *INFORMS journal on Computing*, 21(4):599–613, 2009.
- [22] Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.
- [23] James T Wilson, Riccardo Moriconi, Frank Hutter, and Marc Peter Deisenroth. The reparameterization trick for acquisition functions. *arXiv*, 2017.
- [24] James T Wilson, Frank Hutter, and Marc Peter Deisenroth. Maximizing acquisition functions for Bayesian optimization. *arXiv*, 2018.
- [25] Amar Shah and Zoubin Ghahramani. Parallel Predictive Entropy Search for Batch Global Optimization of Expensive Objective Functions. Advances in neural information processing systems, 28, 2015.
- [26] David Ginsbourger, Janis Janusevskis, and Rodolphe Le Riche. Dealing with asynchronicity in parallel Gaussian process based global optimization. PhD thesis, Mines Saint-Etienne, 2011.
- [27] Enrico Crovini, Simon L Cotter, Konstantinos Zygalakis, and Andrew B Duncan. Batch bayesian optimization via particle gradient flows. arXiv preprint arXiv:2209.04722, 2022.
- [28] J. Dick, F. Y. Kuo, and I. H. Sloan. High-dimensional integration: the quasi-monte carlo way. *Acta Numerica*, 22:133–288, 2013.
- [29] R. E. Caflisch. Monte carlo and quasi-monte carlo methods. Acta Numerica, 7:1-49, 1998.
- [30] Durk P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. *Advances in neural information processing systems*, 28, 2015.

- [31] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pages 1278–1286. PMLR, 2014.
- [32] András Sóbester, Stephen J Leary, and Andy J Keane. On the design of optimization strategies based on global response surface approximation models. *Journal of Global Optimization*, 33:31–59, 2005.
- [33] David Ginsbourger, Rodolphe Le Riche, and Laurent Carraro. *Kriging Is Well-Suited to Parallelize Optimization*, pages 131–162. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- [34] Thomas Desautels, Andreas Krause, and Joel W. Burdick. Parallelizing exploration-exploitation tradeoffs in gaussian process bandit optimization. *Journal of Machine Learning Research*, 15(119):4053–4103, 2014.
- [35] Javier Gonzalez, Zhenwen Dai, Philipp Hennig, and Neil Lawrence. Batch bayesian optimization via local penalization. In Arthur Gretton and Christian C. Robert, editors, *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 648–657, Cadiz, Spain, 09–11 May 2016. PMLR.
- [36] Emile Contal, David Buffoni, Alexandre Robicquet, and Nicolas Vayatis. Parallel gaussian process optimization with upper confidence bound and pure exploration. In Hendrik Blockeel, Kristian Kersting, Siegfried Nijssen, and Filip Železný, editors, *Machine Learning and Knowledge Discovery in Databases*, pages 225–240, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [37] Javier Gonzalez, Michael Osborne, and Neil Lawrence. Glasses: Relieving the myopia of bayesian optimisation. In Arthur Gretton and Christian C. Robert, editors, *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 790–799, Cadiz, Spain, 09–11 May 2016. PMLR.
- [38] Philipp Hennig and Christian J. Schuler. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13(57):1809–1837, 2012.
- [39] José Miguel Hernández-Lobato, Matthew W. Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. ArXiv, abs/1406.2541, 2014.
- [40] Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient bayesian optimization. In *International Conference on Machine Learning*, pages 3627–3635. PMLR, 2017.
- [41] Henry B Moss, David S Leslie, and Paul Rayson. Mumbo: Multi-task max-value bayesian optimization. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2020, Ghent, Belgium, September 14–18, 2020, Proceedings, Part III,* pages 447–462. Springer, 2021.
- [42] Henry B. Moss, David S. Leslie, Javier I. González, and Paul Rayson. Gibbon: General-purpose information-based bayesian optimisation. *ArXiv*, abs/2102.03324, 2021.
- [43] Alex Kulesza and Ben Taskar. Determinantal point processes for machine learning. *Found. Trends Mach. Learn.*, 5:123–286, 2012.
- [44] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3-4):285–294, 1933.
- [45] Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 844–853. PMLR, 06–11 Aug 2017.
- [46] Kirthevasan Kandasamy, Akshay Krishnamurthy, Jeff Schneider, and Barnabás Póczos. Parallelised bayesian optimisation via thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, pages 133–142. PMLR, 2018.

- [47] Vahab Mirrokni, Mohammad Shadravan, et al. Parallelizing thompson sampling. In *Advances in Neural Information Processing Systems*, 2021.
- [48] Amin Karbasi, Vahab Mirrokni, and Mohammad Shadravan. Parallelizing thompson sampling. *Advances in Neural Information Processing Systems*, 34:10535–10548, 2021.
- [49] Masaki Adachi, Satoshi Hayakawa, Saad Hamid, Martin Jørgensen, Harald Oberhauser, and Micheal A Osborne. Sober: Scalable batch bayesian optimization and quadrature using recombination constraints. arXiv preprint arXiv:2301.11832, 2023.
- [50] David Eriksson and Martin Jankowiak. High-dimensional Bayesian optimization with sparse axis-aligned subspaces. In Cassio de Campos and Marloes H. Maathuis, editors, *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, volume 161 of *Proceedings of Machine Learning Research*, pages 493–503. PMLR, 27–30 Jul 2021.
- [51] Zeji Yi, Yunyue Wei, Chu Xin Cheng, Kaibo He, and Yanan Sui. Improving sample efficiency of high dimensional bayesian optimization with mcmc. *arXiv preprint arXiv:2401.02650*, 2024.
- [52] Maximilian Balandat, Brian Karrer, Daniel R. Jiang, Samuel Daulton, Benjamin Letham, Andrew Gordon Wilson, and Eytan Bakshy. Botorch: a framework for efficient monte-carlo bayesian optimization. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS'20, Red Hook, NY, USA, 2020. Curran Associates Inc.
- [53] Ziyu Wang, Frank Hutter, Masrour Zoghi, David Matheson, and Nando De Freitas. Bayesian optimization in a billion dimensions via random embeddings. *J. Artif. Int. Res.*, 55(1):361–387, jan 2016.
- [54] Leonard Papenmeier, Luigi Nardi, and Matthias Poloczek. Increasing the scope as you learn: Adaptive bayesian optimization in nested subspaces. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, Advances in Neural Information Processing Systems, 2022.
- [55] Zi Wang, Clement Gehring, Pushmeet Kohli, and Stefanie Jegelka. Batched large-scale bayesian optimization in high-dimensional spaces. In Amos Storkey and Fernando Perez-Cruz, editors, Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics, volume 84 of Proceedings of Machine Learning Research, pages 745–754. PMLR, 09–11 Apr 2018.
- [56] Zi Wang, Chengtao Li, Stefanie Jegelka, and Pushmeet Kohli. Batched high-dimensional bayesian optimization via structural kernel learning. In *International Conference on Machine Learning*, 2017.
- [57] Jacob R. Gardner, Geoff Pleiss, David Bindel, Kilian Q. Weinberger, and Andrew Gordon Wilson. Gpytorch: blackbox matrix-matrix gaussian process inference with gpu acceleration. In Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18, page 7587–7597, Red Hook, NY, USA, 2018. Curran Associates Inc.
- [58] Sebastian Ament, Sam Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Unexpected improvements to expected improvement for bayesian optimization. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [59] Geoff Pleiss, Jacob Gardner, Kilian Weinberger, and Andrew Gordon Wilson. Constanttime predictive distributions for Gaussian processes. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4114–4123. PMLR, 10–15 Jul 2018.
- [60] Shali Jiang, Daniel R. Jiang, Maximilian Balandat, Brian Karrer, Jacob R. Gardner, and R. Garnett. Efficient nonmyopic bayesian optimization via one-shot multi-step trees. *ArXiv*, abs/2006.15779, 2020.
- [61] George De Ath, Richard M. Everson, Alma A. M. Rahat, and Jonathan E. Fieldsend. Greed is good: Exploration and exploitation trade-offs in bayesian optimisation. *ACM Trans. Evol. Learn. Optim.*, 1(1), apr 2021.

- [62] Edwin V Bonilla, Kian Chai, and Christopher Williams. Multi-task gaussian process prediction. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007.
- [63] Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(8):235–284, 2008.
- [64] Benjamin Charlier, Jean Feydy, Joan Alexis Glaunès, François-David Collin, and Ghislain Durif. Kernel operations on the gpu, with autodiff, without memory overflows. *Journal of Machine Learning Research*, 22(74):1–6, 2021.

Appendix

Table of Contents

A	Approximating the expectation of the softmax weighted sum	15
	A.1 Motivation	. 15
	A.2 Derivation	. 16
	A.3 Practical considerations	. 17
	A.4 Number of effective points	. 18
В	Relationship to other acquisition strategies	18
	B.1 Relationship of BEEBO T and UCB κ hyperparameters	. 18
	B.2 GIBBON	. 19
	B.3 Determinantal Point Processes	. 20
	B.4 Local penalization	. 21
	B.5 RAHBO	. 21
C	Implementation details	22
	C.1 Acquisition function optimization	. 22
	C.2 Benchmark BO methods	. 22
	C.3 Test problems	. 23
	C.4 Heteroskedastic noise	. 23
D	Extended results	24
	D.1 Results including additional baselines	. 24
	D.2 Results for batch sizes 5 and 10	. 26
	D.3 Control problems	. 28
	D.4 Run time	. 28
	D.5 Results with random initialization in round 0	. 29
	D.6 BO curves for all experiments in Table 2 and Table A1	. 30

A Approximating the expectation of the softmax weighted sum

A.1 Motivation

We are free to choose any energy function \tilde{E} in BEEBO, the only requirement being that we are able to compute an expectation of \tilde{E} in order to obtain the scalar summary $E = \mathbb{E}[-\tilde{E}(\mathbf{f})]$. Of particular interest is the softmax weighted sum,

$$\tilde{E}(\mathbf{f}) = \sum_{i=1}^{Q} \operatorname{softmax}(\beta \mathbf{f})_i f_i,$$
 (A1)

where β is the softmax inverse temperature. The softmax weight vector ω computed as

$$\omega_i = \frac{\exp(\beta f_i)}{\sum_{j=1}^{Q} \exp(\beta f_j) + \exp(\beta_y y_{max})}$$
(A2)

where $\exp(\beta_y y_{max})$ is an optional reference threshold value, as in expected improvement, which we set to 0 if not used (β_y is either simply β or a dynamically scaled value that ensures $\tilde{E}(\mathbf{f})$ does

not become 0, see Equation A28). The parameter β allows us to interpolate between two extreme regimes,

$$\tilde{E}(\mathbf{f}) = \frac{1}{Q} \sum_{i=1}^{Q} f_i \qquad \text{for } \beta \to 0$$
 (A3)

$$\tilde{E}(\mathbf{f}) = \max(\mathbf{f})$$
 for $\beta \to \infty$ (A4)

In the first regime, all points of a batch equally contribute to the energy, whereas in the second regime only the single point "responsible" for the maximum is controlling the energy. Note that for numerical reasons, operating towards the $\beta \to \infty$ limit is impractical, as it will lead to zero gradients for all but one point, preventing optimization. We can set $\beta = A^{-1/2}$, with A being the prior uncertainty scale of the GP kernel. Since A represents the expected energy fluctuations for points far from data, this weighting scheme will reflect a natural compromise between Equation A3 and Equation A4.

As opposed to the mean, in the general case, the expectation of the maximum of a Q-dimensional multivariate normal is not available in closed form. To our best knowledge, this is also the case for the softmax weighted sum. In the following, we derive a closed-form approximation of the expectation of the softmax of βf that can be used for gradient-based optimization.

A.2 Derivation

Consider the softmax denominator

$$d(\mathbf{f}) = \sum_{j=1}^{Q} \exp(\beta f_j) + \exp(\beta_y y_{\text{max}}). \tag{A5}$$

We will Taylor expand ln(d) to the second order, using

$$\frac{\partial \ln(d)}{\partial f_i} = \beta \frac{1}{d} \exp(\beta f_i) = \beta \omega_i \tag{A6}$$

$$\frac{\partial^2 \ln(d)}{\partial f_i \partial f_j} = \beta^2 \left(-\frac{1}{d^2} \exp(\beta(f_i + f_j) + \frac{1}{d^2} \delta_{ij} \exp(\beta f_i)) \right)$$
(A7)

$$= \beta^2(\omega_i \delta_{ij} - \omega_i \omega_j), \tag{A8}$$

where δ_{ij} is the Kronecker delta. So

$$\ln(d) \approx \ln(d_{\mathbf{a}}) + \beta \mathbf{w}^T \cdot \Delta \mathbf{f} + \frac{\beta^2}{2} \Delta \mathbf{f}^T \cdot W \cdot \Delta \mathbf{f}, \tag{A9}$$

where \mathbf{a} is the Taylor expansion point, $d_{\mathbf{a}}$ is d evaluated at \mathbf{a} , $\Delta \mathbf{f} = \mathbf{f} - \mathbf{a}$, \mathbf{w} is the Q-dimensional vector $\boldsymbol{\omega}$ evaluated at \mathbf{a} , and W is the $Q \times Q$ matrix $W = \operatorname{diag}(\mathbf{w}) - \mathbf{w}\mathbf{w}^T$. Inserted into Equation A2 we have

$$\omega(\mathbf{f})_i \approx w_i * \exp(\beta \Delta f_i - \beta \mathbf{w}^T \cdot \Delta \mathbf{f} - \frac{\beta^2}{2} \Delta \mathbf{f}^T \cdot W \cdot \Delta \mathbf{f})$$
(A10)

With this approximation we can calculate the expectation value

$$\mathbb{E}[\omega(\mathbf{f})_i * f_i] = \frac{w_i \sqrt{\det(C(\mathbf{x})^{-1})}}{(2\pi)^{Q/2}} \int \exp(\lambda_i(\mathbf{f})) * f_i df$$
(A11)

$$\lambda_i(\mathbf{f}) = \beta \Delta f_i - \beta \mathbf{w}^T \cdot \Delta \mathbf{f}^T - \frac{\beta^2}{2} \Delta \mathbf{f}^T \cdot W \cdot \Delta \mathbf{f} - \frac{1}{2} (\mathbf{f} - \boldsymbol{\mu})^T \cdot C(\mathbf{x})^{-1} \cdot (\mathbf{f} - \boldsymbol{\mu})$$
(A12)

$$= c^{(i)} - \frac{1}{2} (\mathbf{f} - \boldsymbol{\nu}^{(i)})^T \cdot C(\mathbf{x})_{\text{softmax}}^{-1} \cdot (\mathbf{f} - \boldsymbol{\nu}^{(i)}), \tag{A13}$$

Where $c^{(i)}$, $\boldsymbol{\nu}^{(i)}$ and $C(\mathbf{x})_{\text{softmax}}$ are defined as follows:

$$C(\mathbf{x})_{\text{softmax}} = \left(C(\mathbf{x})^{-1} + \beta^2 W\right)^{-1} \tag{A14}$$

$$\mathbf{b}^{(i)} = \mathbf{e}^{(i)} - \mathbf{w} \tag{A15}$$

$$\boldsymbol{\nu}^{(i)} = C(\mathbf{x})_{\text{softmax}} \cdot \left(\beta \mathbf{b}^{(i)} + C(\mathbf{x})^{-1} \cdot \boldsymbol{\mu} + \beta^2 W \cdot \mathbf{a}\right)$$
(A16)

$$c^{(i)} = \frac{1}{2} (\boldsymbol{\nu}^{(i)})^T \cdot C(\mathbf{x})_{\text{softmax}}^{-1} \cdot \boldsymbol{\nu}^{(i)} - \beta (\mathbf{b}^{(i)})^T \cdot \mathbf{a}$$
(A17)

$$-\frac{\beta^2}{2}\mathbf{a}^T \cdot W \cdot \mathbf{a} - \frac{1}{2}\boldsymbol{\mu}^T \cdot C(\mathbf{x})^{-1} \cdot \boldsymbol{\mu}$$
 (A18)

and $\mathbf{e}^{(i)}$ is the *i*'th basis vector with components $e_j^{(i)} = \delta_{ij}$. We can avoid the explict use of the precision matrix by rewriting the updated covariance matrix as

$$C(\mathbf{x})_{\text{softmax}} = (C(\mathbf{x})^{-1} \cdot (I + \beta^2 C(\mathbf{x}) \cdot W))^{-1} = U(\mathbf{x}) \cdot C(\mathbf{x}),$$

where we have defined $U(\mathbf{x}) = (I + \beta^2 C(\mathbf{x}) \cdot W)^{-1}$. The updated mean vectors can conveniently be expressed as

$$\boldsymbol{\nu}^{(i)} = C(\mathbf{x})_{\text{softmax}} \cdot (\beta \mathbf{b}^{(i)} + C(\mathbf{x})^{-1} \cdot \boldsymbol{\mu} + \beta^2 W \cdot \mathbf{a})$$
(A19)

$$\nu^{(i)} = \beta C(\mathbf{x})_{\text{softmax}} \cdot \mathbf{e}^{(i)} + C(\mathbf{x})_{\text{softmax}} \cdot \left(-\beta \mathbf{w} + C(\mathbf{x})^{-1} \cdot \boldsymbol{\mu} + \beta^2 W \cdot \mathbf{a}\right)$$
(A20)

$$\boldsymbol{\nu}^{(i)} = \beta C(\mathbf{x})_{\text{softmax}} \cdot \mathbf{e}^{(i)} + \boldsymbol{\nu}', \tag{A21}$$

where ν' is a constant vector for all (i). Similarly

$$c^{(i)} = -\beta a_i + \frac{1}{2}\beta^2 (C(\mathbf{x})_{\text{softmax}})_{i,i} + \beta \nu_i' + c'$$
(A22)

$$c' = \beta \mathbf{w}^T \cdot \mathbf{a} + \frac{1}{2} \boldsymbol{\nu}'^T \cdot C(\mathbf{x})_{\text{softmax}}^{-1} \cdot \boldsymbol{\nu}' - \frac{\beta^2}{2} \mathbf{a}^T \cdot W \cdot \mathbf{a} - \frac{1}{2} \boldsymbol{\mu}^T \cdot C(\mathbf{x})^{-1} \cdot \boldsymbol{\mu}$$
(A23)

with c' again being a constant for all (i). The expectation of the softmax weighted summary is given by

$$\mathbb{E}[\tilde{E}] = K * \sum_{i=1}^{Q} w_i * \exp(c^{(i)}) * \nu_i^{(i)}$$
(A24)

where $K = \frac{\sqrt{\det(C(\mathbf{x})^{-1})}}{\sqrt{\det(C(\mathbf{x})^{-1} + \beta^2 W)}} = \sqrt{\det(U(\mathbf{x}))}$. The most natural choice for the expansion point is $\mathbf{a} = \boldsymbol{\mu}$ in which case $\boldsymbol{\nu}^{(i)}$ and $c^{(i)}$ reduces to

$$\nu^{(i)} = \beta C(\mathbf{x})_{\text{softmax}} \cdot (\mathbf{e}^{(i)} - \mathbf{w}) + \mu$$
(A25)

$$c^{(i)} = \frac{\beta^2}{2} (\mathbf{e}^{(i)} - \mathbf{w})^T \cdot C(\mathbf{x})_{\text{softmax}} \cdot (\mathbf{e}^{(i)} - \mathbf{w})$$
(A26)

A.3 Practical considerations

Linear algebra For numerical reasons, we avoid computing $C(\mathbf{x})_{\text{softmax}}$ explicitly, and instead use the $U(\mathbf{x}) \cdot C(\mathbf{x})$ factorization to compute solutions with $U(\mathbf{x})^{-1} = I + \beta^2 C(\mathbf{x}) \cdot W$.

$$C(\mathbf{x})_{\text{softmax}} \cdot \left(\mathbf{e}^{(i)} - \mathbf{w}\right) = U(\mathbf{x}) \cdot \left(C(\mathbf{x}) \cdot \mathbf{e}^{(i)} - C(\mathbf{x}) \cdot \mathbf{w}\right)$$
(A27)

Following GPyTorch practices, we make use of the LinearOperator package to exploit the structure of $U(\mathbf{x})^{-1}$ as an AddedDiagLinearOperator when solving. For determinants, we find that Linear-Operator's logdet implementation gives nondeterministic results, and we therefore perform a dense cast before computing K using default Pytorch.

While the factorization is numerically advantageous, it is still limited with regards to β . We find that at $\beta > 5$, numerical errors prevent a reliable calculation of the expectation. In practice, $A^{-1/2}$ lies in a range that allows numerically accurate solutions.

93134

Softmax When y_{max} grows much larger than the softmax input vector \mathbf{f} - a situation that can arise easily when initializing with random points for gradient-based optimization - the softmax weights ω can become numerically zero for all "real" points, thus leading to $E(\mathbf{x})=0$, and vanishing gradients. As we always wish to preserve a minimal energy contribution from the real points, we parametrize the inverse temperature applied to y_{max} , β_y , using a hyperparameter α that denotes the minimal fraction of probability mass pertaining to real points. This parametrization resembles the LogEI version of the expected improvement acquisition function [58] to address the problem of vanishing EI-gradients.

Let N denote the softmax denominator excluding y_{max} , $N = \sum_{i=1}^{Q} \exp(\beta \Delta f_i)$. We define

$$\exp(\beta_y \Delta y_{\text{max}}) = \min\left(\frac{1 - \alpha}{\alpha} N, \exp(\beta \Delta y_{\text{max}})\right)$$
(A28)

We used $\alpha = 0.05$ as a default in all our experiments.

A.4 Number of effective points

We can interpret the softmax as the number of effective points contributing to the energy of the batch. The entropy H of the softmax is given by

$$H(\boldsymbol{\omega}) = -\sum_{i=1}^{Q} \omega_i \ln(\omega_i), \tag{A29}$$

and the number of effective points, D_{eff} , is $\exp(H(\omega))$, so that

$$D_{\text{eff}} = \exp\left(-\sum_{i=1}^{Q} \omega_i \ln(\omega_i)\right). \tag{A30}$$

 D_{eff} is bounded by 1 (approaching the maximum) and Q (approaching the mean). Note that if we include y_{max} in the softmax denominator, we add $\omega_{y_{\mathrm{max}}} \ln(\omega_{y_{\mathrm{max}}})$ to $H(\boldsymbol{\omega})$, and the resulting number becomes bounded by 1 and Q+1.

B Relationship to other acquisition strategies

In the following section, we will discuss how BEEBO is related to UCB, GIBBON, Determinantal Point Processes (DPP), the Local Penalization heuristic and RAHBO. We will base our analysis on meanBEEBO, as the softmax-mediated interdependency of points in maxBEEBO prevents a simple interpretation of the objective in a single-point stepwise manner and does not allow for the same direct analogies to other strategies.

B.1 Relationship of BEEBO T and UCB κ hyperparameters

BEEBO bears some resemblance to the UCB acquisition function, which in the single particle mode, Q=1, reads

$$a_{\text{UCB}}(x) = \mu(x) + \sqrt{\kappa}\sqrt{C(x)},\tag{A31}$$

where the parameter κ controls the balance between exploitation and exploration and $\mu(x)$ and C(x) are respectively the mean and variance of the posterior distribution, $P(f \mid x, D)$, as before. We note that a_{UCB} does not account for the uncertainty of the measurement at x, and therefore remains risk-neutral under heteroskedastic noise [11]. To understand the relationship between BEEBO and UCB, we will therefore limit ourselves to the homoskedastic case and furthermore assume that measurement variances σ^2 are much smaller than the typical prior variance of the GP surrogate, A, of f, e.g. $A \simeq N^{-1}Tr(K)$, so $\sigma^2 \ll A$ and $M^{-1} = (K + \sigma^2)^{-1} \approx K^{-1}$. In this limit, the variance of f(x) after measurement (indexed at i = n, say) reduces to σ^2 :

$$C(x) = ((K^{-1} + \sigma^{-2}I)^{-1})_{nn} = (K(K + \sigma^{2}I)^{-1}\sigma^{2})_{nn} \approx \sigma^{2}$$
(A32)

and the information gain becomes

$$I(x) \approx \frac{1}{2} \ln(C(x)) - \log(\sigma).$$

Consequently, the gradient of the two acquisition functions reads

$$\nabla a_{\text{UCB}}(x) = \nabla \mu(x) + \frac{\sqrt{\kappa}}{2 \cdot \sqrt{C(x)}} \cdot \nabla C(x)$$

$$\nabla a_{\text{BEEBO}}(x) = \nabla \mu(x) + \frac{T}{2 \cdot C(x)} \cdot \nabla C(x).$$

The two gradients will be identical at points x where the posterior uncertainties satisfy $\sqrt{C(x)} = \frac{T}{\sqrt{\kappa}}$. For comparison, we may desire equal gradients at iso-surfaces corresponding to a given fraction, ν , of the prior uncertainty scale \sqrt{A} , by setting T accordingly as $T = \nu \cdot \sqrt{A} \cdot \sqrt{\kappa}$. In our experiments, we use $\nu = \frac{1}{2}$ and configure BEEBO using a dimensionless T' explore-exploit parameter, defined as $T' = \frac{T}{\sqrt{A}}$, and set $T' = \frac{1}{2}\sqrt{\kappa}$ for a given benchmark experiment.

B.2 GIBBON

GIBBON [42] approximates the (intractable) General-purpose max-value Entropy Search acquisition function, which quantifies the mutual information $MI(f^*_{\text{true}};\mathbf{y}|D)$ of a batch of measurements \mathbf{y} and the unknown optimum f^*_{true} . It does so using a lower bound on the information gain and MC estimation of the expectation over f^*_{true} . It can be written as

$$\alpha_{\text{GIBBON}}(\mathbf{x}) = \frac{1}{2} \ln \det(R) - \frac{1}{2|\mathcal{M}|} \sum_{m \in \mathcal{M}} \sum_{i=1}^{Q} \ln \left(1 - \rho_i^2 \frac{\phi(\gamma_i(m))}{\phi(\gamma_i(m))} \left[\gamma_i(m) + \frac{\phi(\gamma_i(m))}{\phi(\gamma_i(m))} \right] \right)$$

$$\alpha_{\text{GIBBON}}(\mathbf{x}) = \frac{1}{2} \ln \det(R) + \sum_{i=1}^{Q} \hat{\alpha}_{\text{GIBBON}}(x_i), \tag{A33}$$

where R is the correlation matrix with entries $R_{ij} = \frac{C(\mathbf{x})_{ij}}{\sqrt{C(\mathbf{x})_{ii}C(\mathbf{x})_{jj}}}$, \mathcal{M} is a set of samples for the max-value f_{true}^* , and ρ_i is the correlation of y_i and $f_{\text{true}}(x_i)$. ϕ and ϕ are the standard normal cumulative distribution and probability density functions, and $\gamma_i(m) = \frac{m-\mu}{\sigma}$.

The definition of BEEBO introduced in Equation 7, with the scalar summarization function set to the expected mean, $E(\mathbf{x}) = -\frac{1}{Q} \sum_{i=1}^{Q} \mu(x_i)$, gives

$$\alpha_{\text{BEEBO}}(\mathbf{x}) = T * \frac{1}{2} \left(\ln \det \left(C(\mathbf{x}) \right) - \ln \det \left(C_{\text{aug}}(\mathbf{x}) \right) \right) + \sum_{i=1}^{Q} \mu(x_i). \tag{A34}$$

From the second formulation of GIBBON, it becomes obvious that although being distinct in their motivation and derivation, BEEBO and GIBBON implement acquisition functions with a similar structure. Taking an information theoretic and multi-fidelity BO standpoint, GIBBON refers to this trade-off as *diversity* against *quality*, whereas in BEEBO we follow the intuitions of UCB, and use *exploration* and *exploitation*.

- Quality Exploitation: GIBBON employs an MC estimate of the lower bound approximation
 of the information gain provided by each point, whereas BEEBO directly summarizes the
 optimality of all points in closed form, either as their mean or an approximated softmax
 weighted sum.
- Diversity Exploration: In GIBBON, the diversity derived from the differential entropy $H(\mathbf{f}|D,\mathbf{x})$ is the entropy of the posterior correlation $\frac{1}{2}\ln\det(R)$. In BEEBO, we employ the reduction of entropy, the information gain $I(\mathbf{x})$. Under homoskedastic noise, $I(\mathbf{x}) \propto \ln\det(C(\mathbf{x}))$. Since $R(\mathbf{x}) = \operatorname{diag}(C(\mathbf{x}))^{-1/2} \cdot C(\mathbf{x}) \cdot \operatorname{diag}(C(\mathbf{x}))^{-1/2}$, we have that $\ln\det(R) = \ln\det(C(\mathbf{x})) \sum_{i}^{Q} \ln(C(\mathbf{x})_{ii})$. Therefore, maximizing the log determinant of R penalizes points that have high variance.

Therefore, while GIBBON presents an attractive approximation of max-value Entropy Search for batched acquisition, BEEBO is an alternative that avoids approximating a quality criterion using MC. Moreover, GIBBON's diversity criterion implicitly penalizes points that have high variance, whereas BEEBO's criterion maximizes the reduction of variance. We find that BEEBO is orders of magnitudes faster to compute than GIBBON (Figure A3).

In the context of large batches (Q >> 10), a modification of GIBBON exists that is further similar to BEEBO. Departing from the strict max-value entropy search derivation, a scaling factor Q^{-2} is introduced to counteract a growing dominance of the diversity term:

$$\alpha_{\text{scaledGIBBON}}(\mathbf{x}) = \frac{1}{2Q^2} * \ln \det(R) + \sum_{i=1}^{Q} \hat{\alpha}_{\text{MES}}(x_i). \tag{A35}$$

This scaling is motivated by the fact that R contains Q^2 elements. However, we note that R is summarized by its log determinant, which scales linearly in Q: As the determinant is the product of the eigenvalues, the log determinant is the sum of the log-eigenvalues. The number of eigenvalues scales linearly with matrix size Q, and so does the log determinant.

B.3 Determinantal Point Processes

A Determinantal Point Process [43] specifies a probability over a set of points, or a "configuration of points" drawn from a ground set. Specifically, the probability of a set of Q points \mathbf{x} is given by

$$P(\mathbf{x}) \propto \det\left(L_{\mathbf{x}}\right),\tag{A36}$$

where $L_{\mathbf{x}}$ is a $Q \times Q$ symmetric matrix. Kulesza et al. [43] provide a decomposition of the general DPP kernel L that makes *quality* and *diversity* components explicit, so that

$$L_{ij} = q(x_i)q(x_j)k(x_i, x_j), \tag{A37}$$

with k being a $\mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^+$ similarity kernel, and q being a unary $\mathbb{R}^d \to \mathbb{R}$ scalar quality function. This framework is naturally amenable to batch BO, as we seek to select a collection of points that trade off quality (optimality) and diversity. Note that both k and q are distinct functions that need to be specified by the user, leading to the practical complication that they must be chosen very carefully so that their scales do not dominate each other, which limits the utility of this decomposition in practice [42].

In the following, we show how BEEBO is equivalent to a DPP, and derive the necessary k and q. Again, we consider BEEBO

$$\alpha_{\text{BEEBO}}(\mathbf{x}) = -E(\mathbf{x}) + T * I(\mathbf{x}), \tag{A38}$$

with the scalar summarization function set to $E(\mathbf{x}) = -\sum_{i=1}^{Q} f(x_i)$. We will first focus on the information gain term $I(\mathbf{x})$, which we can rearrange as

$$I(\mathbf{x}) = \frac{1}{2} \ln \det(C(\mathbf{x})) - \frac{1}{2} \ln \det(C_{\text{aug}}(\mathbf{x})) = \frac{1}{2} \ln \det\left(C(\mathbf{x}) \cdot C_{\text{aug}}^{-1}(\mathbf{x})\right). \tag{A39}$$

Our similarity kernel k is therefore given by the entries of the matrix $S = C(\mathbf{x}) \cdot C_{\text{aug}}(\mathbf{x})^{-1}$, so that $k(x_i, x_j) = S_{ij}$. Note that due to the augmented covariance term, the implied k also depends on all other currently selected points in \mathbf{x} , and $L_{\mathbf{x}}$ is not a submatrix of an all-sample L. Therefore, BEEBO does not implement a DPP under heteroskedastic noise. However, if we only consider homoskedastic noise, BEEBO's $I(\mathbf{x})$ simplifies to the posterior entropy [63], and therefore $S = C(\mathbf{x})$. As $C(\mathbf{x})$ can be accessed as a submatrix of an all-sample C, this permits a DPP.

Given the choice of $E(\mathbf{f})$, we can rewrite BEEBO as

$$\alpha_{\text{BEEBO}}(\mathbf{x}) = \ln \det (S) * T * \frac{1}{2} + \sum_{i=1}^{Q} \mu_{i}$$

$$\frac{2}{T} * \alpha_{\text{BEEBO}}(\mathbf{x}) = \ln \det (S) + \sum_{i=1}^{Q} \frac{2}{T} * \mu_{i}$$

$$\frac{2}{T} * \alpha_{\text{BEEBO}}(\mathbf{x}) = \ln \det (S) + \ln \det (D))$$

$$\frac{2}{T} * \alpha_{\text{BEEBO}}(\mathbf{x}) = \ln \det \left(D^{\frac{1}{2}} \cdot S \cdot D^{\frac{1}{2}}\right)$$

$$\alpha_{\text{BEEBO}}(\mathbf{x}) = \ln \det \left(D^{\frac{1}{2}} \cdot S \cdot D^{\frac{1}{2}}\right) * T * \frac{1}{2}$$

$$\alpha_{\text{BEEBO}}(\mathbf{x}) = \ln \det \left(L\right) * T * \frac{1}{2}$$

$$\alpha_{\text{BEEBO}}(\mathbf{x}) = \ln \det (L) * T * \frac{1}{2}$$
(A40)

where L is a matrix with entries $L_{ij} = S_{ij} \exp(\frac{1}{T} * \mu_i) \exp(\frac{1}{T} * \mu_j)$. BEEBO therefore uses the DPP quality function $q(x_i) = \exp(\frac{1}{T} * \mu_i)$, and, like proven previously for GIBBON, a batch \mathbf{x} with maximal α_{BEEBO} corresponds to the MAP of a DPP.

B.4 Local penalization

Local penalization (LP) is a greedy batch selection strategy that given any arbitrary single-point acquisition function, ensures diversity by applying a *penalization function* $\psi(x, x')$ that downweights the acquisition value of candidate locations x based on their proximity to already selected points. The criterion for selecting x_i is given by

$$x_i = \arg\max\alpha(x) \prod_{j=1}^{i-1} \psi(x, x_j). \tag{A41}$$

Note that in this formulation, the product includes all previously selected points, not just the current batch. The penalization function ψ may in principle be chosen freely. Gonzalez et al. [35] propose exploiting the fact that f_{true} is Lipschitz continuous in order to bound the position of the unknown optimum and penalize accordingly. The Lipschitz constant L is inferred from the GP surrogate and used to parametrize ψ . In LP, acquisition function optimization proceeds iteratively. After an x_i is chosen, the corresponding penalizing multiplier is added to the objective before optimizing for the next x_{i+1} .

While BEEBO enables optimization to proceed in parallel, it is of course possible to also optimize BEEBO greedily (under homoskedastic noise, I is submodular). In this case, it implements an LP strategy where $\alpha(x) = \mu(x)$. Rather than a product of individual $\mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ function evaluations, the penalizer implied by BEEBO is the information gain $I(\mathbf{x}): \mathbb{R}^{i \times d} \to \mathbb{R}$ that we evaluate by concatenating a candidate point to the already acquired x at each iteration. Like in GIBBON, this constitutes an LP strategy that does not require estimation of any properties of f_{true} beyond learning the GP surrogate.

B.5 RAHBO

Risk-averse Heteroskedastic Bayesian Optimization (RAHBO) [11] is a UCB-derived single-point acquisition function that avoids heteroskedastic risk, preferentially selecting points with low noise. While it is not applicable to batched acquisition directly, we here compare it to single-sample BEEBO to highlight different ways of addressing noise. Given a heteroskedastic surrogate model that learns an additional GP for the noise, the *variance proxy*, RAHBO reads

$$\begin{array}{ll} \alpha_{\rm RAHBO}(x) = & {\rm UCB}_f(x) - \alpha * {\rm LCB}_{\rm var}(x) \\ \alpha_{\rm RAHBO}(x) = & \mu_f(x) + \beta_f * \sigma_f(x) - \alpha(\mu_{\rm var}(x) - \beta_{\rm var} * \sigma_{\rm var}(x)) \end{array} , \tag{A42}$$

where μ_f and σ_f are the posterior mean and variance of the surrogate model and β_f is the standard UCB trade-off hyperparameter, yielding the standard upper confidence bound UCB_f. α is the chosen risk tolerance, and LCB is the lower confidence bound of the variance GP with posterior mean μ_{var} and variance σ_{var} traded off using β_{var} .

At Q = 1, BEEBO can be expressed as

$$\alpha_{\text{BEEBO}} = \mu_f(x) + T * \frac{1}{2} \ln(\sigma_f(x)) - T * \frac{1}{2} \ln(\sigma_f^{\text{aug}}(x))$$

$$\alpha_{\text{BEEBO}} = \mu_f(x) + T * \frac{1}{2} \ln\left(\frac{\sigma_f(x)}{\sigma_f^{\text{aug}}(x)}\right), \tag{A43}$$

where the variance proxy at x is considered via the augmented posterior variance σ_f^{aug} .

While RAHBO penalizes risk on an absolute scale, subject to α , BEEBO optimizes for high uncertainty reduction, quantified as the log ratio of the variance before and after making measurements.

Moreover, RAHBO differentiates between known and unknown variance proxies, and uses the LCB_{var} term to discount the predicted variance according to its uncertainty. In its closed-form analytical expression, BEEBO does not permit for the uncertainty of the variance proxy to be taken into account, being more similar to the known variance RAHBO

$$\alpha_{\text{RAHBO}}(x) = \mu_f(x) + \beta_f * \sigma_f(x) - \alpha \mu_{\text{var}}(x)$$
(A44)

where μ_{var} is a noise-free proxy. Either a sampling-based approach, or approximations to I(x) would need to be introduced to handle variance proxy uncertainty in BEEBO.

C Implementation details

C.1 Acquisition function optimization

BEEBO was implemented for full compatibility with the BoTorch framework (version 0.9.4) [52] as an AnalyticAcquistionFunction. Standard BoTorch utilities for initializing and training GPs, initializing q-batches and performing gradient descent optimization of the acquisition function are used. We trained GPyTorch (version 1.11) [57] GP models with KeOps [64] Matérn 5/2 kernels (following BoTorch defaults with a separate length scale for each input dimension, and Gamma priors on the length and output scales). Log determinants for the information gain were computed using singular value decomposition for numerical stability.

GPyTorch provides a get_fantasy_model method that allows for the efficient augmentation of the training data of a GP with a set of points, as done in BEEBO. However, we observed that GPyTorch's implementation suffers from GPU memory leaks when used with automatic differentiation enabled. We therefore instantiate augmented models explicitly, not making use of the (more efficient) augmentation strategy.

All experiments were performed with double precision. SobolQMCNormalSampler was used for acquisition functions making use of the reparametrization trick. Experiments were run on individual Nvidia RTX 6000 and V100 GPUs. Five replicates for the benchmarking experiments required a total of approx. 5,000 RTX 6000 GPU hours, with the majority of the run time dedicated to the GIBBON baseline, rather than BEEBO itself (Figure A3, Table A9).

C.2 Benchmark BO methods

All methods were benchmarked in BoTorch. For q-EI, we used LogEI [58]. For TS, 10.000 base Sobol samples were drawn and sampled with MaxPosteriorSampling using the Cholesky decomposition

of the covariance matrix. GIBBON was optimized using sequential optimization following the BoTorch tutorial. We additionally implemented a custom version of GIBBON that applies the Q^{-2} scaling factor to the diversity term, as proposed in GIBBON's supplementary material. We used 100,000 random discretized candidates for max-value sampling. In a few iterations, optimizing GIBBON seemed challenging, with BoTorch reporting that no nonzero initialization candidate could be identified. KB was optimized using a custom greedy optimization loop with fantasized observations, using (single-sample) LogEI as the underlying acquisition function. TuRBO-1 was optimized following its BoTorch tutorial. None of the methods use a hyperparameter for controlling their explore-exploit trade-off. The results are therefore based on 10 iterations at defaults.

C.3 Test problems

Test functions All test functions were used in their BoTorch implementations. As done in previous work, the embedded Hartmann function was created by appending all-0 dummy dimensions to the original six dimensions [53, 54, 50].

Control problems We consider two control problems from previous work: A 14-dimensional parameter tuning task for controlling robot arms pushing two objects to a target location [55], and a 60-dimensional trajectory planning task for a rover navigating through a maze of obstacles [56]. Instead of converting the problem objectives into rewards as in the original work, we operate on the actual minimization objectives directly (distance to target, navigation loss), and follow BoTorch's approach of simply inverting the objective in order to yield maximization problems. Both problems were adapted from their available implementations in Wang et al. [56] to follow the BoTorch test problem API.

C.4 Heteroskedastic noise

The (inverted) Branin function has three global optima $f(x^*) = -0.397887$ at $x_1^* = (9.42478, 2.475), x_2^* = (-\pi, 12.275)$ and $x_3^* = (\pi, 2.275)$. We define heteroskedastic noise so that the variance is maximal at x_2^* and x_3^* . The noise decays exponentially with the distance from any of the two noised optima at a rate λ .

$$\sigma^{2}(x) = \sigma_{max}^{2} * \exp(-\lambda * \min(\|x - x_{2}^{*}\|_{2}, \|x - x_{3}^{*}\|_{2})$$
(A45)

For our experiments, we set $\sigma_{max}^2=100$ and $\lambda=0.05$. As the surrogate function, we use a HeteroskedasticSingleTaskGP provided in BoTorch. This model learns two GPs simultaneously, one for the function f(x) and one for the (also unknown) variance function $\sigma^2(x)$. When querying the oracle with a batch of points, noised observations of f(x) are provided together with the true σ^2 at each point. The homoskedastic control experiment uses a SingleTaskGP with inferred noise level. The homoskedastic noise is set to $\sigma^2=77.5$, which is the average noise level of the heteroskedastic function over the whole domain.

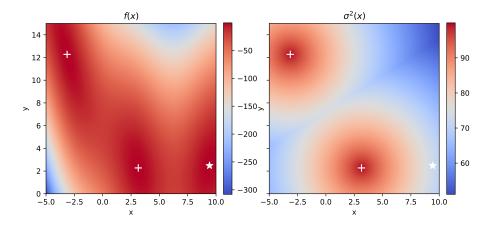


Figure A1: The Branin function with added heteroskedastic noise following Equation A45. $\sigma_{max}^2 = 100, \lambda = 0.05.$

D Extended results

D.1 Results including additional baselines

Table A1: BO on noise-free synthetic test problems. The normalized highest observed value after 10 rounds of BO with q=100 is shown. Colors are normalized row-wise. The BEE-BO and q-UCB columns are equivalent to Table 2. Higher means better. Results are means over five replicate runs.

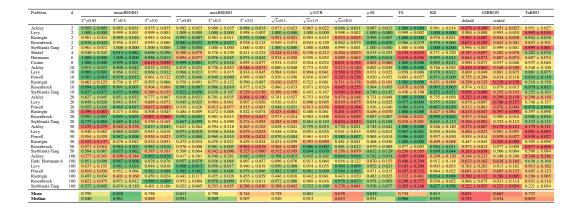


Table A2: BO on noise-free synthetic test problems. The relative batch instantaneous regret of the last, exploitative batch is shown. Colors are normalized row-wise. The BEEBO and q-UCB columns are equivalent to Table 3. Lower means better. Results are means over five replicate runs.



Table A3: Paired t-test p-values for the results of meanBEEBO in Table 2. The combined p-value was computed using Fisher's method. P-values smaller than 0.05 are indicated in bold.

				п	neanBEEE	O T'=0.05						meanBEE	BO T'=0.5						neanBEE	BO T'=5.0		
Problem	d	q-UCB	q-EI	TS	KB	GIBBON	GIBBON (s)	TuRBO	q-UCB	q-EI	TS	KB	GIBBON	GIBBON (s)	TuRBO	q-UCB	q-EI	TS	KB	GIBBON	GIBBON (s)	TuRBO
Ackley	2	8E-03	7E-02	1E+00	9E-03	3E-03	8E-04	1E-01	1E-01	6E-01	9E-01	3E-01	8E-03	1E-02	4E-01	9E-01	8E-01	1E+00	7E-01	8E-03	7E-02	6E-01
Levy	2	4E-01	5E-01	2E-01	1E+00	6E-04	4E-02	7E-02	7E-01	1E+00	3E-01	1E+00	5E-04	5E-02	7E-02	2E-01	1E+00	7E-01	1E+00	9E-04	7E-02	7E-02
Rastrigin	2	7E-03	9E-01	1E+00	3E-01	1E-02	1E-02	9E-01	2E-01	8E-01	1E+00	4E-02	7E-03	2E-03	5E-01	5E-04	1E+00	1E+00	2E-01	4E-03	9E-03	8E-01
Rosenbrock	2	1E-01	8E-01	2E-01	5E-01	6E-03	1E-01	3E-02	3E-01	9E-01	8E-01	9E-01	6E-03	1E-01	1E-01	8E-01	9E-01	9E-01	1E+00	9E-03	1E-01	8E-02
Styblinski-Tang	2	9E-01	9E-01	9E-01	9E-01	9E-01	9E-01	2E-05	3E-01	5E-01	1E-01	9E-01	2E-03	5E-02	5E-06	1E-01	6E-01	6E-02	1E+00	2E-03	1E-02	5E-06
Shekel	4	5E-03	4E-01	4E-04	4E-02	4E-04	6E-03	9E-01	4E-06	4E-04	8E-10	2E-05	1E-08	7E-09	5E-03	6E-04	9E-02	1E-04	1E-02	2E-04	6E-04	4E-01
Hartmann	6	8E-04	7E-02	1E-06	7E-02	4E-05	5E-04	1E-04	7E-03	7E-02	1E-06	7E-02	4E-05	5E-04	1E-04	3E-04	9E-01	2E-06	9E-01	1E-04	1E-03	3E-04
Cosine	8	6E-05	1E-06	1E+00	1E-03	8E-04	1E-03	6E-05	2E-05	1E-06	1E+00	2E-03	9E-04	1E-03	8E-05	5E-01	1E+00	1E+00	1E+00	1E+00	1E+00	1E+00
Ackley	10	4E-05	3E-05	1E+00	3E-06	5E-09	1E-06	8E-05	6E-05	2E-04	1E+00	5E-06	5E-10	8E-07	1E-04	8E-05	1E-01	1E+00	8E-04	4E-09	1E-05	1E-02
Levy	10	5E-05	1E-05	1E-07	6E-03	5E-05	7E-04	3E-03	1E-03	3E-03	1E-01	9E-01	8E-04	5E-03	7E-01	2E-04	1E-03	2E-01	8E-01	1E-04	5E-04	7E-01
Powell	10	3E-03	1E-03	1E-04	3E-04	3E-03	1E-03	1E-01	2E-03	9E-03	6E-04	6E-01	4E-03	2E-03	9E-01	1E-05	9E-01	7E-01	1E+00	4E-02	3E-01	1E+00
Rastrigin	10	2E-01	4E-01	1E+00	2E-01	6E-02	6E-05	1E+00	4E-01	6E-02	1E+00	5E-03	5E-03	4E-05	1E+00	1E-04	7E-05	1E+00	3E-04	8E-06	4E-05	1E+00
Rosenbrock	10	5E-03	2E-03	6E-06	1E-01	1E-03	8E-04	3E-01	5E-03	3E-03	5E-06	1E+00	4E-03	4E-03	8E-01	4E-03	1E+00	1E-01	1E+00	1E+00	1E+00	1E+00
Styblinski-Tang	10	6E-06	1E-02	3E-04	8E-01	2E-09	2E-06	3E-03	1E-05	1E-02	2E-04	8E-01	3E-09	3E-06	3E-03	8E-04	1E+00	1E+00	1E+00	2E-04	3E-01	1E+00
Ackley	20	8E-04	7E-04	1E+00	9E-03	4E-08	1E-07	3E-06	7E-05	1E-04	1E+00	1E-03	6E-09	2E-07	1E-06	9E-06	4E-02	1E+00	2E-01	2E-08	1E-06	8E-06
Levy	20	6E-02	2E-01	1E+00	8E-01	2E-02	1E-03	1E-01	1E-02	2E-01	1E+00	1E+00	3E-02	1E-03	2E-01	2E-02	1E+00	1E+00	1E+00	4E-01	1E-02	9E-01
Powell	20	3E-01	1E-02	1E+00	1E+00	2E-02	7E-03	1E+00	2E-03	1E-02	6E-01	6E-01	9E-03	5E-03	8E-01	2E-01	1E+00	1E+00	1E+00	1E+00	6E-01	1E+00
Rastrigin	20	8E-01	1E+00	1E+00	9E-01	7E-01	5E-01	1E+00	6E-01	2E-01	1E+00	1E-01	2E-02	2E-02	9E-01	1E-04	1E-02	1E+00	2E-02	2E-03	2E-03	6E-01
Rosenbrock	20	6E-04	1E-02	2E-05	1E+00	9E-03	1E-02	3E-01	4E-05	3E-04	1E-05	8E-01	8E-03	6E-03	1E-01	8E-01	1E+00	1E+00	1E+00	1E+00	1E+00	1E+00
Styblinski-Tang	20	1E-08	5E-04	2E-05	2E-03	3E-10	8E-07	1E-04	2E-02	6E-02	6E-05	1E-01	3E-08	2E-05	3E-02	5E-05	1E+00	1E-01	1E+00	5E-04	5E-01	1E+00
Ackley	50	1E+00	1E+00	1E+00	1E+00	6E-01	3E-01	4E-01	1E+00	1E+00	1E+00	1E+00	3E-01	6E-02	8E-02	2E-04	4E-03	1E+00	2E-04	1E-09	2E-07	6E-10
Levy	50	6E-01	7E-01	1E+00	7E-01	3E-02	8E-02	9E-05	2E-01	4E-02	1E+00	8E-02	8E-06	2E-02	3E-08	4E-01	9E-01	1E+00	9E-01	6E-06	1E-01	2E-06
Powell	50	9E-01	8E-01	1E+00	6E-01	5E-01	1E-01	1E-02	2E-04	3E-04	1E+00	8E-03	2E-06	9E-05	1E-06	1E+00	1E+00	1E+00	7E-01	8E-01	9E-02	1E-02
Rastrigin	50	1E+00	1E+00	1E+00	1E+00	1E+00	4E-01	1E+00	5E-04	3E-01	1E+00	3E-02	5E-02	1E-04	1E-01	1E+00	1E+00	1E+00	2E-01	8E-01	4E-03	8E-01
Rosenbrock	50	3E-01	9E-01	9E-01	1E+00	6E-01	8E-01	2E-02	9E-01	8E-03	3E-05	4E-01	1E-02	2E-02	4E-04	3E-01	2E-01	4E-02	6E-01	2E-02	1E-01	3E-04
Styblinski-Tang	50	1E-04	9E-01	2E-04	1E+00	3E-07	3E-02	3E-03	9E-01	1E-01	4E-05	8E-01	9E-10	3E-04	1E-05	2E-01	1E+00	8E-02	1E+00	2E-03	9E-01	9E-01
Ackley	100	1E+00	1E+00	2E-02	6E-01	5E-01	7E-01	7E-02	1E+00	1E+00	6E-02	8E-01	6E-01	7E-01	2E-01	4E-02	5E-04	3E-15	3E-07	2E-05	2E-04	2E-15
Emb. Hartmann 6	100	7E-02	1E-02	4E-04	1E-01	2E-05	3E-05	3E-01	2E-03	2E-02	1E-04	4E-02	4E-05	5E-05	3E-02	4E-01	2E-01	8E-05	4E-01	6E-05	8E-05	7E-01
Levy	100	1E+00	1E+00	2E-05	9E-01	3E-01	9E-01	6E-03	2E-02	2E-01	6E-06	1E-03	3E-02	2E-01	3E-02	1E+00	8E-01	9E-06	3E-02	4E-02	3E-01	5E-10
Powell	100	1E+00	1E+00	1E-03	1E+00	3E-03	5E-03	1E+00	9E-01	8E-01	2E-04	8E-01	4E-07	7E-05	1E-03	7E-01	3E-02	2E-04	2E-02	1E-05	2E-05	2E-08
Rastrigin	100	2E-03	1E-01	1E-01	1E+00	5E-04	4E-06	6E-05	8E-01	9E-01	3E-01	1E+00	2E-03	2E-02	4E-01	7E-02	1E+00	2E-01	1E+00	3E-04	9E-05	6E-05
Rosenbrock	100	1E+00	1E+00	3E-04	1E+00	1E+00	1E+00	9E-01	2E-01	7E-01	2E-04	2E-02	4E-01	2E-02 2E-01	6E-08	4E-01	5E-02	1E-04	2E-02	3E-04 3E-01	1E-01	5E-08
Styblinski-Tang	100	3E-06	5E-01	8E-04	1E+00	4E-08	4E-08	3E-06	9E-01	9E-01	9E-03	1E+00	3E-05	3E-05	3E-02	3E-03	1E+00	2E-02	1E+00	4E-04	1E-03	4E-01
	100																					
Combined		4E-16	2E-02	8E-02	1E+00	6E-47	1E-40	2E-18	1E-20	2E-12	7E-06	1E-03	2E-63	3E-51	2E-36	2E-21	1E+00	1E+00	1E+00	4E-44	4E-27	2E-23

Table A4: Paired t-test p-values for the results of maxBEEBO in Table 2. The combined p-value was computed using Fisher's method. P-values smaller than 0.05 are indicated in bold.

					maxBEEB	O $T'=0.05$						maxBEEI	SO T'=0.5						maxBEEE	3O $T'=5.0$		
Problem	d	q-UCB	q-EI	TS	KB	GIBBON	GIBBON (s)	TuRBO	q-UCB	q-EI	TS	KB	GIBBON	GIBBON (s)	TuRBO	q-UCB	q-EI	TS	KB	GIBBON	GIBBON (s)	TuRBO
Ackley	2	2E-01	7E-01	1E+00	5E-01	3E-03	2E-03	4E-01	2E-01	7E-01	9E-01	5E-01	1E-02	4E-02	5E-01	5E-01	5E-01	1E+00	1E-01	4E-03	5E-03	3E-01
Levy	2	8E-01	9E-01	3E-01	1E+00	7E-04	5E-02	7E-02	3E-01	9E-01	2E-01	1E+00	5E-04	4E-02	7E-02	5E-01	1E+00	9E-01	1E+00	3E-03	9E-02	7E-02
Rastrigin	2	2E-04	7E-01	1E+00	8E-03	5E-03	2E-03	3E-01	5E-01	1E+00	1E+00	2E-01	1E-02	4E-03	9E-01	3E-05	7E-01	1E+00	2E-02	6E-03	1E-03	2E-01
Rosenbrock	2	8E-02	5E-01	2E-01	3E-01	6E-03	9E-02	3E-02	2E-01	6E-01	3E-01	4E-01	5E-03	9E-02	4E-02	9E-01	9E-01	1E+00	1E+00	1E-02	2E-01	2E-01
Styblinski-Tang	2	4E-01	2E-01	7E-04	1E+00	2E-03	1E-02	5E-06	5E-01	8E-01	2E-01	1E+00	2E-03	3E-02	5E-06	4E-01	8E-01	6E-01	9E-01	4E-03	1E-01	5E-06
Shekel	4	1E-01	1E+00	2E-04	8E-01	5E-04	3E-01	1E+00	3E-01	9E-01	5E-03	5E-01	2E-02	1E-01	1E+00	3E-02	9E-01	4E-03	3E-01	7E-03	4E-02	1E+00
Hartmann	6	9E-01	1E+00	2E-05	1E+00	4E-03	3E-01	1E+00	1E-01	8E-01	1E-05	8E-01	2E-04	2E-03	2E-02	9E-04	9E-01	7E-06	9E-01	2E-04	3E-03	1E-02
Cosine	8	6E-05	1E-06	1E+00	1E-03	8E-04	1E-03	7E-05	1E-03	2E-05	1E+00	1E+00	4E-03	2E-02	8E-04	2E-03	1E-02	1E+00	1E+00	6E-01	9E-01	9E-01
Ackley	10	2E-01	2E-01	1E+00	2E-04	2E-08	1E-05	2E-02	9E-01	1E+00	1E+00	7E-01	6E-08	6E-05	9E-01	2E-01	1E+00	1E+00	1E+00	1E-04	5E-01	1E+00
Levy	10	5E-04	3E-03	9E-02	9E-01	6E-04	5E-03	8E-01	3E-03	4E-03	9E-01	1E+00	7E-04	5E-03	1E+00	1E-03	9E-01	1E+00	1E+00	1E-01	1E-01	1E+00
Powell	10	8E-02	4E-02	1E-03	9E-01	2E-03	1E-03	1E+00	4E-02	4E-02	2E-03	9E-01	3E-03	2E-03	1E+00	9E-06	8E-01	7E-02	1E+00	1E-02	5E-02	1E+00
Rastrigin	10	3E-04	1E-02	1E+00	3E-03	6E-04	9E-05	1E+00	2E-01	2E-03	1E+00	1E-04	1E-03	1E-04	1E+00	2E-04	5E-03	1E+00	4E-03	1E-04	7E-05	9E-01
Rosenbrock	10	1E-02	4E-03	6E-06	1E+00	6E-03	6E-03	8E-01	1E-02	3E-01	4E-06	1E+00	9E-03	2E-02	1E+00	1E-03	8E-01	4E-06	1E+00	4E-01	7E-01	1E+00
Styblinski-Tang	10	7E-06	5E-02	2E-04	1E+00	5E-10	4E-06	2E-03	7E-03	9E-01	8E-03	1E+00	2E-07	8E-05	6E-01	1E-02	1E+00	1E+00	1E+00	9E-03	6E-01	1E+00
Ackley	20	1E-04	6E-04	1E+00	9E-03	1E-08	3E-07	2E-06	1E-02	2E-02	1E+00	9E-02	4E-08	7E-07	2E-06	9E-01	1E+00	1E+00	1E+00	6E-04	2E-02	1E+00
Levy	20	1E-01	1E-01	1E+00	8E-01	6E-02	9E-04	1E-01	3E-01	1E+00	1E+00	1E+00	3E-01	2E-03	8E-01	2E-02	9E-01	1E+00	1E+00	2E-01	3E-03	7E-01
Powell	20	9E-01	4E-01	1E+00	1E+00	6E-02	2E-02	1E+00	5E-01	9E-01	1E+00	1E+00	5E-01	2E-02	1E+00	3E-03	9E-01	1E+00	1E+00	4E-01	1E-02	1E+00
Rastrigin	20	5E-03	7E-02	1E+00	3E-02	1E-02	3E-03	9E-01	6E-01	1E-01	1E+00	1E-01	4E-02	1E-02	9E-01	9E-05	1E-02	1E+00	2E-02	1E-03	5E-03	4E-01
Rosenbrock	20	5E-04	1E-02	1E-05	1E+00	8E-03	9E-03	4E-01	4E-01	9E-01	4E-06	1E+00	2E-02	2E-01	1E+00	3E-01	1E+00	9E-01	1E+00	9E-01	1E+00	1E+00
Styblinski-Tang	20	5E-08	1E-01	9E-05	3E-01	2E-09	2E-06	1E-01	1E+00	1E+00	1E-02	1E+00	4E-06	7E-02	1E+00	4E-07	1E+00	1E-01	1E+00	1E-05	9E-01	1E+00
Ackley	50	6E-01	9E-01	1E+00	2E-01	2E-02	9E-04	8E-04	1E+00	1E+00	1E+00	3E-01	7E-04	3E-03	4E-04	1E+00	1E+00	1E+00	9E-01	5E-04	3E-04	5E-05
Levy	50	2E-02	2E-02	1E+00	8E-03	3E-06	1E-02	1E-07	3E-01	1E-01	1E+00	4E-01	8E-07	4E-02	4E-07	1E+00	1E+00	1E+00	1E+00	5E-01	7E-01	5E-05
Powell	50	4E-02	3E-02	1E+00	5E-02	2E-04	7E-04	8E-06	8E-02	2E-01	1E+00	9E-02	6E-03	1E-03	5E-06	1E+00	1E+00	1E+00	1E+00	1E+00	5E-01	4E-03
Rastrigin	50	1E-02	4E-01	1E+00	2E-02	4E-02	6E-04	2E-01	2E-04	5E-01	1E+00	1E-02	8E-02	3E-04	2E-01	1E+00	1E+00	1E+00	1E-01	7E-01	5E-04	7E-01
Rosenbrock	50	7E-03	9E-01	7E-01	1E+00	2E-01	7E-01	5E-03	1E+00	2E-01	2E-02	8E-01	4E-02	2E-01	7E-04	1E+00	1E+00	1E+00	1E+00	1E+00	1E+00	3E-01
Styblinski-Tang	50	5E-05	1E+00	2E-03	1E+00	3E-07	5E-01	2E-01	1E+00	1E+00	8E-02	1E+00	4E-05	1E+00	1E+00	8E-01	1E+00	1E-01	1E+00	6E-05	1E+00	1E+00
Ackley	100	1E+00	1E+00	3E-03	2E-01	6E-02	1E-01	1E-02	9E-01	1E+00	9E-05	2E-02	1E-02	1E-01	4E-04	1E+00	9E-01	1E-09	8E-07	2E-04	2E-03	1E-08
Emb. Hartmann 6	100	3E-02	4E-02	7E-05	9E-02	2E-05	3E-05	2E-01	3E-01	1E-01	2E-04	2E-01	2E-04	2E-04	5E-01	6E-01	2E-01	3E-04	6E-01	3E-04	4E-04	9E-01
Levy	100	5E-03	7E-02	6E-06	1E-03	2E-02	2E-01	5E-10	1E-01	6E-01	7E-06	1E-02	4E-02	3E-01	1E-09	1E+00	9E-01	7E-06	2E-02	4E-02	3E-01	2E-09
Powell	100	6E-02	6E-03	2E-04	9E-03	1E-05	2E-05	7E-09	7E-01	7E-02	2E-04	2E-02	1E-05	2E-05	9E-09	1E+00	1E-01	2E-04	3E-02	1E-05	2E-05	5E-08
Rastrigin	100	6E-01	8E-01	2E-01	1E+00	1E-02	4E-03	6E-02	2E-01	1E+00	2E-01	1E+00	1E-03	9E-05	3E-04	2E-01	1E+00	2E-01	1E+00	2E-03	1E-04	2E-04
Rosenbrock	100	8E-01	8E-01	1E-04	2E-01	6E-01	4E-01	4E-03	1E-01	2E-01	1E-04	1E-02	3E-01	1E-01	3E-08	9E-01	9E-01	1E-04	4E-02	4E-01	2E-01	2E-07
Styblinski-Tang	100	3E-04	1E+00	6E-03	1E+00	2E-05	5E-06	1E-01	1E+00	1E+00	4E-02	1E+00	5E-04	5E-04	9E-01	1E-01	1E+00	1E-01	1E+00	4E-03	2E-03	1E+00
Combined		1E-27	3E-04	1E+00	1E+00	6E-69	1E-44	1E-24	9E-02	1E+00	1E+00	1E+00	3E-51	9E-36	3E-16	1E-04	1E+00	1E+00	1E+00	8E-31	6E-16	3E-03

93142

D.2 Results for batch sizes 5 and 10

Table A5: BO on noise-free synthetic test problems. The normalized highest observed value after 10 rounds of BO with q=5 is shown. Colors are normalized row-wise. Higher means better. Results are means over ten replicate runs.

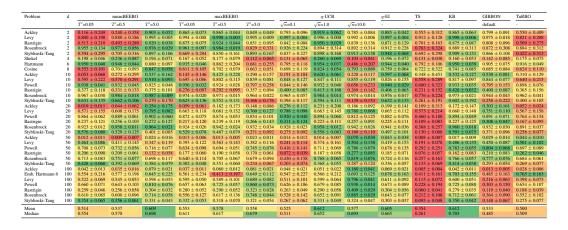


Table A6: BO on noise-free synthetic test problems. The normalized highest observed value after 10 rounds of BO with q=10 is shown. Colors are normalized row-wise. Higher means better. Results are means over ten replicate runs.

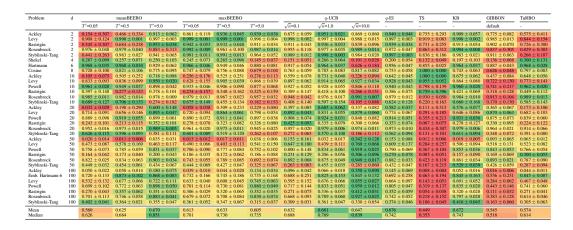


Table A7: BO on noise-free synthetic test problems. The relative batch instantaneous regret of the last, exploitative batch with q=5 is shown. Colors are normalized row-wise. Lower means better. Results are means over ten replicate runs.

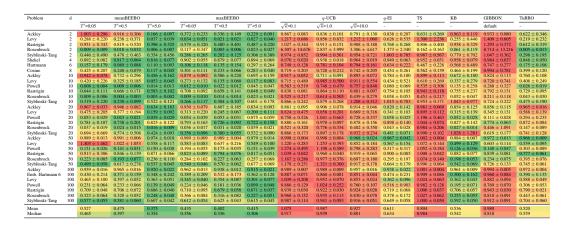


Table A8: BO on noise-free synthetic test problems. The relative batch instantaneous regret of the last, exploitative batch with q=10 is shown. Colors are normalized row-wise. Lower means better. Results are means over ten replicate runs.

Problem	d		meanBEEBO			maxBEEBO			q-UCB		q-EI	TS	KB	GIBBON	TuRBO
		T'=0.05	T'=0.5	T'=5.0	T'=0.05	T'=0.5	T'=5.0	$\sqrt{\kappa}$ =0.1	$\sqrt{\kappa}$ =1.0	$\sqrt{\kappa}$ =10.0		-	-	default	-
Ackley	2	0.935 ± 0.250	0.582 ± 0.306	0.213 ± 0.092	0.333 ± 0.167	0.221 ± 0.113	0.181 ± 0.081	0.907 ± 0.036	0.887 ± 0.061	0.885 ± 0.053	0.749 ± 0.093	0.622 ± 0.333	0.846 ± 0.240	0.958 ± 0.029	0.530 ± 0.402
Levy	2	0.157 ± 0.099	0.060 ± 0.075	0.063 ± 0.052	0.033 ± 0.027	0.087 ± 0.039	0.077 ± 0.067	1.166 ± 0.735	1.137 ± 1.309	1.049 ± 0.691	0.638 ± 0.350	0.407 ± 0.453	0.073 ± 0.113	1.996 ± 1.164	0.138 ± 0.148
Rastrigin	2	0.631 ± 0.269	0.676 ± 0.294	0.404 ± 0.209	0.566 ± 0.189	0.482 ± 0.399	0.662 ± 0.296	0.994 ± 0.255	0.927 ± 0.213	1.016 ± 0.234	0.608 ± 0.109	0.831 ± 0.427	0.708 ± 0.121	1.217 ± 0.214	0.292 ± 0.299
Rosenbrock	2	0.002 ± 0.002	0.002 ± 0.004	0.005 ± 0.004	0.003 ± 0.003	0.002 ± 0.002	0.007 ± 0.012	1.054 ± 0.835	1.335 ± 1.611	2.055 ± 4.821	0.471 ± 0.644	0.116 ± 0.228	0.009 ± 0.011	2.095 ± 3.809	0.001 ± 0.001
Styblinski-Tang	2	0.233 ± 0.080	0.124 ± 0.071	0.371 ± 0.396	0.190 ± 0.181	0.290 ± 0.353	0.234 ± 0.259	1.129 ± 0.298	1.104 ± 0.304	0.861 ± 0.295	0.930 ± 0.436	0.359 ± 0.312	0.294 ± 0.400	1.273 ± 0.469	0.261 ± 0.105
Shekel	4	0.825 ± 0.067	0.827 ± 0.039	0.783 ± 0.095	0.862 ± 0.050	0.827 ± 0.078	0.871 ± 0.039	0.983 ± 0.014	0.965 ± 0.015	0.978 ± 0.012	0.936 ± 0.031	0.968 ± 0.034	1.007 ± 0.013	0.990 ± 0.012	0.720 ± 0.138
Hartmann	6	0.189 ± 0.110	0.118 ± 0.070	0.106 ± 0.096	0.059 ± 0.080	0.061 ± 0.064	0.126 ± 0.098	0.861 ± 0.097	0.853 ± 0.079	0.775 ± 0.092	0.834 ± 0.169	0.347 ± 0.157	0.748 ± 0.293	0.876 ± 0.085	0.039 ± 0.027
Cosine	8	0.217 ± 0.084	0.159 ± 0.100	0.167 ± 0.073	0.213 ± 0.066	0.188 ± 0.064	0.270 ± 0.092	0.831 ± 0.118	0.917 ± 0.166	0.797 ± 0.215	1.055 ± 0.570	0.570 ± 0.194	0.793 ± 0.604	1.003 ± 0.159	0.200 ± 0.067
Ackley	10	0.866 ± 0.093	0.548 ± 0.223	0.277 ± 0.090	0.734 ± 0.162	0.531 ± 0.226	0.707 ± 0.123	0.833 ± 0.061	0.842 ± 0.063	0.924 ± 0.044	0.891 ± 0.102	0.859 ± 0.026	0.965 ± 0.095	0.972 ± 0.018	0.445 ± 0.082
Levy	10	0.196 ± 0.167	0.143 ± 0.123	0.024 ± 0.010	0.123 ± 0.098	0.086 ± 0.055	0.118 ± 0.064	0.936 ± 0.297	0.845 ± 0.284	0.866 ± 0.284	0.504 ± 0.386	0.591 ± 0.168	0.305 ± 0.230	0.747 ± 0.286	0.139 ± 0.060
Powell	10	0.012 ± 0.016	0.017 ± 0.015	0.018 ± 0.017	0.027 ± 0.026	0.032 ± 0.017	0.056 ± 0.033	0.885 ± 0.507	0.948 ± 0.604	0.985 ± 0.461	0.213 ± 0.204	0.682 ± 0.268	0.322 ± 0.709	0.442 ± 0.327	0.009 ± 0.007
Rastrigin	10	0.773 ± 0.187	0.775 ± 0.116	0.615 ± 0.122	0.714 ± 0.177	0.593 ± 0.117	0.598 ± 0.169	0.869 ± 0.097	0.909 ± 0.084	0.856 ± 0.128	0.719 ± 0.042	0.828 ± 0.090	0.738 ± 0.149	0.791 ± 0.067	0.423 ± 0.073
Rosenbrock	10	0.008 ± 0.008	0.006 ± 0.005	0.004 ± 0.002	0.011 ± 0.008	0.013 ± 0.007	0.034 ± 0.017	0.550 ± 0.287	0.684 ± 0.259	0.744 ± 0.178	0.155 ± 0.115	0.303 ± 0.136	0.087 ± 0.072	0.633 ± 0.537	0.005 ± 0.006
Styblinski-Tang	10	0.276 ± 0.155	0.389 ± 0.329	0.523 ± 0.135	0.257 ± 0.086	0.419 ± 0.077	0.545 ± 0.070	1.036 ± 0.164	1.079 ± 0.255	1.074 ± 0.198	1.060 ± 0.498	0.907 ± 0.125	1.064 ± 0.761	0.926 ± 0.186	0.256 ± 0.084
Ackley	20	0.957 ± 0.040	0.858 ± 0.204	0.392 ± 0.146	0.923 ± 0.048	0.674 ± 0.233	0.757 ± 0.071	0.850 ± 0.045	0.822 ± 0.088	0.899 ± 0.044	0.778 ± 0.141	0.984 ± 0.034	0.492 ± 0.078	0.978 ± 0.012	0.871 ± 0.062
Levy	20	0.232 ± 0.066	0.685 ± 0.872	0.065 ± 0.033	0.262 ± 0.209	0.099 ± 0.039	0.135 ± 0.034	0.683 ± 0.291	0.970 ± 0.274	0.794 ± 0.390	0.348 ± 0.134	0.950 ± 0.119	0.425 ± 0.183	0.431 ± 0.394	0.327 ± 0.102
Powell	20	0.078 ± 0.144	0.025 ± 0.017	0.029 ± 0.022	0.038 ± 0.027	0.030 ± 0.012	0.036 ± 0.011	0.568 ± 0.241	0.748 ± 0.472	0.603 ± 0.397	0.082 ± 0.039	0.937 ± 0.209	0.035 ± 0.014	0.126 ± 0.054	0.098 ± 0.057
Rastrigin	20	0.739 ± 0.102	0.730 ± 0.076	0.707 ± 0.074	0.732 ± 0.101	0.660 ± 0.102	0.660 ± 0.099	0.843 ± 0.119	0.865 ± 0.055	0.824 ± 0.059	0.712 ± 0.035	0.999 ± 0.053	0.747 ± 0.093	0.724 ± 0.061	0.783 ± 0.081
Rosenbrock	20	0.019 ± 0.007	0.013 ± 0.008	0.004 ± 0.002	0.018 ± 0.009	0.011 ± 0.005	0.031 ± 0.030	0.461 ± 0.335	0.784 ± 0.288	0.410 ± 0.200	0.036 ± 0.008	0.788 ± 0.119	0.018 ± 0.007	0.095 ± 0.039	0.074 ± 0.045
Styblinski-Tang	20	0.268 ± 0.074	0.288 ± 0.069	0.445 ± 0.058	0.311 ± 0.082	0.410 ± 0.117	0.679 ± 0.152	1.104 ± 0.212	1.104 ± 0.216	1.080 ± 0.296	0.628 ± 0.224	0.985 ± 0.078	0.852 ± 1.063	0.859 ± 0.657	0.523 ± 0.131
Ackley	50	0.981 ± 0.018	0.990 ± 0.005	0.960 ± 0.034	0.988 ± 0.015	0.987 ± 0.009	0.961 ± 0.011	0.960 ± 0.046	0.994 ± 0.007	0.862 ± 0.058	0.893 ± 0.085	1.001 ± 0.005	0.982 ± 0.005	0.934 ± 0.064	0.949 ± 0.010
Levy	50	0.464 ± 0.115	0.965 ± 1.096	0.445 ± 0.107	0.482 ± 0.119	0.489 ± 0.120	0.417 ± 0.127	1.023 ± 0.305	1.373 ± 0.357	0.611 ± 0.174	0.516 ± 0.309	1.032 ± 0.194	1.034 ± 1.514	0.554 ± 0.194	0.487 ± 0.055
Powell	50	0.127 ± 0.030	0.109 ± 0.018	0.076 ± 0.032	0.147 ± 0.031	0.113 ± 0.037	0.177 ± 0.076	1.177 ± 0.598	1.302 ± 0.455	0.522 ± 0.332	0.269 ± 0.418	0.972 ± 0.192	0.082 ± 0.031	0.398 ± 0.913	0.218 ± 0.062
Rastrigin	50	0.889 ± 0.124	0.858 ± 0.050	0.764 ± 0.111	0.804 ± 0.065	0.789 ± 0.045	0.776 ± 0.085	0.943 ± 0.093	0.940 ± 0.037	0.829 ± 0.063	0.826 ± 0.133	0.997 ± 0.046	1.000 ± 0.323	0.832 ± 0.041	0.696 ± 0.035
Rosenbrock	50	0.131 ± 0.030	0.123 ± 0.054	0.075 ± 0.035	0.197 ± 0.043	0.159 ± 0.073	0.249 ± 0.103	0.784 ± 0.390	1.066 ± 0.234	0.417 ± 0.180	0.107 ± 0.023	1.039 ± 0.137	0.117 ± 0.042	0.135 ± 0.056	0.285 ± 0.085
Styblinski-Tang	50	0.450 ± 0.039	0.500 ± 0.148	0.534 ± 0.238	0.467 ± 0.089	0.487 ± 0.068	0.593 ± 0.053	1.050 ± 0.210	1.313 ± 0.213	0.797 ± 0.117	0.520 ± 0.073	0.990 ± 0.055	0.633 ± 0.631	0.804 ± 0.821	0.697 ± 0.068
Ackley	100	0.950 ± 0.023	0.948 ± 0.016	0.815 ± 0.075	0.961 ± 0.010	0.954 ± 0.019	0.861 ± 0.033	0.968 ± 0.022	0.992 ± 0.010	0.914 ± 0.056	0.905 ± 0.057	0.999 ± 0.002	0.949 ± 0.015	0.992 ± 0.002	0.959 ± 0.007
Emb. Hartmann 6	100	0.265 ± 0.179	0.118 ± 0.092	0.107 ± 0.085	0.220 ± 0.151	0.199 ± 0.155	0.208 ± 0.125	0.784 ± 0.109	0.815 ± 0.127	0.756 ± 0.133	0.543 ± 0.163	0.870 ± 0.048	0.157 ± 0.175	0.941 ± 0.065	0.301 ± 0.120
Levy	100	0.402 ± 0.108	0.291 ± 0.069	0.242 ± 0.083	0.328 ± 0.058	0.351 ± 0.077	0.247 ± 0.053	1.017 ± 0.196	1.154 ± 0.164	0.976 ± 0.179	0.442 ± 0.386	0.999 ± 0.040	0.318 ± 0.120	0.821 ± 0.054	0.515 ± 0.037
Powell	100	0.189 ± 0.063	0.141 ± 0.031	0.066 ± 0.024	0.190 ± 0.067	0.166 ± 0.047	0.083 ± 0.027	0.865 ± 0.237	1.232 ± 0.274	0.861 ± 0.315	0.141 ± 0.029	0.916 ± 0.092	0.112 ± 0.020	0.727 ± 0.063	0.226 ± 0.039
Rastrigin	100	0.726 ± 0.068	0.702 ± 0.077	0.628 ± 0.042	0.709 ± 0.084	0.666 ± 0.085	0.636 ± 0.046	0.946 ± 0.057	0.946 ± 0.051	0.806 ± 0.090	0.788 ± 0.119	1.002 ± 0.027	0.730 ± 0.058	0.927 ± 0.024	0.715 ± 0.037
Rosenbrock	100	0.227 ± 0.076	0.200 ± 0.059	0.093 ± 0.040	0.255 ± 0.080	0.229 ± 0.080	0.153 ± 0.080	0.880 ± 0.167	1.060 ± 0.106	0.823 ± 0.223	0.225 ± 0.047	0.988 ± 0.070	0.164 ± 0.016	0.766 ± 0.057	0.366 ± 0.047
Styblinski-Tang	100	0.527 ± 0.052	0.568 ± 0.036	0.598 ± 0.054	0.566 ± 0.044	0.581 ± 0.036	0.625 ± 0.029	1.051 ± 0.145	1.267 ± 0.286	0.965 ± 0.017	0.667 ± 0.271	0.987 ± 0.039	0.522 ± 0.020	0.898 ± 0.033	0.677 ± 0.041
Mean		0.422	0.410	0.322	0.386	0.360	0.387	0.909	1.005	0.867	0.581	0.813	0.525	0.844	0.401
Median		0.268	0.291	0.242	0.262	0.290	0.249	0.936	0.965	0.861	0.628	0.937	0.522	0.859	0.327

D.3 Control problems

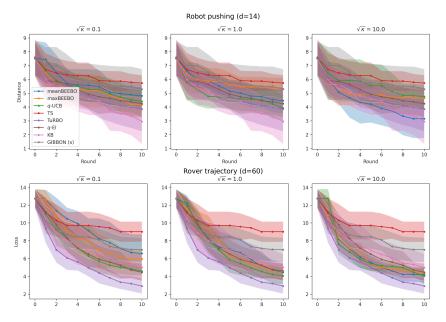


Figure A2: Experiments on the 14D robot arm pushing and 60D rover trajectory planning control problems. 10 replicates each. GIBBON (s) refers to the scaled larged-batch variant of GIBBON.

D.4 Run time

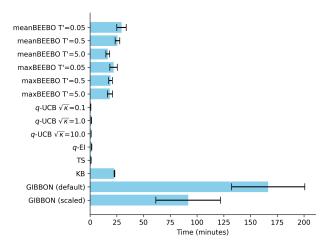


Figure A3: Example run times for the 10-round BO experiment on the 6D Hartmann problem with Q=100. Error bars are over 5 replicate runs. Run times vary depending on the test problem, with GIBBON appearing especially sensitive, becoming e.g. 10x slower on the 50D Ackley problem.

Table A9: Total run times for five replicates of the experiments presented in Table A1. We sum over all test problems.

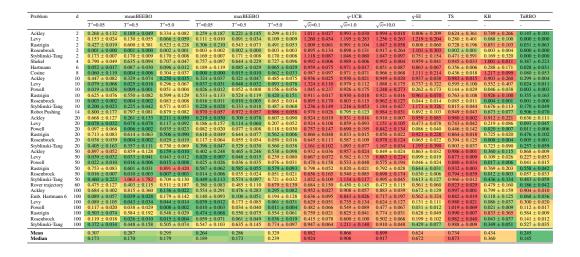
Method	Configuration	Total time [h]
meanBEEBO	T' = 0.05	66.12
meanBEEBO	T' = 0.5	47.08
meanBEEBO	T' = 5.0	37.13
maxBEEBO	T' = 0.05	54.85
maxBEEBO	T' = 0.5	44.20
maxBEEBO	T' = 5.0	47.63
q-UCB	$\sqrt{\kappa} = 0.1$	3.33
q-UCB	$\sqrt{\kappa} = 1.0$	3.70
q-UCB	$\sqrt{\kappa} = 10.0$	4.41
q-EI	-	24.33
TS	-	6.56
KB	-	223.78
GIBBON	default	3380.48
GIBBON	scaled	1055.93

D.5 Results with random initialization in round 0

Table A10: BO with random initialization on noise-free synthetic test problems. The normalized highest observed value after 10 rounds of BO with q=100 is shown. Colors are normalized row-wise. Higher means better. Results are means over five replicate runs.

Ackley 2 0 Levy 2 0 Rastrigin 2 0 Rosenbrock 2 0 Styblinski-Tang 2 1 Sheke 1 0 Hartmann 6 1 Hartmann 6 1 Levy 10 0 Powell 10 0 Rosenbrock 10 0 Rosenbrock 10 0 Rosenbrock 10 0 Rosenbrock 2 Styblinski-Tang 10 0 Robo Pushing 14 0 R	T'=0.05 0.994 ± 0.011 0.995 ± 0.008 0.838 ± 0.190 0.893 ± 0.080 1.000 ± 0.001 0.839 ± 0.311 1.000 ± 0.000 0.935 ± 0.026	T'=0.5 0.971 ± 0.041 0.997 ± 0.002 0.674 ± 0.433 0.900 ± 0.103 1.000 ± 0.000 0.829 ± 0.360 1.000 ± 0.000	T'=5.0 0.988 ± 0.006 0.995 ± 0.005 0.836 ± 0.138 0.605 ± 0.395 1.000 ± 0.000	T'=0.05 0.985 ± 0.012 0.997 ± 0.002 0.784 ± 0.350 0.525 ± 0.496	T'=0.5 0.977 ± 0.029 0.998 ± 0.002 0.756 ± 0.356	T'=5.0 0.912 ± 0.069 0.966 ± 0.044	$\sqrt{\kappa}$ =0.1 0.793 ± 0.255	$\sqrt{\kappa}$ =1.0 0.904 ± 0.059	$\sqrt{\kappa}$ =10.0 0.968 ± 0.019	0.957 ± 0.056	1.000 ± 0.000	0.928 ± 0.076	
Levy 2 0 Rastrigin 2 0 Rosenbrock 2 0 Rosenbrock 2 1 Shekel 4 0 Hartmann 6 1 Cosine 8 1 Ackley 10 0 Levy 10 0 Rosenbrock 10 0 Rosenbrock 10 0 Rosenbrock 2 10 0 Robor Pushing 14 0 Robor Pushing 14 0 Robor Pushing 14 0 Robor Pushing 14 0	0.995 ± 0.008 0.838 ± 0.190 0.893 ± 0.080 1.000 ± 0.001 0.839 ± 0.311 1.000 ± 0.000 1.000 ± 0.000	0.997 ± 0.002 0.674 ± 0.433 0.900 ± 0.103 1.000 ± 0.000 0.829 ± 0.360	0.995 ± 0.005 0.836 ± 0.138 0.605 ± 0.395 1.000 ± 0.000	0.997 ± 0.002 0.784 ± 0.350	0.998 ± 0.002			0.904 ± 0.059	0.968 ± 0.019	0.957 ± 0.056	1.000 ± 0.000	0.029 ± 0.076	0.072 0.027
Rastrigin 2 0 Nosenbrock 2 0 Styblinski-Tang 2 1 Shekel 4 0 Cosine 8 1 Cosine 8 Ackley 10 0 Nosenbrock 10 Nose	0.838 ± 0.190 0.893 ± 0.080 1.000 ± 0.001 0.839 ± 0.311 1.000 ± 0.000 1.000 ± 0.000	0.674 ± 0.433 0.900 ± 0.103 1.000 ± 0.000 0.829 ± 0.360	0.836 ± 0.138 0.605 ± 0.395 1.000 ± 0.000	0.784 ± 0.350		0.966 ± 0.044							0.972 ± 0.027
Rosenbrock 2 0 1 1 1 1 1 1 1 1 1	0.893 ± 0.080 1.000 ± 0.001 0.839 ± 0.311 1.000 ± 0.000 1.000 ± 0.000	0.900 ± 0.103 1.000 ± 0.000 0.829 ± 0.360	0.605 ± 0.395 1.000 ± 0.000		0.756 ± 0.256		0.983 ± 0.023	0.985 ± 0.010	0.954 ± 0.047	0.994 ± 0.004	0.993 ± 0.004	0.997 ± 0.003	0.992 ± 0.017
Styblinski-Tang 2 1	1.000 ± 0.001 0.839 ± 0.311 1.000 ± 0.000 1.000 ± 0.000	1.000 ± 0.000 0.829 ± 0.360	1.000 ± 0.000	0.525 ± 0.496		0.847 ± 0.199	0.540 ± 0.359	0.675 ± 0.250	0.296 ± 0.255	0.873 ± 0.173	0.800 ± 0.447	0.691 ± 0.299	0.902 ± 0.220
Shekel 4 0 Hartmann 6 1 Cosine 8 1 Cosine 8 1 Cosine 10 0 Cevy 10 Cevy	0.839 ± 0.311 1.000 ± 0.000 1.000 ± 0.000	0.829 ± 0.360			0.705 ± 0.413	0.242 ± 0.306		0.475 ± 0.452	0.469 ± 0.485	0.728 ± 0.304	0.888 ± 0.249	0.753 ± 0.423	0.875 ± 0.266
Hartmann	1.000 ± 0.000 1.000 ± 0.000			0.998 ± 0.002	0.999 ± 0.002	0.999 ± 0.002	0.998 ± 0.001	0.999 ± 0.002	0.993 ± 0.008	0.999 ± 0.001	0.998 ± 0.002	1.000 ± 0.000	1.000 ± 0.000
Cosine	1.000 ± 0.000	1.000 ± 0.000	0.706 ± 0.326	0.390 ± 0.353	0.476 ± 0.274	0.388 ± 0.316	0.183 ± 0.068	0.376 ± 0.308	0.266 ± 0.073	0.530 ± 0.350	0.087 ± 0.060	0.178 ± 0.071	0.824 ± 0.247
Ackley 10 0 Levy 10 0 Powell 10 0 Rastrigin 10 0 Rosenbrock 10 0 Styblinski-Tang 10 0 Robot Pushing 14 0 Ackley 20 0			0.984 ± 0.016	0.955 ± 0.063	0.989 ± 0.024	0.992 ± 0.010	0.960 ± 0.052	0.998 ± 0.002	0.919 ± 0.057	0.991 ± 0.020	0.715 ± 0.150	0.998 ± 0.001	0.965 ± 0.045
Levy 10 0 0 10 0 10 10 10 10 10 10 10 10 10	0.935 ± 0.026	0.997 ± 0.002	0.377 ± 0.144	0.999 ± 0.000	0.968 ± 0.023	0.871 ± 0.057	0.906 ± 0.079	0.903 ± 0.050	0.390 ± 0.242	0.787 ± 0.118	1.000 ± 0.000	0.987 ± 0.009	0.912 ± 0.066
Powell		0.904 ± 0.039	0.820 ± 0.037	0.816 ± 0.037	0.742 ± 0.052	0.507 ± 0.137	0.790 ± 0.015	0.776 ± 0.038	0.561 ± 0.170	0.782 ± 0.036	1.000 ± 0.000	0.789 ± 0.023	0.784 ± 0.026
Rastrigin 10 0 Rosenbrock 10 0 Styblinski-Tang 10 0 Robot Pushing 14 0 Ackley 20 0	0.981 ± 0.013	0.957 ± 0.028	0.930 ± 0.034	0.956 ± 0.026	0.926 ± 0.024	0.919 ± 0.022	0.866 ± 0.036	0.839 ± 0.068	0.813 ± 0.126	0.920 ± 0.064	0.942 ± 0.026	0.949 ± 0.041	0.941 ± 0.072
Rosenbrock 10 0 Styblinski-Tang 10 0 Robot Pushing 14 0 Ackley 20 0	0.971 ± 0.029	0.953 ± 0.030	0.672 ± 0.387	0.923 ± 0.109	0.900 ± 0.103	0.702 ± 0.398	0.886 ± 0.061	0.822 ± 0.161	0.219 ± 0.302	0.851 ± 0.191	0.833 ± 0.169	0.939 ± 0.058	0.985 ± 0.008
Styblinski-Tang 10 0 Robot Pushing 14 0 Ackley 20 0	0.465 ± 0.144	0.495 ± 0.180	0.570 ± 0.091	0.526 ± 0.087	0.516 ± 0.140	0.642 ± 0.083	0.394 ± 0.143	0.564 ± 0.152	0.193 ± 0.172	0.441 ± 0.065	1.000 ± 0.000	0.394 ± 0.088	0.672 ± 0.165
Robot Pushing 14 0 Ackley 20 0	0.992 ± 0.006	0.990 ± 0.002	0.865 ± 0.068	0.990 ± 0.003	0.986 ± 0.007	0.966 ± 0.037	0.965 ± 0.043	0.952 ± 0.019	0.220 ± 0.373	0.975 ± 0.017	0.820 ± 0.029	0.992 ± 0.003	0.993 ± 0.004
Ackley 20 0	0.815 ± 0.051	0.814 ± 0.062	0.245 ± 0.050	0.784 ± 0.025	0.643 ± 0.124	0.217 ± 0.157	0.165 ± 0.115	0.584 ± 0.177	0.028 ± 0.062	0.619 ± 0.093	0.399 ± 0.163	0.827 ± 0.066	0.654 ± 0.085
	0.350 ± 0.121	0.377 ± 0.124	0.560 ± 0.172	0.425 ± 0.107	0.310 ± 0.140	0.395 ± 0.131	0.424 ± 0.154	0.522 ± 0.170	0.379 ± 0.093	0.417 ± 0.160	0.247 ± 0.118	0.694 ± 0.255	0.518 ± 0.149
Levy 20 0	0.843 ± 0.016	0.857 ± 0.028	0.789 ± 0.048	0.819 ± 0.030	0.788 ± 0.040	0.390 ± 0.106	0.706 ± 0.063	0.775 ± 0.039	0.460 ± 0.052	0.740 ± 0.054	1.000 ± 0.000	0.763 ± 0.048	0.438 ± 0.103
	0.936 ± 0.035	0.939 ± 0.019	0.896 ± 0.044	0.953 ± 0.027	0.901 ± 0.046	0.911 ± 0.038	0.928 ± 0.016	0.929 ± 0.029	0.768 ± 0.048	0.921 ± 0.024	0.979 ± 0.003	0.956 ± 0.013	0.925 ± 0.041
Powell 20 0	0.947 ± 0.036	0.966 ± 0.013	0.840 ± 0.111	0.936 ± 0.019	0.880 ± 0.100	0.908 ± 0.076	0.946 ± 0.007	0.928 ± 0.050	0.819 ± 0.122	0.926 ± 0.047	0.964 ± 0.014	0.969 ± 0.016	0.957 ± 0.036
Rastrigin 20 0	0.373 ± 0.042	0.462 ± 0.049	0.518 ± 0.054	0.514 ± 0.059	0.480 ± 0.077	0.491 ± 0.087	0.450 ± 0.115	0.463 ± 0.069	0.383 ± 0.094	0.451 ± 0.053	1.000 ± 0.000	0.481 ± 0.068	0.523 ± 0.064
Rosenbrock 20 0	0.992 ± 0.004	0.993 ± 0.004	0.920 ± 0.041	0.991 ± 0.005	0.982 ± 0.013	0.923 ± 0.054	0.967 ± 0.018	0.984 ± 0.005	0.915 ± 0.044	0.984 ± 0.007	0.939 ± 0.020	0.994 ± 0.003	0.993 ± 0.002
Styblinski-Tang 20 0	0.706 ± 0.061	0.669 ± 0.113	0.305 ± 0.198	0.607 ± 0.086	0.417 ± 0.122	0.279 ± 0.077	0.204 ± 0.130	0.524 ± 0.184	0.054 ± 0.074	0.639 ± 0.082	0.271 ± 0.262	0.665 ± 0.129	0.604 ± 0.088
Ackley 50 0	0.221 ± 0.293	0.146 ± 0.116	0.842 ± 0.010	0.622 ± 0.243	0.705 ± 0.042	0.457 ± 0.098	0.627 ± 0.053	0.739 ± 0.034	0.736 ± 0.020	0.727 ± 0.051	1.000 ± 0.000	0.683 ± 0.122	0.175 ± 0.022
Levy 50 0	0.976 ± 0.010	0.978 ± 0.012	0.943 ± 0.021	0.977 ± 0.012	0.955 ± 0.018	0.867 ± 0.013	0.952 ± 0.016	0.966 ± 0.025	0.933 ± 0.007	0.943 ± 0.017	0.987 ± 0.002	0.926 ± 0.041	0.793 ± 0.054
Powell 50 0	0.940 ± 0.036	0.978 ± 0.010	0.959 ± 0.025	0.976 ± 0.009	0.970 ± 0.014	0.929 ± 0.024	0.965 ± 0.014	0.958 ± 0.016	0.978 ± 0.007	0.964 ± 0.013	0.985 ± 0.004	0.957 ± 0.022	0.920 ± 0.039
Rastrigin 50 0	0.273 ± 0.156	0.505 ± 0.031	0.453 ± 0.040	0.473 ± 0.042	0.466 ± 0.016	0.445 ± 0.027	0.466 ± 0.073	0.418 ± 0.023	0.503 ± 0.047	0.468 ± 0.012	1.000 ± 0.000	0.423 ± 0.085	0.459 ± 0.070
Rosenbrock 50 0	0.976 ± 0.012	0.985 ± 0.003	0.988 ± 0.005	0.978 ± 0.011	0.981 ± 0.005	0.968 ± 0.016	0.974 ± 0.004	0.987 ± 0.003	0.984 ± 0.013	0.981 ± 0.003	0.979 ± 0.003	0.987 ± 0.003	0.968 ± 0.018
Styblinski-Tang 50 0	0.605 ± 0.067	0.693 ± 0.038	0.417 ± 0.163	0.536 ± 0.072	0.415 ± 0.055	0.332 ± 0.037	0.341 ± 0.082	0.716 ± 0.031	0.371 ± 0.040	0.690 ± 0.069	0.254 ± 0.220	0.720 ± 0.041	0.499 ± 0.073
Rover trajectory 60 0	0.448 ± 0.209	0.678 ± 0.029	0.708 ± 0.060	0.533 ± 0.066	0.657 ± 0.104	0.629 ± 0.040	0.626 ± 0.076	0.613 ± 0.070	0.665 ± 0.079	0.635 ± 0.074	0.265 ± 0.069	0.616 ± 0.043	0.764 ± 0.074
Ackley 100 0	0.310 ± 0.408	0.347 ± 0.452	0.864 ± 0.023	0.526 ± 0.335	0.536 ± 0.293	0.707 ± 0.086	0.696 ± 0.015	0.721 ± 0.080	0.850 ± 0.007	0.747 ± 0.071	0.007 ± 0.005	0.289 ± 0.081	0.110 ± 0.014
Emb. Hartmann 6 100 0	0.980 ± 0.009	0.988 ± 0.008	0.916 ± 0.101	0.982 ± 0.016	0.933 ± 0.057	0.914 ± 0.051	0.941 ± 0.035	0.915 ± 0.038	0.913 ± 0.116	0.922 ± 0.110	0.554 ± 0.315	0.949 ± 0.065	0.931 ± 0.084
Levy 100 0	0.890 ± 0.150	0.966 ± 0.024	0.943 ± 0.019	0.962 ± 0.012	0.942 ± 0.017	0.946 ± 0.013	0.952 ± 0.005	0.937 ± 0.029	0.964 ± 0.014	0.943 ± 0.030	0.310 ± 0.382	0.908 ± 0.056	0.692 ± 0.013
	0.786 ± 0.051	0.929 ± 0.099	0.985 ± 0.004	0.985 ± 0.002	0.981 ± 0.009	0.981 ± 0.003	0.983 ± 0.004	0.978 ± 0.013	0.983 ± 0.008	0.963 ± 0.018	0.288 ± 0.165	0.967 ± 0.021	0.860 ± 0.027
Rastrigin 100 0	0.522 ± 0.027	0.367 ± 0.194	0.467 ± 0.027	0.481 ± 0.057	0.479 ± 0.007	0.469 ± 0.037	0.442 ± 0.021	0.442 ± 0.047	0.432 ± 0.020	0.493 ± 0.014	0.238 ± 0.426	0.674 ± 0.126	0.394 ± 0.022
	0.810 ± 0.029	0.972 ± 0.012	0.976 ± 0.008	0.928 ± 0.119	0.978 ± 0.006	0.975 ± 0.008	0.977 ± 0.008	0.968 ± 0.012	0.985 ± 0.008	0.974 ± 0.003	0.270 ± 0.406	0.943 ± 0.058	0.857 ± 0.010
	0.564 ± 0.034	0.470 ± 0.152	0.396 ± 0.079	0.432 ± 0.043	0.331 ± 0.061	0.309 ± 0.022	0.321 ± 0.027	0.542 ± 0.130	0.280 ± 0.017	0.591 ± 0.023	0.198 ± 0.273	0.593 ± 0.063	0.412 ± 0.034
Median 0	0.776	0.793 0.929	0.751	0.779	0.762	0.697 0.847	0.714 0.793	0.768 0.822	0.618	0.788	0.692 0.888	0.788	0.750 0.857

Table A11: BO with random initialization on noise-free synthetic test problems. The relative batch instantaneous regret of the last, exploitative batch is shown. Colors are normalized row-wise. Lower means better. Results are means over five replicate runs.



D.6 BO curves for all experiments in Table 2 and Table A1

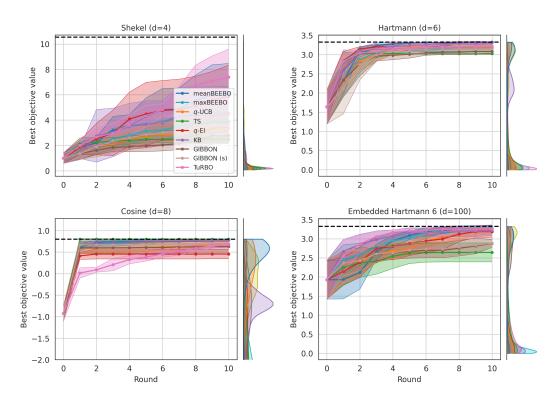


Figure A4: Experiments on the Shekel, Hartmann, Cosine and embedded Hartmann test functions with $\sqrt{\kappa}=0.1$ for BEEBO and q-UCB.

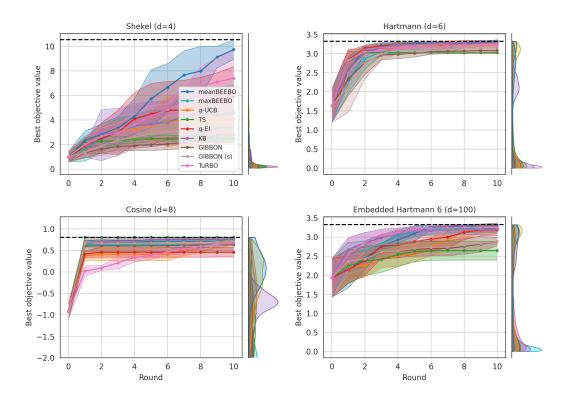


Figure A5: Experiments on the Shekel, Hartmann, Cosine and embedded Hartmann test functions with $\sqrt{\kappa} = 1.0$ for BEEBO and q-UCB.

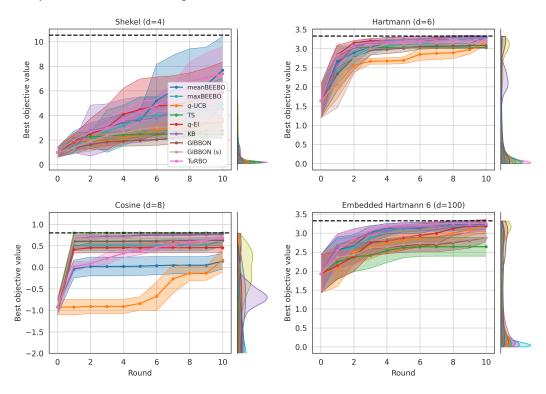


Figure A6: Experiments on the Shekel, Hartmann, Cosine and embedded Hartmann test functions with $\sqrt{\kappa}=10.0$ for BEEBO and q-UCB.

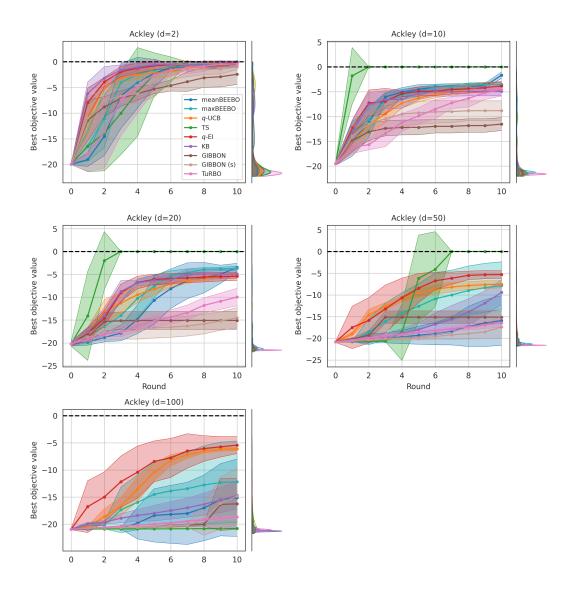


Figure A7: Experiments on the Ackley test function with $\sqrt{\kappa}=0.1$ for BEEBO and q-UCB.

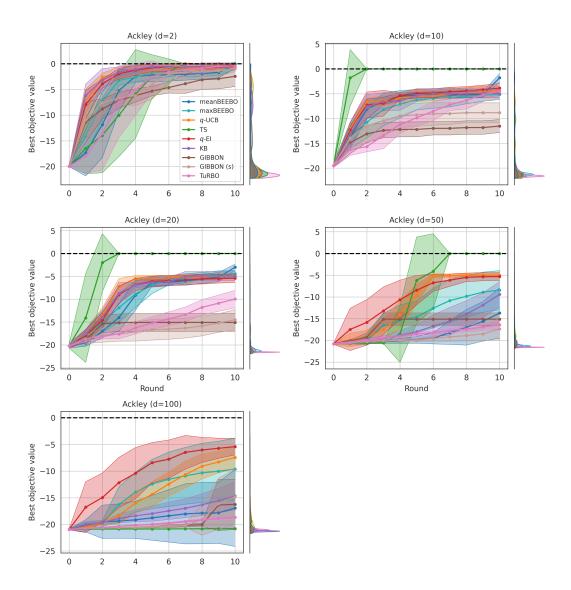


Figure A8: Experiments on the Ackley test function with $\sqrt{\kappa}=1.0$ for BEEBO and q-UCB.

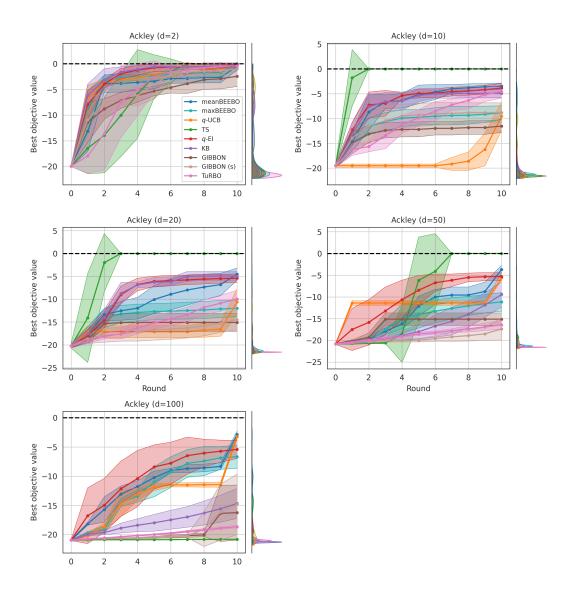


Figure A9: Experiments on the Ackley test function with $\sqrt{\kappa}=10.0$ for BEEBO and q-UCB.

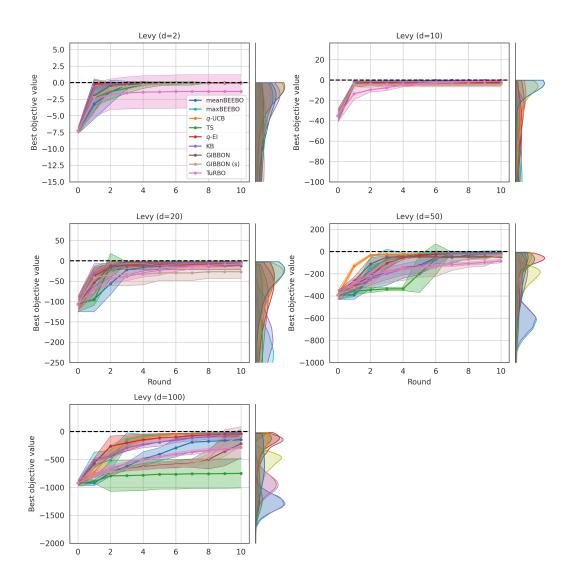


Figure A10: Experiments on the Levy test function with $\sqrt{\kappa}=0.1$ for BEEBO and q-UCB.

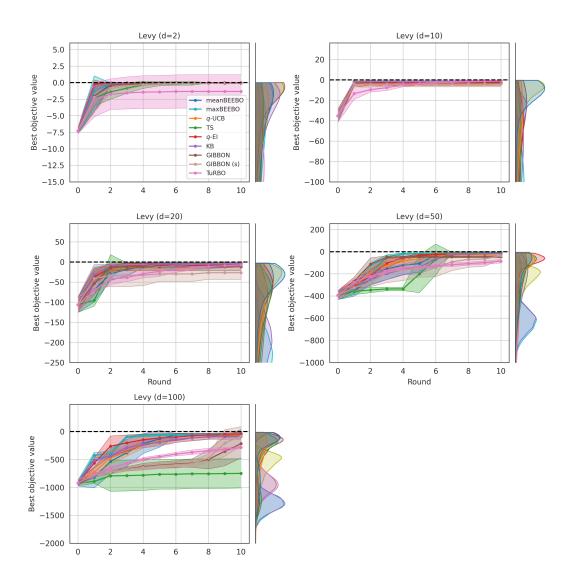


Figure A11: Experiments on the Levy test function with $\sqrt{\kappa}=1.0$ for BEEBO and q-UCB.

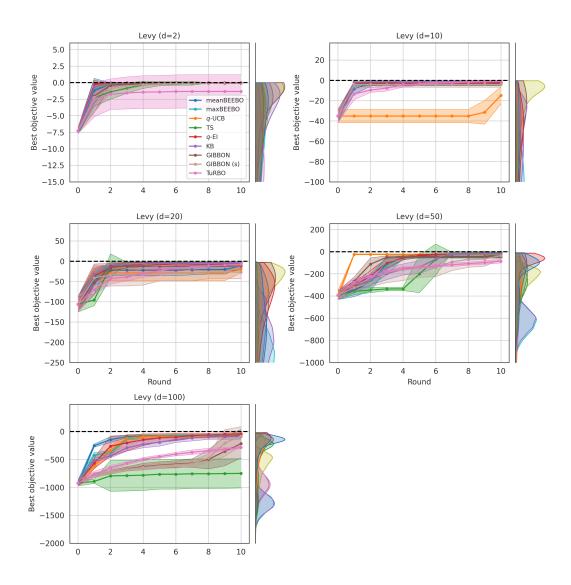


Figure A12: Experiments on the Levy test function with $\sqrt{\kappa}=10.0$ for BEEBO and q-UCB.

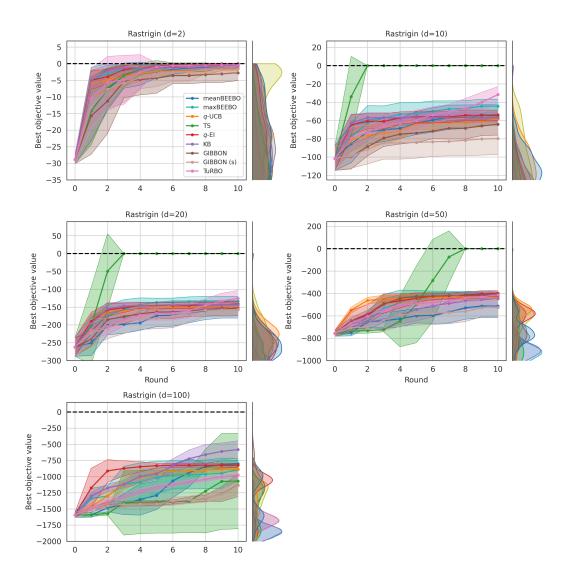


Figure A13: Experiments on the Rastrigin test function with $\sqrt{\kappa}=0.1$ for BEEBO and q-UCB.

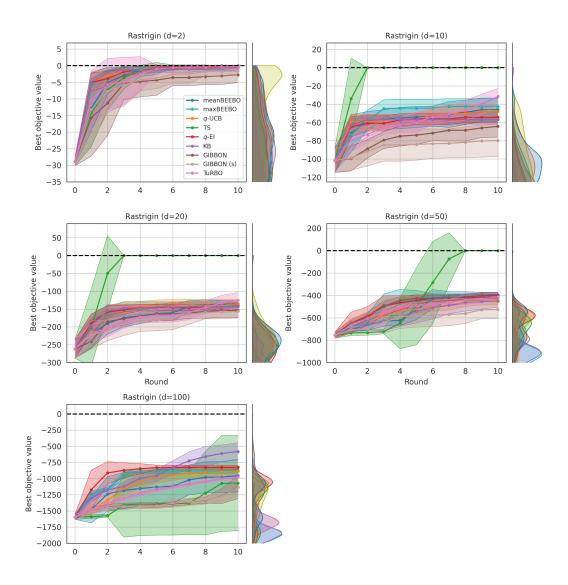


Figure A14: Experiments on the Rastrigin test function with $\sqrt{\kappa}=1.0$ for BEEBO and q-UCB.

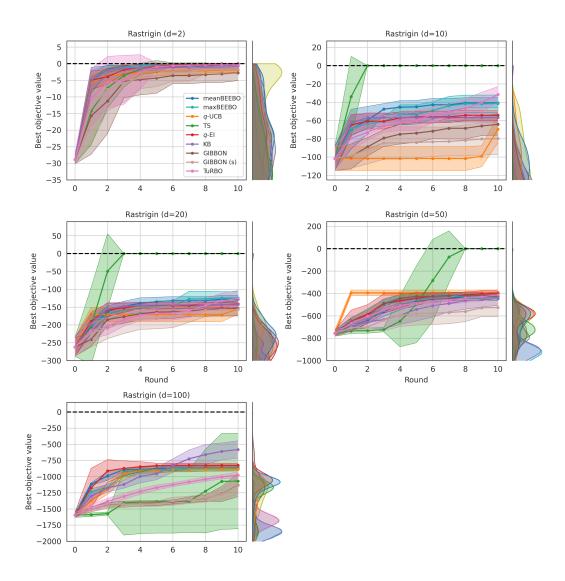


Figure A15: Experiments on the Rastrigin test function with $\sqrt{\kappa}=10.0$ for BEEBO and q-UCB.

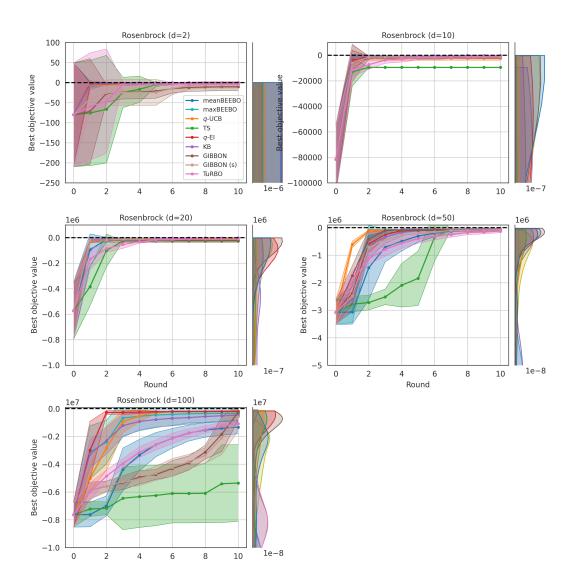


Figure A16: Experiments on the Rosenbrock test function with $\sqrt{\kappa}=0.1$ for BEEBO and q-UCB.

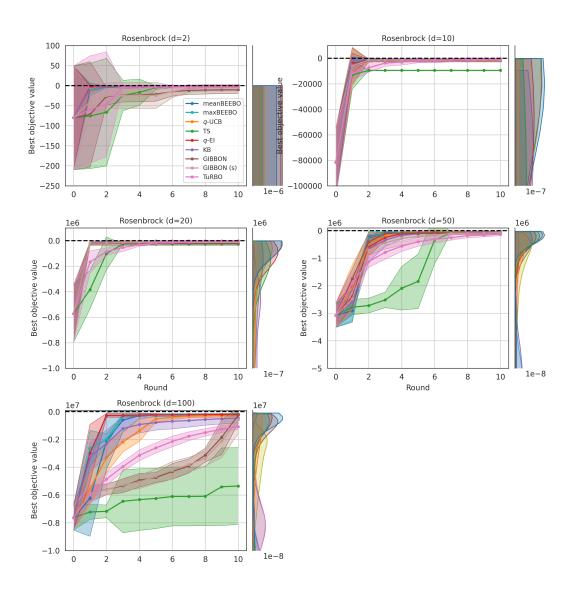


Figure A17: Experiments on the Rosenbrock test function with $\sqrt{\kappa}=1.0$ for BEEBO and q-UCB.

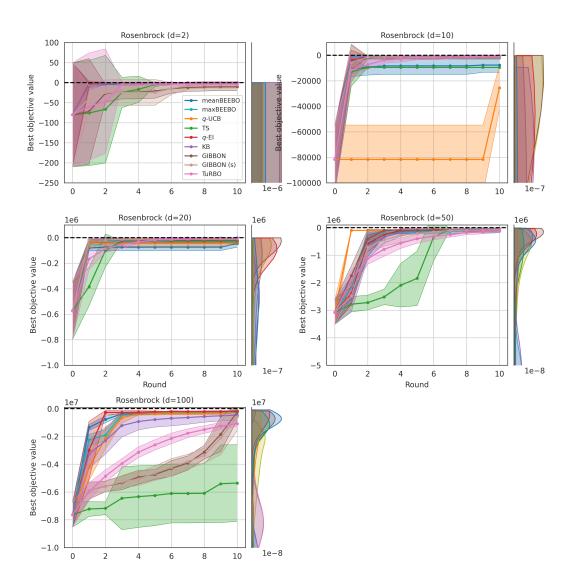


Figure A18: Experiments on the Rosenbrock test function with $\sqrt{\kappa}=10.0$ for BEEBO and q-UCB.

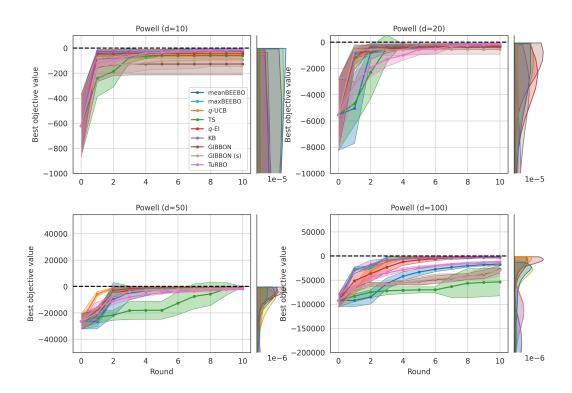


Figure A19: Experiments on the Powell test function with $\sqrt{\kappa}=0.1$ for BEEBO and q-UCB.

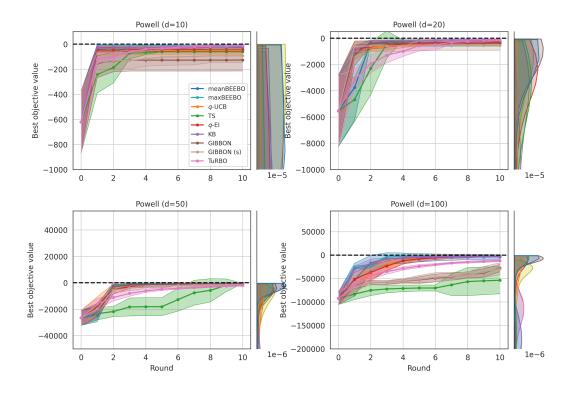


Figure A20: Experiments on the Powell test function with $\sqrt{\kappa}=1.0$ for BEEBO and q-UCB.

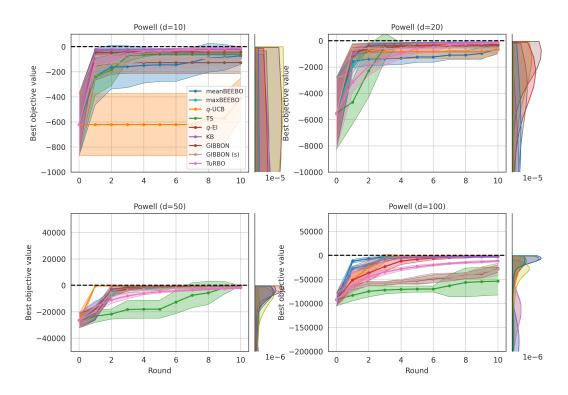


Figure A21: Experiments on the Powell test function with $\sqrt{\kappa}=10.0$ for BEEBO and q-UCB.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: As claimed, we experimentally demonstrate a) the controllability of the acquisition strategy, b) competitive performance on 33 test problems compared to q-UCB, q-EI, Thompson sampling, GIBBON, TuRBO and Kriging Believer, and c) behaviour under heteroskedastic noise.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the
 contributions made in the paper and important assumptions and limitations. A No or
 NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We address limitations in our discussion section, highlighting computational complexity constraints in exact GP inference as well as challenges under heteroskedastic noise.

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was
 only tested on a few datasets or with a few runs. In general, empirical results often
 depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not make use of any theoretical results. All reported results are based on empirical experiments. All underlying assumptions are standard in research on BO with GPs.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: As described in the methods section, we use standard BoTorch and GPyTorch utilities for all our experiments, and provide extended details on the technical implementation in the supplementary section. Our repository includes the full benchmarking setup with appropriate run scripts and instructions.

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
- (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The repository includes the implementation of the proposed method as well as the benchmarking setup with alternative methods. No additional data is required for reproduction.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how
 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We follow GPyTorch and BoTorch for all hyperparameters pertaining to GPs, and describe this accordingly. Our appendix includes additional details on method hyperparameters to ensure reproducibility.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We include full BO curves with standard deviations over five replicates for all quantitative experiments in the appendix. These detailed curves are referenced in the main text at the appropriate place.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We list the used hardware and total GPU hours in the supplement and provide example timings for experiment runtimes.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: The paper does not make use of human participants or datasets. To the best of our understanding, there are no potential harmful consequences and wider negative societal impact expected from the proposed method.

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a
 deviation from the Code of Ethics.

The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: The paper introduces a method for Bayesian optimization (BO). While BO has widespread applications in the sciences and engineering, there is no direct societal impact expected from this contribution.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper does not introduce any trained models or novel data.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We credit the GPyTorch and BoTorch packages that our codebase builds upon. The packages are used as dependencies, and as such are not included directly as assets.

Guidelines

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the
 package should be provided. For popular datasets, paperswithcode.com/datasets
 has curated licenses for some datasets. Their licensing guide can help determine the
 license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The implementation of BEEBO constitutes the only asset, which follows BoTorch APIs and has a README file demonstrating its application when working in BoTorch.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: None of the above are included in this paper.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The presented paper does not involve any human subjects.

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.