
Parallel Bayesian Optimization of Multiple Noisy Objectives with Expected Hypervolume Improvement

Samuel Daulton
Facebook, University of Oxford
sdaulton@fb.com

Maximilian Balandat
Facebook
balandat@fb.com

Eytan Bakshy
Facebook
ebakshy@fb.com

Abstract

Optimizing multiple competing black-box objectives is a challenging problem in many fields, including science, engineering, and machine learning. Multi-objective Bayesian optimization (MOBO) is a sample-efficient approach for identifying the optimal trade-offs between the objectives. However, many existing methods perform poorly when the observations are corrupted by noise. We propose a novel acquisition function, NEHVI, that overcomes this important practical limitation by applying a Bayesian treatment to the popular expected hypervolume improvement (EHVI) criterion and integrating over this uncertainty in the Pareto frontier. We argue that, even in the noiseless setting, generating multiple candidates in parallel is an incarnation of EHVI with uncertainty in the Pareto frontier and therefore can be addressed using the same underlying technique. Through this lens, we derive a natural parallel variant, q NEHVI, that reduces computational complexity of parallel EHVI from *exponential* to *polynomial* with respect to the batch size. q NEHVI is one-step Bayes-optimal for hypervolume maximization in both noisy and noiseless environments, and we show that it can be optimized effectively with gradient-based methods via sample average approximation. Empirically, we demonstrate not only that q NEHVI is substantially more robust to observation noise than existing MOBO approaches, but also that it achieves state-of-the-art optimization performance and competitive wall-times in large-batch environments.

1 Introduction

Black-box optimization problems that involve multiple competing noisy objectives are ubiquitous in science and engineering. For example, a real-time communications service may be interested in tuning the parameters of a control policy to adapt video quality in real time in order to maximize video quality and minimize latency [10, 17]. In robotics, scientists may seek to design hardware components that maximize locomotive speed and minimize energy expended [8, 38]. In agriculture, development agencies may seek to balance crop yield and environmental impact [28]. For such multi-objective optimization (MOO) problems, there typically is no single solution that is best with respect to all objectives. Rather, the goal is to identify the *Pareto frontier*: a set of optimal trade-offs such that improving one objective means deteriorating another. In many cases, the objectives are expensive to evaluate. For instance, randomized trials used in agriculture and the internet industry may take weeks or months to conduct and incur opportunity costs, and manufacturing and testing hardware is both costly and time-consuming. Therefore, it is imperative to be able to identify good trade-offs with as few objective evaluations as possible.

Bayesian optimization (BO), a method for efficient global black-box optimization, is often used to tackle such problems. BO employs a probabilistic surrogate model in conjunction with an acquisition function to navigate the trade-off between exploration (evaluating designs with high uncertainty) and exploitation (evaluating designs that are believed to be optimal). Although a significant number of works have explored multi-objective Bayesian optimization (MOBO), most available methods

[3, 39, 51, 60] do not take into account the fact that, in practice, observations are often subject to noise. For example, results of an A/B test are highly variable due to heterogeneity in the underlying user population and other factors. Agricultural trials are affected by the stochastic nature of plant growth and environmental factors such as soil composition or wind currents. In robotics, devices are subject to manufacturing tolerances, and observations of quantities such as locomotive speed and efficiency may be corrupted by measurement error from noisy sensors and environmental factors such as temperature or surface friction. While previous work has shown that a principled treatment of noisy observations can significantly improve optimization performance in the single-objective case [24, 37], this issue is understudied in the multi-objective setting. Furthermore, many applications in which evaluations take a long time require evaluating large *batches* of candidates in parallel in order to achieve reasonable throughput. For example, when firms optimize systems via A/B tests, it may take several weeks to test any particular configuration. Because of this, large batches of candidate policies are tested simultaneously [36]. In biochemistry and materials design, dozens of tests can be conducted parallel on a single microplate [63]. Even in sophisticated high throughput chemistry settings, these batches may take several hours or days to set up and evaluate [42]. Most existing MOBO methods, however, are either designed for purely sequential optimization [3, 51] or do not scale well to large batch sizes [11].

Contributions: In this work, we propose a novel MOBO algorithm, based on expected hypervolume improvement (EHVI), that scales to highly parallel evaluations of noisy objectives. Our approach is made possible by a general-purpose, differentiable, cached box decomposition (CBD) implementation that dramatically speeds up critical computations needed to account for uncertainty introduced by noisy observations and generate new candidate points for highly parallel batch or asynchronous evaluation. In particular, our CBD-based approach solves the fundamental problem of scaling parallel EHVI-based methods to large batch sizes, *reducing time and space complexity from exponential to polynomial*. Our proposed algorithm, *noisy* expected hypervolume improvement (NEHVI), is the one-step Bayes-optimal policy for hypervolume improvement and provides state-of-the-art performance across a variety of benchmarks. To our knowledge, our work provides the most extensive evaluation of noisy parallel MOBO to date. A high-quality implementation of q NEHVI, as well as many of the baselines considered here, will be made available as open-source software upon publication.

2 Preliminaries

Our goal is to find the set of optimal designs \mathbf{x} over a bounded set $\mathcal{X} \subset \mathbb{R}^d$ that maximize one or more objectives $\mathbf{f}(\mathbf{x}) \in \mathbb{R}^M$, with no known analytical expression nor gradient information of \mathbf{f} .

Multi-Objective Optimization (MOO) aims to identify the set of *Pareto optimal* objective trade-offs. We say a solution $\mathbf{f}(\mathbf{x}) = [f^{(1)}(\mathbf{x}), \dots, f^{(M)}(\mathbf{x})]$ *dominates* another solution $\mathbf{f}(\mathbf{x}) \succ \mathbf{f}(\mathbf{x}')$ if $f^{(m)}(\mathbf{x}) \geq f^{(m)}(\mathbf{x}')$ for $m = 1, \dots, M$ and $\exists m \in \{1, \dots, M\}$ s.t. $f^{(m)}(\mathbf{x}) > f^{(m)}(\mathbf{x}')$. We define the *Pareto frontier* as $\mathcal{P}^* = \{\mathbf{f}(\mathbf{x}) : \mathbf{x} \in \mathcal{X}, \nexists \mathbf{x}' \in \mathcal{X} \text{ s.t. } \mathbf{f}(\mathbf{x}') \succ \mathbf{f}(\mathbf{x})\}$, and denote the set of Pareto optimal designs as $\mathcal{X}^* = \{\mathbf{x} : \mathbf{f}(\mathbf{x}) \in \mathcal{P}^*\}$. Since the Pareto frontier (PF) is often an infinite set of points, MOO algorithms usually aim to identify a finite approximate PF \mathcal{P} . A natural measure of the quality of a PF is the hypervolume of the region of objective space that is dominated by the PF and bounded from below by a reference point. Provided with the approximate PF, the decision-maker can select a particular Pareto optimal trade-off according to their preferences.

Bayesian Optimization (BO) is a sample-efficient optimization method that leverages a probabilistic surrogate model to make principled decisions to balance exploration and exploitation [19, 50]. Typically, the surrogate is a Gaussian Process (GP), a flexible, non-parametric model known for its well-calibrated predictive uncertainty [47]. To decide which points to evaluate next, BO employs an acquisition function $\alpha(\cdot)$ that specifies the value of evaluating a set of new points \mathbf{x} based on the surrogate’s predictive distribution at \cdot . While evaluating the true black-box function \mathbf{f} is time-consuming or costly, evaluating the surrogate is cheap and relatively fast; therefore, numerical optimization can be used to find the maximizer of the acquisition function $\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} \alpha(\mathbf{x})$ to evaluate next on the black-box function. BO sequentially selects new points to evaluate and updates the model to incorporate the new observations.

Evolutionary algorithms (EAs) such as NSGA-II [12] are a popular choice for solving MOO problems (see Zitzler et al. [67] for a review of various other approaches). However, EAs generally suffer from high sample complexity, rendering them infeasible for optimizing expensive-to-evaluate black-box

functions. Multi-objective Bayesian optimization (MOBO), which combines a Bayesian surrogate with an acquisition function designed for MOO, provides a much more sample-efficient alternative.

3 Related Work

Methods based on hypervolume improvement (HVI) seek to expand the volume of the objective space dominated by the Pareto frontier. Expected hypervolume improvement (EHVI) [16] is a natural extension of the popular expected improvement (EI) [29] acquisition function to the MOO setting. Recent work has led to efficient computational paradigms using box decomposition algorithms [59] and practical enhancements such as support for parallel candidate generation and gradient-based acquisition optimization [11, 58]. However, EHVI still suffers from some limitations, including (i) the assumption that observations are noise-free, and (ii) the exponential scaling of its batch variant, q EHVI, in the batch size q , which precludes large-batch optimization. DGEMO [39] is a recent method for parallel MOBO that greedily maximizes HVI while balancing the diversity of the design points being sampled. Although DGEMO scales well to large batch sizes, it does not account for noisy observations. TSEMO [5] is a Thompson sampling (TS) heuristic that can acquire batches of points by optimizing a random fourier feature (RFF) [46] approximation of a GP surrogate using NSGA-II and selecting a subset of points from the EA’s population to sequentially greedily maximize HVI. This heuristic approach for maximizing HVI currently has no theoretical guarantees and relies on zeroth-order optimization methods, which tend to be slower and exhibit worse optimization performance than gradient-based approaches.

Entropy-based methods such as PESMO [25], MESMO [3], and PFES [51] are an alternative to EHVI. Of these three methods, PESMO is the only one that accounts for observation noise. However, PESMO involves intractable entropy computations and therefore relies on complex approximations, as well as challenging and time-consuming numerical optimization procedures [25]. Garrido-Merchán & Hernández-Lobato [21] recently proposed an extension to PESMO that supports parallel candidate generation. However, the authors of this work provide limited evaluation and have not provided code to reproduce their results.¹

MOO can also be cast into a single-objective problem by applying a random scalarization of the objectives. ParEGO maximizes the expected improvement using random augmented Chebyshev scalarizations [32]. MOEA/D-EGO [64] extends ParEGO to the batch setting using multiple random scalarizations and the genetic algorithm MOEA/D [65] to optimize these scalarizations in parallel. Recently, q ParEGO, another batch variant of ParEGO was proposed that uses compositional Monte Carlo objectives and sequential greedy candidate selection [11]. Additionally, the authors proposed a noisy variant, q NParEGO, but the empirical evaluation of that variant was limited. TS-TCH [45] combines random Chebyshev scalarizations with Thompson sampling [54], which is naturally robust to noise when the objective is scalarized. Golovin & Zhang [23] propose to use a hypervolume scalarization with the property that the expected value of the scalarization over a specific distribution of weights is equivalent to the hypervolume indicator. The authors propose an upper confidence bound algorithm using randomly sampled weights, but provide a very limited empirical evaluation.

Many prior attempts by the simulation community to handle MOO with noisy observations found that accounting for the noise did not improve optimization performance: Horn et al. [26] suggest that the best approach is to ignore noise, and Koch et al. [33] concluded that further research was needed to determine if modeling techniques such as re-interpolation could improve BO performance with noisy observations. In contrast, we find that accounting for noise *does substantially* improve performance in noisy settings.

Lastly, previous works have considered methods for quantifying and monitoring uncertainty in the Pareto frontiers during the optimization [4, 7]. In contrast, we provide a solution to performing MOBO in noisy settings, rather than purely reasoning about the uncertainty in the Pareto frontier.

4 Background on Expected Hypervolume Improvement

In this section, we review hypervolume, hypervolume improvement, and expected hypervolume improvement as well as efficient methods for computing these metrics using box decompositions.

¹We contacted the authors twice asking for code to reproduce their results, but they graciously declined.

Definition 1. The hypervolume indicator (HV) of a finite approximate Pareto frontier \mathcal{P} is the M -dimensional Lebesgue measure λ_M of the space dominated by \mathcal{P} and bounded from below by a reference point. $\mathbf{r} \in \mathbb{R}^M$: $\text{HV}(\mathcal{P}|\mathbf{r}) = \lambda_M(\bigcup_{\mathbf{v} \in \mathcal{P}} [\mathbf{r}, \mathbf{v}])$, where $[\mathbf{r}, \mathbf{v}]$ denotes the hyper-rectangle bounded by vertices \mathbf{r} and \mathbf{v} .

As in previous work, we assume that the reference point \mathbf{r} is known and specified by the decision maker [58].

Definition 2. The hypervolume improvement (HVI) of a set of points \mathcal{P}' w.r.t. an existing approximate Pareto frontier \mathcal{P} and reference point \mathbf{r} is defined as² $\text{HVI}(\mathcal{P}'|\mathcal{P}, \mathbf{r}) = \text{HV}(\mathcal{P} \cup \mathcal{P}'|\mathbf{r}) - \text{HV}(\mathcal{P}|\mathbf{r})$.

Computing HV requires calculating the volume of a typically non-rectangular polytope and is known to have time complexity that is super-polynomial in the number of objectives [59]. An efficient approach for computing HV is to (i) decompose the region that is dominated by the Pareto frontier \mathcal{P} and bounded from below by the reference point \mathbf{r} into disjoint axis-aligned hyperrectangles [34], (ii) compute the volume of each hyperrectangle in the decomposition, and (iii) sum over all hyperrectangles. So-called box decomposition algorithms have also been applied to partition the region that is *not dominated* by the Pareto frontier \mathcal{P} , which can be used to compute the HVI from a set of new points [15, 59]. See Appendix B for further details.

Expected Hypervolume Improvement: Since function values at unobserved points are unknown in black-box optimization, so is the HVI of an out-of-sample point. However, in BO the probabilistic surrogate model provides a posterior distribution $p(\mathbf{f}(\mathbf{x})|\mathcal{D})$ over the function values for each \mathbf{x} , which can be used to compute the expected hypervolume improvement (EHVI) acquisition function: $\alpha_{\text{EHVI}}(\mathbf{x}|\mathcal{P}) = \mathbb{E}[\text{HVI}(\mathbf{f}(\mathbf{x})|\mathcal{P})]$. Although α_{EHVI} can be expressed analytically when (i) the objectives are assumed to be conditionally independent given \mathbf{x} and (ii) the candidates are generated and evaluated sequentially [58], Monte Carlo (MC) integration is commonly used since it does not require either assumption [16]. The more general parallel variant using MC integration is given by

$$\alpha_{q\text{EHVI}}(\mathcal{X}_{\text{cand}}|\mathcal{P}) \approx \hat{\alpha}_{q\text{EHVI}}(\mathcal{X}_{\text{cand}}|\mathcal{P}) = \frac{1}{N} \sum_{t=1}^N \text{HVI}(\tilde{\mathbf{f}}_t(\mathcal{X}_{\text{cand}})|\mathcal{P}), \quad (1)$$

where $\tilde{\mathbf{f}}_t \sim p(\mathbf{f}|\mathcal{D})$ for $t = 1, \dots, N$ and $\mathcal{X}_{\text{cand}} = \{\mathbf{x}_i\}_{i=1}^q$ [11]. The same box decomposition algorithms used to compute HVI can be used to compute EHVI (either analytic or via MC) using piece-wise integration. EHVI computation is agnostic to the choice of box decomposition algorithm (and can also use approximate methods [9]). Similar to EI in the single-objective case, EHVI is a one-step Bayes-optimal algorithm for maximizing hypervolume in the MOO setting under the following assumptions: (i) only a single design will be generated and evaluated, (ii) the observations are noise-free, (iii) the final approximate Pareto frontier (and final design that will be deployed) will be drawn from the set of observed points [19].

5 Expected Hypervolume Improvement with Noisy Observations

We consider the case that frequently arises in practice where we only receive noisy observations $\mathbf{y}_i = \mathbf{f}(\mathbf{x}_i) + \boldsymbol{\epsilon}_i$, $\boldsymbol{\epsilon}_i \sim \mathcal{N}(0, \Sigma_i)$, where Σ_i is the noise covariance. In this setting, EHVI is no longer (one-step) Bayes-optimal. This is because we can no longer compute the true Pareto frontier $\mathcal{P}_n = \{\mathbf{f}(\mathbf{x}) \mid \mathbf{x} \in X_n, \nexists \mathbf{x}' \in X_n \text{ s.t. } \mathbf{f}(\mathbf{x}') \succ \mathbf{f}(\mathbf{x})\}$ over the previously evaluated points $X_n = \{\mathbf{x}_i\}_{i=1}^n$. Simply using the observed Pareto frontier, $\mathcal{Y}_n = \{\mathbf{y} \mid \mathbf{y} \in Y_n, \nexists \mathbf{y}' \in Y_n \text{ s.t. } \mathbf{y}' \succ \mathbf{y}, \mathbf{y}\}$ where $Y_n = \{\mathbf{y}_i\}_{i=1}^n$, can have strong detrimental effects on optimization performance. This is illustrated in Figure 1, which shows how EHVI is misled by noisy observations that appear to be Pareto optimal. EHVI proceeds to spend its evaluation budget trying to optimize noise, resulting in a clumped Pareto frontier that lacks diversity. Although the posterior mean could serve as a "plug-in" estimate of the true function values at the observed points and provide some regularization [61], we find that this heuristic also leads to clustered Pareto frontiers (EHVI-PM in Fig. 1). Similar patterns emerge with DGEMO (which does not account for noise), and other baselines that utilize the posterior mean rather than the observed values when computing hypervolume improvement (see Appendix H). To our knowledge, all previous work on EHVI assumes that observations are noiseless [16, 58] or imputes the unknown true function values with the posterior mean.

²For brevity we omit the reference point \mathbf{r} when referring to HVI.




Figure 1: An illustration of the effect of noisy observations on the true noiseless Pareto frontiers identified by NEHVI (our proposed algorithm), EHVI, and EHVI-PM, which uses the modeled posterior mean as point estimate of the true in-sample function values. All algorithms are tested on a BraninCurrin synthetic problem, where observations are corrupted with zero-mean, additive Gaussian noise with a standard deviation of 5% of the range of respective objective. All methods use sequential ($q = 1$) optimization. See Appendix G for details.

5.1 A Bayes-optimal algorithm for hypervolume maximization in noisy environments

In contrast with EHVI(-PM), we instead approach the problem of hypervolume maximization under noisy observations from a Bayesian perspective and derive a novel one-step Bayes-optimal expected hypervolume improvement criterion that iterates the expectation over the posterior distribution $p(\mathbf{f}(X_n)|\mathcal{D}_n)$ of the function values at the previously evaluated points X_n given noisy observations $\mathcal{D}_n = \{\mathbf{x}_i, \mathbf{y}_i, (\Sigma_i)\}_{i=1}^n$. Our acquisition function, *noisy expected hypervolume improvement* (NEHVI), is defined as

$$\alpha_{\text{NEHVI}}(\mathbf{x}) = \int \alpha_{\text{EHVI}}(\mathbf{x}|\mathcal{P}_n)p(\mathbf{f}|\mathcal{D}_n)d\mathbf{f} \quad (2)$$

where \mathcal{P}_n denotes the Pareto frontier over $\mathbf{f}(X_n)$.

By integrating over the uncertainty in the function values at the observed points, NEHVI retains one-step Bayes-optimality in noisy environments (in noiseless environments, NEHVI is equivalent to EHVI). Empirically, Figure 1 shows that NEHVI is robust to noise and identifies a well-distributed Pareto frontier with no signs of clumping, even under very noisy observations.³

The integral in (2) is analytically intractable, but can easily be approximated using MC integration. Let $\tilde{\mathbf{f}}_t \sim p(\mathbf{f}|\mathcal{D}_n)$ for $t = 1, \dots, N$ be samples from the posterior, and let $\mathcal{P}_t = \{\tilde{\mathbf{f}}_t(\mathbf{x}) \mid \mathbf{x} \in X_n, \tilde{\mathbf{f}}_t(\mathbf{x}) \succ \tilde{\mathbf{f}}_t(\mathbf{x}') \forall \mathbf{x}' \in X_n\}$ be the Pareto frontier over the previously evaluated points under the sampled function $\tilde{\mathbf{f}}_t$. Then, $\alpha_{\text{NEHVI}}(\mathbf{x}) \approx \frac{1}{N} \sum_{t=1}^N \alpha_{\text{EHVI}}(\mathbf{x}|\mathcal{P}_t)$. Using MC integration, we can compute the inner expectation in α_{EHVI} simultaneously using samples from the joint posterior $\tilde{\mathbf{f}}_t(X_n, \mathbf{x}) \sim p(\mathbf{f}(X_n, \mathbf{x})|\mathcal{D}_n)$ over \mathbf{x} and X_n :

$$\hat{\alpha}_{\text{NEHVI}}(\mathbf{x}) = \frac{1}{N} \sum_{t=1}^N \text{HVI}(\tilde{\mathbf{f}}_t(\mathbf{x})|\mathcal{P}_t). \quad (3)$$

See Appendix B for details on computing (3) using box decompositions. Note that this “full-MC” variant of NEHVI does not require objectives to be modeled independently, and supports multi-task covariance functions across correlated objectives.

5.2 Parallel Noisy Expected Hypervolume Improvement

Generating and evaluating batches of candidates is imperative to achieving adequate throughput in many real-world scenarios. qNEHVI can naturally be extended to the parallel (asynchronous or batch) setting by evaluating HVI with respect to a batch of q points $\mathcal{X}_{\text{cand}} = \{\mathbf{x}_i\}_{i=1}^q$

$$\alpha_{q\text{NEHVI}}(\mathcal{X}_{\text{cand}}) = \int \alpha_{q\text{EHVI}}(\mathcal{X}_{\text{cand}}|\mathcal{P}_n)p(\mathbf{f}|\mathcal{D}_n)d\mathbf{f} \approx \hat{\alpha}_{q\text{NEHVI}}(\mathcal{X}_{\text{cand}}) = \frac{1}{N} \sum_{t=1}^N \text{HVI}(\tilde{\mathbf{f}}_t(\mathcal{X}_{\text{cand}})|\mathcal{P}_t) \quad (4)$$

Since optimizing q candidates jointly is a difficult numerical optimization problem over a qd -dimensional domain, we use a sequential greedy approximation in the parallel setting and solve a sequence of q simpler optimization problems with d dimensions, which been shown empirically to improve optimization performance [57]. While selecting candidates according to a “sequential greedy” policy does not guarantee that the selected batch of candidates is a maximizer of the $\alpha_{q\text{NEHVI}}$, the submodularity of $\alpha_{q\text{NEHVI}}$ allows us to *bound the regret* of this approximation to be no more than $\frac{1}{e}\alpha_{q\text{NEHVI}}^*$, where $\alpha_{q\text{NEHVI}}^* = \max_{\mathcal{X}_{\text{cand}} \in \mathcal{X}} \alpha_{q\text{NEHVI}}(\mathcal{X}_{\text{cand}})$ (see Appendix F).

³This noise level is 5x greater than the ones considered by previous works that evaluate noisy MOBO [25].

6 Efficient Evaluation with Cached Box Decompositions

Although $\hat{\alpha}_{\text{NEHVI}}(\mathbf{x})$ in (3) has a concise mathematical form, computing it requires determining the Pareto frontier \mathcal{P}_t under each sample $\tilde{\mathbf{f}}_t$ for $t = 1, \dots, N$ and then partitioning the region that is not dominated by \mathcal{P}_t into disjoint hyperrectangles $\{S_{k_t}\}_{k_t=1}^{K_t}$. Optimizing the unbiased MC estimator of α_{NEHVI} would require re-sampling $\{\tilde{\mathbf{f}}_t\}_{t=1}^N$ at each evaluation of α_{NEHVI} . However, computing the Pareto frontier and performing a box decomposition under each of the N samples during every evaluation of α_{NEHVI} in the inner optimization loop ($\mathbf{x}^* = \arg \max_{\mathbf{x}} \alpha_{\text{NEHVI}}(\mathbf{x}|\mathcal{D}_n)$) would be prohibitively expensive. This is because box decomposition algorithms have super-polynomial time complexity in the number of objectives [59]. We instead propose an efficient alternative computational technique for repeated evaluations of EHVI with uncertain Pareto frontiers.

Cached Box Decompositions: For repeated evaluations of the integral in (2), we use a set of fixed samples $\{\tilde{\mathbf{f}}_t(X_n)\}_{t=1}^N$, which allows us to compute the Pareto frontiers and box decompositions once, and cache them for the entirety of the acquisition function optimization, thereby making those two computationally intensive operations a one-time cost per BO iteration.⁴ We refer to this approach as using *cached box decompositions* (CBD). The method of optimizing over fixed random samples is known as sample average approximation (SAA) [2].

Conditional Posterior Sampling: Under the CBD formulation, computing $\hat{\alpha}_{\text{NEHVI}}(\mathbf{x})$ with joint samples from $\tilde{\mathbf{f}}_t(X_n, \mathbf{x}) \sim p(\mathbf{f}(X_n, \mathbf{x})|\mathcal{D}_n)$ requires sampling from the conditional distributions

$$\tilde{\mathbf{f}}_t(\mathbf{x}) \sim p(\mathbf{f}(\mathbf{x})|\mathbf{f}(X_n) = \tilde{\mathbf{f}}_t(X_n), \mathcal{D}_n), \quad (5)$$

where $t = 1, \dots, N$ and $\{\tilde{\mathbf{f}}_t(X_n)\}_{t=1}^N$ are the realized samples at the previously evaluated points. For multivariate Gaussian posteriors (as is the case with GP surrogates), we can sample from $p(\mathbf{f}(X_n)|\mathcal{D}_n)$ via the reparameterization trick [30] by evaluating $\tilde{\mathbf{f}}_t(\mathbf{x}) = \boldsymbol{\mu}_n + L_n^T \boldsymbol{\zeta}_{n,t}$, where $\boldsymbol{\zeta}_{n,t} \sim \mathcal{N}(\mathbf{0}, I_{nM})$, $\boldsymbol{\mu}_n \in \mathbb{R}^{nM}$ is the posterior mean, and $L_n \in \mathbb{R}^{nM \times nM}$ is a lower triangular root decomposition of the posterior covariance matrix, typically a Cholesky decomposition. Given L_n , we can obtain a root decomposition L'_n of the covariance matrix of the joint posterior $p(\mathbf{f}(X_n, \mathbf{x})|\mathcal{D}_n)$ by performing efficient low-rank updates [44]. Given L'_n and the posterior mean of $p(\mathbf{f}(X_n, \mathbf{x})|\mathcal{D}_n)$, we can sample from (5) via the reparameterization trick by augmenting the existing base samples $\boldsymbol{\zeta}_{n,t}$ with M new base samples for the new point.

6.1 Efficient Sequential Greedy Batch Selection using CBD

The CBD technique addresses the general problem of inefficient repeated evaluations of EHVI with uncertain Pareto frontiers. In this section, we show that sequential greedy batch selection (with both q EHVI and q NEHVI) is an incarnation of EHVI with uncertain Pareto frontiers.

The original formulation of parallel EHVI in Daulton et al. [11] uses the inclusion-exclusion principle (IEP), which involves computing the volume jointly dominated by each of the $2^q - 1$ nonempty subsets of points in $\mathcal{X}_{\text{cand}}$. However, using large batch sizes is *not computationally feasible* under this formulation because time and space complexity are exponential in q and multiplicative in the number of hyperrectangles in the box decomposition [11] (see Appendix D for a complexity analysis). Although q EHVI is optimized using sequential greedy batch selection, the IEP is used over all candidates $\mathbf{x}_1, \dots, \mathbf{x}_i$ when selecting candidate i . Although the IEP could similarly be used to compute q NEHVI, we instead leverage CBD, which yields a sequential greedy approximation of the joint (noisy) EHVI that is *mathematically equivalent* to the IEP formulation, but *significantly reduces computational overhead*. That is, the IEP and CBD approaches produce exactly the same acquisition value for a given set of points $\mathcal{X}_{\text{cand}}$, but the IEP and the CBD approaches have *exponential* and *polynomial* time complexities in q , respectively.

When selecting \mathbf{x}_i for $i \in \{2, \dots, q\}$, all \mathbf{x}_j for which $j < i$ have already been selected and are therefore held constant. Thus, we can decompose q NEHVI into the q NEHVI from the previously selected candidates $\mathbf{x}_1, \dots, \mathbf{x}_{i-1}$ and NEHVI from \mathbf{x}_i given the previously selected candidates

⁴For greater efficiency, we may also prune X_n to remove points that are dominated with high probability, which we estimate via MC.




Figure 2: Acquisition optimization wall time under a sequential greedy approximation using L-BFGS-B. CBD enables scaling to much larger batch sizes q than using the IEP and avoids running out-of-memory (OOM) on a GPU. Independent GPs are used for each outcome. The Pareto frontier of the 2-objective, 6-dimensional DTLZ2 problem [13] is initialized with 20 points. Wall times were measured on a Tesla V100 SXM2 GPU (16GB RAM) and a 2x Intel Xeon 6138 CPU @ 2GHz (251GB RAM). See Appendix H.2 for results with more objectives.

$$\hat{\alpha}_{q\text{NEHVI}}(\{\mathbf{x}_j\}_{j=1}^i) = \frac{1}{N} \sum_{t=1}^N \text{HVI}(\{\tilde{\mathbf{f}}_t(\mathbf{x}_j)\}_{j=1}^{i-1} | \mathcal{P}_t) + \frac{1}{N} \sum_{t=1}^N \text{HVI}(\tilde{\mathbf{f}}_t(\mathbf{x}_i) | \mathcal{P}_t \cup \{\tilde{\mathbf{f}}_t(\mathbf{x}_j)\}_{j=1}^{i-1}) \quad (6)$$

Note that the first term on the right hand side is constant, since $\{\mathbf{x}_j\}_{j=1}^{i-1}$ and $\{\tilde{\mathbf{f}}_t(\mathbf{x}_j)\}_{j=1}^{i-1}$ are fixed for all $t = 1, \dots, N$. The second term is $\hat{\alpha}_{\text{NEHVI}}(\mathbf{x}_i)$, where the NEHVI is taken with respect to the Pareto frontier across $\mathbf{f}(X_n, \mathbf{x}_1, \dots, \mathbf{x}_{i-1})$ and computed using the fixed samples $\{\tilde{\mathbf{f}}_t(X_n, \mathbf{x}_1, \dots, \mathbf{x}_{i-1})\}_{t=1}^N$. To compute the second term when selecting candidate \mathbf{x}_i , the N Pareto frontiers and CBDs are updated to include $\{\tilde{\mathbf{f}}_t(X_n, \mathbf{x}_1, \dots, \mathbf{x}_{i-1})\}_{t=1}^N$. As in the sequential $q = 1$ setting, the box decompositions are only computed and cached while selecting each candidate point. See Appendix C.2 for a derivation of (6). Although we have focused on $q\text{NEHVI}$ in the above, the CBD formulation for $q\text{EHVI}$ is obtained by simply replacing \mathcal{P}_t with the Pareto frontier over the observed values \mathcal{Y}_n .

Despite computing N box decompositions when selecting each candidate \mathbf{x}_i for $i = 2, \dots, q$, the CBD approach reduces the time and space complexity from exponential (under the IEP) to polynomial in q (see Appendix D for details on time and space complexity). Figure 2 shows the total acquisition optimization time (including box decompositions) for various batch sizes and demonstrates that using CBD allows to scale to batch sizes that are *completely infeasible* when using IEP.

7 Optimizing NEHVI

Differentiability: Importantly, $\hat{\alpha}_{\text{NEHVI}}(\mathbf{x})$ is differentiable w.r.t. \mathbf{x} . Although determining the Pareto frontier and computing the box decompositions are non-differentiable operations, these operations do not involve \mathbf{x} , even when re-sampling from the joint posterior $p(\mathbf{f}(X_n, \mathbf{x}) | \mathcal{D}_n)$. Exact sample-path gradients of $\nabla_{\mathbf{x}} \hat{\alpha}_{\text{NEHVI}}(\mathbf{x})$ can easily be computed using auto-differentiation in modern computational frameworks. This enables efficient gradient-based optimization of $q\text{NEHVI}$.⁵

SAA Convergence Results: In addition to approximating the outer expectation over $\mathbf{f}(X_n)$ with fixed posterior samples, we can similarly fix the base samples used for the new candidate point \mathbf{x} . This approach yields a deterministic acquisition function, which enables using (quasi-) higher-order optimization methods to obtain fast convergence rates for acquisition optimization [2]. Importantly, we prove that the theoretical convergence guarantees on acquisition optimization under the SAA approach proposed by Balandat et al. [2] also hold for NEHVI.

Theorem 1. *Suppose \mathcal{X} is compact and \mathbf{f} has a multi-output GP prior with continuously differentiable mean and covariance functions. Let $X_n = \{\mathbf{x}_i\}_{i=1}^n$ denote the previously evaluated points and $\{\zeta\}_{i=1}^N$ be base samples $\zeta \sim \mathcal{N}(\mathbf{0}, I_{(n+1)M})$. Let $\hat{\alpha}_{\text{NEHVI}}$ denote the deterministic acquisition function computed using $\{\zeta\}_{i=1}^N$ as $\hat{\alpha}_{\text{NEHVI}}^N$ and define $S^* := \arg \max_{\mathbf{x} \in \mathcal{X}} \alpha_{\text{NEHVI}}(\mathbf{x})$ to be the set of maximizers of $\alpha_{\text{NEHVI}}(\mathbf{x})$ over \mathcal{X} . Suppose $\hat{\mathbf{x}}_N^* \in \arg \max_{\mathbf{x} \in \mathcal{X}} \hat{\alpha}_{\text{NEHVI}}^N(\mathbf{x})$. Then (1) $\hat{\alpha}_{\text{NEHVI}}^N(\hat{\mathbf{x}}_N^*) \rightarrow \alpha_{\text{NEHVI}}(\mathbf{x}_N^*)$ almost surely, and (2) $\text{dist}(\hat{\mathbf{x}}_N^*, S^*) \rightarrow 0$, where $\text{dist}(\hat{\mathbf{x}}_N^*, S^*) := \inf_{\mathbf{x} \in S^*} \|\hat{\mathbf{x}}_N^* - \mathbf{x}\|$ is the Euclidean distance between $\hat{\mathbf{x}}_N^*$ and the set S^* .*

Theorem 1 also holds in the parallel setting, so $q\text{NEHVI}$ enjoys the same convergence guarantees as NEHVI on acquisition optimization under the SAA. See Appendix E for further details and proof.

⁵One can also show that the gradient of the full MC estimator $\hat{\alpha}_{q\text{NEHVI}}$ is an unbiased estimator of the gradient of the true joint noisy expected hypervolume improvement $\alpha_{q\text{NEHVI}}$. However, this result is not necessary for our SAA approach.

8 Approximation of q NEHVI using Approximate GP Sample Paths

Although CBD yields polynomial complexity of q NEHVI with respect to q (rather than exponential complexity with the IEP), it still requires computing N box decompositions and repeatedly evaluating the joint posterior over $f(X_n, \{\mathbf{x}_j\}_{j=1}^{i-1})$ for selecting each candidate \mathbf{x}_i for $i = 1, \dots, q$. A cheaper alternative is to approximate the integral in (4) using a single approximate GP sample path $\tilde{\mathbf{f}}_i$ using RFFs when optimizing candidate \mathbf{x}_i . A single-sample approximation of q NEHVI, which we refer to as q NEHVI-1, can be computed by using $\tilde{\mathbf{f}}_i$ as the sampled GP in (6). Since the RFF is a deterministic model, it is much less computationally expensive to evaluate than the GP posterior on out-of-sample points, and exact gradients of q NEHVI-1 with respect to current candidate \mathbf{x}_i can be computed and used for efficient multi-start optimization of q NEHVI-1 using second-order gradient methods. q NEHVI-1 requires CBD for efficient sequential greedy batch selection and gradient-based optimization, but does not use a sample average approximation for optimizing a new candidate \mathbf{x}_i ; instead, it uses an approximate *sample path*. See Rahimi & Recht [46] for details on RFFs.

q NEHVI-1 is related to TSEMO in that both use sequential greedy batch selection using HVI based on RFF samples. However, TSEMO does not directly maximize HVI when selecting candidate \mathbf{x}_i , where $i = 1, \dots, q$; rather, it relies on a heuristic approach of running NSGA-II on an RFF sample of each objective to create a discrete population of candidates and then selecting the point from the discrete population that maximizes HVI under the RFF sample. In contrast, q NEHVI-1 *directly optimizes* HVI under the RFF using exact sample-path gradients, which leads to improved optimization performance (see Appendix H). Furthermore, we find that q NEHVI-1 is significantly faster than TSEMO, because rather than using NSGA-II it uses second order gradient methods to optimize HVI (see Appendix H). Gradient-based optimization is only possible because CBD enables scalable, differentiable HVI computation. While the primary goal of this work is to develop a principled, scalable method for parallel EHVI in noisy environments, we include empirical comparisons with q NEHVI-1 throughout the appendix to demonstrate the generalizability of the CBD approach and practical performance of the q NEHVI-1 approximation. q NEHVI-1 achieves the fastest batch selection times of any method tested on a GPU on every problem; in many cases, this is an order of magnitude speed-up over q NEHVI. Moreover, q NEHVI-1 has a remarkable ability to scale to large batch sizes when the dimensionality of optimization problem is modest. Further investigation of q NEHVI-1 is needed, but we hope that the readers can recognize the ways in which q NEHVI can create broader opportunities for research into hypervolume improvement based acquisition functions.

9 Experiments

We empirically evaluate q NEHVI on a set of synthetic and real-world benchmark problems. We compare it against the following recently proposed methods from the literature: PESMO, MESMO (which we extend to handle noisy observations using the noisy information gain from Takeno et al. [52]), PFES, DGEMO, MOEA/D-EGO, TSEMO, TS-TCH, q EHVI (and q EHVI-PM-CBD, which uses the posterior mean as a plug-in estimate for the function values at the in-sample points, along with CBD to scale to large batch sizes), and q NParEGO. We optimize all methods using multi-start L-BFGS-B with exact gradients (except for PFES, which uses gradients approximated via finite differences), including TS-TCH where we optimize approximate function samples using RFFs with 500 basis functions. We model each outcome with an independent GP with a Matérn 5/2 ARD kernel and infer the GP hyperparameters via maximum a posteriori (MAP) estimation. For all problems, we assume that the noise variances are observed (except ABR, where we infer the noise level). See Appendix G for more details on the experiments and acquisition function implementations.

We evaluate all methods using the logarithm of the difference in hypervolume between the true Pareto frontier and the approximate Pareto frontier recovered by the algorithm. Since evaluations are noisy, we compute the hypervolume dominated by the noiseless Pareto frontier across the observed points for each method.

Synthetic Problems: We consider a noisy variant of the *BraninCurrin* problem ($M = 2, d = 2$) and the *DTLZ2* problem ($M = 2, d = 6$) [13], in which observations are corrupted with zero-mean additive Gaussian noise with standard deviation of 5% of the range of each objective for *BraninCurrin* and 10% for *DTLZ2*.

Adaptive Bitrate (ABR) Control Policy Optimization: ABR controllers are used for real-time communication and media streaming applications. Policies for these controllers must be tuned to




Figure 3: Sequential optimization performance. The shaded region indicates two standard errors of the mean over 100 replications (only 20 replications were feasible for PESMO due to large runtimes).

deliver a high quality of experience with respect to multiple objectives [40]. In industry settings, A/B tests with dozens of policies are tested simultaneously since each policy may take days or weeks to evaluate, producing noisy measurements across multiple objectives. In this experiment, we tune policies to maximize video quality (bitrate) and minimize stall time. The policy has $d = 4$ parameters, which are detailed in Appendix G. We use the Park simulator [41] and sample a random set of 100 traces to obtain noisy measurements of the objectives under a given policy. For comparing the performance of different methods, we estimate the true noiseless objective using mean objectives across 300 traces. We infer a homoskedastic noise level jointly with the GP hyperparameters via MAP estimation.

Vehicle Design Optimization: Optimizing the design of the frame an automobile is important to maximizing passenger safety, vehicle durability and fuel efficiency. Evaluating a vehicle design is time-consuming, since either a vehicle must be manufactured and crashed, or a nonlinear finite element-based crash analysis must be run to simulate a collision (which can take over 20 hours per run) [62]. Hence, evaluating many designs in parallel is critical for reducing end-to-end optimization time. Observations are often noisy due to manufacturing imperfections, measurement error, or non-deterministic simulations. In this experiment, we tune the $d = 5$ widths of various components of a vehicle’s frame to minimize proxy metrics for (1) fuel consumption, (2) passenger trauma in a full frontal collision, and (3) vehicle fragility [53]. See Appendix G for details. For this demonstration, we add zero-mean Gaussian noise with a standard deviation of 1% of the objective range, which roughly corresponds to the manufacturing noise level used in previous work [62].

9.1 Summary of Results:

We find that $qNEHVI$ and $qNEHVI-1$ outperform all other methods on the noisy benchmarks, both in the sequential and parallel setting. In the sequential setting (Fig 3), $qNEHVI$ and $qNEHVI-1$ are followed closely by $qEHVI-PM$, and in some cases, even $qEHVI$. $TS-TCH$ is firmly in the middle of the pack, while information-theoretic acquisition functions appear to perform the worst. This is consistent across noise levels; for experiments where we add noise to the objectives, we consider noise levels ranging from 1% to 10% of the range of each objective (these are magnitudes of the noise often seen in practice). Previous works have only evaluated MOBO algorithms with noise levels of 1% [25]. In Appendix H, we perform a study showing that $qNEHVI$ consistently performs best with increasing noise levels up to 30% of the range of each objective.

While parallel evaluation can provide optimization speedups on order of the batch size q , these evaluations do affect the overall sample complexity of the algorithm, since less information is available within the synchronous batch setting compared with fully sequential optimization. We find that, by and large, $qNEHVI$ achieves the greatest hyper-volume for increasingly large batch sizes, and scales more elegantly relative to $TS-TCH$ and the ParEGO variants (Fig 4). $qNEHVI$ also consistently outperforms $qEHVI-PM-CBD$. In Appendix H, we observe that $qNEHVI$ and $qNEHVI-1$ provides excellent anytime performance all values of q that we tested. We provide results on 4 additional test problems in Appendix H.3, and in Appendix H.8, we demonstrate that leveraging CBD and a single sample path approximation, $qNEHVI-1$ enables scaling to 5-objective problems, which is a first for an HVI-based method, to our knowledge.




Figure 4: The quality of the final Pareto frontier identified by each method with increasing batch sizes q given a budget of 224 function evaluations. q EHVI is only included for $q = 1$ and $q = 8$ because the IEP scales exponential with q . DGEMO is omitted on the ABR problem because it was prohibitively slow with time-consuming ABR simulations and on the VehicleSafety problem because DGEMO consistently crashed in the graph cutting algorithm.

In our experiments, we find that q NEHVI-1 is among the top performers on relatively low-dimensional problems. Given the strong performance of q NEHVI-1, we examine its performance as the dimensionality of the search space increases in Appendix H.5. We find that q NEHVI is more robust than q NEHVI-1 in higher-dimensional search spaces, but further investigation is needed into how the number of the Fourier basis functions affects the performance of q NEHVI-1 in high-dimensional search spaces.

Optimization wall time: Across all experiments, we observe competitive wall times for optimizing q NEHVI and q NEHVI-1 (all wall time comparisons are provided in Appendix H). On a GPU, optimizing q NEHVI-1 incurs the lowest wall time of any method that we tested on every single problem and optimizing q NEHVI is faster than optimizing information-theoretic methods on all problems. Using efficient low-rank Cholesky updates, q NEHVI is often faster than the q NParEGO implementation in BoTorch on a GPU.

10 Discussion

We proposed NEHVI, a novel acquisition function that provides a principled approach to parallel and noisy multi-objective Bayesian optimization. NEHVI is a one-step Bayes-optimal policy for maximizing the hypervolume dominated by the Pareto frontier in noisy and noise-free settings. NEHVI is made feasible by a new approach to computing joint hypervolumes (CBD), and we demonstrated that CBD enables scalable, parallel candidate generation with both noiseless q EHVI and q NEHVI. We provide theoretical results on optimizing a MC estimator of q NEHVI using sample average approximation and demonstrate significant improvements in optimization performance over state-of-the-art MOBO algorithms.

Yet, our work has some limitations. While the information-theoretic acquisition functions tested here perform poorly on our benchmarks, they do allow for decoupled evaluations of different objectives in cases where querying one objective may be more resource-intensive than querying other objectives. Optimizing such acquisition functions is a non-trivial task, and it is possible that with improved procedures, such acquisition functions could yield improved performance and provide a principled approach to selecting evaluation sources on a budget. Although practically fast enough for most Bayesian optimization tasks, exact hypervolume computation has super-polynomial complexity in the number of objectives. Combining q NEHVI with differentiable approximate methods for computing hypervolume (e.g. Couckuyt et al. [9], Golovin & Zhang [23]) could lead to further speed-ups.

We hope that the core ideas presented in this work, including the CBD approach, can provide a framework to support the development of new computationally efficient MOBO methods.

References

- [1] Asadpour, A., Nazerzadeh, H., and Saberi, A. Stochastic submodular maximization. In Papadimitriou, C. and Zhang, S. (eds.), *Internet and Network Economics*. Springer Berlin Heidelberg, 2008.
- [2] Balandat, M., Karrer, B., Jiang, D. R., Daulton, S., Letham, B., Wilson, A. G., and Bakshy, E. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In *Advances in Neural Information Processing Systems 33*, 2020.
- [3] Belakaria, S., Deshwal, A., and Doppa, J. R. Max-value entropy search for multi-objective bayesian optimization. In *Advances in Neural Information Processing Systems 32*, 2019.
- [4] Binois, M., Ginsbourger, D., and Roustant, O. Quantifying uncertainty on pareto fronts with gaussian process conditional simulations. *Eur. J. Oper. Res.*, 243:386–394, 2015.
- [5] Bradford, E., Schweidtmann, A. M., and Lapkin, A. Efficient multiobjective optimization employing gaussian processes, spectral sampling and a genetic algorithm. *Journal of global optimization*, 71(2):407–438, 2018.
- [6] Brockhoff, D., Tusar, T., Auger, A., and Hansen, N. Using well-understood single-objective functions in multiobjective black-box optimization test suites, 2019.
- [7] Calandra, R. and Peters, J. Pareto front modeling for sensitivity analysis in multi-objective bayesian optimization. 2014.
- [8] Calandra, R., Seyfarth, A., Peters, J., and Deisenroth, M. P. Bayesian optimization for learning gaits under uncertainty. *Annals of Mathematics and Artificial Intelligence*, 76(1):5–23, Feb 2016. ISSN 1573-7470. doi: 10.1007/s10472-015-9463-9.
- [9] Couckuyt, I., Deschrijver, D., and Dhaene, T. Towards efficient multiobjective optimization: Multiobjective statistical criterions. In *2012 IEEE Congress on Evolutionary Computation*, pp. 1–8, 2012.
- [10] Daulton, S., Singh, S., Avadhanula, V., Dimmery, D., and Bakshy, E. Thompson sampling for contextual bandit problems with auxiliary safety constraints. In *NeurIPS Workshop on Safety and Robustness in Decision Making*, 2019.
- [11] Daulton, S., Balandat, M., and Bakshy, E. Differentiable expected hypervolume improvement for parallel multi-objective Bayesian optimization. In *Advances in Neural Information Processing Systems 33*, NeurIPS, 2020.
- [12] Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.
- [13] Deb, K., Thiele, L., Laumanns, M., and Zitzler, E. Scalable multi-objective optimization test problems. volume 1, pp. 825–830, 06 2002. ISBN 0-7803-7282-4. doi: 10.1109/CEC.2002.1007032.
- [14] Deb, K., Gupta, S., Daum, D., Branke, J., Mall, A. K., and Padmanabhan, D. Reliability-based optimization using evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 13(5):1054–1074, 2009. doi: 10.1109/TEVC.2009.2014361.
- [15] Dächert, K., Klamroth, K., Lacour, R., and Vanderpooten, D. Efficient computation of the search region in multi-objective optimization. *European Journal of Operational Research*, 260(3):841 – 855, 2017.
- [16] Emmerich, M. T. M., Giannakoglou, K. C., and Naujoks, B. Single- and multiobjective evolutionary optimization assisted by gaussian random field metamodels. *IEEE Transactions on Evolutionary Computation*, 10(4):421–439, 2006.
- [17] Feng, Q., Letham, B., Bakshy, E., and Mao, H. High-Dimensional Contextual Policy Search with Unknown Context Rewards using Bayesian Optimization. In *Advances in Neural Information Processing Systems 33*, 2020.
- [18] Fisher, M. L., Nemhauser, G. L., and Wolsey, L. A. *An analysis of approximations for maximizing submodular set functions—II*, pp. 73–87. Springer Berlin Heidelberg, Berlin, Heidelberg, 1978.
- [19] Frazier, P. I. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

- [20] Garrido-Merchán, E. C. and Hernández-Lobato, D. Predictive entropy search for multi-objective bayesian optimization with constraints. *Neurocomputing*, 361:50–68, 2019.
- [21] Garrido-Merchán, E. C. and Hernández-Lobato, D. Parallel predictive entropy search for multi-objective bayesian optimization with constraints, 2020.
- [22] Gelbart, M. A., Snoek, J., and Adams, R. P. Bayesian optimization with unknown constraints. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, UAI, 2014.
- [23] Golovin, D. and Zhang, Q. Random hypervolume scalarizations for provable multi-objective black box optimization, 2020.
- [24] Hernández-Lobato, J. M., Hoffman, M. W., and Ghahramani, Z. Predictive entropy search for efficient global optimization of black-box functions. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’14, pp. 918–926, Cambridge, MA, USA, 2014. MIT Press.
- [25] Hernández-Lobato, D., Hernández-Lobato, J. M., Shah, A., and Adams, R. P. Predictive entropy search for multi-objective bayesian optimization, 2015.
- [26] Horn, D., Dagge, M., Sun, X., and Bischl, B. First investigations on noisy model-based multi-objective optimization. volume 10173, pp. 298–313, 02 2017. ISBN 978-3-319-54156-3. doi: 10.1007/978-3-319-54157-0_21.
- [27] Igel, C., Hansen, N., and Roth, S. Covariance matrix adaptation for multi-objective optimization. *Evolutionary Computation*, 15(1):1–28, 2007. doi: 10.1162/evco.2007.15.1.1.
- [28] Jiang, S., Zhang, H., Cong, W., Liang, Z., Ren, Q., Wang, C., Zhang, F., and Jiao, X. Multi-objective optimization of smallholder apple production: Lessons from the bohai bay region. *Sustainability*, 12(16):6496, 2020.
- [29] Jones, D. R., Schonlau, M., and Welch, W. J. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13:455–492, 1998.
- [30] Kingma, D. P. and Welling, M. Auto-Encoding Variational Bayes. *arXiv e-prints*, pp. arXiv:1312.6114, Dec 2013.
- [31] Klamroth, K., Lacour, R., and Vanderpooten, D. On the representation of the search region in multi-objective optimization. *European Journal of Operational Research*, 245(3):767–778, Sep 2015. ISSN 0377-2217. doi: 10.1016/j.ejor.2015.03.031. URL <http://dx.doi.org/10.1016/j.ejor.2015.03.031>.
- [32] Knowles, J. Parego: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1): 50–66, 2006.
- [33] Koch, P., Wagner, T., Emmerich, M. T., Back, T., and Konen, W. Efficient multi-criteria optimization on noisy machine learning problems. *Appl. Soft Comput.*, 29(C):357–370, April 2015. ISSN 1568-4946. doi: 10.1016/j.asoc.2015.01.005. URL <https://doi.org/10.1016/j.asoc.2015.01.005>.
- [34] Lacour, R., Klamroth, K., and Fonseca, C. M. A box decomposition algorithm to compute the hypervolume indicator. *Computers & Operations Research*, 79:347 – 360, 2017.
- [35] LeCun, Y., Cortes, C., and Burges, C. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- [36] Letham, B. and Bakshy, E. Bayesian optimization for policy search via online-offline experimentation. *Journal of Machine Learning Research*, 20(145):1–30, 2019. URL <http://jmlr.org/papers/v20/18-225.html>.
- [37] Letham, B., Karrer, B., Ottoni, G., and Bakshy, E. Constrained bayesian optimization with noisy experiments. *Bayesian Analysis*, 14(2):495–519, 06 2019. doi: 10.1214/18-BA1110.
- [38] Liao, T., Wang, G., Yang, B., Lee, R., Pister, K., Levine, S., and Calandra, R. Data-efficient learning of morphology and controller for a microrobot. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 2488–2494. IEEE, 2019.
- [39] Lukovic, K. M., Tian, Y., and Matusik, W. Diversity-guided multi-objective bayesian optimization with batch evaluations. *Advances in Neural Information Processing Systems*, 33, 2020.

- [40] Mao, H., Chen, S., Dimmery, D., Singh, S., Blaisdell, D., Tian, Y., Alizadeh, M., and Bakshy, E. Real-world video adaptation with reinforcement learning. 2019.
- [41] Mao, H., Negi, P., Narayan, A., Wang, H., Yang, J., Wang, H., Marcus, R., Addanki, R., Shirkoohi, M. K., He, S., Nathan, V., Cangialosi, F., Venkatakrisnan, S. B., Weng, W.-H., Han, S.-W., Kraska, T., and Alizadeh, M. Park: An open platform for learning-augmented computer systems. In *NeurIPS*, 2019.
- [42] Mennen, S. M., Alhambra, C., Allen, C. L., Barberis, M., Berritt, S., Brandt, T. A., Campbell, A. D., Castañón, J., Cherney, A. H., Christensen, M., Damon, D. B., Eugenio de Diego, J., García-Cerrada, S., García-Losada, P., Haro, R., Janey, J., Leitch, D. C., Li, L., Liu, F., Lobben, P. C., MacMillan, D. W. C., Magano, J., McInturff, E., Monfette, S., Post, R. J., Schultz, D., Sitter, B. J., Stevens, J. M., Strambeanu, I. I., Twilton, J., Wang, K., and Zajac, M. A. The evolution of high-throughput experimentation in pharmaceutical development and perspectives on the future. *Organic Process Research & Development*, 23(6):1213–1242, 2019. doi: 10.1021/acs.oprd.9b00140.
- [43] Namkoong, H., Daulton, S., and Bakshy, E. Distilled thompson sampling: Practical and efficient thompson sampling via imitation learning. In *NeurIPS Offline Reinforcement Learning Workshop*, 2020.
- [44] Osborne, M. A. Bayesian gaussian processes for sequential prediction, optimisation and quadrature. 2010.
- [45] Paria, B., Kandasamy, K., and Póczos, B. A Flexible Multi-Objective Bayesian Optimization Approach using Random Scalarizations. *ArXiv e-prints*, May 2018.
- [46] Rahimi, A. and Recht, B. Random features for large-scale kernel machines. In *Proceedings of the 20th International Conference on Neural Information Processing Systems, NIPS’07*, pp. 1177–1184, Red Hook, NY, USA, 2007. Curran Associates Inc. ISBN 9781605603520.
- [47] Rasmussen, C. E. *Gaussian Processes in Machine Learning*, pp. 63–71. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004.
- [48] Real, E., Aggarwal, A., Huang, Y., and Le, Q. V. Regularized evolution for image classifier architecture search. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 4780–4789, Jul. 2019. doi: 10.1609/aaai.v33i01.33014780. URL <https://ojs.aaai.org/index.php/AAAI/article/view/4405>.
- [49] Schuster, M. Speech recognition for mobile devices at google. In Zhang, B.-T. and Orgun, M. A. (eds.), *PRICAI 2010: Trends in Artificial Intelligence*, pp. 8–10, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. ISBN 978-3-642-15246-7.
- [50] Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and de Freitas, N. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.
- [51] Suzuki, S., Takeno, S., Tamura, T., Shitara, K., and Karasuyama, M. Multi-objective Bayesian optimization using pareto-frontier entropy. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9279–9288. PMLR, 13–18 Jul 2020. URL <http://proceedings.mlr.press/v119/suzuki20a.html>.
- [52] Takeno, S., Fukuoka, H., Tsukada, Y., Koyama, T., Shiga, M., Takeuchi, I., and Karasuyama, M. Multi-fidelity Bayesian optimization with max-value entropy search and its parallelization. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9334–9345. PMLR, 13–18 Jul 2020. URL <http://proceedings.mlr.press/v119/takeno20a.html>.
- [53] Tanabe, R. and Ishibuchi, H. An easy-to-use real-world multi-objective optimization problem suite. *Applied Soft Computing*, 89:106078, 2020. ISSN 1568-4946. doi: <https://doi.org/10.1016/j.asoc.2020.106078>.
- [54] Thompson, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- [55] Touré, C., Hansen, N., Auger, A., and Brockhoff, D. Uncrowded hypervolume improvement: Como-cma-es and the sofomore framework. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO ’19*, pp. 638–646, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450361118. doi: 10.1145/3321707.3321852. URL <https://doi.org/10.1145/3321707.3321852>.

- [56] Wang, R., Xiong, J., Ishibuchi, H., Wu, G., and Zhang, T. On the effect of reference point in moea/d for multi-objective optimization. *Applied Soft Computing*, 58:25–34, 2017. ISSN 1568-4946. doi: <https://doi.org/10.1016/j.asoc.2017.04.002>. URL <https://www.sciencedirect.com/science/article/pii/S1568494617301722>.
- [57] Wilson, J., Hutter, F., and Deisenroth, M. Maximizing acquisition functions for bayesian optimization. In *Advances in Neural Information Processing Systems 31*, pp. 9905–9916. 2018.
- [58] Yang, K., Emmerich, M., Deutz, A., and Bäck, T. Multi-objective bayesian global optimization using expected hypervolume improvement gradient. *Swarm and Evolutionary Computation*, 44: 945 – 956, 2019. ISSN 2210-6502. doi: <https://doi.org/10.1016/j.swevo.2018.10.007>.
- [59] Yang, K., Emmerich, M., Deutz, A. H., and Bäck, T. Efficient computation of expected hypervolume improvement using box decomposition algorithms. *CoRR*, abs/1904.12672, 2019.
- [60] Yang, K., Palar, P., Emmerich, M., Shimoyama, K., and Bäck, T. A multi-point mechanism of expected hypervolume improvement for parallel multi-objective bayesian global optimization. pp. 656–663, 07 2019. doi: 10.1145/3321707.3321784.
- [61] Yang, K., Palar, P. S., Emmerich, M., Shimoyama, K., and Bäck, T. A multi-point mechanism of expected hypervolume improvement for parallel multi-objective bayesian global optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '19*, pp. 656–663, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450361118. doi: 10.1145/3321707.3321784. URL <https://doi.org/10.1145/3321707.3321784>.
- [62] Youn, B. D., Choi, K., Yang, R.-J., and Gu, L. Reliability-based design optimization for crashworthiness of vehicle side impact. *Structural and Multidisciplinary Optimization*, 26: 272–283, 02 2004. doi: 10.1007/s00158-003-0345-0.
- [63] Zhang, G. and Block, D. E. Using highly efficient nonlinear experimental design methods for optimization of lactococcus lactis fermentation in chemically defined media. *Biotechnology progress*, 25(6):1587–1597, 2009.
- [64] Zhang, Q., Liu, W., Tsang, E., and Virginas, B. Expensive multiobjective optimization by moea/d with gaussian process model. *IEEE Transactions on Evolutionary Computation*, 14(3): 456–474, 2010. doi: 10.1109/TEVC.2009.2033671.
- [65] Zhou, A., Zhang, Q., and Zhang, G. A multiobjective evolutionary algorithm based on decomposition and probability model. In *2012 IEEE Congress on Evolutionary Computation*, pp. 1–8, 2012. doi: 10.1109/CEC.2012.6252954.
- [66] Zitzler, E., Deb, K., and Thiele, L. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evol. Comput.*, 8(2):173–195, June 2000. ISSN 1063-6560. doi: 10.1162/106365600568202. URL <https://doi.org/10.1162/106365600568202>.
- [67] Zitzler, E., Deb, K., and Thiele, L. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary computation*, 8(2):173–195, 2000.