
Differentiable Expected Hypervolume Improvement for Parallel Multi-Objective Bayesian Optimization

Samuel Daulton
Facebook
sdaulton@fb.com

Maximilian Balandat
Facebook
balandat@fb.com

Eytan Bakshy
Facebook
ebakshy@fb.com

Abstract

In many real-world scenarios, decision makers seek to efficiently optimize multiple competing objectives in a sample-efficient fashion. Multi-objective Bayesian optimization (BO) is a common approach, but many of the best-performing acquisition functions do not have known analytic gradients and suffer from high computational overhead. We leverage recent advances in programming models and hardware acceleration for multi-objective BO using Expected Hypervolume Improvement (EHVI)—an algorithm notorious for its high computational complexity. We derive a novel formulation of q -Expected Hypervolume Improvement (q EHVI), an acquisition function that extends EHVI to the parallel, constrained evaluation setting. q EHVI is an exact computation of the joint EHVI of q new candidate points (up to Monte-Carlo (MC) integration error). Whereas previous EHVI formulations rely on gradient-free acquisition optimization or approximated gradients, we compute exact gradients of the MC estimator via auto-differentiation, thereby enabling efficient and effective optimization using first-order and quasi-second-order methods. Our empirical evaluation demonstrates that q EHVI is computationally tractable in many practical scenarios and outperforms state-of-the-art multi-objective BO algorithms at a fraction of their wall time.

1 Introduction

The problem of optimizing multiple competing objectives is ubiquitous in scientific and engineering applications. For example in automobile design, an automaker will want to maximize vehicle durability and occupant safety, while using lighter materials that afford increased fuel efficiency and lower manufacturing cost [44, 72]. Evaluating the crash safety of an automobile design experimentally is expensive due to both the manufacturing time and the destruction of a vehicle. In such a scenario, sample efficiency is paramount. For a different example, video streaming web services commonly use adaptive control policies to determine the bitrate as the stream progresses in real time [47]. A decision maker may wish to optimize the control policy to maximize the quality of the video stream, while minimizing the stall time. Policy evaluation typically requires using the suggested policy on segments of live traffic, which is subject to opportunity costs. If long evaluation times are the limiting factor, multiple designs may be evaluated in parallel to significantly decrease end-to-end optimization time. For example, an automaker could manufacture multiple vehicle designs in parallel or a web service could deploy several control policies to different segments of traffic at the same time.

1.1 Background

Multi-Objective Optimization: In this work, we address the problem of optimizing a *vector-valued objective* $\mathbf{f}(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}^M$ with $\mathbf{f}(\mathbf{x}) = (f^{(1)}(\mathbf{x}), \dots, f^{(M)}(\mathbf{x}))$ over a bounded set $\mathcal{X} \subset \mathbb{R}^d$. We consider the scenario in which the $f^{(i)}$ are expensive-to-evaluate black-box functions with

no known analytical expression, and no observed gradients. *Multi-objective* (MO) optimization problems typically do not have a single best solution; rather, the goal is to identify the set of *Pareto optimal* solutions such that any improvement in one objective means deteriorating another. Without loss of generality, we assume the goal is to maximize all objectives. We say a solution $\mathbf{f}(\mathbf{x})$ *Pareto dominates* another solution $\mathbf{f}(\mathbf{x}')$ if $f^{(m)}(\mathbf{x}) \geq f^{(m)}(\mathbf{x}') \forall m = 1, \dots, M$ and there exists $m' \in \{1, \dots, M\}$ such that $f^{(m')}(\mathbf{x}) > f^{(m')}(\mathbf{x}')$. We write $\mathbf{f}(\mathbf{x}) \succ \mathbf{f}(\mathbf{x}')$. Let $\mathcal{P}^* = \{\mathbf{f}(\mathbf{x}) \text{ s.t. } \nexists \mathbf{x}' \in \mathcal{X} : \mathbf{f}(\mathbf{x}') \succ \mathbf{f}(\mathbf{x})\}$ and $\mathcal{X}^* = \{\mathbf{x} \in \mathcal{X} \text{ s.t. } \mathbf{f}(\mathbf{x}) \in \mathcal{P}^*\}$ denote the set of Pareto optimal solutions and Pareto optimal inputs, respectively. Provided with the Pareto set, decision-makers can select a solution with an objective trade-off according to their preferences.

A common approach for solving MO problems is to use evolutionary algorithms (e.g. NSGA-II), which are robust multi-objective optimizers, but require a large number of function evaluations [14]. Bayesian optimization (BO) offers a far more sample-efficient alternative [57].

Bayesian Optimization: BO [38] is an established method for optimizing expensive-to-evaluate black-box functions. BO relies on a probabilistic *surrogate model*, typically a Gaussian Process (GP) [55], to provide a posterior distribution $\mathbb{P}(\mathbf{f}|\mathcal{D})$ over the true function values \mathbf{f} given the observed data $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$. An *acquisition function* $\alpha : \mathcal{X}_{\text{cand}} \mapsto \mathbb{R}$ employs the surrogate model to assign a utility value to a set of candidates $\mathcal{X}_{\text{cand}} = \{\mathbf{x}_i\}_{i=1}^q$ to be evaluated on the true function. While the true \mathbf{f} may be expensive-to-evaluate, the surrogate-based acquisition function is not, and can thus be efficiently optimized to yield a set of candidates $\mathcal{X}_{\text{cand}}$ to be evaluated on \mathbf{f} . If gradients of $\alpha(\mathcal{X}_{\text{cand}})$ are available, gradient-based methods can be utilized. If not, gradients are either approximated (e.g. with finite differences) or gradient-free methods (e.g. DIRECT [37] or CMA-ES [32]) are used.

1.2 Limitations of current approaches

In the single-objective (SO) setting, a large body of work focuses on practical extensions to BO for supporting *parallel* evaluation and outcome constraints [49, 30, 66, 25, 43]. Less attention has been given to such extensions in the MO setting. Moreover, the existing constrained and parallel MO BO options have limitations: 1) many rely on scalarizations to transform the MO problem into a SO one [40]; 2) many acquisition functions are computationally expensive to compute [52, 21, 6, 71]; 3) few have known analytical gradients or are differentiable [19, 62, 33]; 4) many rely on heuristics to extend sequential algorithms to the parallel setting [27, 62].

A natural acquisition function for MO BO is Expected Hypervolume Improvement (EHVI). Maximizing the hypervolume (HV) has been shown to produce Pareto fronts with excellent coverage [73, 12, 69]. However, there has been little work on EHVI in the parallel setting, and the work that has been done resorts to approximate methods [71, 28, 62]. A vast body of literature has focused on efficient EHVI computation [34, 20, 67], but the time complexity for computing EHVI is exponential in the number of objectives—in part due the hypervolume indicator itself incurring a time complexity that scales super-polynomially with the number of objectives [68]. Our core insight is that by exploiting advances in auto-differentiation and highly parallelized hardware [51], we can make EHVI computations fast and practical.

1.3 Contributions

In this work, we derive a novel formulation of the parallel q -Expected Hypervolume Improvement acquisition function (q EHVI) that is exact up to Monte-Carlo (MC) integration error. We compute the *exact* gradient of the MC estimator of q EHVI using auto-differentiation, which allows us to employ efficient and effective gradient-based optimization methods. Rather than using first-order gradient methods, we instead leverage the sample average approximation (SAA) approach from [5] to use higher-order deterministic optimization methods, and we prove theoretical convergence guarantees under the SAA approach. Our formulation of q EHVI is embarrassingly parallel, and despite its computational cost would achieve constant time complexity given infinite processing cores. We demonstrate that, using modern GPU hardware and computing exact gradients, optimizing q EHVI is faster than existing state-of-the-art methods in many practical scenarios. Moreover, we extend q EHVI to support auxiliary outcome constraints, making it practical in many real-world scenarios. Lastly, we demonstrate how modern auto-differentiation can be used to compute exact gradients of analytic EHVI, which has never been done before for $M > 2$ objectives. Our empirical evaluation

shows that q EHVI outperforms state-of-the-art multi-objective BO algorithms while using only a fraction of their wall time.

2 Related Work

Yang et al. [69] is the only previous work to consider exact gradients of EHVI, but the authors only derive an analytical gradient for the unconstrained two-objective, sequential optimization setting. All other works either do not optimize EHVI (e.g. they use it for pre-screening candidates [18]), optimize it with gradient-free methods [68], or using approximate gradients [62]. In contrast, we use exact gradients and demonstrate that optimizing EHVI using this gradient information is far more efficient.

There are many alternatives to EHVI for MO BO. For example, ParEGO [40] and TS-TCH [50] randomly scalarize the objectives and use Expected Improvement [38] and Thompson Sampling [61], respectively. SMS-EGO [53] uses HV in a UCB-based acquisition function and is more scalable than EHVI [54]. ParEGO and SMS-EGO have only been considered for the $q = 1$, unconstrained setting. Predictive entropy search for MO BO (PESMO) [33] has been shown to be another competitive alternative and has been extended to handle constraints [26] and parallel evaluations [27]. MO max-value entropy search (MO-MES) has been shown to achieve superior optimization performance and faster wall times than PESMO, but is limited to $q = 1$.

Wilson et al. [65] empirically and theoretically show that *sequential greedy* selection of q candidates achieves performance comparable to jointly optimizing q candidates for many acquisition functions (including [63, 66]). The sequential greedy approach integrates over the posterior of the unobserved outcomes corresponding to the previously selected candidates in the q -batch. Sequential greedy optimization often yields better empirical results because the optimization problem has a lower dimension: d in each step, rather than qd in the joint problem. Most prior works in the MO setting use a sequential greedy approximation or heuristics [62, 71, 28, 10], but impute the unobserved outcomes with the posterior mean rather than integrating over the posterior [30]. For many joint acquisition functions involving expectations, this shortcut sacrifices the theoretical error bound on the sequential greedy approximation because the exact joint acquisition function over $\mathbf{x}_1, \dots, \mathbf{x}_i$, $1 \leq i \leq q$ requires integration over the joint posterior $\mathbb{P}(\mathbf{f}(\mathbf{x}_1), \dots, \mathbf{f}(\mathbf{x}_q) | \mathcal{D})$ and is not computed for $i > 1$.

Garrido-Merchán and Hernández-Lobato [27] and Wada and Hino [62] jointly optimize the q candidates and, noting the difficulty of the optimization, both papers focus on deriving gradients to aid in the optimization. Wada and Hino [62] defined the q EHVI acquisition function, but after finding it challenging to optimize q candidates jointly (without exact gradients), the authors propose optimizing an alternative acquisition function instead of exact q EHVI. In contrast, our novel q EHVI formulation allows for gradient-based parallel and sequential greedy optimization, with proper integration over the posterior for the latter.

Feliot et al. [22] and Abdolshah et al. [1] proposed extensions of EHVI to the constrained $q = 1$ setting, but neither considers the batch setting and both rely on gradient-free optimization.

3 Differentiable q -Expected Hypervolume Improvement

In this section, we review HVI and EHVI computation by means of box decompositions, and explain our novel formulation for the parallel setting.

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

Definition 2. Given a Pareto set \mathcal{P} and reference point \mathbf{r} , the hypervolume improvement (HVI) of a set of points \mathcal{Y} is: $\text{HVI}(\mathcal{Y}, \mathcal{P}, \mathbf{r}) = \text{HV}(\mathcal{P} \cup \mathcal{Y}, \mathbf{r}) - \text{HV}(\mathcal{P}, \mathbf{r})$.¹

EHVI is the expectation of HVI over the posterior $\mathbb{P}(\mathbf{f}, \mathcal{D})$: $\alpha_{\text{EHVI}}(\mathcal{X}_{\text{cand}}) = \mathbb{E}[\text{HVI}(\mathbf{f}(\mathcal{X}_{\text{cand}}))]$. In the sequential setting, and assuming the objectives are independent and modeled with independent

¹In this work, we omit the arguments \mathcal{P} and \mathbf{r} when referring to HVI for brevity.

GPs, EHVI can be expressed in closed form [69]. In other settings, EHVI can be approximated with MC integration. Following previous work, we assume that the reference point is known and specified by the decision maker [69] (see Appendix E.1.1 for additional discussion).

3.1 A review of hypervolume improvement computation using box decompositions

Definition 3. For a set of objective vectors $\{\mathbf{f}(\mathbf{x}_i)\}_{i=1}^q$, a reference point $\mathbf{r} \in \mathbb{R}^M$, and a non-dominated set \mathcal{P} , let $\Delta(\{\mathbf{f}(\mathbf{x}_i)\}_{i=1}^q, \mathcal{P}, \mathbf{r}) \subset \mathbb{R}^M$ denote the set of points (i) are dominated by $\{\mathbf{f}(\mathbf{x}_i)\}_{i=1}^q$, dominate \mathbf{r} , and are not dominated by \mathcal{P} .

Given \mathcal{P}, \mathbf{r} , the HVI of a new point $\mathbf{f}(\mathbf{x})$ is the HV of the intersection of space dominated by $\mathcal{P} \cup \{\mathbf{f}(\mathbf{x})\}$ and the non-dominated space. Figure 1b illustrates this for one new point $\mathbf{f}(\mathbf{x})$ for $M = 2$. The yellow region is $\Delta(\{\mathbf{f}(\mathbf{x})\}, \mathcal{P}, \mathbf{r})$ and the hypervolume improvement is the volume covered by $\Delta(\{\mathbf{f}(\mathbf{x})\}, \mathcal{P}, \mathbf{r})$. Since $\Delta(\{\mathbf{f}(\mathbf{x})\}, \mathcal{P}, \mathbf{r})$ is often a non-rectangular polytope, HVI is typically computed by partitioning the non-dominated space into disjoint axis-parallel rectangles [12, 68] (see Figure 1a) and using piece-wise integration [18].

Let $\{S_k\}_{k=1}^K$ be a partitioning the of non-dominated space into disjoint hyper-rectangles, where each S_k is defined by a pair of lower and upper vertices $\mathbf{l}_k \in \mathbb{R}^M$ and $\mathbf{u}_k \in \mathbb{R}^M \cup \{\infty\}$. The high level idea is to sum the HV of $S_k \cap \Delta(\{\mathbf{f}(\mathbf{x})\}, \mathcal{P}, \mathbf{r})$ over all S_k . For each hyper-rectangle S_k , the intersection of S_k and $\Delta(\{\mathbf{f}(\mathbf{x})\}, \mathcal{P}, \mathbf{r})$ is a hyper-rectangle where the lower bound vertex is \mathbf{l}_k and the upper bound vertex is the component-wise minimum of \mathbf{u}_k and the new point $\mathbf{f}(\mathbf{x})$: $\mathbf{z}_k := \min[\mathbf{u}_k, \mathbf{f}(\mathbf{x})]$.

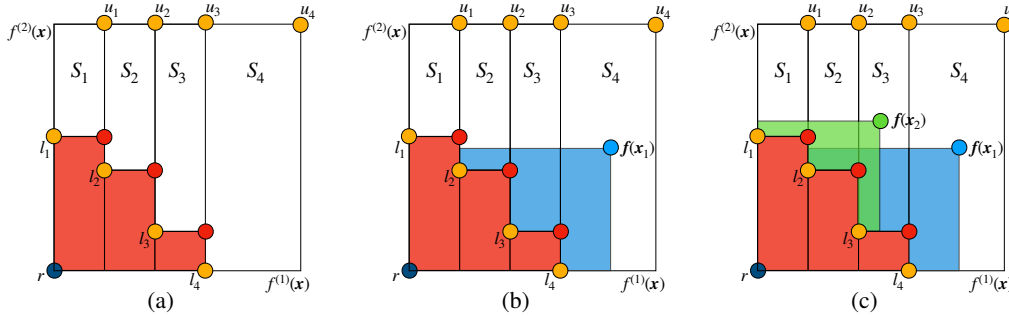


Figure 1: For $M=2$, (a) the dominated space (red) and the non-dominated space partitioned into disjoint boxes (white), (b) the HVI of one new point $\mathbf{f}(\mathbf{x})$, and (c) the HVI of two new points $\mathbf{f}(\mathbf{x}_1), \mathbf{f}(\mathbf{x}_2)$.

Hence, the HVI of a single outcome vector $\mathbf{f}(\mathbf{x})$ within S_k is given by $\text{HVI}_k(\mathbf{f}(\mathbf{x}), \mathbf{l}_k, \mathbf{u}_k) = \lambda_M(S_k \cap \Delta(\{\mathbf{f}(\mathbf{x})\}, \mathcal{P}, \mathbf{r})) = \prod_{m=1}^M [z_k^{(m)} - l_k^{(m)}]_+$, where $u_k^{(m)}, l_k^{(m)}, f^{(m)}(\mathbf{x})$, and $z_k^{(m)}$ denote the m^{th} component of the corresponding vector and $[\cdot]_+$ denotes the $\min(\cdot, 0)$ operation. Summing over rectangles yields

$$\text{HVI}(\mathbf{f}(\mathbf{x})) = \sum_{k=1}^K \text{HVI}_k(\mathbf{f}(\mathbf{x}), \mathbf{l}_k, \mathbf{u}_k) = \sum_{k=1}^K \prod_{m=1}^M [z_k^{(m)} - l_k^{(m)}]_+ \quad (1)$$

3.2 Computing q -Hypervolume Improvement via the Inclusion-Exclusion Principle

Figure 1c illustrates the HVI in the $q = 2$ setting. Given q new points $\{\mathbf{f}(\mathbf{x}_i)\}_{i=1}^q$, let $A_i := \Delta(\{\mathbf{f}(\mathbf{x}_i)\}, \mathcal{P}, \mathbf{r})$ for $i = 1, \dots, q$ be the space dominated by $\mathbf{f}(\mathbf{x}_i)$ but not dominated by \mathcal{P} , independently of the other $q - 1$ points. Note that $\lambda_M(A_i) = \text{HVI}(\mathbf{f}(\mathbf{x}_i))$. The union of the subsets A_i is the space dominated jointly by the q new points: $\bigcup_{i=1}^q A_i = \bigcup_{i=1}^q \Delta(\{\mathbf{f}(\mathbf{x}_i)\}, \mathcal{P}, \mathbf{r})$, and the Lebesgue measure $\lambda_M(\bigcup_{i=1}^q A_i)$ is the joint HVI from the q new points. Since each subspace A_i is bounded, the restricted Lebesgue measure is finite and we may compute $\lambda_M(\bigcup_{i=1}^q A_i)$ using the inclusion-exclusion principle [13, 59]:

$$\text{HVI}(\{\mathbf{f}(\mathbf{x}_i)\}_{i=1}^q) = \lambda_M\left(\bigcup_{i=1}^q A_i\right) = \sum_{j=1}^q (-1)^{j+1} \sum_{1 \leq i_1 \leq \dots \leq i_j \leq q} \lambda_M(A_{i_1} \cap \dots \cap A_{i_j}) \quad (2)$$

Since $\{S_k\}_{k=1}^K$ is a disjoint partition, $\lambda_M(A_{i_1} \cap \dots \cap A_{i_j}) = \sum_{k=1}^K \lambda_M(S_k \cap A_{i_1} \cap \dots \cap A_{i_j})$, we can compute $\lambda_M(A_{i_1} \cap \dots \cap A_{i_j})$ in a piece-wise fashion across the K hyper-rectangles $\{S_k\}_{k=1}^K$ as the HV of the intersection of $A_{i_1} \cap \dots \cap A_{i_j}$ with each hyper-rectangle S_k . The inclusion-exclusion principle has been proposed for computing HV (not HVI) [45], but it is rarely used because complexity scales exponentially with the number of elements. However, the inclusion-exclusion principle is practical for computing the joint HVI of q points since typically $q \ll |\mathcal{P}|$.

This formulation has three advantages. First, while the new dominated space A_i can be a non-rectangular polytope, the intersection $A_i \cap S_k$ is a *rectangular* polytope, which simplifies computation of overlapping hypervolume. Second, the vertices defining the hyper-rectangle $S_k \cap A_{i_1} \cap \dots \cap A_{i_j}$ are easily derived. The lower bound is simply the l_k lower bound of S_k , and the upper bound is the component-wise minimum $z_{k,i_1,\dots,i_j} := \min[\mathbf{u}_k, \mathbf{f}(\mathbf{x}_{i_1}), \dots, \mathbf{f}(\mathbf{x}_{i_j})]$. Third, computation can be across all intersections of subsets $A_{i_1} \cap \dots \cap A_{i_j}$ for $1 \leq i_1 \leq \dots \leq i_j \leq q$ and across all K hyper-rectangles can be performed in parallel. Explicitly, the HVI is computed as:

$$\text{HVI}(\{\mathbf{f}(\mathbf{x}_i)\}_{i=1}^q) = \sum_{k=1}^K \sum_{j=1}^q \sum_{X_j \in \mathcal{X}_j} (-1)^{j+1} \prod_{m=1}^M [z_{k,X_j}^{(m)} - l_k^{(m)}]_+ \quad (3)$$

where $\mathcal{X}_j := \{X_j \subset \mathcal{X}_{\text{cand}} : |X_j| = j\}$ is the superset of all subsets of $\mathcal{X}_{\text{cand}}$ of size j , and $z_{k,X_j}^{(m)} := z_{k,i_1,\dots,i_j}^{(m)}$ for $X_j = \{\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_j}\}$. See Appendix A for further details of the derivation.

3.3 Computing Expected q -Hypervolume Improvement

The above approach for computing HVI assumes that we know the true objective values $\mathbf{f}(\mathcal{X}_{\text{cand}}) = \{\mathbf{f}(\mathbf{x}_i)\}_{i=1}^q$. In BO, we instead compute $q\text{EHVI}$ as the expectation over the posterior model posterior:

$$\alpha_{q\text{EHVI}}(\mathcal{X}_{\text{cand}}) = \mathbb{E}[\text{HVI}(\mathbf{f}(\mathcal{X}_{\text{cand}}))] = \int_{-\infty}^{\infty} \text{HVI}(\mathbf{f}(\mathcal{X}_{\text{cand}})) d\mathbf{f}. \quad (4)$$

Since no known analytical form is known [70] for $q > 1$ (or in the case of correlated outcomes), we estimate (4) using MC integration with samples from the joint posterior $\{\mathbf{f}_t(\mathbf{x}_i)\}_{i=1}^q \sim \mathbb{P}(\mathbf{f}(\mathbf{x}_1), \dots, \mathbf{f}(\mathbf{x}_q) | \mathcal{D}), t = 1, \dots, N$. Let $z_{k,X_j,t}^{(m)} := \min[\mathbf{u}_k, \min_{\mathbf{x}' \in X_j} \mathbf{f}_t(\mathbf{x}')]_+$. Then,

$$\hat{\alpha}_{q\text{EHVI}}^N(\mathcal{X}_{\text{cand}}) = \frac{1}{N} \sum_{t=1}^N \text{HVI}(\mathbf{f}_t(\mathcal{X}_{\text{cand}})) = \frac{1}{N} \sum_{t=1}^N \sum_{k=1}^K \sum_{j=1}^q \sum_{X_j \in \mathcal{X}_j} (-1)^{j+1} \prod_{m=1}^M [z_{k,X_j,t}^{(m)} - l_k^{(m)}]_+ \quad (5)$$

Provided that $\{S_k\}_{k=1}^K$ is an exact partitioning, (5) is an *exact* computation of $q\text{EHVI}$ up to the MC estimation error, which scales as $1/\sqrt{N}$ when using *iid* MC samples regardless of the dimension of the search space [18]. In practice, we use randomized quasi MC methods [8] to reduce the variance and empirically observe low estimation error (see Figure 5a in the Appendix for a comparison of analytic EHVI and (quasi-)MC-based $q\text{EHVI}$).

$q\text{EHVI}$ requires computing the volume of $2^q - 1$ hyper-rectangles (the number of subsets of q) for each of K hyper-rectangles and N MC samples. Given posterior samples, the time complexity on a single-threaded machine is: $T_1 = O(MNK(2^q - 1))$. In the two-objective case, $K = |\mathcal{P}| + 1$, but K is super-polynomial in M [68]. The number of boxes required for a decomposition of the non-dominated space is unknown for $M \geq 4$ [68]. $q\text{EHVI}$ is agnostic to the partitioning algorithm used, and in F.4 we demonstrate using $q\text{EHVI}$ in higher-dimensional objective spaces using an approximate box decomposition algorithm [11]. Despite the daunting workload, the critical work path—the time complexity of the smallest non-parallelizable unit—is constant: $T_\infty = O(1)$.² On highly-threaded many-core hardware (e.g. GPUs), our formulation achieves tractable wall times in many practical scenarios: as is shown in Figure 11 in the Appendix, the computation time is nearly constant with increasing q until an inflection point at which the workload saturates the available cores. For additional discussion of both time and memory complexity of $q\text{EHVI}$ see Appendix A.4.

3.4 Outcome Constraints

Our proposed $q\text{EHVI}$ acquisition function is easily extended to constraints on auxiliary outcomes. We consider the scenario where we receive observations of M objectives $\mathbf{f}(\mathbf{x}) \in \mathbb{R}^M$ and V constraints

²As evident from (5), the critical path consists of 3 multiplications and 5 summations.

$\mathbf{c}^{(v)} \in \mathbb{R}^V$, all of which are assumed to be “black-box”. We assume w.l.o.g. that $\mathbf{c}^{(v)}$ is feasible iff $\mathbf{c}^{(v)} \geq \mathbf{0}$. In the constrained optimization setting, we aim to identify the feasible Pareto set: $\mathcal{P}_{\text{feas}} = \{\mathbf{f}(\mathbf{x}) \text{ s.t. } \mathbf{c}(\mathbf{x}) \geq \mathbf{0}, \nexists \mathbf{x}' : \mathbf{c}(\mathbf{x}') \geq \mathbf{0}, \mathbf{f}(\mathbf{x}') \succ \mathbf{f}(\mathbf{x})\}$. The natural improvement measure in the constrained setting is *feasible* HVI, which we define for a single candidate point \mathbf{x} as $\text{HVI}_c(\mathbf{f}(\mathbf{x}), \mathbf{c}(\mathbf{x})) := \text{HVI}[\mathbf{f}(\mathbf{x})] \cdot \mathbb{1}[\mathbf{c}(\mathbf{x}) \geq \mathbf{0}]$. Taking expectations, the constrained expected HV can be seen to be the HV weighted by the probability of feasibility. In Appendix A.3, we detail how performing feasibility-weighting on the sample-level allows us to include such auxiliary outcome constraints into our MC formulation in a straightforward way.

4 Optimizing q -Expected Hypervolume Improvement

4.1 Differentiability

While an analytic formula for the gradient of EHVI exists for the $M = 2$ objective case in the unconstrained, sequential ($q = 1$) setting, no such formula is known in 1) the case of $M > 2$ objectives, 2) the constrained setting, and 3) for $q > 1$. Leveraging the re-parameterization trick [39, 64] and auto-differentiation, we are able to automatically compute exact gradients of the MC-estimator $q\text{EHVI}$ in *all* of the above settings, as well as the gradient of analytic EHVI for $M \geq 2$ (see Figure 5b in the Appendix for a comparison of the exact gradients of EHVI and the sample average gradients of $q\text{EHVI}$ for $M = 3$).^{3,4}

4.2 Optimization via Sample Average Approximation

We show in Appendix C that if mean and covariance function of the GP are sufficiently regular, the gradient of the MC estimator (5) is an unbiased estimate of the gradient of the exact acquisition function (4). To maximize $q\text{EHVI}$, we could therefore directly apply stochastic optimization methods, as has previously been done for single-outcome acquisition functions [64, 66]. Instead, we opt to use the sample average approximation (SAA) approach from Balandat et al. [5], which allows us to employ deterministic, higher-order optimizers to achieve faster convergence rates. Informally (see Appendix C for the formal statement), if $\hat{\mathbf{x}}_N^* \in \arg \max_{\mathbf{x} \in \mathcal{X}} \hat{\alpha}_{q\text{EHVI}}^N(\mathbf{x})$, we can show under some regularity conditions that, as $N \rightarrow \infty$, (i) $\hat{\alpha}_{q\text{EHVI}}^N(\hat{\mathbf{x}}_N^*) \rightarrow \max_{\mathbf{x} \in \mathcal{X}} \alpha_{q\text{EHVI}}(\mathbf{x})$ *a.s.*, and (ii) $\text{dist}(\hat{\mathbf{x}}_N^*, \arg \max_{\mathbf{x} \in \mathcal{X}} \alpha_{q\text{EHVI}}(\mathbf{x})) \rightarrow 0$ *a.s.*. These results hold for any covariance function satisfying the regularity conditions, including such ones that model correlation between outcomes. In particular, our results do not require the outputs to be modeled by independent GPs.

Figure 2a demonstrates the importance of using exact gradients for efficiently and effectively optimizing EHVI and $q\text{EHVI}$ by comparing the following optimization methods: L-BFGS-B with exact gradients, L-BFGS-B with gradients approximated via finite differences, and CMA-ES (without gradients). The cumulative time spent optimizing the acquisition function is an order of magnitude less when using exact gradients rather than approximate gradients or zeroth order methods.

4.3 Sequential Greedy and Joint Batch Optimization

Jointly optimizing q candidates increases in difficulty with q because the problem dimension is dq . An alternative is to sequentially and greedily select candidates and condition the acquisition function on the previously selected pending points when selecting the next point [65]. Using a submodularity argument similar to that in Wilson et al. [64], the sequential greedy approximation of $q\text{EHVI}$ enjoys regret of no more than $\frac{1}{e} \alpha_{q\text{EHVI}}^*$, where $\alpha_{q\text{EHVI}}^*$ is the optima of $\alpha_{q\text{EHVI}}$ [23] (see Appendix B).

Although sequential greedy approaches have been considered for many acquisition functions [65], no previous work has proposed a proper sequential greedy approach (with integration over the posterior) for parallel EHVI, as this would require computing the Pareto front under each sample \mathbf{f}_i from the joint posterior before computing the hypervolume improvement. These operations would be computationally expensive for even modest N and non-differentiable. $q\text{EHVI}$ avoids determining the Pareto set for each sample by using inclusion-exclusion principle to compute the joint HVI over the pending points $\mathbf{x}_1, \dots, \mathbf{x}_{i-1}$ and the new candidate \mathbf{x}_i for each MC sample. Figure 2b empirically

³Technically, min and max are only sub-differentiable, but are known to be well-behaved [64]. In our MC setting with GP posteriors, $q\text{EHVI}$ is differentiable w.p. 1 if \mathbf{x} contains no repeated points.

⁴For the constrained case, we replace the indicator with a differentiable sigmoid approximation.

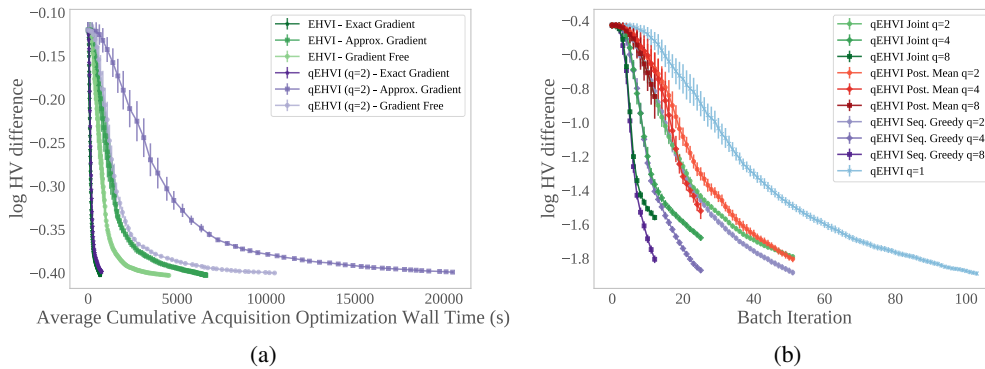


Figure 2: (a) A comparison of EHVI and q EHVI ($q = 2$) optimized with L-BFGS-B using exact gradients, L-BFGS-B using gradients approximated using finite differences, and CMA-ES, a gradient-free method. (b) A comparison of joint optimization, sequential greedy optimization with proper integration at the pending points, and sequential greedy using the posterior mean. Both plots show optimization performance on a DTLZ2 problem ($d = 6, M = 2$) with a budget of 100 evaluations (plus the initial quasi-random design). We report means and 2 standard errors across 20 trials.

demonstrates the improved optimization performance from properly integrating over the unobserved outcomes rather than using the posterior mean or jointly optimizing the q candidates.

5 Benchmarks

We empirically evaluate q EHVI on synthetic and real world optimization problems. We compare q EHVI⁵ against existing state-of-the-art methods including SMS-EGO⁶, PESMO⁶, TS-TCH⁵, and analytic EHVI [68] with gradients⁵. Additionally, we compare against a novel extension of ParEGO [40] that supports parallel evaluation and constraints (neither of which have been done before to our knowledge); we call this method q PAREGO⁵. Additionally, we include a quasi-random baseline that selects candidates from a scrambled Sobol sequence. See Appendix E.1 for details on all baseline algorithms.

Synthetic Benchmarks We evaluate optimization performance on four benchmark problems in terms of log hypervolume difference, which is defined as the difference between the hypervolume of the true (feasible) Pareto front and the hypervolume of the approximate (feasible) Pareto front based on the observed data; in the case that the true Pareto front is unknown (or not easily approximated), we evaluate the hypervolume indicator. All reference points and search spaces are provided in Appendix E.2. For synthetic problems, we consider the Branin-Currin problem ($d = 2, M = 2$, convex Pareto front) [6] and the C2-DTLZ2 ($d = 12, M = 2, V = 1$, concave Pareto front), which is a standard constrained benchmark from the MO literature [16] (see Appendix F.1 for additional synthetic benchmarks).

Real-World Benchmarks *Structural Optimization in Automobile Safety Design* (VEHICLESAFETY): Vehicle crash safety is an important consideration in the structural design of automobiles. A lightweight car is preferable because of its potentially lower manufacturing cost and better fuel economy, but lighter material can fare worse than sturdier alternatives in a collision, potentially leading to increased vehicle damage and more severe injury to the vehicle occupants [72]. We consider the problem designing the thickness of 5 reinforced parts of the frontal frame of a vehicle that considerably affect crash safety. The goal is to minimize: 1) the *mass* of the vehicle; 2) the *collision acceleration* in a full frontal crash—a proxy for bio-mechanical trauma to the vehicle occupants from the acceleration; and 3) the *toe-board intrusion*—a measure of the most extreme

⁵Acquisition functions are available as part of the open-source library BoTorch [5]. Code is available at <https://github.com/pytorch/botorch>

⁶We leverage existing implementations from the Spearmint library. The code is available at <https://github.com/HIPS/Spearmint/tree/PESM>

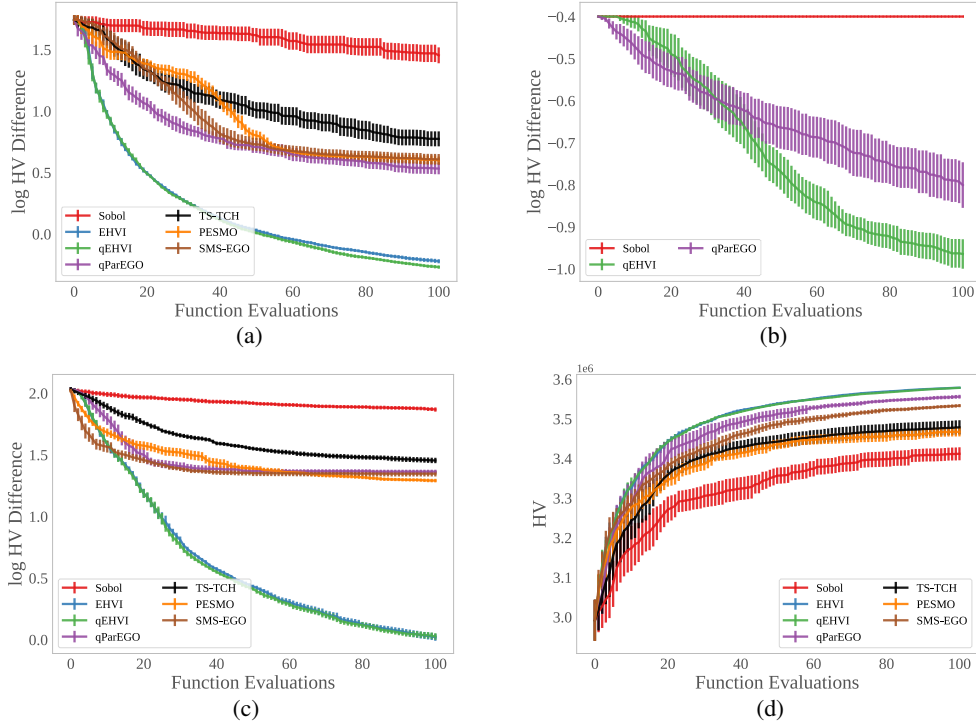


Figure 3: Sequential optimization performance on (a) on the Branin-Currin problem ($q = 1$), (b) the C2-DTLZ2 problem, (c) the vehicle crash safety problem ($q = 1$), and (d) the ABR control problem ($q = 1$). We report the means and 2 standard errors across 20 trials.

mechanical damage to the vehicle in an off-frontal collision [44]. For this problem, we optimize the surrogate from Tanabe and Ishibuchi [60].

Policy Optimization for Adaptive Bitrate Control (ABR): Many web services adapt video playback quality adaptively based on the receiver’s network bandwidth to maintain steady, high quality stream with minimal stalls and buffer periods [47]. Previous works have proposed controllers with different scalarized objective functions [46], but in many cases, engineers may prefer to learn the set of optimal trade-offs between their metrics of interest, rather than specifying a scalarized objective in advance. In this problem, we decompose the objective function proposed in Mao et al. [46] into its constituent metrics and optimize 4 parameters of an ABR control policy on the Park simulator [48] to maximize video quality (bitrate) and minimize stall time. See Appendix E.2 for details.

5.1 Results

Figure 3 shows that $qEHVI$ outperforms all baselines in terms of sequential optimization performance on all evaluated problems. Table 1 shows that $qEHVI$ achieves wall times that are an order of magnitude smaller than those of PESMO on a CPU in sequential optimization, and maintains competitive wall times even relative to $qPAREGO$ (which has a significantly smaller workload) for large q on a GPU. TS-TCH has by far the fastest wall time, but this comes at the cost of inferior optimization performance.

Figure 4 illustrates optimization performance of *parallel* acquisition functions for varying batch sizes. Increasing the level of parallelism leads to faster convergence for all algorithms (Figure 4a). In contrast with other algorithms, $qEHVI$ ’s sample complexity does not deteriorate substantially when high levels of parallelism are used (Figure 4b).

Table 1: Acquisition Optimization wall time in seconds on a CPU (2x Intel Xeon E5-2680 v4 @ 2.40GHz) and a GPU (Tesla V100-SXM2-16GB). We report the mean and 2 standard errors across 20 trials. NA indicates that the algorithm does not support constraints.

CPU	BRANINCURRIN	C2DTLZ2	ABR	VEHICLESAFETY
PESMO ($q=1$)	249.16 (± 19.35)	NA	214.16 (± 18.38)	492.64 (± 58.98)
SMS-EGO ($q=1$)	146.1 (± 8.57)	NA	89.54 (± 5.79)	115.11 (± 8.21)
TS-TCH ($q=1$)	2.82 (± 0.03)	NA	17.22 (± 0.04)	47.46 (± 0.05)
q PAREGO ($q=1$)	1.56 (± 0.16)	4.01 (± 0.77)	7.47 (± 0.67)	1.74 (± 0.27)
EHVI ($q=1$)	3.04 (± 0.16)	NA	2.48 (± 0.19)	15.18 (± 2.24)
q EHVI ($q=1$)	3.63 (± 0.23)	5.4 (± 1.18)	6.15 (± 0.71)	67.54 (± 10.45)
GPU	BRANINCURRIN	C2DTLZ2	ABR	VEHICLESAFETY
TS-TCH ($q=1$)	0.07 (± 0.00)	NA	0.16 (± 0.00)	0.32 (± 0.0)
TS-TCH ($q=2$)	0.07 (± 0.00)	NA	0.15 (± 0.00)	0.34 (± 0.01)
TS-TCH ($q=4$)	0.09 (± 0.01)	NA	0.15 (± 0.00)	0.31 (± 0.01)
TS-TCH ($q=8$)	0.08 (± 0.00)	NA	0.16 (± 0.00)	0.34 (± 0.01)
q PAREGO ($q=1$)	3.2 (± 0.37)	3.85 (± 0.91)	9.64 (± 0.96)	3.44 (± 0.51)
q PAREGO ($q=2$)	7.12 (± 0.81)	12.1 (± 2.77)	21.19 (± 1.53)	7.32 (± 0.97)
q PAREGO ($q=4$)	15.34 (± 1.69)	39.71 (± 7.40)	35.46 (± 2.32)	17.2 (± 2.29)
q PAREGO ($q=8$)	32.11 (± 4.14)	99.58 (± 15.20)	72.52 (± 5.04)	39.72 (± 7.13)
EHVI ($q=1$)	4.53 (± 0.23)	NA	6.82 (± 0.55)	8.95 (± 0.64)
q EHVI ($q=1$)	5.98 (± 0.28)	3.36 (± 0.94)	7.71 (± 0.67)	10.43 (± 0.64)
q EHVI ($q=2$)	11.37 (± 0.56)	21.56 (± 3.45)	18.32 (± 1.48)	17.67 (± 1.54)
q EHVI ($q=4$)	25.29 (± 1.51)	89.18 (± 10.86)	44.44 (± 3.53)	54.25 (± 4.17)
q EHVI ($q=8$)	102.46 (± 9.22)	215.74 (± 15.85)	100.64 (± 7.22)	255.72 (± 23.73)

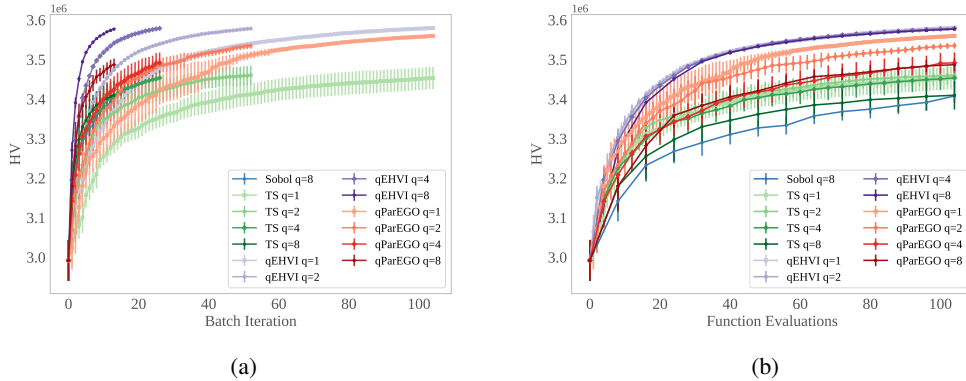


Figure 4: Parallel optimization performance on the ABR problem with varying batch sizes (q) by (a) *batch BO iterations* and (b) *function evaluations*.

6 Discussion

We present a practical and efficient acquisition function, q EHVI, for parallel, constrained multi-objective Bayesian optimization. Leveraging differentiable programming, modern parallel hardware, and the Sample Average Approximation, we efficiently optimize q EHVI via quasi second-order methods and provide theoretical convergence guarantees for our approach. Empirically, we demonstrate that our method out-performs state-of-the-art multi-objective Bayesian optimization methods.

One limitation of our approach is that it currently assumes noiseless observations, which, to our knowledge, is the case with all formulations of EHVI. Integrating over the uncertainty around the previous observations [43] by using MC samples over the new candidates and the training points, one may be able to account for the noise. Another limitation of q EHVI is that its scalability is limited the partitioning algorithm, precluding its use in high-dimensional objective spaces. More scalable partitioning algorithms, either approximate algorithms (e.g. the algorithm proposed by Couckuyt et al. [11], which we examine briefly in Appendix F.4) or more efficient exact algorithms that result in fewer disjoint hyper-rectangles (e.g. [41, 17, 69]), will improve the scalability and computation time of q EHVI. We hope this work encourages researchers to consider more improvements from applying modern computational paradigms and tooling to Bayesian optimization.

7 Statement of Broader Impact

Optimizing a single outcome commonly comes at the expense of other secondary outcomes. In some cases, decision makers may be able to form a scalarization of their objectives in advance, but in the researcher’s experience, formulating such trade-offs in advance is difficult for most. Improvements to the optimization performance and practicality of multi-objective Bayesian optimization have the potential to allow decision makers to better understand and make more informed decisions across multiple trade-offs. We expect these directions to be particularly important as Bayesian optimization is increasingly used for applications such as recommender systems [42], where auxiliary goals such as fairness must be accounted for. Of course, at the end of the day, exactly what objectives decision makers choose to optimize, and how they balance those trade-offs (and whether that is done in equitable fashion) is up to the individuals themselves.

Acknowledgments

We would like to thank Daniel Jiang for helpful discussions around our theoretical results.

References

- [1] M. Abdolshah, A. Shilton, S. Rana, S. Gupta, and S. Venkatesh. Expected hypervolume improvement with constraints. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 3238–3243, 2018.
- [2] Arash Asadpour, Hamid Nazerzadeh, and Amin Saberi. Stochastic submodular maximization. In Christos Papadimitriou and Shuzhong Zhang, editors, *Internet and Network Economics*. Springer Berlin Heidelberg, 2008.
- [3] R. Astudillo and P. Frazier. Bayesian optimization of composite functions. *Forthcoming, in Proceedings of the 35th International Conference on Machine Learning*, 2019.
- [4] Anne Auger, Johannes Bader, Dimo Brockhoff, and Eckart Zitzler. Theory of the hypervolume indicator: Optimal mu-distributions and the choice of the reference point. In *Proceedings of the Tenth ACM SIGEVO Workshop on Foundations of Genetic Algorithms*, FOGA ’09, page 87–102, New York, NY, USA, 2009. Association for Computing Machinery.
- [5] 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 *Advances in Neural Information Processing Systems 33*, 2020.
- [6] Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Max-value entropy search for multi-objective bayesian optimization. In *Advances in Neural Information Processing Systems 32*, 2019.
- [7] Eric Bradford, Artur Schweidtmann, and Alexei Lapkin. Efficient multiobjective optimization employing gaussian processes, spectral sampling and a genetic algorithm. *Journal of Global Optimization*, 71, 02 2018. doi: 10.1007/s10898-018-0609-2.
- [8] Russel E Cafilisch. Monte carlo and quasi-monte carlo methods. *Acta numerica*, 7:1–49, 1998.
- [9] Mauro Cerasoli and Aniello Fedullo. The inclusion-exclusion principle. *Journal of Interdisciplinary Mathematics*, 5(2):127–141, 2002.
- [10] Anirban Chaudhuri, Raphael Haftka, Peter Ifju, Kelvin Chang, Christopher Tyler, and Tony Schmitz. Experimental flapping wing optimization and uncertainty quantification using limited samples. *Structural and Multidisciplinary Optimization*, 51, 11 2014. doi: 10.1007/s00158-014-1184-x.
- [11] I. Couckuyt, D. Deschrijver, and T. Dhaene. Towards efficient multiobjective optimization: Multiobjective statistical criterions. In *2012 IEEE Congress on Evolutionary Computation*, pages 1–8, 2012.
- [12] Ivo Couckuyt, Dirk Deschrijver, and Tom Dhaene. Fast calculation of multiobjective probability of improvement and expected improvement criteria for pareto optimization. *J. of Global Optimization*, 60(3): 575–594, November 2014.
- [13] Daniel A. da Silva. Proprietades geraes. *J. de l’Ecole Polytechnique*, cah. 30. I, 1854.

- [14] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.
- [15] Kalyan Deb, L. Thiele, Marco Laumanns, and Eckart Zitzler. Scalable multi-objective optimization test problems. volume 1, pages 825–830, 06 2002. ISBN 0-7803-7282-4. doi: 10.1109/CEC.2002.1007032.
- [16] Kalyanmoy Deb. *Constrained Multi-objective Evolutionary Algorithm*, pages 85–118. Springer International Publishing, Cham, 2019.
- [17] Kerstin Dächert, Kathrin Klamroth, Renaud Lacour, and Daniel Vanderpooten. Efficient computation of the search region in multi-objective optimization. *European Journal of Operational Research*, 260(3):841 – 855, 2017.
- [18] M. T. M. Emmerich, K. C. Giannakoglou, and B. Naujoks. Single- and multiobjective evolutionary optimization assisted by gaussian random field metamodels. *IEEE Transactions on Evolutionary Computation*, 10(4):421–439, 2006.
- [19] M. T. M. Emmerich, A. H. Deutz, and J. W. Klinkenberg. Hypervolume-based expected improvement: Monotonicity properties and exact computation. In *2011 IEEE Congress of Evolutionary Computation (CEC)*, pages 2147–2154, 2011.
- [20] Michael Emmerich, Kaifeng Yang, André Deutz, Hao Wang, and Carlos M. Fonseca. *A Multicriteria Generalization of Bayesian Global Optimization*, pages 229–242. Springer International Publishing, 2016.
- [21] Michael T. M. Emmerich and Carlos M. Fonseca. Computing hypervolume contributions in low dimensions: Asymptotically optimal algorithm and complexity results. In Ricardo H. C. Takahashi, Kalyanmoy Deb, Elizabeth F. Wanner, and Salvatore Greco, editors, *Evolutionary Multi-Criterion Optimization*, pages 121–135, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- [22] Paul Feliot, Julien Bect, and Emmanuel Vazquez. A bayesian approach to constrained single- and multi-objective optimization. *Journal of Global Optimization*, 67(1-2):97–133, Apr 2016. ISSN 1573-2916. doi: 10.1007/s10898-016-0427-3. URL <http://dx.doi.org/10.1007/s10898-016-0427-3>
- [23] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey. *An analysis of approximations for maximizing submodular set functions—II*, pages 73–87. Springer Berlin Heidelberg, Berlin, Heidelberg, 1978.
- [24] Tobias Friedrich and Frank Neumann. Maximizing submodular functions under matroid constraints by multi-objective evolutionary algorithms. In Thomas Bartz-Beielstein, Jürgen Branke, Bogdan Filipič, and Jim Smith, editors, *Parallel Problem Solving from Nature – PPSN XIII*, pages 922–931, Cham, 2014. Springer International Publishing. ISBN 978-3-319-10762-2.
- [25] Jacob Gardner, Matt Kusner, Zhixiang, Kilian Weinberger, and John Cunningham. Bayesian optimization with inequality constraints. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 937–945, Beijing, China, 22–24 Jun 2014. PMLR.
- [26] Eduardo C Garrido-Merchán and Daniel Hernández-Lobato. Predictive entropy search for multi-objective bayesian optimization with constraints. *Neurocomputing*, 361:50–68, 2019.
- [27] Eduardo C Garrido-Merchán and Daniel Hernández-Lobato. Parallel predictive entropy search for multi-objective bayesian optimization with constraints, 2020.
- [28] David Gaudrie, Rodolphe Le Riche, Victor Picheny, Benoît Enaux, and Vincent Herbert. Targeting solutions in bayesian multi-objective optimization: sequential and batch versions. *Annals of Mathematics and Artificial Intelligence*, 88(1-3):187–212, Aug 2019. ISSN 1573-7470. doi: 10.1007/s10472-019-09644-8. URL <http://dx.doi.org/10.1007/s10472-019-09644-8>
- [29] Michael A. Gelbart, Jasper Snoek, and Ryan P. Adams. Bayesian optimization with unknown constraints. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, UAI, 2014.
- [30] David Ginsbourger, Rodolphe Le Riche, and Laurent Carraro. *Kriging Is Well-Suited to Parallelize Optimization*, pages 131–162. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- [31] P. Glasserman. Performance continuity and differentiability in monte carlo optimization. In *1988 Winter Simulation Conference Proceedings*, pages 518–524, 1988.
- [32] Nikolaus Hansen. *The CMA Evolution Strategy: A Comparing Review*, volume 192, pages 75–102. 06 2007. doi: 10.1007/3-540-32494-1_4.

- [33] Daniel Hernández-Lobato, José Miguel Hernández-Lobato, Amar Shah, and Ryan P. Adams. Predictive entropy search for multi-objective bayesian optimization, 2015.
- [34] Iris Hupkens, Andre Deutz, Kaifeng Yang, and Michael Emmerich. Faster exact algorithms for computing expected hypervolume improvement. In Antonio Gaspar-Cunha, Carlos Henggeler Antunes, and Carlos Coello Coello, editors, *Evolutionary Multi-Criterion Optimization*, pages 65–79. Springer International Publishing, 2015.
- [35] Hisao Ishibuchi, Naoya Akedo, and Yusuke Nojima. A many-objective test problem for visually examining diversity maintenance behavior in a decision space. In *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation, GECCO '11*, page 649–656, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450305570. doi: 10.1145/2001576.2001666. URL <https://doi.org/10.1145/2001576.2001666>
- [36] Hisao Ishibuchi, Ryo Imada, Yu Setoguchi, and Yusuke Nojima. How to specify a reference point in hypervolume calculation for fair performance comparison. *Evol. Comput.*, 26(3):411–440, September 2018.
- [37] Donald Jones, C. Perttunen, and B. Stuckman. Lipschitzian optimisation without the lipschitz constant. *Journal of Optimization Theory and Applications*, 79:157–181, 01 1993. doi: 10.1007/BF00941892.
- [38] 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.
- [39] Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. *arXiv e-prints*, page arXiv:1312.6114, Dec 2013.
- [40] J. Knowles. Parego: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1):50–66, 2006.
- [41] Renaud Lacour, Kathrin Klamroth, and Carlos M. Fonseca. A box decomposition algorithm to compute the hypervolume indicator. *Computers & Operations Research*, 79:347 – 360, 2017.
- [42] Benjamin Letham and Eytan Bakshy. 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>
- [43] Benjamin Letham, Brian Karrer, Guilherme Ottoni, and Eytan Bakshy. Constrained bayesian optimization with noisy experiments. *Bayesian Analysis*, 14(2):495–519, 06 2019. doi: 10.1214/18-BA1110.
- [44] Xingtao Liao, Qing Li, Xujing Yang, Weigang Zhang, and Wei Li. Multiobjective optimization for crash safety design of vehicles using stepwise regression model. *Structural and Multidisciplinary Optimization*, 35:561–569, 06 2008. doi: 10.1007/s00158-007-0163-x.
- [45] Edgar Manoatl Lopez, Luis Miguel Antonio, and Carlos A. Coello Coello. A gpu-based algorithm for a faster hypervolume contribution computation. In António Gaspar-Cunha, Carlos Henggeler Antunes, and Carlos Coello Coello, editors, *Evolutionary Multi-Criterion Optimization*, pages 80–94. Springer International Publishing, 2015.
- [46] Hongzi Mao, Ravi Netravali, and Mohammad Alizadeh. Neural adaptive video streaming with pensieve. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication, SIGCOMM '17*, page 197–210, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450346535. doi: 10.1145/3098822.3098843. URL <https://doi.org/10.1145/3098822.3098843>
- [47] Hongzi Mao, Shannon Chen, Drew Dimmery, Shaun Singh, Drew Blaisdell, Yuandong Tian, Mohammad Alizadeh, and Eytan Bakshy. Real-world video adaptation with reinforcement learning. 2019.
- [48] Hongzi Mao, Parimarjan Negi, Akshay Narayan, Hanrui Wang, Jiacheng Yang, Haonan Wang, Ryan Marcus, Ravichandra Addanki, Mehrdad Khani Shirkoohi, Songtao He, Vikram Nathan, Frank Cangialosi, Shaileshh Bojja Venkatakrishnan, Wei-Hung Weng, Shu-Wen Han, Tim Kraska, and Mohammad Alizadeh. Park: An open platform for learning-augmented computer systems. In *NeurIPS*, 2019.
- [49] Sébastien Marmin, Clément Chevalier, and David Ginsbourger. Differentiating the multipoint expected improvement for optimal batch design. In Panos Pardalos, Mario Pavone, Giovanni Maria Farinella, and Vincenzo Cutello, editors, *Machine Learning, Optimization, and Big Data*, pages 37–48, Cham, 2015. Springer International Publishing.
- [50] B. Paria, K. Kandasamy, and B. Póczos. A Flexible Multi-Objective Bayesian Optimization Approach using Random Scalarizations. *ArXiv e-prints*, May 2018.

- [51] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. 2017.
- [52] Victor Picheny. Multiobjective optimization using gaussian process emulators via stepwise uncertainty reduction. *Statistics and Computing*, 25, 10 2013. doi: 10.1007/s11222-014-9477-x.
- [53] Wolfgang Ponweiser, Tobias Wagner, Dirk Biermann, and Markus Vincze. Multiobjective optimization on a limited budget of evaluations using model-assisted s-metric selection. In Günter Rudolph, Thomas Jansen, Nicola Beume, Simon Lucas, and Carlo Poloni, editors, *Parallel Problem Solving from Nature – PPSN X*, pages 784–794, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
- [54] Alma A. M. Rahat, Richard M. Everson, and Jonathan E. Fieldsend. Alternative infill strategies for expensive multi-objective optimisation. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '17*, page 873–880, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450349208.
- [55] Carl Edward Rasmussen. *Gaussian Processes in Machine Learning*, pages 63–71. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004.
- [56] Jerry Seeger. Inclusion-exclusion and characteristic functions. *Mathematics Magazine*, 71(3):216–218, 1998. ISSN 0025570X, 19300980. URL <http://www.jstor.org/stable/2691209>
- [57] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.
- [58] Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on International Conference on Machine Learning, ICML'10*, page 1015–1022, Madison, WI, USA, 2010. Omnipress. ISBN 9781605589077.
- [59] J. Sylvester. Note sur la théorème de legendre. *Comptes Rendus Acad. Sci.*, 96:463–465, 1883.
- [60] Ryoji Tanabe and Hisao Ishibuchi. 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>.
- [61] 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.
- [62] Takashi Wada and Hideitsu Hino. Bayesian optimization for multi-objective optimization and multi-point search, 2019.
- [63] Jialei Wang, Scott C. Clark, Eric Liu, and Peter I. Frazier. Parallel bayesian global optimization of expensive functions, 2016.
- [64] J. T. Wilson, R. Moriconi, F. Hutter, and M. P. Deisenroth. The reparameterization trick for acquisition functions. *ArXiv e-prints*, December 2017.
- [65] James Wilson, Frank Hutter, and Marc Deisenroth. Maximizing acquisition functions for bayesian optimization. In *Advances in Neural Information Processing Systems 31*, pages 9905–9916. 2018.
- [66] Jian Wu and Peter I. Frazier. The parallel knowledge gradient method for batch bayesian optimization. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, page 3134–3142, Red Hook, NY, USA, 2016. Curran Associates Inc. ISBN 9781510838819.
- [67] Kaifeng Yang, Michael Emmerich, André Deutz, and Carlos M. Fonseca. Computing 3-d expected hypervolume improvement and related integrals in asymptotically optimal time. In *9th International Conference on Evolutionary Multi-Criterion Optimization - Volume 10173, EMO 2017*, page 685–700, Berlin, Heidelberg, 2017. Springer-Verlag.
- [68] Kaifeng Yang, Michael Emmerich, André H. Deutz, and Thomas Bäck. Efficient computation of expected hypervolume improvement using box decomposition algorithms. *CoRR*, abs/1904.12672, 2019.
- [69] Kaifeng Yang, Michael Emmerich, André Deutz, and Thomas Bäck. 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>. URL <http://www.sciencedirect.com/science/article/pii/S2210650217307861>

- [70] Kaifeng Yang, Pramudita Palar, Michael Emmerich, Koji Shimoyama, and Thomas Bäck. A multi-point mechanism of expected hypervolume improvement for parallel multi-objective bayesian global optimization. pages 656–663, 07 2019. doi: 10.1145/3321707.3321784.
- [71] Kaifeng Yang, Pramudita Satria Palar, Michael Emmerich, Koji Shimoyama, and Thomas Bäck. 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*, page 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>.
- [72] R. J. Yang, N. Wang, C. H. Tho, J. P. Bobineau, and B. P. Wang. Metamodeling Development for Vehicle Frontal Impact Simulation. *Journal of Mechanical Design*, 127(5):1014–1020, 01 2005.
- [73] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. G. da Fonseca. Performance assessment of multiobjective optimizers: an analysis and review. *IEEE Transactions on Evolutionary Computation*, 7(2): 117–132, 2003.