
Random Hypervolume Scalarizations for Provable Multi-Objective Black Box Optimization

Daniel Golovin¹ Qiuyi (Richard) Zhang¹

Abstract

Single-objective black box optimization (also known as zeroth-order optimization) is the process of minimizing a scalar objective $f(x)$, given evaluations at adaptively chosen inputs x . In this paper, we consider multi-objective optimization, where $f(x)$ outputs a vector of possibly competing objectives and the goal is to converge to the Pareto frontier. Quantitatively, we wish to maximize the standard *hypervolume indicator* metric, which measures the dominated hypervolume of the entire set of chosen inputs. In this paper, we introduce a novel scalarization function, which we term the *hypervolume scalarization*, and show that drawing random scalarizations from an appropriately chosen distribution can be used to efficiently approximate the *hypervolume indicator* metric. We utilize this connection to show that Bayesian optimization with our scalarization via common acquisition functions, such as Thompson Sampling or Upper Confidence Bound, provably converges to the whole Pareto frontier by deriving tight *hypervolume regret* bounds on the order of $\tilde{O}(\sqrt{T})$. Furthermore, we highlight the general utility of our scalarization framework by showing that any provably convergent single-objective optimization process can be effortlessly converted to a multi-objective optimization process with provable convergence guarantees.

1. Introduction

Single-objective optimization has traditionally been the focus in the field of machine learning for many practical and interesting applications, from standard regression objectives (Kutner et al., 2005) to ever-increasingly compli-

cated losses used in deep learning (LeCun et al., 2015) and reinforcement learning (Sutton et al., 1998). However, in the recent years, there has been a growing need to care about multiple objectives in learning and understanding the inherent tradeoffs between these conflicting objectives. Examples include classical tradeoffs such as bias vs variance (Neal et al., 2018), and accuracy vs calibration (Guo et al., 2017) but there are increasingly complex tradeoffs between accuracy and robustness to attack (Zhang et al., 2019), accuracy and fairness (Zliobaite, 2015), between multiple correlated tasks in multi-task learning (Kendall et al., 2018), between network adaptations in meta-learning (Finn et al., 2017), or any combination of the above.

To understand and visualize a growing number of complex tradeoffs, many turn to multi-objective optimization, which is the maximization of k objectives $F(x) := (f_1(x), \dots, f_k(x))$ over the parameter space $x \in \mathcal{X} \subseteq \mathbb{R}^n$. Since one often cannot simultaneously maximize all f_i , multi-objective optimization aims to find the entire *Pareto frontier* \mathcal{F} of the objective space, where intuitively x is on the Pareto frontier if there is no way to improve on all objectives simultaneously. To measure progress, a natural and widely used metric to compare Pareto sets is the *hypervolume indicator*, which is the volume of the dominated portion of the Pareto set (Zitzler & Thiele, 1999). The hypervolume metric is especially desirable because it has strict Pareto compliance meaning that if a set $A \subseteq B \subset \mathbb{R}^k$, then the hypervolume of B is greater than that of A . While the hypervolume indicator satisfies strict Pareto compliance, almost all other unary metrics do not, thereby explaining the popularity of using hypervolume as the predominant measure of progress (Zitzler et al., 2003).

While the hypervolume indicator remains the gold standard in evaluating multi-objective algorithms, note that it is a function of a set of objective vectors and not just a single evaluation. This means evolutionary algorithms often cannot use the hypervolume as a fitness score although some discretized version exist (Emmerich et al., 2005). Furthermore, it is computationally inefficient: the best current asymptotic runtime calculating the hypervolume indicator in \mathbb{R}^k is $O(n^{k/2})$ via a reduction to Klee’s Measure Problem (Beume & Rudolph, 2006). This cannot be sub-

^{*}Equal contribution ¹Google Brain, Pittsburgh, Pennsylvania, USA. Correspondence to: Richard Zhang <qiuyiz@google.com>.

stantially sped up as hypervolume calculation is shown to be $\#P$ -hard and even approximating the hypervolume improvement by a $O(2^{d^{1-\epsilon}})$ factor is \mathbf{NP} -hard for any $\epsilon > 0$ (Bringmann & Friedrich, 2010). For $k \leq 3$, fast albeit complicated algorithms exist (Yang et al., 2019).

Due to the difficulty of using the hypervolume directly, many multi-objective optimization problems use a heuristic-based scalarization strategy, which splits the multi-objective optimization into numerous single "scalarized" objectives (Rojers et al., 2013). For some weights $\lambda \in \mathbb{R}^k$, we have scalarization functions $s_\lambda(y) : \mathbb{R}^k \rightarrow \mathbb{R}$ that convert multi-objective outputs into a single-objective scalars. Bayesian optimization is then applied to this family of single-objective functions $s_\lambda(F(x))$ for various λ and if we construct s_λ to be monotonically increasing in all coordinates, then

$$x_\lambda = \arg \max_{x \in \mathcal{X}} s_\lambda(F(x))$$

is on the Pareto frontier (Paria et al., 2018).

Early works on scalarization include heuristic-based algorithms such as ParEgo (Knowles, 2006) and MOEAD (Zhang & Li, 2007). The most popular scalarizations are the linear scalarization $s_\lambda(y) = \sum_i \lambda_i y_i$ and the Chebyshev scalarization $s_\lambda(y) = \min_i \lambda_i (y_i - z_i)$ for some reference point z and some distribution over λ (Nakayama et al., 2009). Using the linear scalarization is especially popular among researchers and practitioners, as it can be viewed as using f_i as a regularization function and λ_i as the corresponding Lagrange multipliers. When f_i are all convex, classical Lagrange duality then guarantees that varying λ will explore the entire Pareto frontier.

However, the choices of weight distribution, reference point, and even the scalarization functions themselves are crucial to the success of multi-objective optimization, leading researchers to put forth a diverse set of options for each of these choices between different papers (Paria et al., 2018; Nakayama et al., 2009; Zhang & Li, 2007). Recently, some works have come up with novel scalarizations that perform better empirically (Aliano Filho et al., 2019; Schmidt et al., 2019) and others have tried to do comparisons between different scalarizations with varying conclusions (Kasimbeyli et al., 2019). Some have also proposed adaptively weighted approaches that have connections to gradient-based multi-objective optimization (Lin et al., 2019).

In addition to scalarization, there are a large diverse array of multi-objective optimization rules that are often heuristic-based and lack theoretical guarantees, such as aggregation-based, decomposition-based, diversity-based, elitism-based, gradient-based, and hybrid approaches (Emmerich & Deutz, 2018; Zitzler et al., 2004; Konak et al.,

2006). Furthermore, many Bayesian optimization approaches have been extended to the multi-objective setting, such as predictive entropy search (Hernández-Lobato et al., 2016) or uncertainty measures (Picheny, 2015). To our knowledge, there have been limited theoretical works that give regret and convergence bounds for multi-objective optimization. Some works show that the algorithms achieve small Pareto regret, which only guarantees that one recovers a single point close to the Pareto frontier (Lu et al., 2019; Öner et al., 2018; Turğay et al., 2018). For many scalarization functions, Paria et al provides a Bayes regret bound with respect to a scalarization-induced regret (Paria et al., 2018) but small Bayes regret with a monotone scalarization only guarantees recovery of a single Pareto optimal point. Zuluaga et al provides sub-linear hypervolume regret bounds; however, they are exponential in k and its analysis only applies to a specially tailored algorithm that requires an unrealistic classification step (Zuluaga et al., 2013).

1.1. Our contributions

We introduce a scalarization function, which we term the hypervolume scalarization, and present a connection, independently discovered by (Deng & Zhang, 2019), between this scalarization function and the hypervolume indicator. We provide the first hypervolume regret bounds that are valid for common Bayesian optimization algorithms such as UCB and Thompson Sampling and can be polynomial in the number of objectives k when \mathcal{X} is suitably compact. This implies that our choice of scalarization functions and weight distribution \mathcal{D}_λ is theoretically sound and leads to provable convergence to the whole Pareto frontier. Furthermore, we can utilize the connection of scalarization to hypervolume indicator to provide a computationally-efficient and accurate estimator of the hypervolume, providing a simple algorithm for approximating the hypervolume that easily generalizes to higher dimensions. The first lemma that we rely on is the following:

Lemma 1 (Hypervolume Scalarization: Informal Restatement of Lemma 5). *Let $Y = \{y_1, \dots, y_m\}$ be a set of m points in \mathbb{R}^k and let $z \in \mathbb{R}^k$ be a reference point such that $y_i \geq z$ for all i . Then, the hypervolume $\mathcal{HV}_z(Y)$ of Y with respect to z is given by:*

$$\mathcal{HV}_z(Y) := c_k \mathbb{E}_{\lambda \sim \mathcal{D}} \left[\max_{y \in Y} s_\lambda(y - z) \right]$$

for some scalarization functions $s_\lambda(y)$ and fixed weight distribution \mathcal{D} .

Note our resulting hypervolume scalarization functions, which mirrors Equation 6 of (Deng & Zhang, 2019), are similar to the Chebyshev scalarization approach $s_\lambda(y) = \min_i \lambda_i (y_i - z_i)$ but they come with provable guarantees.

Specifically, using these scalarizations, we can first easily extend a single-objective optimization algorithm into the multi-objective setting: at each step, we first apply a random scalarization, then we attempt to maximize along that scalarization function via any single-objective optimization procedure to find the next suggested point. We show novel regret bounds for combining our hypervolume scalarization with popular and standard Bayesian optimization algorithms, such as Thompson Sampling (TS) and Upper Confidence Bound (UCB) methods.

If we let $x_t \in \mathcal{X}$ denote the point chosen in step t , and let $Y_t := \{F(x_i) : i = 1, 2, \dots, t\}$, then the *hypervolume regret* at time t is defined as $r_t := \mathcal{HV}_z(Y^*) - \mathcal{HV}_z(Y_t)$, where Y^* is the Pareto frontier. We show that our cumulative hypervolume regret at step T for TS and UCB is bounded by $O(k^2 \sqrt{\gamma_T T \ln T})$, where γ_T is a kernel-dependent quantity known as the maximum information gain. Intuitively, the maximum information gain allows us to quantify the increase in certainty after evaluating T points and is usually mild. For example, $\gamma_T = O(\text{poly} \log(T))$ for the squared exponential kernel and since r_t is by definition monotonically decreasing, our regret bound translates immediately into a convergence bound of $\tilde{O}(T^{-1/2})$ and we provably converge rapidly to the Pareto frontier. We note that the bound for the single-objective case can be recovered by setting $k = 1$ and matches the bounds of previous work (Russo & Van Roy, 2014; Srinivas et al., 2010). Furthermore, since the single-objective bounds are shown to be tight up to poly-logarithmic factors (Scarlett et al., 2017), our regret bounds are thus also tight for our dependence on T . Our main theorem can be stated as follows:

Theorem 2 (Convergence of Bayesian Optimization with Hypervolume Scalarization: Informal Restatement of Theorem 8). *The cumulative hypervolume regret for using random hypervolume scalarization with UCB or TS after T observations is upper bounded as*

$$\sum_{t=1}^T (\mathcal{HV}_z(Y^*) - \mathcal{HV}_z(Y_t)) \leq O(k^2 n^{1/2} [\gamma_T T \ln(T)]^{1/2})$$

Furthermore, $\mathcal{HV}_z(Y_T) \geq \mathcal{HV}_z(Y^*) - \epsilon_T$, where $\epsilon_T = O(k^2 n^{1/2} [\gamma_T \ln(T)/T]^{1/2})$.

Lastly, we show that if the single-objective optimization procedure admits good convergence properties, then we can solve multi-objective optimization by simply applying the single-objective procedure on sufficiently large number of randomly chosen hypervolume scalarizations. Furthermore, this method is provably correct as we can deduce hypervolume error bounds via concentration properties. Specifically, by using the single-objective procedure for T iterations on $l = O(1/\epsilon^k)$ different randomly chosen scalarizations, we can bound the hypervolume error by ϵ

after the total number of observations of $T \cdot l = O(T/\epsilon^k)$. Therefore, we present a novel generic framework that extends any single-objective convergence bound into a convergence bound for the multi-objective case that provably demonstrates convergence to the Pareto frontier. Note that these bounds hold for any algorithm but are not tight and can likely be improved, unlike those presented above in the Bayesian optimization setting.

Theorem 3 (Convergence of General Optimization with Hypervolume Scalarization: Informal Restatement of Theorem 9). *Let \mathcal{A} be any single-objective optimization algorithm with the guarantee that it converges to within ϵ_T of the maximum after T iterations and observations.*

Then, running \mathcal{A} with random hypervolume scalarization converges to the Pareto frontier and admits an hypervolume error of $O(\epsilon_T)$ after $O(T/\epsilon_T^k)$ observations.

We empirically validate our theoretical contributions by running our multiobjective algorithms with hypervolume scalarizations on the Black-Box Optimization Benchmark (BBOB) functions, which can be used for bi-objective optimization problems (Tušar et al., 2016). We see that our multi-objective Bayesian optimization algorithms, which admit strong regret bounds, consistently outperforms the multi-objective evolutionary algorithms. Furthermore, we observe the superior performance of the hypervolume scalarization functions over other scalarizations, although that difference is less pronounced when the Pareto frontier is even somewhat convex.

Guided by our theory, we believe that researchers and practitioners alike should slowly move away from using linear scalarizations for multi-objective optimization and move towards hypervolume scalarizations. We summarize our contributions as follows:

- Introduction of the theoretically guided hypervolume scalarization function, with connections to hypervolume indicator and provable guarantees.
- Development of simple algorithm for approximating hypervolume indicator with hypervolume scalarizations that is accurate due to good smoothness and concentration properties.
- Novel hypervolume regret bounds for Bayesian optimization with UCB or TS when using hypervolume scalarization.
- Derivation of convergence hypervolume error bounds for any single-objective optimization procedure when using hypervolume scalarization.

2. Preliminaries

We first define a few notations for the rest of the paper. For two vectors x, y , we define $x \leq y$ and similarly all other comparisons/operators element-wise and $\|\cdot\|$ is the Euclidean norm unless specified otherwise. We define $a + b := a + b1$ for $a \in \mathbb{R}^k$ and $b \in \mathbb{R}$. For a function $f(x) : \mathbb{R}^n \rightarrow \mathbb{R}$, we say f is L -Lipschitz if $|f(x) - f(x')| \leq L\|x - x'\|_1$.

Let \mathcal{X} be a compact subset of \mathbb{R}^n and let $F(x) : \mathbb{R}^n \rightarrow \mathbb{R}^k$ be our multi-objective function with k objective functions f_i . For two points $x_1, x_2 \in \mathcal{X}$, we say that x_1 is *Pareto-dominated* by x_2 if $f_i(x_1) \leq f_i(x_2)$ for all i and there exists j such that $f_j(x_1) < f_j(x_2)$. A point is *Pareto-optimal* in \mathcal{X} if no point in \mathcal{X} dominates it. Let X^* denote the set of Pareto-optimal points in \mathcal{X} . A *Pareto set* for input points $X = \{x_1, \dots, x_m\}$ is the set of points in \mathbb{R}^k is $\{F(x)|x \in X\}$, denoted as simply $F(X)$. The *Pareto frontier* $Y^* := \{F(x)|x \in X^*\}$ is the Pareto set of X^* .

Our main progress indicators are given by the hypervolume indicator. For $S \subseteq \mathbb{R}^k$ compact, let $\text{vol}(S)$ be the hypervolume of S .

Definition 4. For $Y \subseteq \mathbb{R}^k$, we define the (dominated) **hypervolume indicator** of Y with respect to reference point z as:

$$\mathcal{HV}_z(Y) = \text{vol}(\{x | x \geq z, x \text{ is dominated by some } y \in Y\})$$

Therefore, for a finite set Y , $\mathcal{HV}_z(Y)$ can be viewed as the hypervolume of the union of the dominated hyperrectangles for each point $y_i \geq z$ that has one corner at z and the other corner at y_i . Note that our definition also holds for non-finite set as a limiting integral in the Lebesgue measure. We let $\mathcal{S}_+^{k-1} = \{y \in \mathbb{R}^k | \|y\| = 1, y \geq 0\}$ and let $y \sim \mathcal{S}_+^{k-1}$ denote that y is drawn uniformly on \mathcal{S}_+^{k-1} .

2.1. Scalarization with Bayesian Optimization

Bayesian optimization uses a probabilistic model to fit to the blackbox function. Gaussian processes (GP) are a standard way to model distributions over functions and are commonly used to derive good regret bounds (Williams & Rasmussen, 2006). We will begin with a brief review of GPs and their role in Bayesian Optimization.

A Gaussian process, $\mathcal{GP}(\mu, \kappa)$, is a distribution over functions. In a GP, the similarity between points x_i, x_j are determined by the kernel function $\kappa(x_i, x_j)$ and for some finite set of points $X = \{x_1, \dots, x_m\} \in \mathcal{X}$, the distribution of $f(x_1), \dots, f(x_m)$ over f is modeled as multivariate Gaussian whose covariance matrix is $\Sigma_{ij} = \kappa(x_i, x_j)$ and mean is given by $\mu_i = \mu(x_i)$. Examples of popular kernels are the squared exponential, $\kappa(x_i, x_j) = \exp(-\gamma\|x_i - x_j\|^2)$, and the Matérn kernel. The mean function, prior to receiving

data, is often assumed to be zero.

When datapoints for values of $y_i = f(x_i) + \epsilon_i$ are received with $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$, the GP is then updated into a posterior distribution over functions that attempts to fit the datapoints values. This is done by applying conditioning and because all relevant distribution are Gaussian, the resulting distribution is still a GP. This posterior GP induces an unique mean function $\mu(x)$ and standard deviation function $\sigma(x)$ given by the standard formulas:

$$\mu(x) = \kappa(x, X)^\top (\Sigma + \sigma^2 I)^{-1} y$$

$$\sigma(x) = \kappa(x, x) - \kappa(x, X)^\top [\Sigma + \sigma^2 I]^{-1} \kappa(x, X)$$

where $\kappa(x, X)$ denotes the vector with entries as $\kappa(x, x_i)$.

Bayesian optimization then uses the probabilistic model to optimize for the best inputs for the blackbox function. Since we have a distribution over functions, the optimization is done via an *acquisition function*, $a(x) : \mathbb{R}^n \rightarrow \mathbb{R}$, that assigns a scalar value to each point based on the aforementioned distribution. A popular and natural method is Thompson Sampling (TS), which draws a random sample from the posterior $f \sim \mathcal{GP}(\mu, \kappa)$, and simply uses that as a non-deterministic acquisition function $a(x) = f(x)$ (Thompson, 1933). Another popular method is the Upper Confidence Bound (UCB) acquisition that combines the mean and variance function into $a(x) = \mu(x) + \sqrt{\beta}\sigma(x)$, where β is a function of m and n (e.g., $\beta = \Theta(n \ln(m))$) in (Srinivas et al., 2010). In both cases, the acquisition $A(x)$ is maximized as well to suggest the next point to explore.

For multi-objective Bayesian optimization, we assume that all objectives f_i are samples from known GP priors $\mathcal{GP}(0, \kappa_i)$ with input domain \mathcal{X} . Therefore, we will have a separate GP and acquisition function $A_i(x)$ for each f_i and let the acquisition vector be $A(x) = (a_1(x), \dots, a_k(x))$. It then becomes unclear if there is a single acquisition function that can be used to explore the Pareto frontier well in this setting. Scalarization is a well-known means to combine multiple single-objective acquisition functions for this task. Given a distribution over scalarization functions $s_\lambda(y)$ for $y \in \mathbb{R}^k$, we simply draw a random scalarization function and optimize according to $s_\lambda(A(x))$. The full algorithm is outline in Algorithm 1.

Examples of scalarizations are the simple linear scalarization $s_\lambda(y) = \sum_i \lambda_i y_i$ and the Chebyshev scalarization $s_\lambda(y) = \min_i \lambda_i (y_i - z_i)$ for some input distribution \mathcal{D}_λ . However, these scalarizations lack provable guarantees and it is in fact known that the linear scalarization can only provide solutions on the convex part of the Pareto frontier (Boyd & Vandenberghe, 2004). Furthermore, the Chebyshev scalarization is criticized for lacking diversity and uniformity in the Pareto frontier (Das & Dennis, 1998).

Algorithm 1 Scalarization for Multi-Objective Bayesian Optimization

Input: $F : \mathbb{R}^n \rightarrow \mathbb{R}^k$: multi-objective function, $T \in \mathbb{Z}_+$: number of iterations, $A(x) : \mathbb{R}^n \rightarrow \mathbb{R}^k$ acquisition function, \mathcal{D}_λ : distribution to sample λ for known $s_\lambda(x)$

for $t = 1$ **to** T **do**

Sample: Draw independently $\lambda_t \sim \mathcal{D}_\lambda$.

Optimize: $x_t = \arg \max_{x \in \mathcal{X}} s_{\lambda_t}(A(x))$ where $A = (A_1, \dots, A_k)$ is evaluated with $\mathcal{GP}^{(t-1)}$

Evaluate: $y_t = F(x_t)$

Update: Incorporate (x_t, y_t) into $\mathcal{GP}^{(t-1)}$ to obtain $\mathcal{GP}^{(t)}$, where the posterior update is done independently for each f_i and $\mathcal{GP}(\mu_i, \kappa_i)$

end for

return $\{x_t\}_{t=1}^T$

3. Hypervolume Scalarization

In this section, we introduce our hypervolume scalarization function and demonstrate that the expected scalarization value under a certain distribution of weights will give the dominated hypervolume, up to a constant factor difference. Our proof technique relies on a volume integration argument in spherical coordinates and mirrors an independent proof of (Deng & Zhang, 2019). Later, we prove concentration of the empirical mean, allowing for an accurate estimator of the dominated hypervolume.

Lemma 5 (Also in (Deng & Zhang, 2019)). *Let $Y = \{y_1, \dots, y_m\}$ be a set of m points in \mathbb{R}^k . Then, the hypervolume of Y with respect to a reference point z is given by:*

$$\mathcal{HV}_z(Y) = c_k \mathbb{E}_{\lambda \sim \mathcal{S}_+^{k-1}} \left[\max_{y \in Y} s_\lambda(y - z) \right]$$

where $s_\lambda(y) = \min_i (\max(0, y_i/\lambda_i))^k$ and $c_k = \frac{\pi^{k/2}}{2^k \Gamma(k/2+1)}$ is a dimension-independent constant.

Intuitively, this lemma says that although the maximization of any specific scalarization will bias the optimization to a certain point on the Pareto frontier, there is a specific combination of scalarizations with different weights such that the sum of the maximization of these scalarizations over Y will give the hypervolume of Y . It is important for the maximization to be inside the expectation, implying that maximizing various randomized single-objective scalarizations could serve to also maximize the hypervolume, which is our ultimate objective. In fact, it is easy to see that the hypervolume of a set cannot be written solely as a maximization over the expectation of any scalarizations.

Furthermore, note that our scalarization function, similar to the Chebyshev scalarization, is a minimum over coordinates $s_\lambda(y) = \min_i (\max(0, y_i/\lambda_i))^k$. Intuitively, this min-

imization captures the notion that Pareto dominance is a coordinate-wise optimality criterion, as we can redefine x_1 being Pareto dominated by x_2 as $\min_i f_i(x_2) - f_i(x_1) \geq 0$.

Now that we can rewrite the hypervolume as a specific expectation of maximization of random scalarizations, we proceed to show that our random estimator has controlled variance and therefore concentrates. Specifically, we show that the hypervolume indicator can be computed efficiently via this integral formulation, providing a simple and fast implementation to give an approximation to the hypervolume indicator. Our argument relies on proving smoothness properties of our hypervolume scalarizations for any $\lambda > 0$ and then applying standard concentration inequalities. We note that it is non-obvious why $s_\lambda(y)$ is smooth, since $s_\lambda(y)$ depends inversely on λ_i so when λ_i is small, s_λ might change very fast. The full proof is in the supplementary material.

Lemma 6 (Hypervolume Concentration). *Let $Y = \{y_1, \dots, y_m\} \subseteq \mathbb{R}^k$ and z be a reference point and assume for some $B \geq 1$, we can bound $y_i \leq z + B$. Then, $s_\lambda(y)$ is $O(B^k k^{1+k/2})$ -Lipschitz for any λ and we can independently draw $\lambda_1, \dots, \lambda_s$ such that*

$$\left| \frac{1}{c_k} \mathcal{HV}_z(Y) - \frac{1}{s} \sum_j \max_{y \in Y} s_{\lambda_j}(y - z) \right| \leq \epsilon$$

holds with probability at least $1 - \delta$ with $s = O(B^{2k} k^k \log(1/\delta)/\epsilon^2)$ samples.

4. Hypervolume Regret Bounds

In this section, we derive novel hypervolume regret bounds for Bayesian optimization with UCB and TS acquisition functions under the hypervolume scalarizations. Furthermore, we introduce an algorithmic framework to turn any single-objective optimization algorithm to a multi-objective optimization algorithm via scalarizations and show that single-objective convergence guarantees can provably translate into multi-objective convergence bounds.

For a fixed weight λ , if $X = \{x_1, \dots, x_m\}$ is our dataset and recall X^* is the maximal input set corresponding to the Pareto frontier Y^* , we define the scalarized regret to be

$$r_\lambda(X) = \max_{x \in X^*} s_\lambda(F(x)) - \max_{x \in X} s_\lambda(F(x))$$

The scalarized regret only captures the regret for a specific weighting λ and we emphasize that scalarized regret bounds will guarantee convergence to only one point in X^* . To prove convergence to the whole frontier, we consider the following *Bayes regret* given by $R(X) = \mathbb{E}_\lambda [r_\lambda(X)]$, which integrates over all λ for some distribution \mathcal{D}_λ .

Note that Bayes regret cannot be minimized by a single point, rather it requires a set of points across the Pareto frontier to achieve a small Bayes regret. Let $X_T = \{x_1, \dots, x_T\}$ be the set of chosen inputs up to time T . Then, we wish to bound $R(X_T)$. We cannot analyze the Bayes regret directly. Rather, we bound it via a slightly surrogate regret measure. Let us define the instantaneous regret at time step t as:

$$r(x_t, \lambda_t) = \max_{x \in X^*} s_{\lambda_t}(F(x)) - s_{\lambda_t}(F(x_t))$$

The cumulative regret at time step T is then $R_C(T) = \sum_{t=1}^T r(x_t, \lambda_t)$. We will first bound the cumulative regret and then use the cumulative regret to bound the Bayes regret, which is a scaled version of our hypervolume regret when using hypervolume scalarizations.

Regret bounds are derived using a metric known as the *maximum information gain* (MIG), which captures an information-theoretic notion of uncertainty reduction of our blackbox function. For any $A \subseteq \mathcal{X}$, we define the random set $y_A = \{y_a = f(a) + \epsilon_a | a \in A\}$ be our noisy evaluations. The reduction in uncertainty of the distribution over functions f induced by the GP by observing y_A is given by the mutual information $\mathbf{I}(y_A; f) = \mathbf{H}(f) - \mathbf{H}(f|y_A) = \mathbf{H}(y_A) - \mathbf{H}(y_A|f)$, where \mathbf{H} denotes the Shannon entropy. The maximum information gain after T observations is defined as :

$$\gamma_T = \max_{A \subseteq \mathcal{X}: |A|=T} \mathbf{I}(y_A; f)$$

The mutual information can also be explicitly calculated via a useful formula: $\mathbf{I}(y_A; f) = \frac{1}{2} \log |I + \sigma^{-2} K_A|$, where K_A is the T -by- T covariance matrix of dataset A and $|\cdot|$ is the determinant operator. Using the formula, one can derive bounds of $\gamma_T = O((\log T)^{n+1})$ for the squared exponential kernel and similar bounds for the Matérn and linear/polynomial kernels for any A (Srinivas et al., 2010). When the maximum information gain is small, the \mathcal{GP} function distribution induced by the corresponding is relatively easier to model and regret bounds are therefore tighter.

Theorem 7 (Theorem 1 in (Paria et al., 2018)). *Let each objective $f_i(x)$ for $x \in [0, 1]^n$ follow a Gaussian distribution with marginal variances bounded by 1 and observation noises $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2)$ are independent with $\sigma_i^2 \leq \sigma^2 \leq 1$. Let $\gamma_{T,k} \leq \gamma_T$, where $\gamma_{T,k}$ is the MIG for the k -th objective. Running Algorithm 1 with L -Lipschitz scalarizations on either UCB or TS acquisition function produces an expected cumulative regret after T steps that is bounded by:*

$$\mathbb{E}[R_C(T)] = O(Lkn^{1/2}[\gamma_T T \ln(T)]^{1/2})$$

where the expectation is over choice of λ_t and \mathcal{GP} measure.

Theorem 8. *Assume the conditions in Theorem 7 holds and let $F(\mathcal{X}) \subseteq [0, 2/5]^k$ and $z \geq 0$ and $Y_t = F(X_t)$. Running Algorithm 1 with hypervolume scalarizations with reference point z and $D_\lambda = \mathcal{S}_+^{k-1}$ on either UCB or TS acquisition function produces hypervolume regret after T observations that is bounded by:*

$$\sum_{t=1}^T (\mathcal{HV}_z(Y^*) - \mathcal{HV}_z(Y_t)) \leq O(k^2 n^{1/2} [\gamma_T T \ln(T)]^{1/2})$$

Furthermore, $\mathcal{HV}_z(Y_T) \geq \mathcal{HV}_z(Y^*) - \epsilon_T$, where $\epsilon_T = O(k^2 n^{1/2} [\gamma_T \ln(T)/T]^{1/2})$.

Proof. WLOG, let $z = 0$. From Lemma 5, we see that for our hypervolume scalarization, we can rewrite our regret using the Bayes regret.

$$\begin{aligned} \mathcal{HV}_z(Y^*) - \mathcal{HV}_z(Y_t) &= c_k \mathbb{E}_{\lambda \sim \mathcal{S}_+^{k-1}} \left[\max_{x \in X^*} s_\lambda(F(x)) - \max_{x \in X_t} s_\lambda(F(x)) \right] \\ &= c_k \mathbb{E}_{\lambda \sim \mathcal{S}_+^{k-1}} [r_\lambda(X_t)] \\ &= c_k R(X_t) \end{aligned}$$

Therefore, our Bayes regret is exactly a scaled version of our hypervolume regret with $c_k = \pi^{k/2} / (2^k \Gamma(k/2 + 1)) \leq (\pi^{k/2} e^{k/2}) / (2^k (k/2 + 1)^{k/2})$ with standard bounds on Γ . Next, we use the instantaneous regret to bound Bayes risk and note that $\lambda_t \sim \mathcal{S}_+^{k-1}$.

$$\begin{aligned} R(X_t) &= \mathbb{E}_{\lambda \sim \mathcal{S}_+^{k-1}} \left[\max_{x \in X^*} s_\lambda(F(x)) - \max_{x \in X_t} s_\lambda(F(x)) \right] \\ &\leq \mathbb{E}_{\lambda \sim \mathcal{S}_+^{k-1}} \left[\max_{x \in X^*} s_\lambda(F(x)) - s_\lambda(F(x_t)) \right] \\ &\leq \mathbb{E}_\lambda [r(x_t, \lambda)] \end{aligned}$$

Therefore, by Theorem 7, we conclude that we can bound the Bayes and hypervolume regret by observing the following relations:

$$\begin{aligned} \sum_{i=1}^T R(X_t) &\leq \mathbb{E}_{\lambda_1, \dots, \lambda_t} \left[\sum_{i=1}^T r(x_t, \lambda_t) \right] \\ &\leq \mathbb{E}[R_C(T)] \\ &= O(Lkn^{1/2}[\gamma_T T \ln(T)]^{1/2}) \end{aligned}$$

Lastly, by Lemma 6, we see that $L \leq 2^{-k} k^{1+k/2}$ and note that $c_k L \leq k(\pi^{k/2} e^{k/2}) / (5^k) \leq k$. Together, we finally conclude that

$$\sum_{t=1}^T (\mathcal{HV}_z(Y^*) - \mathcal{HV}_z(Y_t)) \leq O(k^2 n^{1/2} [\gamma_T T \ln(T)]^{1/2})$$

Algorithm 2 Scalarization with General Single-Objective Optimization

Input: $F : \mathbb{R}^n \rightarrow \mathbb{R}^k$: multi-objective function, $T \in \mathbb{Z}_+$: number of iterations, \mathcal{A} : single-objective optimization algorithm, \mathcal{D}_λ : distribution to sample λ for known $s_\lambda(x)$, $l \in \mathbb{Z}_+$: number of scalarization samples

for $i = 1$ **to** l **do**

Sample: Draw independently $\lambda_i \sim \mathcal{D}_\lambda$.

Optimize: Run \mathcal{A} on single-objective function $s_{\lambda_i}(F(x))$ for T iterations to obtain points

$\{x_{\lambda_i}^{(1)}, x_{\lambda_i}^{(2)}, \dots, x_{\lambda_i}^{(T)}\}$

end for

return $\bigcup_i \{x_{\lambda_i}^{(t)}\}_{t=1}^T$

For the final claim, by definition of \mathcal{HV}_z and our equivalence between \mathcal{HV}_z and R , we conclude that $R(X_t)$ must be monotonically decreasing. Our final claim follows from monotonicity and simple algebra. \square

We note that our regret bounds hold for classical Bayesian optimization procedures that are widely used with no artificial modifications. Furthermore, we can generalize our results to show that for any single-objective optimization procedure, one can use hypervolume scalarizations to convert the procedure into a multi-objective optimization procedure via a natural extension (Algorithm 2). More remarkably, if the single-objective optimization procedure admits convergence bounds, we can immediately derive hypervolume convergence bounds by appealing to our previous connections between scalarization and hypervolume via the following theorem. Practically, this implies that any currently used single-objective optimization algorithms can be easily generalized to the multi-objective setting with minimal effort and enjoy provable convergence bounds. The full proof is in the supplementary material.

Theorem 9. *Let $F(\mathcal{X}) \subseteq [0, B]^k$ and $z \geq 0$ and let \mathcal{A} be any single-objective maximization algorithm on objective function $g(x)$ that guarantees that after T iterations, it returns x_T such that $g(x_T) \geq g(x^*) - \epsilon_T$, where x^* is the optima. Then, running Algorithm 2 on hypervolume scalarizations with reference point z and $D_\lambda = \mathcal{S}_+^{k-1}$ with $l = \tilde{O}((2B)^{k^2} k^{k+1} / (c_k^{k-1} \epsilon_T^{k+1}))$ converges to the Pareto frontier and after $l \cdot T$ observations, we have*

$$\mathcal{HV}_z(Y_T) \geq \mathcal{HV}_z(Y^*) - 5c_k \epsilon_T$$

5. Experiments

We empirically demonstrate the utility of hypervolume scalarizations by running our proposed multiobjective algorithms on the Black-Box Optimization Benchmark

(BBOB) functions, which can be paired up into multiple bi-objective optimization problems (Tušar et al., 2016). To emphasize the utility of a new theoretically sound scalarization, we focus on scalarization-related algorithms and comparisons were not made to the vast array of diverse multi-objective blackbox optimization used in certain practical settings. Therefore, our goal is therefore to compare scalarizations on commonly used optimization algorithms, such as UCB and evolutionary strategies. Furthermore, we note that scalarized algorithms have very fast practical runtimes and are especially relevant in settings with a low computational budget.

Our objectives are given by BBOB functions, which are usually non-negative and are minimized. The input space is always a compact hypercube $[-5, 5]^n$ and the global minima is often at the origin. For bi-objective optimization, given two different BBOB functions f_1, f_2 , we attempt to maximize the hypervolume spanned by $(-f_1(x_i), -f_2(x_i))$ over choices of inputs x_i with respect to the reference point $(-5, -5)$. Therefore, all relevant output points are contained in the square between $(-5, -5)$ and $(0, 0)$, giving a maximum hypervolume of 25. Because BBOB functions can drastically different ranges, we first normalize the function by a measure of standard deviation computed by taking the empirical variance over a deterministic set of 30 different inputs. We also apply large random shifts/rotations as well as allow for adding moderate random observation noise to the objective function.

We run each of our algorithms in dimensions $n = 8, 16, 24$ and optimize for 70 iterations with 5 repeats. Our algorithms are the Random algorithm, UCB algorithm, and Evolutionary Strategy (ES). Our scalarizations include the linear and hypervolume scalarization with the weight distribution D_λ as \mathcal{S}_+^1 . Note that for brevity, we do not include the Chebyshev scalarization because it is almost a monotonic transformation of the hypervolume scalarization with a different weight distribution. We run the UCB algorithm via an implementation of Algorithm 1 with a constant standard deviation multiplier of 1.8 and a standard Matérn kernel, while we run the ES algorithms using Algorithm 2 with $T = 1$ and $l = 70$ by relying on a well-known single-objective evolutionary strategy known as Eagle (Yang & Deb, 2010).

From our results, there is a clear trend that UCB algorithms outperform both ES and Random algorithms in all cases, empirically affirming the utility of Bayesian optimization and its strong theoretical regret bounds. We believe that this is because evolutionary strategies tend to follow local descent procedures and get stuck at local minimas, therefore implicitly having less explorative mechanisms. Meanwhile, Bayesian optimization inherently promotes exploration via usage of its standard deviation estimates.

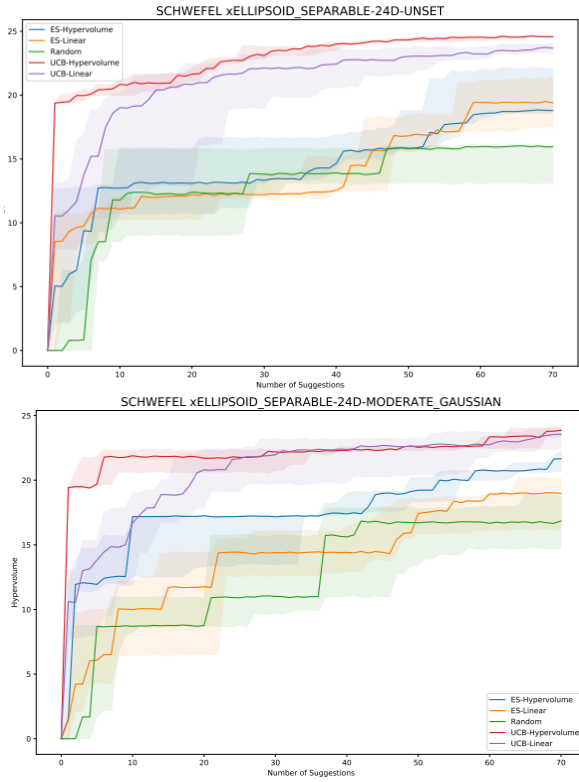


Figure 1. Dominated hypervolume plot of bi-objective optimization of SCHWEFEL and ELLIPSOID BBOB functions in 24D with no and Gaussian observation noise. Notice that Hypervolume scalarization is slightly better than Linear scalarization in both UCB and ES algorithms and this is more stark in the noise-less case.




Figure 2. Dominated hypervolume and Pareto plot of bi-objective optimization of ELLIPSOID and SPHERE BBOB functions in 24D with no (UNSET) observation noise. Notice that the UCB algorithms outperform the ES algorithms and the Hypervolume scalarization is slightly better. Also, when comparing UCB-Hypervolume (blue points) and UCB-Linear (pink points) in the Pareto plot, the Hypervolume scalarization is slightly better and never suggests points that are significantly sub-optimal, although the difference is not very stark because of the convex nature of the Pareto frontier.

Furthermore, we see that the hypervolume scalarization slightly outperforms the linear scalarization, with UCB-Hypervolume being a clear winner in certain cases, even with Gaussian observation noise (see Fig 2). The superior performance of the hypervolume scalarization is seen in both UCB and ES algorithms (see Fig 1), and using the hypervolume scalarization is often never worse than using the linear scalarization. We note that the difference is more stark when there is no noise added, allowing for less variance when calculating the scalarization. Also, we note that when the Pareto frontier is almost convex, the difference in performance becomes hard to observe. As seen from the Pareto plot in Fig 2, we see that although UCB-Hypervolume does produce a better Pareto frontier than UCB-Linear, the convex nature of the Pareto frontier allows linear scalarizations to perform decently. A complete profile of the plots are given in the supplementary material.

6. Conclusion

We introduced the hypervolume scalarization functions and utilized its connection to the hypervolume indicator to derive hypervolume regret bounds for many classes of multi-objective optimization algorithms that rely on single-objective subroutines. Ultimately, we hoped to have presented a convincing case as to the underlying reason why researchers and practitioners should use hypervolume scalarizations, as opposed to the more popular linear scalarizations, which are provably suboptimal. Even though our scalarization function enjoys smoothness and concentration bounds, we note that possible improvements to the scalarizations can be made for variance reduction. We believe that a lower-variance and better-concentrated scalarization can be constructed that also can similar provable hypervolume error guarantees. Furthermore, it is conceivable

that variance reduction techniques during the optimization process can be applied to achieve better concentration and convergence.

Furthermore, our regret bounds can likely be improved, especially the general regret bound for any single-objective optimization procedure given by Theorem 9. We believe the exponential dependence on k^2 can be improved, as well as removing the extra l factor completely. However, to achieve these better convergence bounds, it will most likely require algorithmic changes. Lastly, our experiments could be improved by concocting a specific optimization with a concave Pareto frontier, looking at optimization with more than two objectives, or considering a more diverse set of algorithms. Understanding and quantifying the full impact of using random scalarizations to generalize single-objective algorithms to multi-objective algorithms is an open problem.

References

- Aliano Filho, A., Moretti, A. C., Pato, M. V., and de Oliveira, W. A. An exact scalarization method with multiple reference points for bi-objective integer linear optimization problems. *Annals of Operations Research*, pp. 1–35, 2019.
- Beume, N. and Rudolph, G. *Faster s -metric calculation by considering dominated hypervolume as klee’s measure problem*. Universitätsbibliothek Dortmund, 2006.
- Boyd, S. P. and Vandenberghe, L. *Convex optimization*. Cambridge University Press, Cambridge, 2004. ISBN 0-521-83378-7.
- Bringmann, K. and Friedrich, T. Approximating the volume of unions and intersections of high-dimensional geometric objects. *Computational Geometry*, 43(6-7):601–610, 2010.
- Das, I. and Dennis, J. E. Normal-boundary intersection: A new method for generating the pareto surface in nonlinear multicriteria optimization problems. *SIAM journal on optimization*, 8(3):631–657, 1998.
- Deng, J. and Zhang, Q. Approximating hypervolume and hypervolume contributions using polar coordinate. *IEEE Transactions on Evolutionary Computation*, 23(5):913–918, 2019.
- Emmerich, M., Beume, N., and Naujoks, B. An emo algorithm using the hypervolume measure as selection criterion. In *International Conference on Evolutionary Multi-Criterion Optimization*, pp. 62–76. Springer, 2005.
- Emmerich, M. T. and Deutz, A. H. A tutorial on multi-objective optimization: fundamentals and evolutionary methods. *Natural computing*, 17(3):585–609, 2018.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1126–1135. JMLR. org, 2017.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1321–1330. JMLR. org, 2017.
- Hernández-Lobato, D., Hernandez-Lobato, J., Shah, A., and Adams, R. Predictive entropy search for multi-objective bayesian optimization. In *International Conference on Machine Learning*, pp. 1492–1501, 2016.
- Kasimbeyli, R., Ozturk, Z. K., Kasimbeyli, N., Yalcin, G. D., and Erdem, B. I. Comparison of some scalarization methods in multiobjective optimization. *Bulletin of the Malaysian Mathematical Sciences Society*, 42(5): 1875–1905, 2019.
- Kendall, A., Gal, Y., and Cipolla, R. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7482–7491, 2018.
- 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.
- Konak, A., Coit, D. W., and Smith, A. E. Multi-objective optimization using genetic algorithms: A tutorial. *Reliability Engineering & System Safety*, 91(9):992–1007, 2006.
- Kutner, M. H., Nachtsheim, C. J., Neter, J., Li, W., et al. *Applied linear statistical models*, volume 5. McGraw-Hill Irwin Boston, 2005.
- LeCun, Y., Bengio, Y., and Hinton, G. Deep learning. *nature*, 521(7553):436–444, 2015.
- Lin, X., Zhen, H.-L., Li, Z., Zhang, Q.-F., and Kwong, S. Pareto multi-task learning. In *Advances in Neural Information Processing Systems 32*, pp. 12037–12047. Curran Associates, Inc., 2019. URL <http://papers.nips.cc/paper/9374-pareto-multi-task-learning.pdf>.
- Lu, S., Wang, G., Hu, Y., and Zhang, L. Multi-objective generalized linear bandits. *arXiv preprint arXiv:1905.12879*, 2019.
- Nakayama, H., Yun, Y., and Yoon, M. *Sequential approximate multiobjective optimization using computational intelligence*. Springer Science & Business Media, 2009.

- Neal, B., Mittal, S., Baratin, A., Tantia, V., Scicluna, M., Lacoste-Julien, S., and Mitliagkas, I. A modern take on the bias-variance tradeoff in neural networks. *arXiv preprint arXiv:1810.08591*, 2018.
- Öner, D., Karakurt, A., Eryılmaz, A., and Tekin, C. Combinatorial multi-objective multi-armed bandit problem. *arXiv preprint arXiv:1803.04039*, 2018.
- Paria, B., Kandasamy, K., and Póczos, B. A flexible framework for multi-objective bayesian optimization using random scalarizations. *arXiv preprint arXiv:1805.12168*, 2018.
- Picheny, V. Multiobjective optimization using gaussian process emulators via stepwise uncertainty reduction. *Statistics and Computing*, 25(6):1265–1280, 2015.
- Rojters, D. M., Vamplew, P., Whiteson, S., and Dazeley, R. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113, 2013.
- Russo, D. and Van Roy, B. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4):1221–1243, 2014.
- Scarlett, J., Bogunovic, I., and Cevher, V. Lower bounds on regret for noisy gaussian process bandit optimization. *arXiv preprint arXiv:1706.00090*, 2017.
- Schmidt, M., Schöbel, A., and Thom, L. Min-ordering and max-ordering scalarization methods for multi-objective robust optimization. *European Journal of Operational Research*, 275(2):446–459, 2019.
- Srinivas, N., Krause, A., Kakade, S., and Seeger, M. Gaussian process optimization in the bandit setting: No regret and experimental design. *ICML*, 2010.
- Sutton, R. S., Barto, A. G., et al. *Introduction to reinforcement learning*, volume 2. MIT press Cambridge, 1998.
- 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.
- Turğay, E., Öner, D., and Tekin, C. Multi-objective contextual bandit problem with similarity information. *arXiv preprint arXiv:1803.04015*, 2018.
- Tušar, T., Brockhoff, D., Hansen, N., and Auger, A. Coco: the bi-objective black box optimization benchmarking (bbob-biobj) test suite. *ArXiv e-prints*, 2016.
- Williams, C. K. and Rasmussen, C. E. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- Yang, K., Emmerich, M., Deutz, A., and Bäck, T. Efficient computation of expected hypervolume improvement using box decomposition algorithms. *arXiv preprint arXiv:1904.12672*, 2019.
- Yang, X.-S. and Deb, S. Eagle strategy using lévy walk and firefly algorithms for stochastic optimization. In *Nature Inspired Cooperative Strategies for Optimization (NICSO 2010)*, pp. 101–111. Springer, 2010.
- Zhang, H., Yu, Y., Jiao, J., Xing, E. P., Ghaoui, L. E., and Jordan, M. I. Theoretically principled trade-off between robustness and accuracy. *arXiv preprint arXiv:1901.08573*, 2019.
- Zhang, Q. and Li, H. Moea/d: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on evolutionary computation*, 11(6):712–731, 2007.
- Zitzler, E. and Thiele, L. Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach. *IEEE transactions on Evolutionary Computation*, 3(4):257–271, 1999.
- Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C. M., and Da Fonseca, V. G. Performance assessment of multi-objective optimizers: An analysis and review. *IEEE Transactions on evolutionary computation*, 7(2):117–132, 2003.
- Zitzler, E., Laumanns, M., and Bleuler, S. A tutorial on evolutionary multiobjective optimization. In *Metaheuristics for multiobjective optimisation*, pp. 3–37. Springer, 2004.
- Zliobaite, I. On the relation between accuracy and fairness in binary classification. *arXiv preprint arXiv:1505.05723*, 2015.
- Zuluaga, M., Sergent, G., Krause, A., and Püschel, M. Active learning for multi-objective optimization. In *International Conference on Machine Learning*, pp. 462–470, 2013.