

---

# Robust Multi-Objective Bayesian Optimization Under Input Noise

---

Samuel Daulton<sup>\*12</sup> Sait Cakmak<sup>\*13</sup> Maximilian Balandat<sup>1</sup> Michael A. Osborne<sup>2</sup> Enlu Zhou<sup>3</sup>  
Eytan Bakshy<sup>1</sup>

## Abstract

Bayesian optimization (BO) is a sample-efficient approach for tuning design parameters to optimize expensive-to-evaluate, black-box performance metrics. In many manufacturing processes, the design parameters are subject to random input noise, resulting in a product that is often less performant than expected. Although BO methods have been proposed for optimizing a single objective under input noise, no existing method addresses the practical scenario where there are multiple objectives that are sensitive to input perturbations. In this work, we propose the first multi-objective BO method that is robust to input noise. We formalize our goal as optimizing the multivariate value-at-risk (MVAR), a risk measure of the uncertain objectives. Since directly optimizing MVAR is computationally infeasible in many settings, we propose a scalable, theoretically-grounded approach for optimizing MVAR using random scalarizations. Empirically, we find that our approach significantly outperforms alternative methods and efficiently identifies optimal robust designs that will satisfy specifications across multiple metrics with high probability.

## 1. Introduction

Scientists and engineers frequently face optimization problems where the goal is to tune a set of parameters to optimize multiple competing black-box objective functions. Typically, no single design is best with respect to every objective. Hence, the goal in multi-objective (MO) optimization is to identify the Pareto frontier (PF) of optimal trade-offs between the objectives and the corresponding optimal designs. These problems are ubiquitous in a variety of domains including manufacturing (Liao et al., 2008), materials design (Ashby, 2000), robotics (Calandra and Peters, 2014),

and machine learning (Izquierdo et al., 2021; Eriksson et al., 2021). Obtaining measurements of the objectives often requires resource-intensive simulation or experimentation. Therefore, any practical routine for optimizing such functions must be highly sample-efficient; that is, the method must identify optimal design parameters while only querying the objective functions at a small number of designs. Bayesian optimization (BO) (Shahriari et al., 2016) is a popular technique for addressing this class of problems.

In many real-world applications, the final implementation of the selected design parameters are subject to input noise, resulting from uncontrollable perturbations of the input parameters (Beyer and Sendhoff, 2007). For example, many vaccine production processes involve freeze-drying procedures to stabilize active pharmaceutical ingredients to increase storage lifetime (Mortier et al., 2016; Xie and Schenckendorf, 2019). Tuning the operating parameters, such as shelf temperature and chamber pressure, can significantly improve the efficiency of the drying step with little reduction in product quality. However, shelf temperature and chamber pressure are subject to uncontrollable random input noise around the *nominal*<sup>1</sup> input parameters. Robustness with respect to this input noise is critical. Higher temperatures lead to greater efficiency, but also make the process more sensitive to perturbations in the temperature because if temperature exceeds the critical collapse threshold, there is irreversible product damage. A more conservative temperature that is robust to input noise may be a better choice than a higher temperature that is worse with respect to the *nominal objectives*. In such high-stakes and high-throughput production pipelines, decision-makers seek to identify robust design parameters that ensure the manufactured products will have high objective quality with high probability.

Optimization without consideration of input noise can lead to solutions that are catastrophic when subjected to input noise at the implementation stage (Doltsinis and Kang, 2004). The toy problem that is illustrated in left plot of Figure 1 demonstrates a scenario where a non-robust design (depicted as a purple square) is better than the robust design

---

<sup>\*</sup>Equal contribution <sup>1</sup>Meta <sup>2</sup>University of Oxford <sup>3</sup>Georgia Institute of Technology. Correspondence to: Samuel Daulton <sdaulton@fb.com>, Sait Cakmak <saicakmak@fb.com>.

Proceedings of the 39<sup>th</sup> International Conference on Machine Learning, Baltimore, Maryland, USA, PMLR 162, 2022. Copyright 2022 by the author(s).

<sup>1</sup>We call the design as specified by the decision maker the *nominal design*, and the corresponding value without any perturbations as the *nominal values* for the *nominal objectives*.




Figure 1: A toy problem where the goal is to tune a single parameter  $x$  to maximize two objectives. Left: The nominal values for a non-robust (purple,  $x_1$ ) and a robust design design (green,  $x_2$ ) are indicated using squares. The plus markers illustrate objective values of each design under zero-mean Gaussian input noise with a standard deviation of 0.1. The non-robust design can lead to low objective values under input perturbations. Center: An illustration of the MVAR sets of the non-robust and robust designs. The triangles represent a discrete approximation of the MVAR set of each design under the input noise distribution. For each point  $z$  in the MVAR set of a given design, the objective values for that design subject to input perturbations are  $\geq z$  with probability  $\alpha \geq 0.9$ . In other words, the objectives under input noise for each design will fall in the respective shaded region with probability  $\geq \alpha$ . Under input noise, the non-robust design  $x_1$  may result in poor objective values despite yielding better values (relative to the robust design  $x_2$ ) without perturbations. After accounting for input noise,  $x_2$  is a more robust solution than  $x_1$ . Right: The MVAR set (black stars) across three optimal designs  $x_1^*, x_2^*, x_3^*$  is the set of optimal points across the union of the MVAR sets (colored triangles) of each design.

(marked by a green square) with respect to the nominal values. However, the performance metrics for the non-robust design can be significantly worse when the design is subject to input noise. In contrast, the objectives are far less sensitive to input perturbations around the robust design.

To identify designs that are robust to a noisy performance metric due to input noise or observation noise, the value-at-risk (VAR) is often used as the robust objective because it provides high-probability performance guarantees (Basel Committee on Banking Supervision, 2012). The  $\alpha$  VAR is the  $(1 - \alpha)$ -quantile (where  $0 \leq \alpha \leq 1$ ) of the cumulative distribution function (CDF) of the performance metric. Under input noise, variation in the objective is induced via the uncertain inputs. Intuitively, the VAR is the largest value such that the objective value of a given design subject to input perturbations will be greater than that value with probability at least  $\alpha$ . Thus, in the context of manufacturing, a design with  $\alpha$  VAR exceeding the target specification produces a yield of at least  $\alpha$ .

Many recent works have considered BO methods that are robust to input noise in the single objective setting, e.g., by optimizing VAR (Nguyen et al., 2021b) or other risk measures (Fröhlich et al., 2020; Cakmak et al., 2020; Nguyen et al., 2021a). However, no previous work has considered sample-efficient, generally-applicable multi-objective Bayesian optimization (MOBO) methods that are robust to input perturbations. To our knowledge, the only existing MO methods that are robust to input noise are evolution-

ary algorithms (EA) that often require tens of thousands of evaluations (e.g., Deb and Gupta (2005)).

In the MO setting, high-probability performance guarantees of a single design can be assessed using the multivariate value-at-risk (MVAR) (Prékopa, 2012). MVAR maps a probability value  $\alpha$  to a set of vectors where each element provides a lower bound on the the objectives’ possible values under input noise with probability  $\alpha$ . As illustrated in the center plot of Figure 1, the MVAR set of a robust design ( $x_2$ ) often provides significantly better probabilistic lower bounds than the MVAR set of a non-robust design ( $x_1$ ).

Similar to the PF in the standard MO setting—where the PF is the set of optimal trade-offs between objectives with no input noise—the *global* MVAR set is the set of optimal trade-offs that can be achieved with probability  $\geq \alpha$  under input noise across all possible designs. The MVAR across a set of 3 designs is illustrated in the right plot of Figure 1.

While MVAR provides a natural measure of robust performance guarantees, it is relatively expensive to compute, making it computationally challenging to directly optimize MVAR in the context of MOBO (see Appendix D.1 for a discussion). In this work, we propose a family of novel, theoretically-grounded methods for optimizing MVAR via random scalarizations that mitigates many challenges associated with optimizing MVAR directly.

## Contributions

1. We introduce robust MO optimization under input

noise, formalize the problem in terms of optimizing global MVAR—a novel, probabilistic form of a robust PF—and discuss computational challenges unique to this setting.

2. We derive a novel theoretical connection between the MVAR and the VAR of a particular scalarization of the objectives, which motivates a family of computationally efficient BO methods for identifying MVAR-optimal designs using an MVAR Approximation based on Random Scalarizations (MARS).
3. We demonstrate that MARS vastly outperforms non-robust alternatives on a variety of synthetic and real-world robust MO optimization problems, including a pharmaceutical manufacturing application.
4. We derive and evaluate extensions of our methods to handle expensive-to-evaluate black-box constraints and parallel candidate generation.

## 2. Background

**Multi-Objective Optimization** In MO optimization, the goal is to, without loss of generality, maximize a vector-valued black-box function  $\max_{\mathbf{x} \in \mathcal{X}} \mathbf{f}(\mathbf{x})$  where  $\mathbf{f}(\mathbf{x}) := [f_1(\mathbf{x}), \dots, f_M(\mathbf{x})]$ ,  $M \geq 2$ , and  $\mathcal{X} \subset \mathbb{R}^d$  is a compact search space. Often, there may be additional black-box constraints  $\mathbf{c}(\mathbf{x}) \geq \mathbf{0}$ , where  $\mathbf{c}(\mathbf{x}) \in \mathbb{R}^V$ ,  $V > 0$ , that must be satisfied. We consider the setting where  $\mathbf{f}$  and  $\mathbf{c}$  have no known analytic expressions and no known or observed gradients. The goal is to identify the Pareto frontier (PF) of optimal trade-offs and corresponding Pareto set of optimal designs  $\mathcal{X}^*$ .

**Notation** For vectors  $\mathbf{y}, \mathbf{y}' \in \mathbb{R}^M$ . Let  $\geq$  and  $>$  denote the component-wise extensions of the order notations  $\geq$  and  $>$ , i.e.,  $\mathbf{y} \geq \mathbf{y}' \iff y_i \geq y'_i \forall i$  and  $\mathbf{y} > \mathbf{y}' \iff y_i > y'_i \forall i$ , where  $\cdot_i$  denotes the  $i^{\text{th}}$  element.

**Definition 2.1.** A vector  $\mathbf{f}(\mathbf{x})$  *Pareto dominates*  $\mathbf{f}(\mathbf{x}')$ , denoted by  $\mathbf{f}(\mathbf{x}) \succ \mathbf{f}(\mathbf{x}')$ , if  $\mathbf{f}(\mathbf{x}) \geq \mathbf{f}(\mathbf{x}')$  and  $\exists j \in \{1, \dots, M\}$  such that  $f^{(j)}(\mathbf{x}) > f^{(j)}(\mathbf{x}')$ .

**Definition 2.2.** The *Pareto frontier* over a set of objective vectors  $\mathcal{F} = \{\mathbf{f}(\mathbf{x}) \mid \mathbf{x} \in X \subseteq \mathcal{X}\}$  is  $\text{PARETO}(\mathcal{F}) = \{\mathbf{f}(\mathbf{x}) \in \mathcal{F} : \nexists \mathbf{x}' \in X \text{ s.t. } \mathbf{f}(\mathbf{x}') \succ \mathbf{f}(\mathbf{x})\}$ . If there are constraints  $\mathbf{c}(\mathbf{x})$ , elements of PARETO are subject to the additional membership assumption that  $\mathbf{c}(\mathbf{x}) \geq \mathbf{0}$ . We call the corresponding set of optimal designs the *Pareto set*.

Although the true PF  $\mathcal{P}^* = \text{PARETO}(\{\mathbf{f}(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}})$  is typically unknown, an MO optimization algorithm can be employed to identify a finite approximation. Hypervolume is a commonly used metric to measure the quality of a PF.

**Definition 2.3.** The *hypervolume (HV) indicator* of a set of points  $\mathcal{Y} \subset \mathbb{R}^M$  is the  $M$ -dimensional Lebesgue measure  $\lambda_M$  of the region dominated by  $\mathcal{P} := \text{PARETO}(\mathcal{Y})$  and

bounded from below by a reference point  $\mathbf{r} \in \mathbb{R}^M$ , which we write as  $\text{HV}(\mathcal{Y}, \mathbf{r})$ .

**Definition 2.4.** The *hypervolume improvement (HVI)* of a set of points  $\mathcal{Y}'$  with respect to an existing Pareto frontier  $\mathcal{P}$  and reference point  $\mathbf{r}$  is defined as  $\text{HVI}(\mathcal{Y}' | \mathcal{P}, \mathbf{r}) = \text{HV}(\mathcal{P} \cup \mathcal{Y}' | \mathbf{r}) - \text{HV}(\mathcal{P} | \mathbf{r})$ .<sup>2</sup>

**Bayesian Optimization (BO)** is a sample-efficient technique for optimizing black-box functions (Frazier, 2018). BO relies on a probabilistic surrogate model—typically a Gaussian process (GP), which provides well-calibrated uncertainty estimates—and an acquisition function that leverages the surrogate model to quantify the value of evaluating the objective functions for a design  $\mathbf{x}$ . The acquisition function balances exploring areas with high uncertainty and exploiting regions believed to be optimal. BO selects the next point  $\mathbf{x}$  to evaluate by maximizing the acquisition function (which is cheap to evaluate relative to the objective functions), observes a (potentially noisy) measurement of the metrics  $\mathbf{y}$ , adds the new observation to the dataset  $\mathcal{D} = \{(\mathbf{x}, \mathbf{y})\} \cup \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$ , updates the surrogate model to incorporate the new observation, and repeats this process for a predetermined budget of evaluations. In the MO setting, a common approach is to optimize a random scalarization of the objectives using a single objective acquisition function, such as expected improvement (Knowles, 2006) or Thompson sampling (Paria et al., 2018). Alternatively, a multi-objective acquisition function can be used to directly optimize the PF. For example, expected hypervolume improvement (EHVI) (Emmerich et al., 2006) aims to maximize the HV of the PF under the surrogate model’s posterior distribution.

## 3. Related Work

Many recent works on BO have considered settings where at the time of implementation either (i) the parameters are subject to noise (Bogunovic et al., 2018; Oliveira et al., 2019; Fröhlich et al., 2020)—as we consider in this work—or (ii) the system performance depends on auxiliary unknown environmental variables (Kirschner et al., 2020; Iwazaki et al., 2021). While previous work has focused on optimizing the expected (Toscano-Palmerin and Frazier, 2018) or the worst-case performance (ur Rehman et al., 2014; Sessa et al., 2020), a recent body of work has focused on optimizing more sophisticated risk measures (Torossian et al., 2020; Cakmak et al., 2020; Nguyen et al., 2021a;b). However, despite recent significant interest in non-robust MOBO (Lukovic et al., 2020; Suzuki et al., 2020), to our knowledge, no prior work has studied robust MOBO.

Motivated by practical limitations due to manufacturing tolerances, Malkomes et al. (2021) proposed constraint ac-

<sup>2</sup>Henceforth, we omit  $\mathbf{r}$  for brevity. As in previous work, we assume that the reference point  $\mathbf{r}$  is known and specified by the decision maker (Yang et al., 2019).

tive search (CAS), which aims to identify diverse solutions in the region of the search space that exceeds a minimum threshold on the objectives. However, CAS does not model or account for input noise, and CAS alone cannot produce any guarantees on robustness to input noise. Methods such as CAS would require a post-hoc analysis using the data collected during optimization to analyze the sensitivity of the solutions to input noise (Calandra and Peters, 2014).

Approaching robust design by decoupling data collection and sensitivity analysis is central to the Taguchi method (Taguchi, 1989). Data acquisition often revolves around finding designs that balance the mean and variance of the sensitive objective under input noise (Beyer and Sendhoff, 2007). Do et al. (2021) propose an approach in this vein for the two-objective setting where only one objective is subject to input noise. However, the algorithm does not seek to identify trade-offs with high probability robustness guarantees, and the method does not handle multiple sensitive objectives. In contrast with the Taguchi method, we aim to unify data collection and sensitivity analysis by selecting designs that are believed to yield high-probability performance guarantees.

Outside the BO literature, robust MO optimization has been studied using either EAs (Gupta and Deb, 2005; He et al., 2019) or assuming access to the explicit mathematical programming formulation of the problem (Majewski et al., 2017; Roberts et al., 2018). Those works have focused on finding the Pareto frontier of the expectation or the worst-case objectives or on finding the Pareto frontier of the nominal objectives with additional constraints on the deviation from the nominal values (Deb and Gupta, 2005; Avigad and Branke, 2008). Some works have considered conceptual properties of different scalarization methods (Ide and Köbis, 2014), but not in relation to MVAR. EAs that are robust to input noise are not applicable to the small evaluation budget regime that we consider because they typically require a large number of function evaluations (Deb and Gupta, 2005). Even methods that combine EAs with GPs require thousands of evaluations per design (Zhou et al., 2018).

As a final differentiator from prior work, we consider the practical setting where there are additional black-box constraints that are sensitive to input noise (Marzat et al., 2013; Li and Li, 2015), which is a subject addressed by only a few BO methods (Beland and Nair, 2017) even in the single objective case.

#### 4. Multi-Objective Optimization with Noisy Inputs

In many practical scenarios, the nominal performance of a design can be evaluated by means of a simulation (e.g., by simulating the pharmaceutical process under nominal operating conditions). We consider the setting where we

can simulate  $f(x)$  for any given design  $x \in \mathcal{X}$ , but that the design is subject to noise  $\xi(x)$  from a known noise process  $\xi(x) \sim P(\xi; x)$  at implementation time.<sup>3</sup> The realized system performance is given by the random variable  $f(x \diamond \xi)$ , where  $x \diamond \xi$  denotes any known function  $g(x, \xi)$  (e.g. for additive noise  $\diamond$  is simply  $+$ ). For an extended problem formulation including black-box constraints, see Appendix A.4.

In robust optimization, the goal is often to optimize a *risk measure*  $\rho[f(x \diamond \xi)]$  that maps a random variable to a statistic of its distribution. A common risk measure is the expectation over the input noise distribution (Deb and Gupta, 2005; Toscano-Palmerin and Frazier, 2018; Fröhlich et al., 2020),  $\mathbb{E}_{\xi \sim P(\xi)}[f(x \diamond \xi)]$ , which can be used instead of the random variable  $f(x \diamond \xi)$  and optimized via standard multi-objective optimization methods. We propose the first MOBO methods for optimizing expectation objectives in Appendix F. Despite its widespread use, the expectation risk measure may not always align with the practitioner’s true robustness goals. Often, one would prefer solutions with objectives that are better than some performance specification  $z \in \mathbb{R}^M$  with high probability (e.g. to maximize production yield) (Sarykalin et al., 2008). Hence, in the single-objective setting, probabilistic risk measures such as value-at-risk (VAR) are frequently used.

**Definition 4.1.** Given input noise  $\xi \sim P(\xi)$  where  $\xi \in \mathbb{R}^d$  and a confidence level  $\alpha \in [0, 1]$ , the value-at-risk for a given point  $x$  is:

$$\text{VAR}_\alpha[f(x \diamond \xi)] = \sup\{z \in \mathbb{R} : P[f(x \diamond \xi) \geq z] \geq \alpha\}.$$

Although several BO methods exist for optimizing VAR (Cakmak et al., 2020; Nguyen et al., 2021b), they cannot directly be used in the MO setting because VAR is not defined for multivariate random variables. A naïve way to extend VAR to the MO setting would be to consider the VAR of each objective independently. However, this ignores the fact that all  $M$  objectives are evaluated at the same realization of  $x \diamond \xi$ . Considering the VAR of each objective independently typically leads to overly optimistic risk estimates because objectives under input noise are not typically *simultaneously* greater than or equal to their respective independent VARs (i.e. the  $(1 - \alpha)$ -quantiles) with probability  $\geq \alpha$ . Thus it is important to use risk measures such as multivariate value-at-risk (MVAR) that account for the joint distribution of the objectives (Prékopa, 2012).

This is illustrated in the center plot in Figure 1. Under input noise, the objective values are correlated, which underscores the importance of accounting for the joint distribution of the objectives in measures of robustness. In addition, the center plot in Figure 1 shows that  $f(x \diamond \xi)$  has an asymmetric distribution, even for this very simple and well-behaved

<sup>3</sup>For brevity, we omit the dependency on  $x$  in our notation and write  $\xi$  and  $P(\xi)$  going forward.



toy problem, which highlights how applying VAR to each objective independently or using the expectation risk measure may conceal underlying variation and risk. Indeed, the results in Appendix I show that optimizing an expectation risk measure on this problem results in poor performance.

In contrast with VAR and the expectation risk measure which map a random variable to *single* scalar or vector, MVAR maps a random variable to a non-dominated *set* of vectors in the outcome space that are dominated by  $\alpha$ -fraction of all possible realizations, where  $\alpha \in [0, 1]$  is a hyperparameter set by the practitioner. That is, each vector in the MVAR set corresponds to an objective specification that a design will meet with probability  $\geq \alpha$ . Therefore,  $\alpha$  is an interpretable risk level that can be valuable in manufacturing applications where one wishes to find the PF of all objective specifications with a guaranteed yield fraction ( $\alpha$ ).

**Definition 4.2.** The MVAR of  $\mathbf{f}$  for a given point  $\mathbf{x}$  and confidence level  $\alpha \in [0, 1]$  is:

$$\text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})] = \text{PARETO}(\{\mathbf{z} \in \mathbb{R}^M : P[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq \mathbf{z}] \geq \alpha\}).$$

The MVAR set over  $X$  specifies objective vectors  $\mathbf{z}$  such that there exists a known design  $\mathbf{x} \in X$  with corresponding random objectives  $\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})$  under  $P(\boldsymbol{\xi})$  that dominate  $\mathbf{z}$  with probability  $\geq \alpha$ .

**Definition 4.3.** The MVAR for a set of points  $X$  is:

$$\text{MVAR}_\alpha[\{\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})\}_{\mathbf{x} \in X}] = \text{PARETO}\left(\bigcup_{\mathbf{x} \in X} \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]\right).$$

The global MVAR across the design space,  $\text{MVAR}_\alpha[\{\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})\}_{\mathbf{x} \in \mathcal{X}}]$ , is a robust analogue of the PF in the standard MO setting. The concept of the MVAR of a set of design points  $X$  is a novel contribution of this work.

**Optimization Goal** In this work, our goal is to identify the MVAR set across the design space:  $\text{MVAR}_\alpha[\{\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})\}_{\mathbf{x} \in \mathcal{X}}]$ . Given an approximate MVAR set across the design space, a decision-maker can pick a design according to their preferences. Similar to the standard MO setting, the HV of the MVAR set across the design space can be used to evaluate optimization performance.

## 5. Optimizing MVAR

A natural approach for optimizing MVAR is to directly maximize the HV dominated by the MVAR set. Although MVAR of a given point typically cannot be evaluated directly, it can be approximated using  $n_\xi$  MC samples of  $\boldsymbol{\xi}$ , provided that independent samples can be drawn from the noise process. Thus, evaluating the MVAR set across the previously evaluated designs using the surrogate requires sampling from the posterior of  $P(\mathbf{f}|\mathcal{D})$  evaluated

jointly at  $\mathbf{x}_1 \diamond \boldsymbol{\xi}_1, \dots, \mathbf{x}_n \diamond \boldsymbol{\xi}_n$  for  $i = 1, \dots, n_\xi$ , where  $X_{1:n} := \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$  are the previously evaluated designs. Since  $\{\mathbf{f}(\mathbf{x}' \diamond \boldsymbol{\xi})\}_{\mathbf{x}' \in X_{1:n}}$  is typically not observed, the corresponding posterior predictions may have large uncertainties. In order to get a reliable estimate of MVAR, we would need to integrate over the posterior distribution of  $\{\mathbf{f}(\mathbf{x}' \diamond \boldsymbol{\xi})\}_{\mathbf{x}' \in X_{1:n}}$ .  $q\text{NEHVI}$  (Daulton et al., 2021a) is a variant of EHVI that integrates over the uncertainty in function values at previously evaluated designs. This makes  $q\text{NEHVI}$  suitable for optimizing MVAR.

However, several computational issues—including time complexity that is exponential in the number of objectives and exponential in the size of  $\text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$ —make it infeasible to directly optimize MVAR with  $q\text{NEHVI}$  in many settings. We defer a detailed discussion to Appendix D and present an empirical evaluation in Appendix I.

### 5.1. Relationship between MVAR and Scalarizations

An alternative to direct optimization of the MVAR set is to apply a scalarization to the objectives and use a standard risk measure on the scalarized objective. Unlike the use of independent risk measures on each objective, this approach accounts for the correlation between outcomes induced by the input perturbation. In this section, we present our main theoretical result: under limited assumptions, there exists a bijection, based on VAR, that maps a particular family of scalarizations—Chebyshev scalarizations (Kaisa, 1999)—to points in the MVAR set. In other words, each point in the MVAR set corresponds to a particular set of scalarization weights. This means that we can recover the entire MVAR set using these scalarizations, without any loss. Proofs and additional theoretical results including extensions to the constrained setting are provided in Appendix A.

**Definition 5.1.** Let  $\mathbf{w} \in \Delta_+^{M-1}$ , where  $\Delta_+^{M-1}$  denotes the positive  $(M-1)$ -simplex, and let  $\mathbf{r} \in \mathbb{R}^M$ . The Chebyshev scalarization  $s[\mathbf{y}, \mathbf{w}, \mathbf{r}]$  is given by  $s[\mathbf{y}, \mathbf{w}, \mathbf{r}] = \min_i w_i(y_i - r_i)$ , where  $\cdot_i$  denotes the  $i^{\text{th}}$  dimension.<sup>4</sup>

The contour in the left plot in Figure 2 shows the Chebyshev scalarization for a fixed  $\mathbf{w}$  for the two-objective toy problem from Figure 1 and illustrates a connection between Pareto dominance and the Chebyshev scalarization, which we formalize below. The black points are function values under sampled perturbations for a single design  $\mathbf{x}$ . The center plot in Figure 2 shows the distribution of Chebyshev scalarization values for a given  $\mathbf{w}$  and the black line indicates the  $\alpha$ -level VAR. The right plot in Figure 2 illustrates that using the VAR of a Chebyshev scalarization, we can deduce a point  $\mathbf{z}$  such that the function values under the

<sup>4</sup>Typically,  $\mathbf{f}$  is scaled to the unit cube using the  $\mathbf{r}$  as the lower bound before applying the scalarization. Since the scaled reference point is  $\mathbf{0}$ , hence forth, we omit  $\mathbf{r}$  for brevity. See Appendix G.1 for details.




Figure 2: Construction of MVAR sets via random scalarizations for the 1-d example in Figure 1. Left: The function values for a single design with zero-mean Gaussian input perturbations with a standard deviation of 0.1 are marked by black points. The background is a contour showing the value of a Chebyshev scalarization across the objective space. Center: The probability density of a Chebyshev scalarization over the function values under input noise and the value-at-risk  $v$  of a Chebyshev scalarization for  $\alpha = 0.9$ . The probability mass to the right of the black line is equal to  $\alpha$ . Right: Leveraging Theorem 5.1, VAR,  $w$  can be mapped to point in the MVAR set that is dominated by the objectives under input perturbations 90% of the time. The green triangles represent a discrete approximation of the MVAR set with 64 samples from the input noise distribution. The green area indicates the region that dominates the identified MVAR point, or, equivalently, the area for which the Chebyshev scalarization defined by  $w$  is greater than  $v$ .

input perturbations will dominate  $z$  with probability  $\geq \alpha$ .

**Lemma 5.1** (VAR of Chebyshev scalarization  $\Rightarrow$  Pareto Dominance). *Let  $v = \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \xi), \mathbf{w}])$ . Then,  $P[\mathbf{f}(\mathbf{x} \diamond \xi) \geq \frac{v}{w}] \geq \alpha$ , where  $\frac{v}{w}$  denotes element-wise division.*

The condition in Lemma 5.1 is one criterion for membership in the MVAR set (the other being Pareto efficiency). The shaded region in right plot in Figure 2 illustrates the region that dominates  $z$ . With probability  $\geq \alpha$ ,  $\mathbf{f}(\mathbf{x} \diamond \xi)$  will fall within the shaded region. *This result enables translating the VAR of a Chebyshev scalarization into an interpretable, high-probability guarantee on robust performance in terms of Pareto dominance.*

**Assumption 5.1.**  $\mathbf{f}(\mathbf{x} \diamond \xi)$  has a continuous, strictly increasing CDF  $F$ . I.e., if  $\mathbf{f}(\mathbf{x}) \succ \mathbf{f}(\mathbf{x}')$ , then  $F[\mathbf{f}(\mathbf{x})] > F[\mathbf{f}(\mathbf{x}')]^5$ .

If Assumption 5.1 is met, then for any  $w$ ,  $\frac{v}{w}$  is not dominated by any other point in the MVAR set, and hence,  $\frac{v}{w}$  is an element of the MVAR set. Furthermore, we have the following:

**Theorem 5.1** (MVAR  $\iff$  VAR of Chebyshev scalarization). *Given  $\mathbf{x}$ ,  $\mathbf{f}$ , and  $P(\xi)$ , let  $h : \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \xi)] \rightarrow \Delta_+^{M-1}$  be the function  $h(\mathbf{z}) = \mathbf{w} = \frac{1}{\mathbf{z} \parallel \frac{1}{\mathbf{z}} \parallel}$ . Under*

<sup>5</sup>Assumption 5.1 holds when  $\mathbf{f}$  is a function sampled from a GP prior with many commonly used covariance functions. See Lemma A.5 for a formal statement and Appendix A.3 for proof.

*Assumption 5.1,  $h(\cdot)$  is bijective and  $h^{-1}(\mathbf{w}) = \mathbf{z} = \frac{1}{w} \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \xi), \mathbf{w}])$ .*

Theorem 5.1 provides a technique for generating points in the MVAR set using the VAR of Chebyshev scalarizations with different weights. Importantly, any design that is globally optimal with respect to MVAR is a maximizer of the VAR of a Chebyshev scalarization. This naturally motivates a methodology for identifying the global MVAR set by optimizing the VAR of random Chebyshev scalarizations.

**Corollary 5.1** (Consistent Optimizers). *Suppose  $\mathbf{z}$  is a point in the global MVAR set, i.e.,  $\mathbf{z} \in \text{MVAR}_\alpha[\{\mathbf{f}(\mathbf{x} \diamond \xi)\}_{\mathbf{x} \in \mathcal{X}}]$ . Let  $\mathcal{X}_z^*$  be the set of designs such that for all  $\mathbf{x} \in \mathcal{X}_z^*$ ,  $\mathbf{z} \in \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \xi)]$ . Then every  $\mathbf{x} \in \mathcal{X}_z^*$  is a maximizer of  $\text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \xi), \mathbf{w}])$  for  $\mathbf{w} = (\mathbf{z} \parallel \frac{1}{\mathbf{z}} \parallel)^{-1}$ .*

## 5.2. MARS: MVAR Approximation via Random Scalarizations

The connection between the VAR of a Chebyshev scalarization and MVAR can be exploited to optimize the global MVAR by randomly sampling a Chebyshev scalarization at each BO iteration<sup>6</sup> and optimizing the VAR of the Chebyshev scalarization using a single objective BO algorithm, such as Noisy Expected Improvement (NEI) (Letham et al., 2019)—which is required (rather than expected improvement) since we must integrate over the unknown best

<sup>6</sup>We sample weights uniformly from  $\Delta_+^{M-1}$ , which we find works well empirically.

unknown incumbent value—or Thompson sampling (TS) (Thompson, 1933; Paria et al., 2018). We refer to this technique as MVAR Approximation via Random Scalarizations (MARS). MARS is a simple, theoretically-grounded technique, and we find, in Section 6, that MARS performs well empirically. In the main text, we focus on MARS with NEI (denoted as MARS-NEI), but we derive and empirically evaluate upper confidence bound (UCB) (Srinivas et al., 2010) and TS variants in Appendices B and G.

For MARS-NEI, we use the MC formulation of NEI (Balandat et al., 2020) so that the Chebyshev scalarizations can be computed in the NEI as composite objectives (Astudillo and Frazier, 2019). We optimize the the acquisition function using sample-path gradients that leverage well-studied gradient estimators of VAR (see Appendix C.1 for details). See Appendix G.1 for details on optimization.

## 6. Experiments

In this section, we provide an empirical demonstration of robust MOBO on synthetic and real-world problems. We compare three broad classes of methods: (i) Non-robust methods including NSGA-II (Deb et al., 2002),  $q$ NPAREGO (Daulton et al., 2020),  $q$ NEHVI (Daulton et al., 2021b), (ii) methods for optimizing expectation risk measures, e.g. using  $q$ NPAREGO (denoted as EXP- $q$ NPAREGO), and (iii) methods for optimizing MVAR via MARS. For readability, we only include one expectation and one MVAR optimization method in the main text, both based on NEI. In Appendix G, we evaluate additional methods including MARS with TS and UCB and methods for direct MVAR optimization based on NEHVI. We consider the expectation risk measure because it is simple and performs well in many scenarios. All robust (expectation and MVAR) methods are our novel contributions. Additionally, we compare against a quasi-random policy, which selects the designs to evaluate according to a scrambled Sobol sequence (Owen, 2003).

For all BO-methods, we begin by evaluating  $2(d+1)$  design points from a scrambled Sobol sequence. We use  $n_{\xi} = 32$  samples because we find that setting  $n_{\xi} > 32$  yields little-to-no improvement in optimization performance (see Figure 12). See Appendix G for details on all methods. Our code is open-sourced at [github.com/facebookresearch/robust\\_mobo](https://github.com/facebookresearch/robust_mobo).

### 6.1. Synthetic Problems

**Gaussian Mixture Model (GMM)** ( $d = 2$ ,  $M = 2$ ,  $\alpha = 0.9$ ): This is a variant of the GMM problem from Fröhlich et al. (2020) where each objective is an independent GMM. We use a multiplicative noise model, i.e.,  $\mathbf{x} \diamond \xi := \mathbf{x}\xi$ , where  $\xi \sim \mathcal{N}(\mu = \mathbf{1}, \Sigma = 0.07I_2)$  with  $I_n$  denoting the  $n$ -dimensional identity matrix. In Appendix I.5, we present multiple variations of this problem to demonstrate the consistency of our methods under different noise models.

**Constrained Branin Currin** ( $d = 2$ ,  $M = 2$ ,  $V = 1$ ,  $\alpha = 0.7$ ): We subject this problem, which originates from Daulton et al. (2020), to a heteroskedastic input noise process given by  $P(\xi; \mathbf{x}) = \mathcal{N}(\mu = \mathbf{0}, \Sigma = 0.05(1 + \sigma(1 - 2x_0))I_2)$ , where  $\sigma(x) = \frac{1}{1+e^{-x}}$ . The optimal designs with respect to the nominal objectives are on the boundary of the feasible region with respect to the outcome constraint. Hence, the optimal designs with respect to the nominal metrics often violate the constraint under input perturbations.

### 6.2. Real-World Problems

**Disc Brake** ( $d = 4$ ,  $M = 2$ ,  $V = 4$ ,  $\alpha = 0.95$ ): In this disc brake manufacturing problem, the goal is to minimize the brake’s mass and the stopping time of a vehicle by tuning the inner and outer radii of the disc, the engaging force, and the number of friction surfaces (Ray and Liew, 2002). Following Emch and Parkinson (1994), we use zero-mean uniform input noise with a maximum absolute perturbation value of 5% of the range of each parameter (except for the number of friction surfaces, which is noise-free).

**Penicillin Production** ( $d = 7$ ,  $M = 3$ ,  $\alpha = 0.8$ ): This problem considers optimizing the manufacturing process of penicillin (Liang and Lai, 2021). The goal is to maximize the yield while minimizing time-to-ferment and CO<sub>2</sub> output by tuning 7 initial conditions of the chemical reaction. Each parameter is subject to independent zero-mean Gaussian noise, where the standard deviation ranges from 0.5% to 3% of the parameter’s domain (see Appendix G for details).

### 6.3. Results

In Figure 9, we evaluate all methods in terms of log HV regret, which is the difference in HV between the true global MVAR set and the MVAR of the set of designs evaluated by each method (see Appendix G.3 for details on estimating true global MVAR). Figure 3 shows that MARS is consistently the best performing method. The non-robust methods consistently perform poorly on all problems. On the GMM problem, EXP- $q$ NPAREGO performs worse than MARS-NEI because the optimal designs under expectation risk measure are in a disjoint part of the search space from the MVAR-optimal designs. On the Constrained Branin Currin problem, the nominal methods perform no better than quasi-random Sobol search. On the penicillin problem, MARS vastly outperforms all other methods.

Although the log HV regret highlights the performance of MARS, it does not fully capture the *necessity* of using a robust method in practice. In Figure 4, we analyze the yield (i.e. the probability of the objectives exceeding a performance specification under the input noise distribution) of a robust and a non-robust design on the Disc Brake problem. Using a target performance specification chosen from the

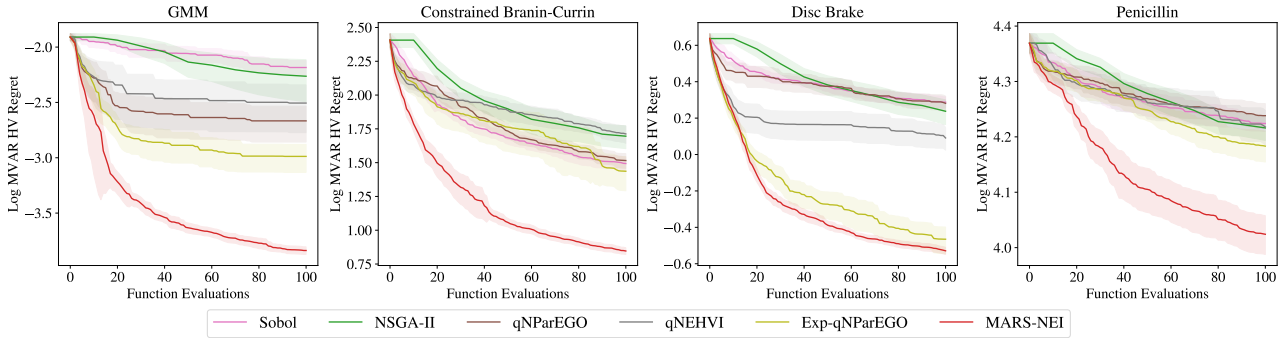


Figure 3: The log MVAR HV regret after the initial space-filling design. For each method, we plot the mean and 2 standard errors of the mean over 20 trials.

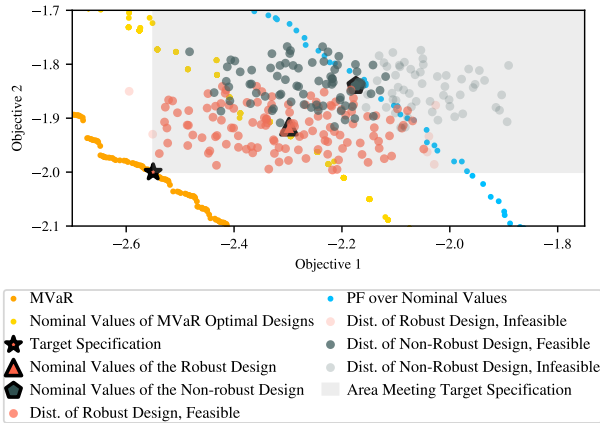


Figure 4: The yield from selecting a robust versus a non-robust design on the Disc Brake problem. Although the non-robust design is feasible under the nominal objectives, it is located near the boundary of the feasible region in design space and violates some of the black-box constraints (not shown) under a large fraction of input perturbations.

MVAR set, we see that if a decision-maker were to select a non-robust design (green pentagon) that is optimal with respect to the nominal objectives and nominal values that meet the target specification, the yield would only be 58.2%. In contrast, an MVAR-optimal solution can be chosen such that it meets target specification with high probability. For example, the robust design marked by the orange triangle enjoys a yield of 95.3%. As shown in Figure 4, this is because the non-robust design often does not satisfy all of the black-box constraints under input perturbations. In this problem, the objectives are relatively robust to noise (much more so than the toy example in Figure 1), but the feasibility of a design (and therefore the yield) is highly sensitive to input noise when the design is near the boundary of the feasible region in design space.

Table 3 reports the wall times for running a single iteration of BO with each algorithm. Not only is MARS-NEI computationally tractable on all problems, but it achieves wall times that are competitive to alternative algorithms.

Evaluation in additional problem settings and comparisons against additional methods in Appendix I further validate that MARS-NEI is consistently a top performer and yields competitive wall times. We find that MARS-TS performs slightly worse, on average, in terms of log HV regret than MARS-NEI, but that it has shorter wall times. Methods for direct MVAR optimization with  $q$ NEHVI perform comparably to MARS-NEI, but the direct MVAR optimization is prohibitively expensive in terms of wall time and memory requirements and was infeasible to run on many problems (including nearly all problems with  $> 2$  objectives). In contrast, Figure 7 shows that MARS-NEI can scale to problems with  $> 2$  objectives and consistently performs best in those settings. Additionally, in Appendix G, we show that MARS-NEI works well under a wide variety of input noise processes, scales well with increasing batch sizes (in the parallel evaluation setting), and is not sensitive to the number of MC samples  $n_\xi$  used to estimate  $P(\xi)$ , for  $n_\xi \geq 32$ .

## 7. Discussion

In this work, we formulate the goal of MO robust optimization under input noise as optimizing the *global* MVAR set—a novel concept that is a robust analogue of the Pareto frontier in the standard MO setting. We derive a correspondence between MVAR and Chebyshev scalarizations based on VAR. This theoretical result naturally motivates a computationally efficient approach (MARS) for using BO to optimize the global MVAR set with high sample-efficiency. Empirically, we find that MARS consistently outperforms alternative approaches and achieves competitive wall times.

Although our focus has been on the small evaluation budget regime and BO, our theoretical results are far more general. The connection between MVAR and Chebyshev scalarizations could be leveraged by gradient-based and evolutionary methods to scale global MVAR optimization to settings with large evaluation budgets. We hope that our contributions serve as a foundation for future advances in methods for robust MO optimization under input noise.



## Acknowledgements

We thank Ben Letham, David Eriksson, James Wilson, Martin Jørgensen, and Raul Astudillo, as well as the members of the Oxford Machine Learning Research Group, for providing insightful feedback. In addition, Sait Cakmak and Enlu Zhou are grateful for support by the Air Force Office of Scientific Research under Grant FA9550-19-1-0283.

## References

- M.F. Ashby. Multi-objective optimization in material design and selection. *Acta Materialia*, 48(1):359–369, 2000. ISSN 1359-6454.
- Raul Astudillo and Peter Frazier. Bayesian optimization of composite functions. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 354–363. PMLR, 09–15 Jun 2019.
- Gideon Avigad and Jürgen Branke. Embedded evolutionary multi-objective optimization for worst case robustness. In *Proceedings of the 10th Annual Conference on Genetic and Evolutionary Computation, GECCO '08*, page 617–624, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781605581309. doi: 10.1145/1389095.1389221.
- 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.
- Basel Committee on Banking Supervision. Fundamental review of the trading book. *Consultative Document*, 2012.
- Justin J. Beland and Prasanth B. Nair. Bayesian optimization under uncertainty. 2017.
- Hans-Georg Beyer and Bernhard Sendhoff. Robust optimization – a comprehensive survey. *Computer Methods in Applied Mechanics and Engineering*, 196(33):3190–3218, 2007. ISSN 0045-7825.
- J. Blank and K. Deb. pymoo: Multi-objective optimization in python. *IEEE Access*, 8:89497–89509, 2020.
- Ilija Bogunovic, Jonathan Scarlett, Stefanie Jegelka, and Volkan Cevher. Adversarially robust optimization with gaussian processes. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Edwin V Bonilla, Kian Chai, and Christopher Williams. Multi-task gaussian process prediction. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2008. URL <https://proceedings.neurips.cc/paper/2007/file/66368270ffd51418ec58bd793f2d9b1b-Paper.pdf>.
- Richard H. Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput.*, 16:1190–1208, 1995.
- Sait Cakmak, Raul Astudillo, Peter Frazier, and Enlu Zhou. Bayesian optimization of risk measures. In *Advances in Neural Information Processing Systems 33*, 2020.
- Roberto Calandra and Jan Peters. Pareto front modeling for sensitivity analysis in multi-objective bayesian optimization. 2014.
- Daniele Calandriello, Luigi Carratino, Alessandro Lazaric, Michal Valko, and Lorenzo Rosasco. Gaussian process optimization with adaptive sketching: Scalable and no regret, 2019.
- Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, page 844–853. JMLR.org, 2017.
- Areski Cousin and Elena Di Bernardino. On multivariate extensions of conditional-tail-expectation. *Insurance: Mathematics and Economics*, 55:272–282, 2014.
- Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Differentiable Expected Hypervolume Improvement for Parallel Multi-Objective Bayesian Optimization. In *Advances in Neural Information Processing Systems 33*, 2020.
- Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Parallel Bayesian Optimization of Multiple Noisy Objectives with Expected Hypervolume Improvement. In *Advances in Neural Information Processing Systems 34*, 2021a.
- Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Multi-objective bayesian optimization over high-dimensional search spaces, 2021b.
- Nando de Freitas, Alex Smola, and Masrour Zoghi. Exponential regret bounds for gaussian process bandits with deterministic observations. In *ICML*, 2012.
- 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.
- Kalyanmoy Deb and Himanshu Gupta. Searching for robust pareto-optimal solutions in multi-objective optimization. In Carlos A. Coello Coello, Arturo Hernández Aguirre, and Eckart Zitzler, editors, *Evolutionary Multi-Criterion Optimization*, pages 150–164, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- Bach Do, Makoto Ohsaki, and Makoto Yamakawa. Bayesian optimization for robust design of steel frames with joint and individual probabilistic constraints. *Engineering Structures*, 245:112859, 2021. ISSN 0141-0296.
- Ioannis Doltsinis and Zhan Kang. Robust design of structures using optimization methods. *Computer Methods in Applied Mechanics and Engineering*, 193(23):2221–2237, 2004. ISSN 0045-7825. doi: <https://doi.org/10.1016/j.cma.2003.12.055>. URL <https://www.sciencedirect.com/science/article/pii/S0045782504000787>.
- G. Emch and A. Parkinson. Robust Optimal Design for Worst-Case Tolerances. *Journal of Mechanical Design*, 116(4):1019–1025, 12 1994. ISSN 1050-0472. doi: 10.1115/1.2919482.
- 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.
- David Eriksson, Pierce I-Jen Chuang, Samuel Daulton, Peng Xia, Akshat Shrivastava, Arun Babu, Shicong Zhao, Ahmed Aly, Ganesh Venkatesh, and Maximilian Balandat. Latency-aware neural architecture search with multi-objective bayesian optimization, 2021.
- Peter I Frazier. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.
- Lukas Fröhlich, Edgar Klenske, Julia Vinogradska, Christian Daniel, and Melanie Zeilinger. Noisy-input entropy search for efficient robust bayesian optimization. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2262–2272. PMLR, 26–28 Aug 2020.
- Himanshu Gupta and Kalyanmoy Deb. Handling constraints in robust multi-objective optimization. In *2005 IEEE Congress on Evolutionary Computation*, volume 1, pages 25–32 Vol.1, 2005. doi: 10.1109/CEC.2005.1554663.
- Nikolaus Hansen. *The CMA Evolution Strategy: A Comparing Review*, volume 192, pages 75–102. 06 2007. doi: 10.1007/3-540-32494-1\_4.
- Zhenan He, Gary G. Yen, and Zhang Yi. Robust multiobjective optimization via evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 23(2):316–330, 2019. doi: 10.1109/TEVC.2018.2859638.
- L. Jeff Hong. Estimating quantile sensitivities. *Operations Research*, 57(1):118–130, 2009. doi: 10.1287/opre.1080.0531.
- L. Jeff Hong, Zhaolin Hu, and Guangwu Liu. Monte carlo methods for value-at-risk and conditional value-at-risk: a review. *ACM Transactions on Modeling and Computer Simulation*, 24(4):1–37, 2014.
- Jonas Ide and Elisabeth Köbis. Concepts of efficiency for uncertain multi-objective optimization problems based on set order relations. *Mathematical Methods of Operations Research*, 80, 08 2014. doi: 10.1007/s00186-014-0471-z.
- 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>.
- 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.
- Shogo Iwazaki, Yu Inatsu, and Ichiro Takeuchi. Mean-variance analysis in bayesian optimization under uncertainty. In Arindam Banerjee and Kenji Fukumizu, editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 973–981. PMLR, 13–15 Apr 2021.
- Sergio Izquierdo, Julia Guerrero-Viu, Sven Hauns, Guilherme Miotto, Simon Schrod, André Biedenkapp, Thomas Elsken, Difan Deng, Marius Lindauer, and Frank Hutter. Bag of baselines for multi-objective joint neural architecture search and hyperparameter optimization. In *8th ICML Workshop on Automated Machine Learning (AutoML)*, 2021.
- Miettinen Kaisa. *Nonlinear Multiobjective Optimization*, volume 12 of *International Series in Operations Research & Management Science*. Kluwer Academic Publishers, Boston, USA, 1999.

- Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. *arXiv e-prints*, page arXiv:1312.6114, Dec 2013.
- Johannes Kirschner, Ilija Bogunovic, Stefanie Jegelka, and Andreas Krause. Distributionally robust bayesian optimization. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2174–2184. PMLR, 26–28 Aug 2020.
- 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.
- Renaud Lacour, Kathrin Klamroth, and Carlos M. Fonseca. A box decomposition algorithm to compute the hypervolume indicator. *Computers & Operations Research*, 79: 347 – 360, 2017.
- Benjamin Letham, Brian Karrer, Guilherme Ottoni, and Eytan Bakshy. Constrained bayesian optimization with noisy experiments. *Bayesian Analysis*, 14(2), 06 2019.
- Zhuangzhi Li and Zukui Li. Optimal robust optimization approximation for chance constrained optimization problem. *Computers & Chemical Engineering*, 74:89–99, 2015.
- Qiaohao Liang and Lipeng Lai. Scalable bayesian optimization accelerates process optimization of penicillin production. In *NeurIPS 2021 AI for Science Workshop*, 2021.
- Xingtiao 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.
- Mina Konakovic Lukovic, Yunsheng Tian, and Wojciech Matusik. Diversity-Guided Multi-Objective Bayesian Optimization With Batch Evaluations. In *Advances in Neural Information Processing Systems 33*, 2020.
- Dinah Elena Majewski, Marco Wirtz, Matthias Lampe, and André Bardow. Robust multi-objective optimization for sustainable design of distributed energy supply systems. *Computers & Chemical Engineering*, 102:26–39, 2017. ISSN 0098-1354. Sustainability & Energy Systems.
- Gustavo Malkomes, Bolong Cheng, Eric H Lee, and Mike Mccourt. Beyond the pareto efficient frontier: Constraint active search for multiobjective experimental design. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 7423–7434. PMLR, 18–24 Jul 2021.
- Julien Marzat, Eric Walter, and H el ene Piet-Lahanier. Worst-case global optimization of black-box functions through kriging and relaxation. *Journal of Global Optimization*, 55:707–727, 04 2013. doi: 10.1007/s10898-012-9899-y.
- Merve Meraklı and Simge K uc ukyavuz. Vector-valued multivariate conditional value-at-risk. *Operations Research Letters*, 46(3):300–305, 2018. ISSN 0167-6377. doi: <https://doi.org/10.1016/j.orl.2018.02.006>. URL <https://www.sciencedirect.com/science/article/pii/S0167637717304145>.
- S erine Th er ese F.C. Mortier, Pieter-Jan Van Bockstal, Jos Corver, Ingmar Nopens, Krist V. Gernaey, and Thomas De Beer. Uncertainty analysis as essential step in the establishment of the dynamic design space of primary drying during freeze-drying. *European Journal of Pharmaceutics and Biopharmaceutics*, 103:71–83, 2016. ISSN 0939-6411.
- Mojmir Mutny and Andreas Krause. Efficient high dimensional bayesian optimization with additivity and quadrature fourier features. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Quoc Phong Nguyen, Zhongxiang Dai, Bryan Kian Hsiang Low, and Patrick Jaillet. Optimizing conditional value-at-risk of black-box functions. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021a.
- Quoc Phong Nguyen, Zhongxiang Dai, Bryan Kian Hsiang Low, and Patrick Jaillet. Value-at-risk optimization with gaussian processes, 2021b.
- Rafael Oliveira, Lionel Ott, and Fabio Ramos. Bayesian optimisation under uncertain inputs. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 1177–1184. PMLR, 16–18 Apr 2019.
- Michael A. Osborne. Bayesian gaussian processes for sequential prediction, optimisation and quadrature. 2010.
- Art B Owen. Quasi-monte carlo sampling. *Monte Carlo Ray Tracing: Siggraph*, 1:69–88, 2003.
- C.J. Paciorek. *Nonstationary Gaussian Processes for Regression and Spatial Modelling*. PhD thesis, Carnegie Mellon University, Pittsburgh, Pennsylvania, 2003.
- Biswajit Paria, Kirthevasan Kandasamy, and Barnab as P ocz os. A flexible framework for multi-objective bayesian optimization using random scalarizations, 2018.

- András Prékopa. Multivariate value at risk and related topics. *Annals of Operations Research*, 193:49–69, 2012.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2008.
- Carl Edward Rasmussen. *Gaussian Processes in Machine Learning*, pages 63–71. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004.
- Tapabrata Ray and K.M. Liew. A swarm metaphor for multi-objective design optimization. *Engineering Optimization*, 34(2):141–153, 2002. doi: 10.1080/03052150210915.
- Justo José Roberts, Agnelo Marotta Cassula, José Luz Silveira, Edson da Costa Bortoni, and Andrés Z. Mendiburu. Robust multi-objective optimization of a renewable based hybrid power system. *Applied Energy*, 223:52–68, 2018. ISSN 0306-2619.
- R. Tyrrell Rockafellar and Stanislav Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7), 2002. ISSN 0378-4266. doi: 10.1016/S0378-4266(02)00271-6.
- Sergey Sarykalin, Gaia Serraino, and Stan Uryasev. Value-at-risk vs conditional value-at-risk in risk management and optimization. 09 2008. ISBN 978-1-877640-23-0. doi: 10.1287/educ.1080.0052.
- Robert J. Serfling. *Approximation Theorems of Mathematical Statistics*. Wiley Series in Probability and Statistics. Wiley-Blackwell, 2008. ISBN 9780470316481. doi: 10.1002/9780470316481.
- Pier Giuseppe Sessa, Ilija Bogunovic, Maryam Kamgarpour, and Andreas Krause. Mixed strategies for robust optimization of unknown objectives, 2020.
- Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104:1–28, 2016.
- 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.
- Shinya Suzuki, Shion Takeno, Tomoyuki Tamura, Kazuki Shitara, and Masayuki Karasuyama. Multi-objective bayesian optimization using pareto-frontier entropy, 2020.
- G. Taguchi. *Einführung in Quality engineering: Minimierung von Verlusten durch Prozeßbeherrschung*. American Supplier Inst., 1989. ISBN 9789283310846.
- 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.
- 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.
- Léonard Torossian, Victor Picheny, and Nicolas Durrande. Bayesian quantile and expectile optimisation. *arXiv: 2001.04833*, 2020.
- Saul Toscano-Palmerin and Peter I. Frazier. Bayesian optimization with expensive integrands. *arXiv: 1803.08661*, 2018.
- Sameer Rehman, Matthijs Langelaar, and Fred van Keulen. Efficient kriging-based robust optimization of unconstrained problems. *Journal of Computational Science*, 5(6):872–881, 2014. ISSN 1877-7503.
- Zi Wang, Clement Gehring, Pushmeet Kohli, and Stefanie Jegelka. Batched large-scale bayesian optimization in high-dimensional spaces, 2018.
- 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.
- James T. Wilson, Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Peter Deisenroth. Efficiently sampling functions from gaussian process posteriors. In *International Conference on Machine Learning*, 2020.
- Xiangzhong Xie and René Schenkendorf. Robust process design in pharmaceutical manufacturing under batch-to-batch variation. *Processes*, 7(8), 2019. ISSN 2227-9717. doi: 10.3390/pr7080509.
- 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.
- Qi Zhou, Ping Jiang, Xiang Huang, Feng Zhang, and Tao Zhou. A multi-objective robust optimization approach based on gaussian process model. *Structural and Multidisciplinary Optimization*, 57, 01 2018. doi: 10.1007/s00158-017-1746-9.



Ciyou Zhu, Richard H. Byrd, Peihuang Lu, and Jorge Nocedal. Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Trans. Math. Softw.*, 23(4):550–560, 1997. ISSN 0098-3500. doi: 10.1145/279232.279236. URL <https://doi.org/10.1145/279232.279236>.

## A. Theory and Proofs

### A.1. Preliminaries

**Definition A.1.** Let  $\mathcal{F} = \{\mathbf{f}(\mathbf{x}) : \mathbf{x} \in X \subseteq \mathcal{X}\}$ . The weakly efficient Pareto frontier is  $\text{WEAKPARETO}(\mathcal{F}) = \{\mathbf{f}(\mathbf{x}) \in \mathcal{F} : \nexists \mathbf{x}' \in X \text{ s.t. } \mathbf{f}(\mathbf{x}') > \mathbf{f}(\mathbf{x})\}$ , and  $\text{PARETO}(\mathcal{F}) \subseteq \text{WEAKPARETO}(\mathcal{F})$ . If there are constraints  $\mathbf{c}(\mathbf{x})$ , elements of  $\text{WEAKPARETO}$  are subject to the additional membership Assumption that  $\mathbf{c}(\mathbf{x}) \geq \mathbf{0}$ . We call the corresponding set of optimal designs the *weak Pareto set*.

**Definition A.2.**  $\text{WEAK-MVAR}$  is defined in the same way as  $\text{MVAR}$ , but only requires that its elements are weakly Pareto optimal.

**Definition A.3** (Prékopa (2012)). If Assumption 5.1 holds, then we can express  $\text{MVAR}$  with an equality with respect to  $\alpha$ :

$$\text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})] = \{\mathbf{z} \in \mathbb{R}^M : P[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq \mathbf{z}] = \alpha\}. \quad (1)$$

### A.2. Proofs

**Lemma A.1.** Let  $\mathbf{y} \in \mathbb{R}^M$  and  $v \in \mathbb{R}$ . Then  $s[\mathbf{y}, \mathbf{w}] \geq v \iff \mathbf{y} \geq \frac{v}{\mathbf{w}}$ .

*Proof.* This follows directly from Definition 5.1.

$$\begin{aligned} s[\mathbf{y}, \mathbf{w}] \geq v &\iff \min_i w_i y_i \geq v \\ &\iff w_i y_i \geq v \quad \forall i \\ &\iff y_i \geq \frac{v}{w_i} \quad \forall i \\ &\iff \mathbf{y} \geq \frac{v}{\mathbf{w}}. \end{aligned}$$

□

Lemma A.1 states that a lower bound  $v$  on the value of a Chebyshev scalarization of an objective vector  $\mathbf{y}$  can be used to define a point,  $\frac{v}{\mathbf{w}}$ , that  $\mathbf{y}$  is greater than or equal to. We can extend this to make a similar statement about the VAR of a Chebyshev scalarization.

**Lemma 5.1** (VAR of Chebyshev scalarization  $\Rightarrow$  Pareto Dominance). Let  $v = \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}])$ . Then,  $P[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq \frac{v}{\mathbf{w}}] \geq \alpha$ , where  $\frac{v}{\mathbf{w}}$  denotes element-wise division.

*Proof.* From Definition 4.1, we have that

$$\text{VAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})] = \sup\{z \in \mathbb{R} : P[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq z] \geq \alpha\}.$$

Hence,  $P(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \geq v) \geq \alpha$ . From Definition 5.1, we have  $P(\min_i [w_i f_i(\mathbf{x} \diamond \boldsymbol{\xi})] \geq v) \geq \alpha$ . By Lemma A.1, the statement  $\min_i [w_i f_i(\mathbf{x} \diamond \boldsymbol{\xi})] \geq v$  is equivalent to  $\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq \frac{v}{\mathbf{w}}$ . Hence,  $P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq \frac{v}{\mathbf{w}}) \geq \alpha$ . □

**Lemma A.2** (VAR of Chebyshev scalarization  $\Rightarrow$  WEAK-MVAR). Let  $v = \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}])$ . Then,  $\frac{v}{\mathbf{w}} \in \text{WEAK-MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$ .

*Proof.* Let  $\mathbf{z} = \frac{v}{\mathbf{w}}$ . Suppose there exists  $\mathbf{z}' \in \text{WEAK-MVAR}_\alpha(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}))$  such that  $\mathbf{z}' > \mathbf{z}$ . Since  $\mathbf{z}' \in \text{WEAK-MVAR}_\alpha(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}))$ , we have that  $P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq \mathbf{z}') \geq \alpha$ . Note that  $\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq \mathbf{z}'$  implies that  $\min_i w_i f_i(\mathbf{x} \diamond \boldsymbol{\xi}) \geq \min_i w_i z'_i$ . Hence,  $P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq \mathbf{z}') \geq \alpha$  implies that  $P(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \geq s[\mathbf{z}', \mathbf{w}]) \geq \alpha$ . Since  $\mathbf{z}' > \mathbf{z}$  and  $\mathbf{w} \in \Delta_+^{M-1}$ , we have that  $s[\mathbf{z}', \mathbf{w}] > s[\mathbf{z}, \mathbf{w}] = v$ . But this contradicts Definition 4.1. Since there does not exist  $\mathbf{z}' \in \text{WEAK-MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$  such that  $\mathbf{z}' > \mathbf{z}$ , we have that  $\mathbf{z} \in \text{WEAK-MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$ . □

If Assumption 5.1 holds, the inequalities with respect to  $\alpha$  in Lemma 5.1 become equalities, and we show that  $\frac{v}{\mathbf{w}}$  is strictly Pareto optimal.

**Lemma A.3** (VAR of Chebyshev scalarization  $\Rightarrow$  MVAR). *Let  $v = \text{VAR}_\alpha[s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}]]$ . If Assumption 5.1 holds, then  $\frac{v}{\mathbf{w}} \in \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$ .*

*Proof.* Let  $z := \frac{v}{\mathbf{w}}$ . Suppose there exists  $z' \in \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$  such that  $z' \succ z$ . Because  $F(\cdot)$  is a strictly increasing CDF,  $P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq z') = F(z') > F(z) = P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq z)$ . From Lemma 5.1, we have that  $P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq z) \geq \alpha$ . Because  $z' \in \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$  from Equation (1), we have that  $P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq z') = \alpha$ . Hence,  $P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq z) \geq P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq z')$ . This is a contradiction.  $\square$

Lemma A.3 provides a technique for generating points in the MVAR set using the VAR of Chebyshev scalarizations with different  $\mathbf{w}$ .

Now, consider the reverse mapping.

**Lemma A.4** (MVAR  $\Rightarrow$  VAR of Chebyshev scalarization). *Suppose that  $z \in \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$ . Then  $z = \frac{v}{\mathbf{w}}$  for  $\mathbf{w} := \frac{1}{z} \|\frac{1}{z}\|_1^{-1} \in \Delta_+^M$  and  $v = \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}])$ .*

*Proof.* Without loss of generality, assume that  $z > \mathbf{0}$ .<sup>7</sup> Let  $v = \|\frac{1}{z}\|_1^{-1} \in \mathbb{R}_+$ . Then,  $z = \frac{v}{\mathbf{w}}$ . Hence, all we need to show is that  $v$  equals  $\text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}])$ . Let us define  $v' := \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}])$ . By definition,

$$v' = \sup\{v'' \in \mathbb{R} : P(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \geq v'') \geq \alpha\}.$$

Since  $z \in \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$ ,  $P[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq z] \geq \alpha$ . Hence,  $P[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq \frac{v}{\mathbf{w}}] \geq \alpha$ . Using Lemma A.1, we have that  $P(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \geq v) \geq \alpha$ . Since  $v'$  is the supremum, we have that  $v \leq v'$ . Suppose now that  $v < v'$ . Note that  $\frac{v}{\mathbf{w}} = z \in \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$ . From Lemma A.2,  $\frac{v'}{\mathbf{w}} \in \text{WEAK-MVAR}_\alpha(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}))$ . Since  $v < v'$ ,  $z = \frac{v}{\mathbf{w}} \prec \frac{v'}{\mathbf{w}}$ . Hence,  $z$  cannot be in  $\text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$  because it is dominated by  $\frac{v'}{\mathbf{w}}$  and  $\frac{v'}{\mathbf{w}} \in \text{WEAK-MVAR}_\alpha(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}))$ . This is a contradiction. Hence  $v \geq v'$  and therefore it follows that  $v = v' = \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}])$ .  $\square$

Let  $h : \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})] \rightarrow \Delta_+^{M-1}$  be given by  $h(z) = \frac{1}{z \|\frac{1}{z}\|}$ . Being the element-wise application of a scalar injective mapping ( $z \mapsto 1/z$ ), it is clear that  $h$  is injective. However,  $h(\cdot)$  is not necessarily bijective, since without Assumption 5.1 there may be weight vectors  $\mathbf{w}$  such that  $\nexists z \in \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$  s.t.  $h(z) = \mathbf{w}$ .

**Theorem 5.1** (MVAR  $\iff$  VAR of Chebyshev scalarization). *Given  $\mathbf{x}$ ,  $\mathbf{f}$ , and  $P(\boldsymbol{\xi})$ , let  $h : \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})] \rightarrow \Delta_+^{M-1}$  be the function  $h(z) = \tilde{\mathbf{w}} = \frac{1}{z \|\frac{1}{z}\|}$ . Under Assumption 5.1,  $h(\cdot)$  is bijective and  $h^{-1}(\mathbf{w}) = z = \frac{1}{\mathbf{w}} \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}])$ .*

*Proof.* This follows directly from Lemma A.3 and Lemma A.4.  $\square$

**Corollary A.1** (MVAR via Scalarization). *WEAK-MVAR enjoys the following scalarization representation:  $\text{WEAK-MVAR}_\alpha(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})) = \{\frac{1}{\mathbf{w}} \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}]) : \mathbf{w} \sim \Delta_+^{M-1}\}$ . If Assumption 5.1 holds, then  $\text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})] = \{\frac{1}{\mathbf{w}} \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}]) : \mathbf{w} \sim \Delta_+^{M-1}\}$ .*

Although the MVAR representation in Corollary A.1 depends on Assumption 5.1, Lemma 5.1 recovers all weakly Pareto optimal points even if this assumption is not met because  $\text{MVAR} \subseteq \text{WEAK-MVAR}$ . Hence, with or without Assumption 5.1, Theorem A.1 can be used to approximate the MVAR set.

**Result A.1** (MVAR Approximation). *MVAR can be approximated with a finite set of weight vectors:  $\widehat{\text{MVAR}}(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})) = \{\frac{1}{\mathbf{w}_i} \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}_i])\}_{i=1}^{N_{\text{MVAR}}}$ .*

### A.3. Discussion of the Assumption of Continuous, Strictly-Increasing CDFs with Gaussian Processes

In Bayesian Optimization with Gaussian Process surrogates, it is assumed that the objective function  $\mathbf{f}$  is sample path from a Gaussian process prior. In this setting, the random variable  $\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})$  for a deterministic sample path  $\mathbf{f}$  (where the

<sup>7</sup>This can easily be guaranteed by shifting the objectives to be strictly positive. If the objectives are not unbounded from below, the reference point, which is commonly supplied by the decision-maker in multi-objective optimization effectively provides a lower bound on the objectives.

stochasticity comes solely from  $\xi$ ) has a continuous, strictly increasing CDF  $F(\cdot)$  for many commonly used covariance functions. Hence, the bijective relationship in Theorem 5.1 holds given a suitable choice of covariance function.<sup>8</sup>

**Lemma A.5.** *If  $f$  is a sample path from a multi-output Gaussian Process prior where the covariance function and the covariance function of the derivative process are strictly positive definite and with sample paths that are differentiable,<sup>9</sup>  $\mathcal{X}$  is a compact set, and  $P(\mathbf{x} \diamond \xi)$  is a continuous distribution with strictly positive support, then Assumption 5.1 holds.*

*Proof.* Consider the case of a scalar function  $f$ . By Lemma 1 of Cakmak et al. (2020), the density of  $f(\mathbf{x} \diamond \xi)$  is strictly positive. Hence, any interval with positive Lebesgue measure has non-zero density. So the cumulative density function of  $f(\mathbf{x} \diamond \xi)$  is strictly increasing. Consider the case of a multi-output sample path  $f$ . Suppose that the joint CDF is not strictly increasing. Then there exist  $\mathbf{y}, \mathbf{y}'$  such that  $\mathbf{y} \geq \mathbf{y}'$  and there exists at least one  $i \in \{1, \dots, M\}$  s.t.  $y_i > y'_i$  and  $F(\mathbf{y}) \leq F(\mathbf{y}')$ . Since  $\mathbf{y} \geq \mathbf{y}'$  and  $F$  is a CDF,  $F(\mathbf{y}) \geq F(\mathbf{y}')$ . Hence,  $F(\mathbf{y}) = F(\mathbf{y}')$ . So, we have  $P(f_1(\mathbf{x} \diamond \xi) \leq y_1, \dots, f_M(\mathbf{x} \diamond \xi) \leq y_M) = P(f_1(\mathbf{x} \diamond \xi) \leq y'_1, \dots, f_M(\mathbf{x} \diamond \xi) \leq y'_M)$ . Suppose  $y'_i = y_i$  for all  $i \neq j$ . Then,  $P(f_1(\mathbf{x} \diamond \xi) \leq y_1, y'_j < f_j(\mathbf{x} \diamond \xi) \leq y_j, \dots, f_M(\mathbf{x} \diamond \xi) \leq y_M) = 0$ . But the hyperrectangle bounded by  $[-\infty, \dots, y'_j, \dots, -\infty]$  and  $\mathbf{y}$  has positive Lebesgue measure. Since the pdf of each of  $f_1, \dots, f_M$  is strictly positive, the cumulative density over the hyperrectangle is greater than zero, which is a contradiction. The argument is easily extended to the case when there exists  $1 \geq k \geq M$  indices  $i_1, \dots, i_k$  such that  $y'_{i_k} < y_{i_k}$ .  $\square$

#### A.4. Extension to Black-Box Constraints Under Input Noise

In this section, we consider the setting where in addition to the objective function  $f$  there is a vector-valued black-box function  $\mathbf{c}(\mathbf{x}) \in \mathbb{R}^V$  specifying the outcome constraint  $\mathbf{c}(\mathbf{x}) > \mathbf{0}$  that is also subject to input noise  $\xi \sim P(\xi)$ . To handle black-box constraints under input noise, we weight the objectives  $f(\mathbf{x} \diamond \xi)$  by a feasibility indicator  $\mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \xi) > \mathbf{0}]$  that is 1 if all constraints are satisfied and 0 otherwise. We define VAR and MVAR for feasibility-weighted objectives and extend the theoretical results from Section 5.1.

The proofs for the results in the constrained setting follow the proofs in Appendix A.2 with slight modifications. 1) Assumption 5.1 does not hold for feasibility-weighted objectives. Hence, the implication is that a random sampled scalarization is only guaranteed to correspond to a point in the WEAK-MVAR set, but importantly any point in the MVAR set does correspond to some scalarization and therefore can be recovered. 2) The proofs handle the special case where some of the constraints are not satisfied and the feasibility-weighted objectives are zero.

**Definition A.4.** The value-at-risk of the feasibility-weighted objective for a given point  $\mathbf{x}$  is:

$$\text{VAR}_\alpha(f(\mathbf{x} \diamond \xi) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \xi) > \mathbf{0}]) = \sup \{z \in \mathbb{R} : P(f(\mathbf{x} \diamond \xi) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \xi) > \mathbf{0}] \geq z) \geq \alpha\}.$$

**Definition A.5.** The MVAR of the feasibility-weighted objectives  $f$  for a given point  $\mathbf{x}$  is:

$$\text{MVAR}_\alpha(f(\mathbf{x} \diamond \xi), \mathbf{c}(\mathbf{x} \diamond \xi)) = \left\{ z \in \mathbb{R}^M : P(f(\mathbf{x} \diamond \xi) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \xi) > \mathbf{0}] \geq z) \geq \alpha, \nexists z' \in \mathbb{R}^M, z' \succ z, P(f(\mathbf{x} \diamond \xi) \mathbb{1}[\mathbf{c} > \mathbf{0}] \geq z') \geq \alpha \right\}.$$

Let  $\mathcal{M}_{\xi, X}^c := \bigcup_{\mathbf{x} \in X} \text{MVAR}_\alpha[f(\mathbf{x} \diamond \xi), \mathbf{c}(\mathbf{x} \diamond \xi)]$ . The *global* MVAR of the feasibility weighted objectives for a set of points  $X$  is defined as

$$\text{MVAR}_{P(\xi), \alpha}(\{f(\mathbf{x} \diamond \xi), \mathbf{c}(\mathbf{x} \diamond \xi)\}_{\mathbf{x} \in X}) = \{z \in \mathcal{M}_{\xi, X}^c : \nexists z' \in \mathcal{M}_{\xi, X}^c \text{ s.t. } z' \succ z\}.$$

WEAK-MVAR of the feasibility-weighted objectives is defined in the same way, but only requires that its elements are weakly Pareto optimal

**Lemma A.6.** *Given a weight vector  $\mathbf{w} \in \Delta_+^{M-1}$ ,  $\mathbf{y} \in \mathbb{R}^M$ ,  $\mathbf{y}_c \in \mathbb{R}^{M'}$  and  $v \in \mathbb{R}$ ,*

$$s[\mathbf{y}, \mathbf{w}] \mathbb{1}[\mathbf{y}_c \succ \mathbf{0}] \geq v \iff \mathbf{y} \mathbb{1}[\mathbf{y}_c \succ \mathbf{0}] \geq \frac{v}{\mathbf{w}}.$$

<sup>8</sup>Using a discrete MC approximation of the input noise distribution means that the CDF is not strictly increasing.

<sup>9</sup>Sample paths of many commonly used covariance functions are differentiable (Paciorek, 2003).



*Proof.* This follows directly from Definition 5.1.

$$\begin{aligned}
 s[\mathbf{y}, \mathbf{w}] \mathbb{1}[\mathbf{y}_c \succ \mathbf{0}] \geq v &\iff \mathbb{1}[\mathbf{y}_c \succ \mathbf{0}] \min_i w_i y_i \geq v \\
 &\iff w_i y_i \mathbb{1}[\mathbf{y}_c \succ \mathbf{0}] \geq v \quad \forall i \\
 &\iff y_i \mathbb{1}[\mathbf{y}_c \succ \mathbf{0}] \geq \frac{v}{w_i} \quad \forall i \\
 &\iff \mathbf{y} \mathbb{1}[\mathbf{y}_c \succ \mathbf{0}] \geq \frac{v}{\mathbf{w}}.
 \end{aligned}$$

□

**Theorem A.1** (VAR of Feasibility-Weighted Chebyshev scalarization  $\Rightarrow$  Pareto Dominance). *Given a weight vector  $\mathbf{w} \in \Delta_+^{M-1}$ , let  $v = \text{VAR}_\alpha \left( s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \right)$ . Then,*

$$P\left(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \geq \frac{v}{\mathbf{w}}\right) \geq \alpha.$$

*Proof.* From Definition A.4, we have that

$$\text{VAR}_\alpha \left( s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \right) = \sup \left\{ z \in \mathbb{R} : s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \geq z \right\} \geq \alpha.$$

Hence,

$$P(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \geq v) \geq \alpha.$$

From Definition 5.1, we have

$$P(\mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \min_i w_i f^{(i)}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq v) \geq \alpha.$$

By Lemma A.6, the following expressions are equivalent

$$\mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \min_i w_i f^{(i)}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq v \iff \mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \geq \frac{v}{\mathbf{w}}.$$

Hence,  $P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \geq \frac{v}{\mathbf{w}}) \geq \alpha$ . □

**Lemma A.7.** *Let  $z \in \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi})]$ . If  $\mathbf{f}(\mathbf{x}) > \mathbf{0}$ , then  $z \neq \mathbf{0}$  if and only if  $z = \mathbf{0}$ .*

*Proof.* The following shows that, since  $\mathbf{f}(\mathbf{x}) > \mathbf{0}$ ,  $z \neq \mathbf{0}$  if and only if  $P(\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}) < \alpha$ . Thus,  $z \neq \mathbf{0}$  if and only if  $z = \mathbf{0}$ .

Since  $\mathbf{f}(\mathbf{x}) > \mathbf{0}$ , we have that  $z_i \geq 0$  for all  $i = 1, \dots, M$  and there exists  $j \in \{1, \dots, M\}$  such that  $z_j = 0$ . Note that since  $\mathbf{f}(\mathbf{x}) > \mathbf{0}$ , the  $i^{\text{th}}$  element  $f^{(i)}(\mathbf{x}) \mathbb{1}[\mathbf{c}(\mathbf{x}) > \mathbf{0}] = 0$  if and only if  $\mathbb{1}[\mathbf{c}(\mathbf{x}) > \mathbf{0}] = 0$ . Hence, either  $\mathbf{f}(\mathbf{x}) \mathbb{1}[\mathbf{c}(\mathbf{x}) > \mathbf{0}] > \mathbf{0}$  or  $\mathbf{f}(\mathbf{x}) \mathbb{1}[\mathbf{c}(\mathbf{x}) > \mathbf{0}] = \mathbf{0}$ .

From Definition A.5, we have that  $P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \geq z) \geq \alpha$ . Since  $\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] > \mathbf{0}$  or  $\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] = \mathbf{0}$ ,

$$P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] > \mathbf{0}) + P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] = \mathbf{0}) = 1.$$

Suppose  $z \neq \mathbf{0}$ . Since  $z$  is not dominated by any other point in the MVAR set,  $P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \succ \mathbf{0}) < \alpha$ . Hence,  $z$  must be  $\mathbf{0}$ . □

**Lemma A.8** (VAR of Feasibility-Weighted Chebyshev scalarization  $\Rightarrow$  Feasibility-Weighted WEAK-MVAR). *Let  $v = \text{VAR}_\alpha \left( s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}], \mathbf{w}] \right)$ . Then,  $\frac{v}{\mathbf{w}} \in \text{WEAK-MVAR}_\alpha(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}))$ .*

*Proof.* Without loss of generality, assume that the objectives  $\mathbf{f}(\mathbf{x}) > \mathbf{0}$ .<sup>7</sup> Let  $\mathbf{z} = \frac{\mathbf{v}}{\mathbf{w}}$ . Suppose there exists  $\mathbf{z}' \in \text{WEAK-MVAR}_\alpha(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}))$  such that  $\mathbf{z}' > \mathbf{z}$ . Since  $\mathbf{z}' \in \text{WEAK-MVAR}_\alpha(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}))$ ,

$$P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \geq \mathbf{z}') \geq \alpha.$$

Note that  $\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \geq \mathbf{z}'$  implies that

$$\mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \min_i w_i f^{(i)}(\mathbf{x} \diamond \boldsymbol{\xi}) \geq \min_i w_i z'_i.$$

Hence,  $P(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}) \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \geq \mathbf{z}') \geq \alpha$  implies that

$$P(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \geq s[\mathbf{z}', \mathbf{w}]) \geq \alpha.$$

Note that  $s[\mathbf{z}', \mathbf{w}] > s[\mathbf{z}, \mathbf{w}] = v$ . But by the Definition A.4,  $v$  is the maximum value such that

$$P(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \geq v) \geq \alpha.$$

This is a contradiction.  $\square$

**Theorem A.2** (Feasibility-Weighted MVAR  $\Rightarrow$  VAR of Feasibility-Weighted Chebyshev scalarization). *For any  $\mathbf{z} \in \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi})]$ , there exists  $\mathbf{w} \in \Delta_+^{M-1}$  such that  $\mathbf{z} = \frac{\mathbf{v}}{\mathbf{w}}$  where  $v = \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}])$ .*

*Proof.* Without loss of generality, assume that the objectives  $\mathbf{f}(\mathbf{x}) > \mathbf{0}$ .<sup>7</sup>

**Case 1:  $\mathbf{z} \not> \mathbf{0}$ .** From Lemma A.7, we have that  $\mathbf{z} = \mathbf{0}$ . Note that since the MVAR set contains only non-dominated points and  $\mathbf{0}$  is a lower bound on  $\mathbf{f}(\mathbf{x}) \mathbb{1}[\mathbf{c}(\mathbf{x}) > \mathbf{0}]$ ,  $\text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi})] = \{\mathbf{0}\}$ . Let  $v := \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}])$ .

Suppose  $v = 0$ . Then, for any  $\mathbf{w} \in \Delta_+^{M-1}$ ,  $\frac{\mathbf{v}}{\mathbf{w}} = \mathbf{0} \in \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi})]$ .

Suppose  $v > 0$ . Then for any  $\mathbf{w} \in \Delta_+^{M-1}$ ,  $\frac{\mathbf{v}}{\mathbf{w}} > \mathbf{0}$ . By Lemma A.8,  $\frac{\mathbf{v}}{\mathbf{w}} \in \text{WEAK-MVAR}_\alpha(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}))$ . But  $\text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi})] = \{\mathbf{0}\}$ , and any point in the MVAR set is non-dominated. So  $\frac{\mathbf{v}}{\mathbf{w}} \not> \mathbf{0}$ . This is a contradiction.

Hence,  $v = 0$ . Therefore, Theorem A.2 holds when  $\mathbf{z} \not> \mathbf{0}$ .

**Case 2:  $\mathbf{z} \succ \mathbf{0}$ .**

Consider the vector  $\frac{1}{\mathbf{z}}$ . If we divide  $\frac{1}{\mathbf{z}}$  by its L1-norm, we obtain a vector  $\mathbf{w} := \frac{1}{\mathbf{z}} \|\frac{1}{\mathbf{z}}\|_1^{-1} \in \Delta_+^M$ , where  $\|\cdot\|_1$  denotes the L1-norm. Let  $v = \frac{1}{\|\frac{1}{\mathbf{z}}\|_1} \in \mathbb{R}_+$ . Then,  $\mathbf{z} = \frac{\mathbf{v}}{\mathbf{w}}$ . Hence, all we need to show is that  $v = \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}])$ . By definition,

$$\text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}]) = \sup \{v'' \in \mathbb{R} : P(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}] \geq v'') \geq \alpha\}.$$

Let us define  $v' := \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}])$ . Suppose that  $v < v'$ . By Lemma A.8,  $\frac{\mathbf{v}'}{\mathbf{w}} \in \text{WEAK-MVAR}_\alpha(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}))$ . Since  $v < v'$ ,  $\mathbf{z} = \frac{\mathbf{v}}{\mathbf{w}} \prec \frac{\mathbf{v}'}{\mathbf{w}}$ . But  $\mathbf{z}$  is in  $\text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$ . So by Definition 4.2,  $\mathbf{z}$  is not dominated by any other vector in  $\text{WEAK-MVAR}_\alpha(\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}))$ . This is a contradiction. Hence,  $v = \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}])$ . Thus, Theorem A.2 holds when  $\mathbf{z} \succ \mathbf{0}$ .  $\square$

**Corollary A.2.** *There is an injective function  $g : \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi})] \rightarrow \Delta_+^{M-1}$  such that  $g(\mathbf{z}) = \frac{1}{\mathbf{z} \|\frac{1}{\mathbf{z}}\|} = \mathbf{w}$  and  $\mathbf{z} = \frac{\mathbf{v}}{\mathbf{w}}$  where  $v = \text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}]) = \frac{1}{\|\frac{1}{\mathbf{z}}\|}$ .*

**Corollary A.3.** *Suppose  $\mathbf{z}$  is a point in the global MVAR set  $\mathbf{z} \in \text{MVAR}_\alpha[\{\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi})\}_{\mathbf{x} \in \mathcal{X}}]$ . Let  $\mathcal{X}_\mathbf{z}^*$  be the set of designs such that for all  $\mathbf{x} \in \mathcal{X}_\mathbf{z}^*$ ,  $\mathbf{z} \in \text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi})]$ . Then every  $\mathbf{x} \in \mathcal{X}_\mathbf{z}^*$  is a maximizer of  $\text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}])$ .*

When  $\text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi})] \neq \{\mathbf{0}\}$ , it follows directly from the injective mapping from  $\mathbf{z}$  to  $\mathbf{w}$  that  $\text{VAR}_\alpha(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}] \mathbb{1}[\mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi}) > \mathbf{0}])$  is the same for all  $\mathbf{x} \in \mathcal{X}_\mathbf{z}^*$ . When  $\text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{c}(\mathbf{x} \diamond \boldsymbol{\xi})] = \{\mathbf{0}\}$ , then all  $\mathbf{x}$  are infeasible designs and  $\mathcal{X}_\mathbf{z}^* = \mathcal{X}$ .

## B. MARS with Alternative Acquisition Functions

In this section, we discuss using MARS with two alternative acquisition functions: Thompson Sampling (TS) and Upper Confidence Bound (UCB).

### B.1. MARS with Thompson Sampling

As discussed in Section 5, direct MVAR optimization with  $q$ NEHVI requires evaluating the joint posterior over  $n_{\xi}(n+1)$  designs. The same is true when using MARS-NEI. Although low-rank Cholesky updates can significantly reduce the complexity (Osborne, 2010), further computational improvements can be obtained by using TS with random Fourier features (RFFs) (Rahimi and Recht, 2008). However, RFFs are approximate GP samples and introduce approximation error (see Appendix D.3 for further discussion).<sup>10</sup> We refer to this method as MARS-TS. MARS-TS naturally supports (i) parallel candidate generation by drawing a new posterior sample and new scalarization weights for each candidate; (ii) constraints, by evaluating the feasibility-weighted objectives under the posterior sample, and (iii) noisy observations.

### B.2. MARS with Upper Confidence Bound

Another computationally efficient approach is to use Upper Confidence Bound (UCB), which does not require the expensive integration over  $\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})$  where  $\mathbf{x} \in X_{1:n}$ . We refer to this method as MARS-UCB. In what follows, we show how to extend the V-UCB algorithm of Nguyen et al. (2021b) to optimize the VAR of Chebyshev scalarizations. The result builds on the following lemma by Chowdhury and Gopalan (2017), which holds under the assumption that the function  $f^{(i)}(\cdot)$  belongs to a reproducing kernel Hilbert space (RKHS)  $\mathcal{F}_{k_i}(B_i)$ , whose RKHS norm is bounded by  $\|f^{(i)}\|_{k_i} \leq B_i$ , where  $\mathbf{f} = [f^{(1)}, \dots, f^{(M)}]$ . We use  $\mu_n^{(i)}(\mathbf{x})$ ,  $\Sigma_n^{(i)}(\mathbf{x}, \mathbf{x})$  to denote the posterior, conditional on observations up to iteration  $n$ , mean and variance of the GP surrogate corresponding to  $i^{\text{th}}$  objective, and use  $\sigma_i^2$  to denote the observation noise for the  $i^{\text{th}}$  objective.

**Lemma B.1.** (Chowdhury and Gopalan, 2017). For  $\delta \in (0, 1)$ ,  $\zeta_{n+1}^{(i)} = B_i + \sigma_i^2 \sqrt{2(\gamma_n + 1 + \log(1/\delta))}$ , the following holds for all  $\mathbf{x} \in \mathcal{X}$  with probability  $\geq 1 - \delta$ :

$$l_n^{(i)}(\mathbf{x}) \leq f^{(i)}(\mathbf{x}) \leq u_n^{(i)}(\mathbf{x}), \quad (2)$$

where  $l_n^{(i)}(\mathbf{x}) := \mu_n^{(i)}(\mathbf{x}) - \zeta_{n+1}^{(i)} (\Sigma_n^{(i)}(\mathbf{x}, \mathbf{x}))^{1/2}$ ,  $u_n^{(i)}(\mathbf{x}) := \mu_n^{(i)}(\mathbf{x}) + \zeta_{n+1}^{(i)} (\Sigma_n^{(i)}(\mathbf{x}, \mathbf{x}))^{1/2}$ , and  $\gamma_n$  denotes the maximum information gain.

Assuming that each objective is modeled using an independent GP surrogate and considering all objectives jointly, we see that (2) holds jointly for all  $i = 1, \dots, m$  with probability at least  $(1 - \delta') := (1 - \delta)^m$ . Applying the Chebyshev scalarization,

$$\begin{aligned} w_i l_n^{(i)}(\mathbf{x}) &\leq w_i f^{(i)}(\mathbf{x}) \leq w_i u_n^{(i)}(\mathbf{x}), \forall i = 1, \dots, M \\ \min_i w_i l_n^{(i)}(\mathbf{x}) &\leq \min_i w_i f^{(i)}(\mathbf{x}) \leq \min_i w_i u_n^{(i)}(\mathbf{x}) \\ s[\mathbf{l}_n(\mathbf{x}), \mathbf{w}] &\leq s[\mathbf{f}(\mathbf{x}), \mathbf{w}] \leq s[\mathbf{u}_n(\mathbf{x}), \mathbf{w}] \end{aligned}$$

holds with probability  $\geq (1 - \delta')$  for all  $\mathbf{x} \in \mathcal{X}$ . Following Lemma 2 of (Nguyen et al., 2021b), we get that

$$\text{VAR}_{\alpha}(s[\mathbf{l}_n(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}]) \leq \text{VAR}_{\alpha}(s[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}]) \leq \text{VAR}_{\alpha}(s[\mathbf{u}_n(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}])$$

holds with probability at least  $(1 - \delta')$ . Thus, the UCB policy for VAR of Chebyshev scalarization is defined as the policy that samples  $\mathbf{x}_{n+1} = \arg \max_{\mathbf{x}} \text{VAR}_{\alpha}(s[\mathbf{u}_n(\mathbf{x} \diamond \boldsymbol{\xi}), \mathbf{w}])$ .

In practice, computing the  $\zeta_{n+1}^{(i)}$  given in Lemma B.1 is impractical, and typically leads to an acquisition function that is overly conservative. Thus, we follow Nguyen et al. (2021b) and use  $\zeta_{n+1}^{(i)} = 2 \log(n^2 \pi^2 / 0.6)$  in the experiments.

MARS-UCB can be extended to support parallel candidate generation by sampling a new scalarization for each candidate and noisy observations. The UCB policy derived above does not hold for feasibility-weighted objectives because Lemma B.1 requires that  $f^{(i)}(\cdot)$  belongs to a RKHS, and this is not the case when  $f^{(i)}(\cdot)$  is weighted by a feasibility indicator because it is no longer continuous (de Freitas et al., 2012).

<sup>10</sup>Alternative approaches for efficient posterior sampling such as decoupled sampling (Wilson et al., 2020) could also be used.

## C. Gradient-based Acquisition Function Optimization

### C.1. Approximate Gradients of VAR

One of the earliest and simplest-to-use gradient estimators for VAR was presented by [Hong \(2009\)](#). Under mild regularity assumptions on the distribution of the random variable, they establish the consistency of the VAR gradient estimator, which can be seen as the sample-path gradient of the well-known estimator of VAR. For this discussion, let  $g(\cdot)$  be a deterministic function of its argument (e.g., a sample path of the GP), let us fix  $\mathbf{x}$ , and let  $\boldsymbol{\xi} \sim P(\boldsymbol{\xi})$  be a continuous random variable. Let  $\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_k$  denote i.i.d. samples from  $P(\boldsymbol{\xi})$ . Define the following ordering of the samples, where the subscript  $(\cdot)$  denote the order statistic:

$$g(\mathbf{x} \diamond \boldsymbol{\xi}_{(1)}) \leq g(\mathbf{x} \diamond \boldsymbol{\xi}_{(2)}) \leq \dots \leq g(\mathbf{x} \diamond \boldsymbol{\xi}_{(k)}).$$

The VAR at risk level  $\alpha$  can be estimated by

$$\text{VAR}_{\boldsymbol{\xi} \sim P(\boldsymbol{\xi})}(g(\mathbf{x} \diamond \boldsymbol{\xi})) \approx g(\mathbf{x} \diamond \boldsymbol{\xi}_{(\lfloor (1-\alpha)k \rfloor)}),$$

where  $\lfloor \cdot \rfloor$  denotes the largest integer less than or equal to  $\cdot$ . It is well known (cf. [Serfling \(2008\)](#)) that this estimator is consistent as  $k \rightarrow \infty$ . [Hong \(2009\)](#) extend this result to show that the corresponding gradient estimator

$$\nabla_{\mathbf{x}} \text{VAR}_{\boldsymbol{\xi} \sim P(\boldsymbol{\xi})}(g(\mathbf{x} \diamond \boldsymbol{\xi})) \approx \nabla_{\mathbf{x}} g(\mathbf{x} \diamond \boldsymbol{\xi}_{(\lfloor (1-\alpha)k \rfloor)})$$

is an asymptotically (as  $k \rightarrow \infty$ ) unbiased estimator of the gradient of VAR. The estimator is also consistent as long as  $\nabla_{\mathbf{x}} g(\mathbf{x} \diamond \boldsymbol{\xi}_{(\lfloor (1-\alpha)k \rfloor)})$  is not a function of  $\boldsymbol{\xi}$ , otherwise, averaging of multiple sample gradients is required to obtain a consistent estimator of the gradient of VAR.

In addition to the sample-path gradient estimator discussed above, there are other estimators of gradients of VAR that are based on, e.g., the likelihood ratio gradient estimation or on the kernel density estimators. A detailed discussion of these can be found in [Hong et al. \(2014\)](#).

### C.2. Approximate Gradients of MVAR

Differentiability of MVAR, more precisely the differentiability of the elements of the MVAR set, is a subject that has not been explored in the literature. Since the computation of the MVAR set corresponding to a set of posterior samples is expensive enough to be the bottleneck during acquisition function optimization, it is highly desirable to avoid the finite-difference gradient estimation, which requires multiple evaluations of the objective and is in general less efficient than the sample-path gradients. Instead, it is preferable to establish a direct connection between the MVAR set and the gradients of the samples on which the MVAR is computed. The method we discuss below is inspired by the gradients of VAR, which correspond to the gradients of the sample that is equal to VAR.

The correspondence between  $\nabla_{\mathbf{x}} \text{VAR}_{\boldsymbol{\xi} \sim P(\boldsymbol{\xi})}(g(\mathbf{x} \diamond \boldsymbol{\xi}))$  and  $\nabla_{\mathbf{x}} g(\mathbf{x} \diamond \boldsymbol{\xi}_{(\lfloor (1-\alpha)k \rfloor)})$  follows from the observation that, since  $g(\cdot)$  is a continuous function, shifting  $\mathbf{x}$  by a sufficiently small  $\epsilon$  should not change the ordering of  $\boldsymbol{\xi}$ 's. We should still have  $\text{VAR}_{\boldsymbol{\xi} \sim P(\boldsymbol{\xi})}(g(\mathbf{x} + \epsilon + \boldsymbol{\xi})) \approx g(\mathbf{x} + \epsilon + \boldsymbol{\xi}_{(\lfloor (1-\alpha)k \rfloor)})$  with the same ordering, as long as  $g(\mathbf{x} \diamond \boldsymbol{\xi}_{(i)}) \neq g(\mathbf{x} \diamond \boldsymbol{\xi}_{(j)})$  for  $i \neq j$ . The same idea extends to the MVAR. Using a finite set of samples to approximate the MVAR set, with  $\mathbf{m}(\mathbf{x} \diamond \boldsymbol{\xi})$  denoting an arbitrary element of the MVAR set, we have that  $\mathbf{m}^{(j)}(\mathbf{x} \diamond \boldsymbol{\xi}) = \mathbf{f}^{(j)}(\mathbf{x} \diamond \boldsymbol{\xi}_{i^{(j)}})$  for some  $i^{(j)} \in \{1, \dots, k\}$ , where  $\mathbf{f} = [f^{(1)}, \dots, f^{(M)}]$  and the  $i^{(j)}$  is dependent on the outcome  $j$  and the particular element of the MVAR set. This can be interpreted as saying that the elements of the MVAR set are constructed by piecing together outcomes from the samples of the random variable.

Similar to what was discussed for VAR, if we perturb  $\mathbf{x}$  by a small  $\epsilon$ , under the assumption that  $f^{(j)}(\mathbf{x} \diamond \boldsymbol{\xi}_i) \neq f^{(j)}(\mathbf{x} \diamond \boldsymbol{\xi}_k)$  for  $i \neq k$ , we should get that  $\mathbf{m}^{(j)}(\mathbf{x} + \epsilon + \boldsymbol{\xi}) = \mathbf{f}^{(j)}(\mathbf{x} + \epsilon + \boldsymbol{\xi}_{i^{(j)}})$  with the same  $i^{(j)}$  as before the perturbation. This, in essence, says that we can calculate the gradients of the elements of the MVAR set as  $\nabla_{\mathbf{x}} \mathbf{m}^{(j)}(\mathbf{x} \diamond \boldsymbol{\xi}) = \nabla_{\mathbf{x}} \mathbf{f}^{(j)}(\mathbf{x} \diamond \boldsymbol{\xi}_{i^{(j)}})$ . Putting all outcomes together, we get

$$\nabla_{\mathbf{x}} \mathbf{m}(\mathbf{x} \diamond \boldsymbol{\xi}) = [\nabla_{\mathbf{x}} f^{(1)}(\mathbf{x} \diamond \boldsymbol{\xi}_{i^{(1)}}), \dots, \nabla_{\mathbf{x}} f^{(M)}(\mathbf{x} \diamond \boldsymbol{\xi}_{i^{(M)}})].$$

A theoretical consistency analysis of these MVAR gradient estimators is beyond the scope of this paper. However, we observe that they do work well in practice, enabling efficient optimization of MVAR-NEHVI-RFF (see Appendix D&I).



## D. Direct MVAR Optimization using $q$ NEHVI

In this section, we discuss direct optimization of MVAR using NEHVI, highlight the computational challenges that come with this approach, and introduce an approximation that mitigates some of these challenges for some problems where the MVAR set for a design is relatively small (e.g. where the number of objectives is small and  $\alpha$  is large). However, we find that these approaches are typically infeasible when  $M \geq 3$  due to GPU memory limits (see for example Table 4).

### D.1. Direct Optimization of MVAR with NEHVI

As described in Section 5, the extension of  $q$ NEHVI to optimize MVAR is conceptually simple.  $q$ NEHVI selects the next point to evaluate by maximizing the expected HVI under the GP posterior. We replace the standard HVI of a new point with respect to the PF with the joint HVI of the MVAR set of a new point with respect to the MVAR set over the previously evaluated designs:

$$\alpha_{\text{MVAR-NEHVI}}(\mathbf{x}) = \mathbb{E}_{\mathbf{f} \sim P(\mathbf{f}|\mathcal{D})} [\text{HVI}(\text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})] \mid \text{MVAR}_\alpha[\{\mathbf{f}(\mathbf{x}' \diamond \boldsymbol{\xi})\}_{\mathbf{x}' \in X_{1:n}}])].$$

However, there are several computational issues that make this approach prohibitively expensive, except for when the objective evaluation takes multiple hours or days. In our experiments, we observe long runtimes, even when using very reasonable parameter values of  $n_\xi = 32$  and  $\alpha = 0.9$ . We discuss the factors contributing to this in the following subsection. Note that all the issues discussed are compounded by the fact that even when using a gradient-based approach to optimize the acquisition function, the acquisition function needs to be evaluated many times.

### D.2. Complexity and Challenges

There are three primary computational bottlenecks, corresponding to three stages of computing  $\alpha_{\text{MVAR-NEHVI}}(\mathbf{x})$ . We discuss each stage and their complexity below.




Figure 5: The maximum size of  $\text{MVAR}_\alpha[\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})]$  across the design space  $\mathbf{x} \in \mathcal{X}$ , which is an estimator of the maximum MVAR set size that will be encountered during numerical optimization, for different  $\alpha$  with  $n_\xi = 32$  on the GMM problem. The size of the MVAR set significantly increases as  $\alpha$  decreases and as  $M$  increases.

1. **Posterior Sampling:** Computing  $\alpha_{\text{MVAR-NEHVI}}(\mathbf{x})$  requires drawing joint posterior samples at the baseline points (points that are already evaluated) and the current candidate(s)  $\mathbf{x}$  under all  $n_\xi$  perturbations. For all methods using GPs, posterior sampling at  $n$  points and  $n_\xi$  perturbations scales as  $O(n_\xi^2 n^3 M)$ . Hence, using GPs is only feasible for modest  $n_\xi$ .
2. **Computing MVAR:** The next stage is computing the MVAR corresponding to each  $\mathbf{x}$  from posterior samples of  $\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})$ . Computing MVAR involves computing the distribution function of  $\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})$ ,<sup>11</sup> which has a time and space complexity of  $\mathcal{O}((1 - \alpha)^M n_\xi^{M+1} M)$ . The MVAR has to be computed for each  $\mathbf{x}$  and each posterior sample, further

<sup>11</sup>The distribution of  $\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})$  is unknown and the GP posterior over  $\mathbf{f}(\mathbf{x} \diamond \boldsymbol{\xi})$  is generally intractable. Hence, the distribution of  $\mathbf{x} \diamond \boldsymbol{\xi}$  is often approximated with a finite set of MC samples (Cakmak et al., 2020) because the GP posterior can be evaluated analytically over a finite set of points.

inflating the computational effort required. The size of the resulting MVAR set is  $\mathcal{O}((1 - \alpha)^M n_\xi^M M)$ . Figure 5 empirically demonstrates that the size of the MVAR set significantly increases as  $\alpha$  decreases, particularly for larger  $M$ .

3. **Computing joint hypervolume improvement:** Given the samples of MVAR corresponding to the baseline points and the candidate(s), the final step of MVAR-NEHVI is to compute the joint HVI of the MVAR set of the candidates over the global MVAR set corresponding to the baseline points. To our knowledge, the only existing differentiable approach for joint HVI computation relies on the inclusion-exclusion principle (IEP, Daulton et al. (2020)). The time and space complexity of computing the joint HVI of a set of  $q'$  points using the IEP is exponential with respect to  $q'$ . To compute the HVI of the MVAR for a set of  $q$  candidates, we must compute the joint HVI of a set of size  $q' = q|\text{MVAR}_\alpha[\{\mathbf{f}(\mathbf{x}_i + \boldsymbol{\xi})\}_{i=1}^q]|$ . Since the size of the MVAR set of a single candidate scales as  $\mathcal{O}((1 - \alpha)^M n_\xi^M M)$ , using IEP quickly becomes infeasible, except for very moderate  $M$ ,  $n_\xi$ , and  $\alpha$ . As shown in Figure 5, even for  $q = 1$ ,  $n_\xi = 32$ , and  $M = 3$ , the size of the MVAR set can be quite large for smaller values of  $\alpha$ , which precludes the use of IEP. Although IEP is necessary to make the joint hypervolume improvement computation differentiable, there are non-differentiable approaches (e.g. Lacour et al. (2017)) that could be used instead to compute the joint hypervolume improvement. However, optimizing  $\alpha_{\text{MVAR-NEHVI}}$  without gradients would be very slow given that Daulton et al. (2020) showed that simply optimizing analytic EHVI (without MVAR) with CMA-ES (Hansen, 2007) or L-BFGS-B (Byrd et al., 1995) with approximate gradients estimated via finite differences is over an order of magnitude slower than when using exact gradients.

A final challenge in using MVAR-NEHVI is that the calculation of MVAR is not differentiable, and there are no known theoretically-grounded gradient estimators of MVAR. Therefore, we use the heuristic approach described in Appendix C.2 for estimating the gradients of MVAR.

### D.3. Approximating qNEHVI with RFF Draws

Random Fourier Features (RFF, Rahimi and Recht (2008)) offer an inexpensive and differentiable approximation of GP sample paths. Daulton et al. (2021a) propose to combine a single RFF draw with NEHVI to obtain a cheap approximation,  $q\text{NEHVI-1}$ . They find that  $q\text{NEHVI-1}$  is competitive with  $q\text{NEHVI}$  in small dimensional search spaces, though its performance degrades as the dimensionality increases.

We follow their approach and extend  $q\text{NEHVI-1}$  to optimize MVAR, and name this method MVAR-NEHVI-RFF. Using RFFs significantly reduces the computational cost of the evaluating the acquisition function because it avoids computationally expensive exact posterior sampling. In addition, since we use a single RFF draw, the MVAR set and its HVI only have to be computed once per acquisition function evaluation rather than for each posterior sample. In the end, for small problem instances, MVAR-NEHVI-RFF ends up with a per-iteration runtime measured in seconds, which is an immense reduction from the time it takes for MVAR-NEHVI using an exact GP model.

Note that MVAR-NEHVI-RFF still requires the use of IEP for differentiable HVI computations, making this approach infeasible in many settings as discussed above. In addition, the performance of RFFs based acquisition functions are known to degrade when the underlying function is difficult to model and variance starvation is a known issue (Wang et al., 2018; Wilson et al., 2020; Mutny and Krause, 2018; Calandriello et al., 2019). Thus, for an acquisition function that works well in all settings, we recommend using MARS.

## E. Pruning for Efficient Joint Posterior Sampling

NEI and  $q\text{NEHVI}$  both require sampling from the joint posterior over function values at the new design  $\mathbf{x}$  and previously evaluated designs  $X_{1:n}: P(\{\mathbf{f}(\mathbf{x})\}_{\mathbf{x} \in X_{1:n} \cup \{\mathbf{x}_1, \dots, \mathbf{x}_q\}} | \mathcal{D})$ . To reduce the cost of posterior sampling, we prune  $X_{1:n}$  to only include the subset of points  $X_{\text{pruned}} \subseteq X_{1:n}$  that have nonzero probability of being optimal. We estimate the probability of being optimal using MC estimation with  $N_{\text{prune}}$  samples from the joint posterior. For NEI, optimality is with respect to a scalar objective and often means  $|X_{\text{pruned}}| \ll n$ . For  $q\text{NEHVI}$ , any design that has nonzero probability of being Pareto optimal is retained; typically, this results in a much larger  $X_{\text{pruned}}$  than we using NEI with a scalar objective. The typically larger size of  $X_{\text{pruned}}$  under  $q\text{NEHVI}$ -based methods has a significant effect when using MVAR- $q\text{NEHVI}$ , where sampling from the joint posterior scales as  $\mathcal{O}(Mn n_\xi)$  and has a very significant effect on runtime. Pruning strategies that leverage techniques from pre-screening and population selection in EAs may further improve computational efficiency.

## F. Optimization of Multi-Objective Expectation Objectives

As noted in the main text, the optimization of expectation of objectives can be achieved via rather straightforward extensions of the existing multi-objective acquisition functions. Here, we discuss the main idea, and show how to extend  $q$ NPAREGO and  $q$ NEHVI (Daulton et al., 2021a).

### F.1. Optimization Expectation Objectives with $q$ NPAREGO

The acquisition function PAREGO is an extension of the well known Expected Improvement acquisition function to the multi-objective setting via augmented Chebyshev scalarizations.  $q$ NPAREGO is an MC-based variant that uses composite objectives with the NEI acquisition function (Daulton et al., 2021a). Given a weight vector  $\mathbf{w} \in \Delta_+^{M-1}$ , it selects the next point to evaluate as follows:

$$\mathbf{x}_{n+1} = \arg \max_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{\mathbf{f} \sim P(\mathbf{f}|\mathcal{D})} \left[ s_a[\mathbf{f}(\mathbf{x}), \mathbf{w}] - \max_{\mathbf{x}' \in X_{1:n}} s_a[\mathbf{f}(\mathbf{x}'), \mathbf{w}] \right]_+, \quad (3)$$

where  $X_{1:n}$  denotes the points evaluated so far, and  $[\cdot]_+$  denotes  $\max(\cdot, 0)$ ,  $s_a[\mathbf{y}, \mathbf{w}] = \min w_i y_i + \beta \sum_i w_i y_i$ , and  $\beta$  is a small positive constant. The expectation in (3) is not available in closed form, and is typically replaced by a (Q)MC approximation obtained by drawing samples of  $\{\mathbf{f}(\mathbf{x}')\}_{\mathbf{x}' \in X_{1:n} \cup \{\mathbf{x}\}}$  from the joint GP posterior. A batch of  $q$  candidates can be selected in a sequential greedy fashion where each point is selected using a different scalarization weight vector and the improvement from the batch of  $q$  points replaces the improvement from a single point in (3).

To extend  $q$ NPAREGO to the expectation objectives, we replace each occurrence of  $s_a[\mathbf{f}(\mathbf{x}), \mathbf{w}]$  in (3) with  $s_a[\mathbb{E}_{\xi \sim P(\xi)}[\mathbf{f}(\mathbf{x} \diamond \xi)], \mathbf{w}]$ . For implementation, we follow the same MC idea, and draw samples from the joint posterior of  $\{\mathbf{f}(\mathbf{x} \diamond \xi)\}_{\mathbf{x} \in X_{1:n} \cup \{\mathbf{x}_1, \dots, \mathbf{x}_q\}, \xi \in \Xi}$  where  $\Xi$  is a set of  $n_\xi$  input noise samples, and approximate (3) using these samples. To improve the computational efficiency, one can also calculate the posterior distribution of  $\{\mathbb{E}_{\xi \sim P(\xi)}[\mathbf{f}(\mathbf{x} \diamond \xi)]\}_{\mathbf{x} \in X_{1:n} \cup \{\mathbf{x}_1, \dots, \mathbf{x}_q\}}$  from the posterior of  $[\mathbf{f}(\mathbf{x} \diamond \xi)]_{\mathbf{x} \in X_{1:n} \cup \{\mathbf{x}_1, \dots, \mathbf{x}_q\}, \xi \in \Xi}$  via a simple matrix-matrix product, and use that to draw the posterior samples. This avoids inverting  $m(n+q)n_\xi \times (n+q)|\Xi|$  matrices, and reduces the cost of posterior sampling from  $\mathcal{O}(m(n+q)^3 n_\xi^3)$  to  $\mathcal{O}(m(n+q)^3)$ . However, this computational technique cannot be used if there are black-box constraints. We refer to this method as EXP- $q$ NPAREGO.

### F.2. Optimization Expectation Objectives with $q$ NEHVI

The other acquisition function we consider is the  $q$ NEHVI:

$$\mathbf{x}_{n+1} = \arg \max_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{\mathbf{f} \sim P(\mathbf{f}|\mathcal{D})} [\text{HVI}(\mathbf{f}(\mathbf{x})|\mathcal{P}_n)], \quad (4)$$

where  $\mathcal{P}_n$  is the PF over  $\mathbf{f}(\mathbf{x})_{\mathbf{x}' \in X_{1:n}}$ . The extension of  $q$ NEHVI to the expectation objectives follows a similar path to that of  $q$ NPAREGO. We replace the hypervolume improvement of  $\mathbf{f}(\mathbf{x})$  in (4) with the hypervolume improvement of  $\mathbb{E}_{\xi \sim P(\xi)}[\mathbf{f}(\mathbf{x})]$  with respect to the PF over  $\{\mathbb{E}_{\xi \sim P(\xi)}[\mathbf{f}(\mathbf{x}')]\}_{\mathbf{x}' \in X_{1:n}}$ . To do so, we replace the posterior samples of  $\mathbf{f}(\mathbf{x})$  and  $\{\mathbf{f}(\mathbf{x}')\}_{\mathbf{x}' \in X_{1:n}}$  that are used in  $q$ NEHVI calculations with the posterior samples of the expectation over  $P(\xi)$ , which can be obtained in the same manner described above for  $q$ NPAREGO. However,  $q$ NEHVI with expectation objective is often prohibitively slow because typically relatively few points from  $X_{1:n}$  can be pruned as discussed in Appendix E. Hence, we only evaluate a single sample approximation using RFFs, analogous to the RFF approximation of MVAR-NEHVI, which we refer to as EXP-NEHVI-RFF.

### F.3. Challenges of Using Expectation with Feasibility-Weighted Objectives

Independently computing the expectation of the objectives and the feasibility and taking the product of the expectations, would ignore the fact that the objective functions and constraint functions are evaluated on the same perturbed designs. To account for the perturbed inputs jointly across in the objectives and constraints, we use feasibility weighted objectives. Feasibility weighting requires penalizing designs that are infeasible such that the feasibility-weighted objectives for an infeasible design are worse than the objectives for any feasible design.

Feasibility weighting can make the expectation sensitive to the range of the objectives. When evaluating a solution near the border of the feasible domain, we end up with a subset of the perturbed solutions evaluating to zero due to infeasibility and others evaluating to their respective objective values. To see how this can affect the performance, consider the following

examples. Suppose that half of the perturbed solutions are infeasible and the objective values are bounded in  $[0, 1]$ . In this case, the feasibility weighted objective take values in  $[0, 0.5]$ , where it will be inferior to some other solutions due to the potential for it to be infeasible. Now, suppose that the objective values are bounded in  $[100, 101]$ . The infeasibility in this case will bring the feasibility weighted objective to the range of  $[50, 50.5]$ , which is strictly worse than any solution that is more feasible, even if by only a small fraction. If we instead set the infeasible solutions to 100 rather than zero, this would lead to the feasibility weighted objective value to  $[100, 100.5]$ . Note that setting the infeasible solutions to 100 is equivalent to normalizing the objectives to  $[0, 1]$  before applying the feasibility and using zero for the infeasible objectives, which in theory should have no effect in the optimization performance. In practice, we typically do not know precise bounds on the objectives, and instead standardize / normalize the objectives during optimization using bounds derived on the go.

The example above highlights the effect of the range of each objective. Often, an infeasibility cost  $\lambda$  is used to penalize for infeasible designs to ensure that infeasible points are worse than any feasible point (for example if the objectives can take negative values) by setting the feasibility-weighted objectives to  $(\mathbf{f}(\mathbf{x}) + \lambda)\mathbb{1}[\mathbf{c} > \mathbf{0}] - \lambda$  for some  $\lambda \geq 0$ . However, the feasibility weighted expectation is typically sensitive to the infeasibility cost, and feasibility weighted objectives give higher value to conservative solutions if the Pareto front lies near the border of the feasible domain. This makes it difficult to determine apriori how conservatively expectation methods will act when using feasibility weighted objectives. In contrast, MVAR avoids this issue by providing high probability guarantees on the value of the feasibility-weighted objectives under input noise. For any that design is feasible with probability  $\alpha$ , the infeasibility cost is in the tail of the multivariate CDF and has no effect on the elements of the MVAR set.

## G. Experiment Details

### G.1. Method Details

We evaluate the following BO methods:

**Methods that optimize the nominal objectives** (see Daulton et al. (2021a)) for details):  $q$ NPAREGO,  $q$ NEHVI, and NEHVI-RFF (referred to as  $q$ NEHVI-1 in Daulton et al. (2021a)), which approximates the expectation in  $q$ NEHVI with a single approximate GP sample using RFFs.

**Methods that optimize the expectation objectives** (Appendix F): EXP- $q$ NPAREGO, and EXP-NEHVI-RFF.

**Methods that optimize MVAR**: MARS-NEI (Section 5.2), MARS-TS (Appendix B), MARS-UCB (Appendix B), MVAR-NEHVI (Appendix D), and MVAR-NEHVI-RFF (Appendix D).

We implemented all methods using the BoTorch library (Balandat et al., 2020) (except for NSGA-II), leveraging the existing implementations of NEI and  $q$ NEHVI available at <https://github.com/pytorch/botorch>. We used the implementation of NSGA-II in the PyMOO library (Blank and Deb, 2020), which is available at <https://github.com/anyoptimization/pymoo>.

For all model-based methods, we model each objective and constraint with an independent GP with a Matérn- $\frac{5}{2}$  ARD kernel (Rasmussen, 2004).<sup>12</sup> For methods that use scalarizations, we use composite objectives (Astudillo and Frazier, 2019). We use maximum a posteriori estimates of the GP hyperparameters using the default priors in BoTorch. For all MC-based acquisition functions, we use  $N_{MC} = 256$  QMC samples from the GP posterior. We use sample-average approximation (Balandat et al., 2020) by using fixed Quasi-MC samples from  $P(\boldsymbol{\xi})$  (for robust methods) and fixed Quasi-MC base samples for all methods to approximate the expectation over the GP posterior.<sup>13,14</sup> This results in an approximation of the acquisition function that is a deterministic function of the input  $\mathbf{x}$ . For RFF-based methods, the approximate GP sample (using 512 random features) is also a deterministic function, which, coupled with fixed samples from  $P(\boldsymbol{\xi})$ , results in a deterministic approximation of the acquisition functions. The deterministic approximations of the acquisition functions enable the use of quasi-Newton methods for optimization. We optimize all acquisition functions using multi-start optimization with L-BFGS-B (Zhu et al., 1997).

For MARS and other MVAR-based methods, we use the known MVAR reference point, which would typically be supplied

<sup>12</sup>All methods except MARS-UCB support using multi-task GPs that model for correlations between objectives (Bonilla et al., 2008).

<sup>13</sup>All fixed samples are re-sampled once per BO iteration.

<sup>14</sup>For heteroskedastic input noise processes, we fix a set base samples and use the reparameterization trick (Kingma and Welling, 2013) to sample from from the heteroskedastic input noise process using the fixed based samples—rather than directly fixing the input noise samples as we do in the case of homoskedastic noise.



by the decision maker. For  $q$ NEHVI and EXP-NEHVI-RFF, we use the heuristic from Daulton et al. (2020) to adaptively infer the reference point during the optimization (the MVAR reference point is not suitable for the nominal and expectation objectives).

For methods involving scalarizations, the objectives are normalized before applying the scalarizations. For MARS methods, the reference point is used as the lower bound and the ideal point (i.e. the component-wise maximum of each objective, Ishibuchi et al. (2018)) across the MVAR set over the previously evaluated designs (estimated using the posterior mean) is used as the upper bound for normalization. For  $q$ NPAREGO, we use the ideal and nadir points (i.e. the component-wise minimum objective values across the PF, Ishibuchi et al. (2018)) across the PF over the previously evaluated designs. Similarly, for EXP- $q$ NPAREGO, we use the ideal and nadir points over the PF expectation objectives (estimated using the posterior mean) over the previously evaluated designs.

For feasibility weighting the objectives, we use a sigmoid function as a differentiable approximation of the indicator function as in Balandat et al. (2020). The infeasibility cost set to be the minimum posterior mean minus six standard deviations.

For methods that use NEI and  $q$ NEHVI with exact posterior sampling, we prune the previously evaluated designs using  $N_{\text{prune}} = 2048$  samples to estimate the probability that a previously evaluated design is optimal. Additionally, we cache the Cholesky decomposition of the posterior covariance matrix over  $\{\mathbf{f}(\mathbf{x}' \diamond \boldsymbol{\xi})\}_{\mathbf{x}' \in X_{\text{pruned}}}$  and use low-rank updates to draw joint samples over  $\{\mathbf{f}(\mathbf{x}' \diamond \boldsymbol{\xi})\}_{\mathbf{x}' \in X_{\text{pruned}} \cup \{\mathbf{x}\}}$  (Osborne, 2010).

For batch (or asynchronous) candidate generation, we use a sequential greedy approach (Wilson et al., 2018), where one new candidate is optimized at time and the joint acquisition value of all candidates  $\mathbf{x}_1, \dots, \mathbf{x}_i$  is optimized to select  $\mathbf{x}_i$ . For methods relying on scalarizations, a new scalarization is sampled for each new candidate.

For NSGA-II, we used the same initial sobol starting points as for the other methods. We used a population size of 10 and adjusted the number of iterations for NSGA-II according to the evaluation budget. The PyMOO default configuration was used for all other settings. In addition to the objectives, observations of the constraints are provided to NSGA-II.

## G.2. Problem Details

In this section, we provide descriptions of the test problems. The reference points used for all problems are provided in Table 2. For each problem, we set the reference point to be slightly worse than the nadir point (using the heuristic from Ishibuchi et al. (2011)) of the MVAR set evaluated over a large grid of design points. For the Penicillin problem, we set the reference point to exclude the region of the objective space with low values of the time objective (following (Liang and Lai, 2021)) as those objective trade-offs are less appealing to decision makers due to providing negligible Penicillin yield.

**Toy Problem** ( $d = 1$ ,  $M = 2$ ,  $\alpha = 0.9$ ): This is the toy problem that was used to highlight the concepts in Figures 1&2. The noise model is given by  $P(\boldsymbol{\xi}) = \mathcal{N}(\boldsymbol{\mu} = \mathbf{0}, \Sigma = 0.1I_2)$ . The first objective is a mixture linear-sinusoidal function, and the second objective is modified from the well-known Levy test function. The exact expressions are given as follows. The function is evaluated on  $x \in [0, 0.7]$ .

$$\begin{aligned}
 f^{(1)}(x) &= 30 - 30 * (p_1(x)p_4(x) + p_2(x)(1 - p_4(x)) + p_3(x)) \\
 p_1(x) &= 2.4 - 10x - 0.1x^2 \\
 p_2(x) &= 2x - 0.1x^2 \\
 p_3(x) &= (x - 0.5)^2 + 0.1 \sin(30x) \\
 p_4(x) &= 1 / (1 + \exp((x - 0.2)/0.005)) \\
 f^{(2)}(x) &= p_5((x * 0.95 + 0.03) * 20 - 10) \\
 p_5(x) &= p_6(1 + (x - 1)/4) - 0.75 * x^2 + 9.0955 \\
 p_6(x) &= (\sin(\pi * x))^2 + (x - 1)^2(1 + 10(\sin(\pi * x))^2)
 \end{aligned}$$

**GMM** ( $d = 2$ ,  $M \in \{2, 3, 4\}$ ,  $\alpha \in \{0.7, 0.8, 0.9\}$ ): In addition to the version presented in the main text, we consider several variations of the GMM problem using different number of objectives, different noise models, and different risk levels to analyze the effects of these factors on the optimization performance of the algorithms. For all GMM problems considered, each objective is a mixture of the probability density function of three Gaussian distributions, modified from the single objective version presented in (Fröhlich et al., 2020). We present the canonical formula of the objectives and

the parameters corresponding to each objective below. In the formula,  $\phi(\mathbf{x}; \mu, \Sigma)$  is used to denote the probability density function of the multivariate Gaussian distribution with mean  $\mu$  and covariance matrix  $\Sigma$ . The search space for all GMM problems is  $\mathcal{X} = [0, 1]^2$ .

The GMM problem used in the main text involves 2 objectives and uses  $\alpha = 0.9$ . Additional experiments in Appendix I.1 use additional independent GMMs to increase the number of objectives to 3 and 4, and evaluate performance with different settings of  $\alpha \in \{0.7, 0.8, 0.9\}$ . In all experiments with 3 and 4 objective GMM, we use additive noise, where  $P(\xi) = \mathcal{N}(\mu = \mathbf{0}, \Sigma = 0.05I_M)$ . Many additional noise processes are discussed and evaluated in Appendix I.5, using the same 2 objective GMM problem from the main text and  $\alpha = 0.9$ .

$$f^{(i)}(\mathbf{x}) = 2\pi \sum_{j=1}^3 \text{var}_j^{(i)} \text{cons}_j^{(i)} \phi(\mathbf{x}; \mu = \text{pos}_j^{(i)}, \Sigma = \text{var}_j^{(i)} I_2)$$

$$\text{pos}_j^{(i)} = \begin{cases} j = 1 & j = 2 & j = 3 \\ [0.2, 0.2] & [0.8, 0.2] & [0.5, 0.7] & \text{if } i = 1 \\ [0.07, 0.2] & [0.4, 0.8] & [0.85, 0.1] & \text{if } i = 2 \\ [0.08, 0.21] & [0.45, 0.75] & [0.86, 0.1] & \text{if } i = 3 \\ [0.09, 0.19] & [0.44, 0.72] & [0.89, 0.13] & \text{if } i = 4 \end{cases}$$

$$\text{var}_j^{(i)} = \begin{cases} j = 1 & j = 2 & j = 3 \\ 0.04 & 0.01 & 0.01 & \text{if } i = 1 \\ 0.04 & 0.01 & 0.0025 & \text{if } i = 2 \\ 0.04 & 0.01 & 0.0049 & \text{if } i = 3 \\ 0.0225 & 0.0049 & 0.0081 & \text{if } i = 4 \end{cases}$$

$$\text{cons}_j^{(i)} = \begin{cases} j = 1 & j = 2 & j = 3 \\ 0.5 & 0.7 & 0.7 & \text{if } i = 1 \\ 0.5 & 0.7 & 0.7 & \text{if } i = 2 \\ 0.5 & 0.7 & 0.9 & \text{if } i = 3 \\ 0.5 & 0.7 & 0.9 & \text{if } i = 4 \end{cases}$$

**Constrained Branin Currin** We use the open source implementation available at <https://github.com/pytorch/botorch>. See Daulton et al. (2020) for details.

**Disc Brake** We use the open source implementation available at <https://github.com/ryojitanabe/reproblems>. See Tanabe and Ishibuchi (2020) for details.

**Penicillin Manufacturing Problem** We use the open-source implementation available at <https://github.com/HarryQL/TuRBO-Penicillin>. See Liang and Lai (2021) for details. We adapt the problem by adding independent zero-mean Gaussian input noise to each parameter. The standard deviation of the input noise distribution for each parameter is listed in Table 1.

Table 1: Standard deviation for independent zero-mean Gaussian input noise for each parameter in the Penicillin Problem (reported as a percentage of the range of each parameter).

Parameter	Noise Level
Culture Volume	3%
Biomass Concentration	3%
Temperature	0.5%
Glucose Concentration	2%
Substrate Feed Rate	1%
Substrate Feed Concentration	1%
H <sup>+</sup> Concentration	1%

Table 2: Reference points for negative versions (i.e. multiplying the objectives by -1 to make the goal maximization of all objectives) of all problems (except the GMM and Toy problems, which are designed for maximization).

Problem	Reference Point
Toy Problem	[-14.1951, -3.1887]
Disc Brake	[-5.89, -3.27]
Constrained Branin Currin (heteroskedastic noise)	[-194.9376, -12.2969]
Constrained Branin Currin (homoskedastic noise)	[-195.4667, -12.4984]
Penicillin	[5.657, -64.1, -340.0]
GMM ( $M = 2, \alpha = 0.9$ , multiplicative noise)	[0.3752, 0.3548]
GMM ( $M = 2, \alpha = 0.9$ , correlated noise)	[0.2727, 0.2583]
GMM ( $M = 2, \alpha = 0.8$ , heteroskedastic noise)	[0.3465, 0.3036]
GMM ( $M = 2, \alpha = 0.9$ , homoskedastic noise, $\sigma = 0.05$ )	[0.2756, 0.2368]
GMM ( $M = 2, \alpha = 0.9$ , homoskedastic noise, $\sigma = 0.1$ )	[0.1047, 0.1112]
GMM ( $M = 2, \alpha = 0.9$ , homoskedastic noise, $\sigma = 0.2$ )	[0.0160, 0.0131]
GMM ( $M = 3, \alpha = 0.9$ , homoskedastic noise, $\sigma = 0.05$ )	[0.2733, 0.0051, 0.1538]
GMM ( $M = 3, \alpha = 0.9$ , homoskedastic noise, $\sigma = 0.05$ )	[0.2733, 0.0051, 0.1538]
GMM ( $M = 3, \alpha = 0.8$ , homoskedastic noise, $\sigma = 0.05$ )	[0.0420, 0.0180, 0.1952]
GMM ( $M = 3, \alpha = 0.7$ , homoskedastic noise, $\sigma = 0.05$ )	[0.0537, -0.0517, -0.0021]
GMM ( $M = 4, \alpha = 0.9$ , homoskedastic noise, $\sigma = 0.05$ )	[0.0264, -0.0396, 0.0619, 0.1689]
GMM ( $M = 4, \alpha = 0.8$ , homoskedastic noise, $\sigma = 0.05$ )	[0.0322, -0.0398, 0.1168, -0.0023]

### G.3. Evaluation Details

The global MVAR set is unknown and is approximated by taking the union of the MVAR sets of all designs evaluated across all methods and all replications. We take this approach because even using an evolutionary algorithm to optimize MVAR is nontrivial, since MVAR maps a single design to a set of points and is relatively computationally intensive to evaluate. To evaluate the performance of a given method, we use  $n_{\xi} = 512$  (except for 4 objective GMM, where we use  $n_{\xi} = 256$ ) to compute a high-fidelity estimate of the MVAR set across the designs selected during optimization by the method. We similarly use the same  $n_{\xi} = 512$  samples to estimate the true MVAR set (by considering all designs evaluated across all methods and all replications).

### H. Wall Times

In Table 3, we present the time it takes to run a single BO iteration using all algorithms we considered in this paper. We include the runtimes for the four problems from the main text. Wall times for additional problems including several problems with 3 and 4 objectives are provided in Table 4 (these problems are described in Appendix I.1). As we discuss in Appendix D, MVAR-NEHVI, when not computationally infeasible, is quite expensive to run, making it an impractical method for most problems. Although MVAR-NEHVI-RFF provides a much cheaper and highly performant approximation, we see that it also runs into computational limitations as the size of the MVAR set grows (e.g. for Penicillin with 3 objectives); this is also pronounced in Table 4 on the problems with 3 and 4 objectives. For the MARS family of methods, we see overall quite reasonable runtimes, with the most expensive one, MARS-NEI, taking on average 41.4 seconds on the most expensive problem instance we considered. MARS-TS offers a cheaper alternative to MARS-NEI, with its average runtime remaining below 10 seconds on all experiments. As we show later in the Appendix, the performance of MARS-TS typically trails closely behind MARS-NEI, making it a strong alternative when the algorithm runtime is of the essence.

Table 3: The wall time (in seconds) per BO iteration. The experiments were timed on a shared cluster using 4 CPU cores, 1 GPU, and 16 GB of RAM. We report the mean and 2 standard errors over 20 trials. An N/A entry denotes that we did not attempt to run a particular experiment (e.g., because the method does not support the problem setting), whereas an OOM entry denotes that we attempted but the experiment did not run due to scalability limitations. The three top-performing algorithms, with respect to the final average MVAR HV regret in each experiment are highlighted using *best*, *second*, *third*, respectively.

Algorithm ( $d, M, V, \alpha$ )	GMM (2, 2, 0, 0.9)	Constrained BC (2, 2, 1, 0.7)	Disc Brake (4, 2, 4, 0.95)	Penicillin (7, 3, 0, 0.8)
Sobol	0.4 ( $\pm 0.6$ )	0.9 ( $\pm 1.0$ )	0.9 ( $\pm 1.0$ )	2.9 ( $\pm 1.5$ )
$q$ NPAREGO	2.5 ( $\pm 2.3$ )	5.6 ( $\pm 4.4$ )	7.9 ( $\pm 15.1$ )	21.2 ( $\pm 25.6$ )
$q$ NEHVI	2.5 ( $\pm 2.0$ )	9.8 ( $\pm 8.9$ )	16.9 ( $\pm 10.3$ )	23.6 ( $\pm 41.0$ )
$q$ NEHVI-RFF	0.8 ( $\pm 0.7$ )	2.1 ( $\pm 3.2$ )	2.7 ( $\pm 5.4$ )	<b>5.0 (<math>\pm 3.9</math>)</b>
Exp- $q$ NPAREGO	3.3 ( $\pm 2.4$ )	10.7 ( $\pm 11.1$ )	23.3 ( $\pm 166.3$ )	121.9 ( $\pm 119.0$ )
EXP-NEHVI-RFF	0.9 ( $\pm 1.0$ )	7.2 ( $\pm 10.8$ )	<b>3.3 (<math>\pm 7.3</math>)</b>	<b>5.1 (<math>\pm 3.3</math>)</b>
MARS-NEI	3.9 ( $\pm 3.0$ )	<b>8.4 (<math>\pm 5.3</math>)</b>	10.3 ( $\pm 45.4$ )	<b>41.4 (<math>\pm 60.1</math>)</b>
MARS-TS	<b>3.3 (<math>\pm 3.6</math>)</b>	<b>3.3 (<math>\pm 3.7</math>)</b>	3.1 ( $\pm 6.5$ )	6.7 ( $\pm 5.8$ )
MARS-UCB	8.2 ( $\pm 11.0$ )	N/A	N/A	12.3 ( $\pm 14.1$ )
MVAR-NEHVI	<b>145.6 (<math>\pm 130.7</math>)</b>	N/A	<b>243.2 (<math>\pm 154.</math>)</b>	N/A
MVAR-NEHVI-RFF	<b>1.8 (<math>\pm 1.7</math>)</b>	<b>10.0 (<math>\pm 11.9</math>)</b>	<b>2.9 (<math>\pm 3.1</math>)</b>	OOM

## I. Additional Experiments

### I.1. Additional Test Problems

In addition to the problems presented in the main text, we studied the performance of the algorithms on the Toy problem used for illustrations in Figures 1 and 2, and the 3 and 4 objective variations of the GMM problem. These problems are described in detail in Appendix G.2. In addition to the acquisition functions presented in the main text, we ran EXP-NEHVI-RFF, MVAR-NEHVI-RFF, MARS-UCB, and MARS-TS. The results of these experiments are presented in Figure 6 for the Toy problem and Figure 7 for the GMM problems, and the runtimes of the algorithms are reported in Table 4. We see that MARS-NEI is overall the best performing method, with MARS-TS typically following closely. MARS-UCB appears to be less reliable, demonstrating significantly worse performance in most experiments. In addition, the MVAR-NEHVI-RFF is missing from all but two of the experiments, which is due to the method running into the scalability limitations discussed in Appendix D.

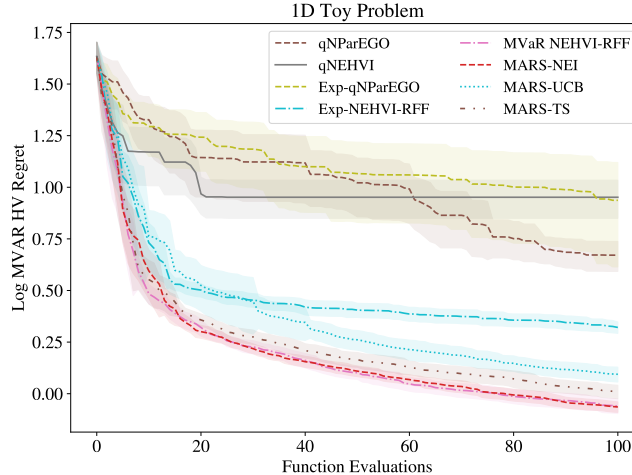


Figure 6: The log MVAR hypervolume regret on the Toy problem. We plot means and 2 standard errors across 20 trials.

Robust Multi-Objective Bayesian Optimization Under Input Noise

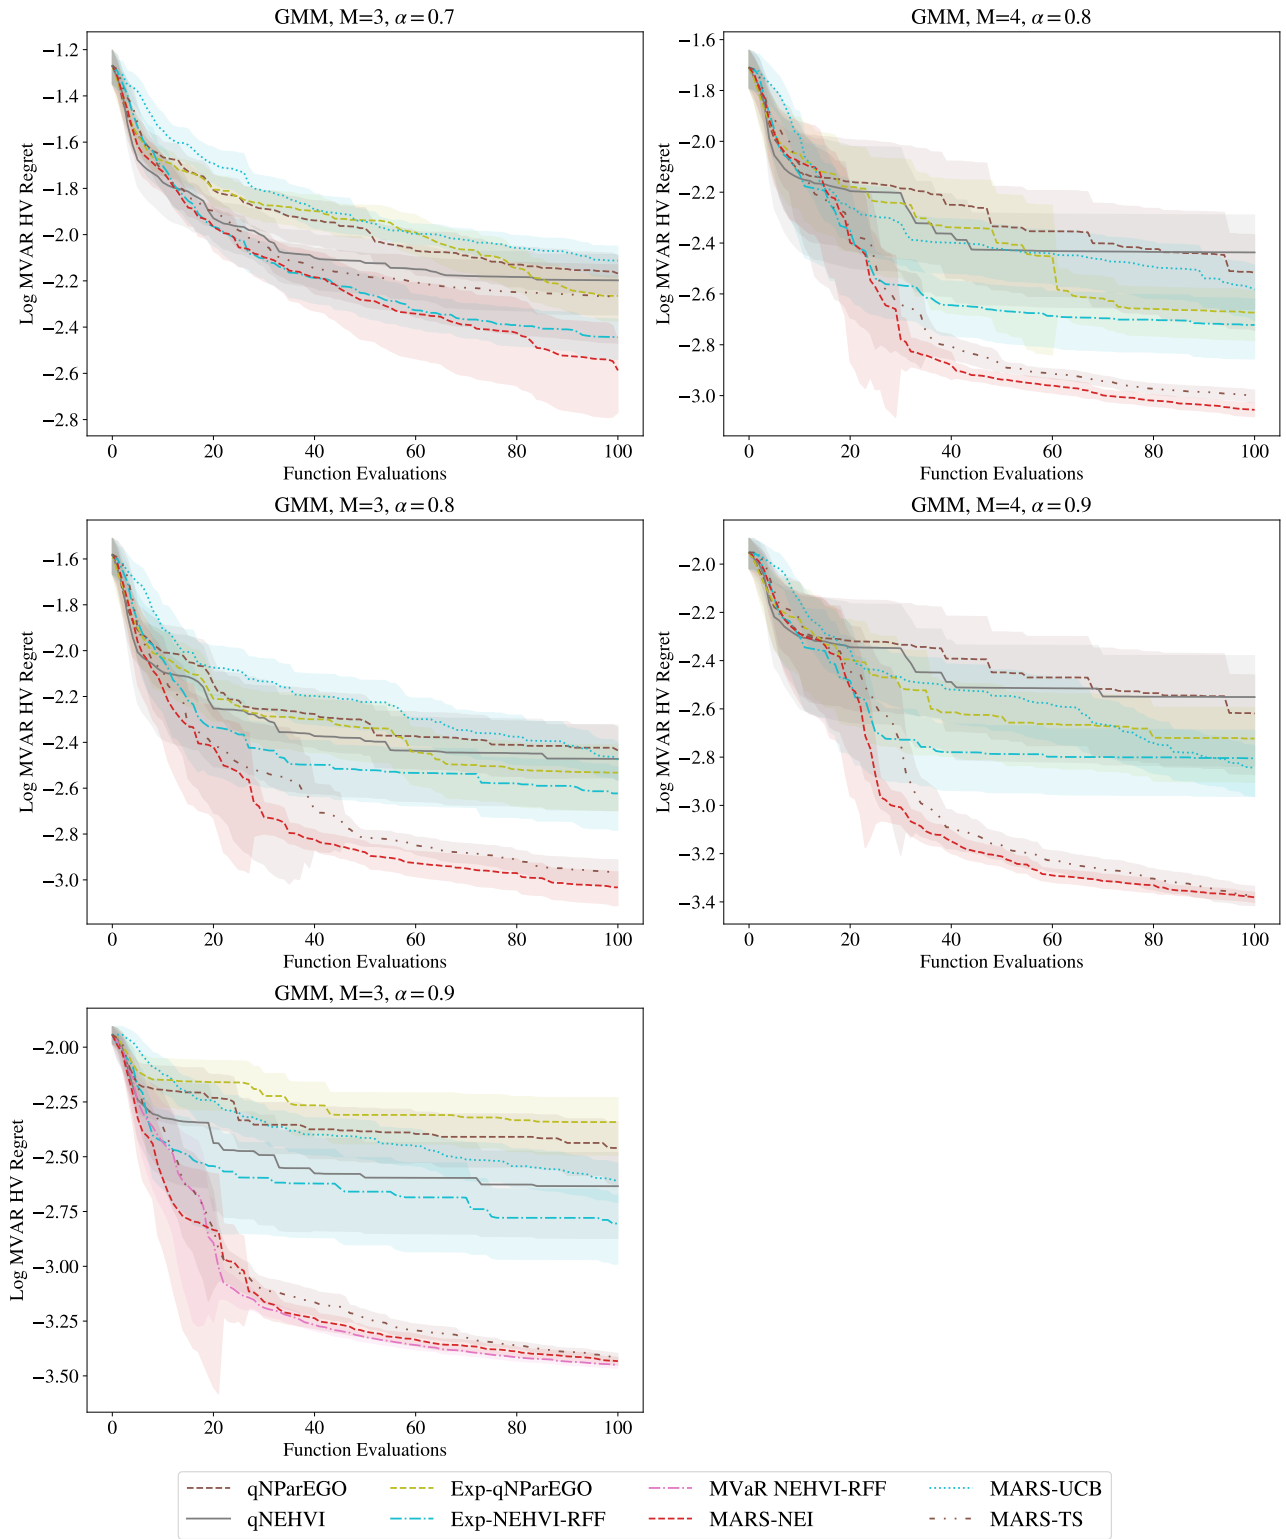


Figure 7: The log MVAR hypervolume regret on 3 and 4-Objective GMM problems. We plot means and 2 standard errors across 20 trials.



Table 4: The wall time (in seconds) per BO iteration for the additional problems. The experiments were timed on a shared cluster using 4 CPU cores, 1 GPU, and 16 GB of RAM. We report the mean and 2 standard errors over 20 trials. An OOM entry denotes that the experiment did not run due to scalability limitations. The three top-performing algorithms, with respect to the final average MVAR HV regret in each experiment are highlighted using *best*, *second*, *third*, respectively.

Algorithm ( $d, M, \alpha$ )	Toy Problem		GMM, M=3		GMM, M=4	
	(1, 2, 0.9)	(2, 3, 0.7)	(2, 3, 0.8)	(2, 3, 0.9)	(2, 4, 0.8)	(2, 4, 0.9)
Sobol	0.4 ( $\pm 0.3$ )	1.9 ( $\pm 2.2$ )	1.2 ( $\pm 1.4$ )	0.5 ( $\pm 0.3$ )	3.3 ( $\pm 3.3$ )	0.8 ( $\pm 0.2$ )
$q$ NPAREGO	2.2 ( $\pm 2.0$ )	4.5 ( $\pm 2.6$ )	3.4 ( $\pm 2.6$ )	1.4 ( $\pm 1.5$ )	8.9 ( $\pm 10.2$ )	4.5 ( $\pm 12.9$ )
$q$ NEHVI	3.0 ( $\pm 2.5$ )	23.3 ( $\pm 26.4$ )	20.2 ( $\pm 22.8$ )	23.8 ( $\pm 43.8$ )	57.3 ( $\pm 75.4$ )	48.0 ( $\pm 55.2$ )
$q$ NEHVI-RFF	0.4 ( $\pm 0.3$ )	<b>3.1 (<math>\pm 5.8</math>)</b>	1.7 ( $\pm 1.0$ )	0.7 ( $\pm 0.5$ )	10.0 ( $\pm 91.8$ )	2.1 ( $\pm 4.0$ )
Exp- $q$ NPAREGO	13.1 ( $\pm 23.5$ )	8.0 ( $\pm 12.8$ )	6.2 ( $\pm 6.4$ )	13.4 ( $\pm 54.2$ )	10.6 ( $\pm 10.6$ )	8.1 ( $\pm 18.3$ )
EXP-NEHVI-RFF	0.4 ( $\pm 0.3$ )	<b>4.0 (<math>\pm 10.3</math>)</b>	<b>2.1 (<math>\pm 6.3</math>)</b>	0.8 ( $\pm 0.8$ )	<b>5.3 (<math>\pm 4.0</math>)</b>	1.9 ( $\pm 1.5$ )
MARS-NEI	<b>10.4 (<math>\pm 21.1</math>)</b>	<b>11.1 (<math>\pm 14.4</math>)</b>	<b>10.7 (<math>\pm 23.0</math>)</b>	<b>8.8 (<math>\pm 13.2</math>)</b>	<b>27.5 (<math>\pm 221.1</math>)</b>	<b>13.5 (<math>\pm 15.9</math>)</b>
MARS-TS	<b>0.9 (<math>\pm 0.7</math>)</b>	5.7 ( $\pm 8.9$ )	<b>3.8 (<math>\pm 2.9</math>)</b>	<b>1.8 (<math>\pm 2.5</math>)</b>	<b>9.9 (<math>\pm 27.1</math>)</b>	<b>4.2 (<math>\pm 4.2</math>)</b>
MARS-UCB	6.7 ( $\pm 7.4$ )	14.4 ( $\pm 22.5$ )	10.4 ( $\pm 11.1$ )	10.3 ( $\pm 12.3$ )	17.6 ( $\pm 19.1$ )	<b>12.1 (<math>\pm 11.9</math>)</b>
MVAR-NEHVI-RFF	<b>2.6 (<math>\pm 2.2</math>)</b>	OOM	OOM	<b>3.0 (<math>\pm 2.1</math>)</b>	OOM	OOM

## I.2. Comparison with $q$ NEHVI Based Methods

Table 5: The final MVAR HV regret obtained using each method. We report the mean and 2 standard errors over 20 trials. An N/A entry denotes that we did not attempt to run a particular experiment (e.g., because the method does not support the problem setting), whereas an OOM entry denotes that we attempted but the experiment did not run due to scalability limitations. The three top-performing algorithms, with respect to the final average MVAR HV regret in each experiment are highlighted using *best*, *second*, *third*, respectively.

Algorithm ( $d, M, V, \alpha$ )	GMM (2, 2, 0, 0.9)	Constrained BC (2, 2, 1, 0.7)	Disc Brake (4, 2, 4, 0.95)	Penicillin (7, 3, 0, 0.8)	Toy Problem (1, 2, 0, 0.9)	GMM, M=4 (2, 4, 0, 0.8)
Scale	$1 \times 10^{-4}$	$1 \times 10^1$	$1 \times 10^0$	$1 \times 10^4$	$1 \times 10^0$	$1 \times 10^{-3}$
Sobol	65.37 ( $\pm 10.23$ )	3.12 ( $\pm 0.36$ )	19.39 ( $\pm 1.35$ )	1.68 ( $\pm 0.05$ )	2.38 ( $\pm 0.27$ )	7.39 ( $\pm 1.35$ )
$q$ NPAREGO	21.54 ( $\pm 7.68$ )	3.28 ( $\pm 0.41$ )	19.10 ( $\pm 1.83$ )	1.73 ( $\pm 0.09$ )	4.69 ( $\pm 0.78$ )	3.04 ( $\pm 1.24$ )
$q$ NEHVI	31.21 ( $\pm 14.34$ )	5.16 ( $\pm 0.76$ )	12.33 ( $\pm 1.77$ )	1.65 ( $\pm 0.12$ )	8.95 ( $\pm 1.87$ )	3.66 ( $\pm 1.47$ )
$q$ NEHVI-RFF	30.54 ( $\pm 12.13$ )	5.60 ( $\pm 0.81$ )	17.11 ( $\pm 1.97$ )	<b>1.38 (<math>\pm 0.09</math>)</b>	5.97 ( $\pm 0.77$ )	2.18 ( $\pm 0.62$ )
Exp- $q$ NPAREGO	10.29 ( $\pm 2.92$ )	2.73 ( $\pm 0.77$ )	3.43 ( $\pm 0.57$ )	1.52 ( $\pm 0.10$ )	8.65 ( $\pm 4.53$ )	2.12 ( $\pm 0.47$ )
EXP-NEHVI-RFF	13.83 ( $\pm 5.95$ )	1.23 ( $\pm 0.05$ )	<b>1.02 (<math>\pm 0.05</math>)</b>	<b>1.46 (<math>\pm 0.06</math>)</b>	2.09 ( $\pm 0.13$ )	<b>1.90 (<math>\pm 0.50</math>)</b>
MARS-NEI	1.45 ( $\pm 0.12$ )	<b>0.70 (<math>\pm 0.04</math>)</b>	2.96 ( $\pm 0.13$ )	<b>1.06 (<math>\pm 0.09</math>)</b>	<b>0.86 (<math>\pm 0.05</math>)</b>	<b>0.88 (<math>\pm 0.05</math>)</b>
MARS-TS	<b>0.92 (<math>\pm 0.07</math>)</b>	<b>0.85 (<math>\pm 0.05</math>)</b>	3.20 ( $\pm 0.21$ )	1.49 ( $\pm 0.05$ )	<b>1.02 (<math>\pm 0.08</math>)</b>	<b>1.00 (<math>\pm 0.06</math>)</b>
MARS-UCB	26.69 ( $\pm 3.28$ )	N/A	N/A	1.77 ( $\pm 0.01$ )	1.24 ( $\pm 0.10$ )	2.61 ( $\pm 0.64$ )
MVAR-NEHVI	<b>0.57 (<math>\pm 0.01</math>)</b>	N/A	<b>1.02 (<math>\pm 0.04</math>)</b>	N/A	1.76 ( $\pm 0.15$ )	N/A
MVAR-NEHVI-RFF	<b>0.64 (<math>\pm 0.03</math>)</b>	<b>0.57 (<math>\pm 0.02</math>)</b>	<b>1.92 (<math>\pm 0.06</math>)</b>	OOM	<b>0.87 (<math>\pm 0.07</math>)</b>	OOM

As discussed in Section 5 and detailed in Appendix D, MVAR can be optimized using  $q$ NEHVI based methods. However, this comes with serious computational challenges, some of which are eased if we use  $q$ NEHVI with RFF draws. In this section, we present results comparing MARS-NEI with EXP-NEHVI-RFF and MVAR-NEHVI-RFF using the test problems from the main text. MVAR-NEHVI was also ran on two of the problems to provide a point of reference for its performance. In addition, we present both the standard  $q$ NEHVI, which uses GPs, as well as its RFF counterpart as a reference point on how the use of RFFs affects the performance of the methods on a given problem.

Figure 8 and Table 5 show that although there are instances where  $q$ NEHVI methods outperform MARS-NEI, MARS-NEI remains competitive throughout. We see that  $q$ NEHVI and its RFF counterpart are competitive, without a clear winner across problems. The expectation  $q$ NEHVI outperforms the expectation  $q$ NPAREGO in all but one problem, highlighting the benefit of using a method that aims to directly maximize the hypervolume in an MO setting. Lastly, it is worth highlighting that the MVAR-NEHVI-RFF is not reported for the Penicillin problem, which is due to the scalability limitations of IEP (discussed in Appendix D), preventing the hypervolume improvement computations from running with the available GPU memory. See Appendix I.1 for additional experiments comparing  $q$ NEHVI based methods.

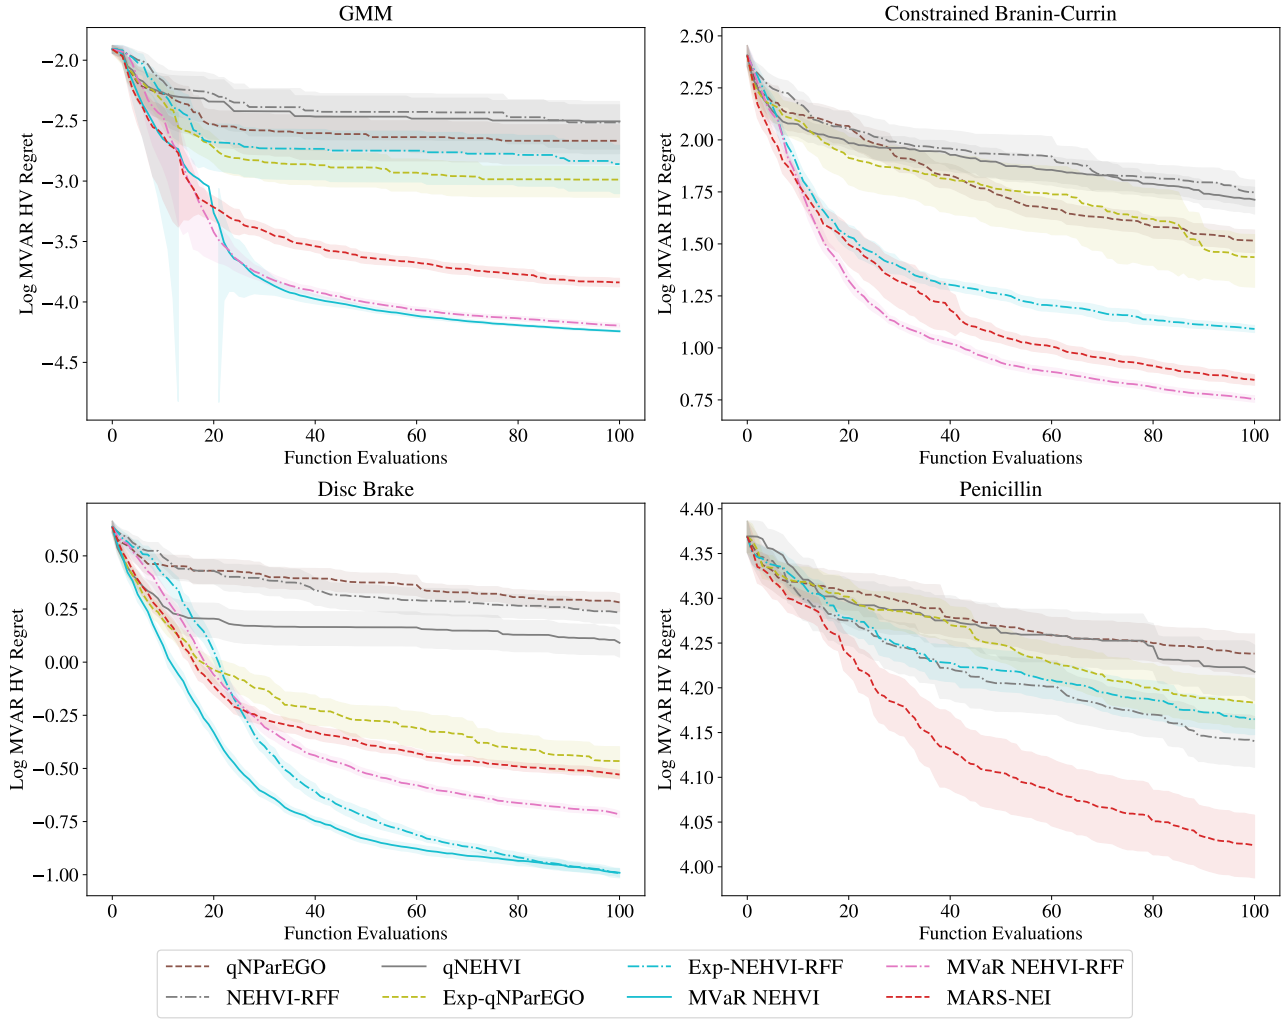


Figure 8: A comparison of the methods from evaluated in the main text with additional  $q$ NEHVI-based methods: NEHVI-RFF, EXP-NEHVI-RFF, MVAR-NEHVI, and MVAR-NEHVI-RFF.

### I.3. Comparison of Methods Optimizing MVAR

In Figure 9 and Table 5, we present results from the test problems from the main text showing the performance of all acquisition functions that we proposed for optimizing MVAR. Results for additional problems comparing are presented in Appendix I.1. We observe that MVAR-NEHVI-RFF is a reasonably cheap method (see Table 3 for runtimes) that performs quite well on smaller problem instances. However, in larger problem instances where the size of the MVAR set is large ( $> 10$ ) it starts running into scalability limitations and no-longer works. Among the class of MARS methods, we find that MARS-NEI consistently performs quite well, with MARS-TS being a close second and a slightly cheaper alternative. On the other hand, MARS-UCB proves to be not as reliable, performing worse than non-robust methods in some problems. We attribute this to its fundamental reliance on the parameter  $\zeta_{n+1}^{(i)}$ , which would have to be tuned on a problem by problem basis to optimize its performance. In light of all the results, we recommend MARS-NEI as a broadly applicable and high performing method for optimizing MVAR, and recommend MVAR-NEHVI-RFF as an alternative when the size of the MVAR set is small.




Figure 9: A comparison of the methods from evaluated in the main text with additional methods for optimizing MVAR: MARS-TS, MARS-UCB, MVAR-NEHVI-RFF, and MVAR-NEHVI.

#### I.4. Parallel Evaluations

In Table 6, we present results demonstrating the effect of varying batch size on the optimization performance of MARS-NEI and MVaR-NEHVI-RFF. As expected, the results show that the performance of both algorithms degrade as the batch size increases. Interestingly, the degradation is rather minimal for MARS-NEI, while the effects of increasing batch size seem to be rather significant for MVaR-NEHVI-RFF. We see that MVaR-NEHVI-RFF underperforms all versions of MARS-NEI even with a batch size of 2. In addition, there are fewer results presented for MVaR-NEHVI-RFF, which is due to the hypervolume improvement computations using IEP running into scalability limits (recall that IEP scales exponentially in the size of the joint MVAR set of the current batch of candidates). These results make a strong case for using MARS-NEI whenever one wished to evaluate candidates in parallel.

#### I.5. Effect of Noise Level

The location of robust designs on a problem depends on many factors, including the magnitude of the input noise. In the edge case where there is no input noise present, the robust designs will be the same as the nominal Pareto optimal designs. As the magnitude of the input noise increases, the robust designs may start to deviate from the nominally optimal designs, with the exact behavior typically being unpredictable and heavily dependent on the problem and noise structure. In Figure 10, we present results on the GMM problem from the main text under various noise models, demonstrating how the

Table 6: Effect of the batch size on optimization performance. We report the final MVAR HV regret and 2 standard errors from 20 trials. OOM denotes that the method ran into scalability issues and did not run.

Algorithm	1D Toy Problem ( $d = 1, M = 2$ )	Constrained Branin-Currin ( $d = 2, M = 2, V = 1$ )
MARS-NEI, $q = 1$	0.86 ( $\pm 0.05$ )	5.43 ( $\pm 0.27$ )
MARS-NEI, $q = 2$	0.92 ( $\pm 0.04$ )	5.20 ( $\pm 0.21$ )
MARS-NEI, $q = 4$	0.92 ( $\pm 0.07$ )	5.58 ( $\pm 0.38$ )
MARS-NEI, $q = 8$	0.93 ( $\pm 0.08$ )	5.59 ( $\pm 0.41$ )
MVAR $q$ NEHVI RFF, $q = 1$	0.87 ( $\pm 0.07$ )	4.36 ( $\pm 0.13$ )
MVAR $q$ NEHVI RFF, $q = 2$	1.03 ( $\pm 0.05$ )	6.16 ( $\pm 0.36$ )
MVAR $q$ NEHVI RFF, $q = 4$	1.21 ( $\pm 0.07$ )	OOM
MVAR $q$ NEHVI RFF, $q = 8$	OOM	OOM

performance of the algorithms change in response to changes in the noise model. The list of noise models considered in this study are as follows:

- Homoscedastic normal noise, std = 0.05:  $P(\xi) = \mathcal{N}(\mu = \mathbf{0}, \Sigma = 0.05I_2)$ .
- Homoscedastic normal noise, std = 0.10:  $P(\xi) = \mathcal{N}(\mu = \mathbf{0}, \Sigma = 0.10I_2)$ .
- Homoscedastic normal noise, std = 0.20:  $P(\xi) = \mathcal{N}(\mu = \mathbf{0}, \Sigma = 0.20I_2)$ .
- Heteroscedastic normal noise, std = 0.2X:  $P(\xi; \mathbf{x}) = \mathcal{N}(\mu = \mathbf{0}, \Sigma = 0.2S)$  with  $S = [x_1, 0; 0, x_2]$  is the  $2 \times 2$  matrix with the given entries.
- Correlated normal noise:  $P(\xi) = \mathcal{N}(\mu = \mathbf{0}, \Sigma = 0.001S)$  with  $S = [2.5, -2; -2, 2.5]$ .
- Multiplicative noise model from the main text:  $\mathbf{x} \diamond \xi := \mathbf{x}\xi'$ , where  $\xi' \sim \mathcal{N}(\mu = \mathbf{1}, \Sigma = 0.07I_2)$ .

We see that MARS-NEI consistently outperforms the alternatives, except for the under homoskedastic noise with a large standard deviation of 0.2. In this setting, the EXP- $q$ NPAREGO is slightly ahead, which is not too surprising since the expectation and MVAR optimal designs happen to be in the same part of the solution space under this noise model.

In addition to the GMM problem, the Constrained Branin Currin problem was also ran under multiple noise models. Along with the heteroscedastic noise model used in the main text, we also studied it using a simple homoscedastic noise model:  $P(\xi) = \mathcal{N}(\mu = \mathbf{0}, \Sigma = 0.05I_2)$ . The plots for these are shown in Figure 11, demonstrating that MARS performs consistently under both noise models.

### I.6. Effect of $n_\xi$ on Optimization Performance

In a final side study, we analyze the effect of varying  $n_\xi$  on the optimization performance of the acquisition functions optimizing expectation and MVAR. We attempted to run the Constrained Branin Currin and Disc Brake experiments with  $n_\xi \in \{8, 16, 32, 64, 96, 128\}$  and included the results of the algorithms that successfully completed without running into scalability issues, such as getting an out-of-memory error.

The results are shown in Figure 12. We see that the methods using GPs (EXP- $q$ NPAREGO and MARS-NEI), do not scale beyond  $n_\xi = 64$  due to the cubic complexity (with respect to the number of points and  $n_\xi$ ) of posterior sampling (see Appendix D.2, the remaining methods avoid this issue since the RFF draws are deterministic functions). Overall, the results for MVAR methods show that too small of an  $n_\xi$  leads to poor performance, while increasing it much beyond our default value of  $n_\xi = 32$  does not yield any a significant benefit, at least in these problems. On Constrained Branin-Currin, we observe that the performance of the expectation methods degrade as  $n_\xi$  increases, which we attribute to these methods becoming more proficient at differentiating the expectation and MVAR optimal regions (which are not co-located), thus focusing their sampling away from the MVAR optimal region.

## J. Efficient Methods for Computing MVAR

Before going into the discussion, we note that the this section presents the MVAR computation for a random variable to be minimized. This simplifies the discussion by enabling the use of common terms such as CDF and quantile, since the

MVAR, as originally defined in Prékopa (2012), corresponds to the  $\alpha$  quantile of a random variable to be minimized. In the maximization setting studied in this paper, the MVAR, as defined in Definition 4.2, can be computed by first computing the MVAR of the negative of the random variable, as discussed here, then negating the result.

The existing algorithms for computing the MVAR (e.g., those presented by Prékopa (2012)) presume the availability of a cheap to evaluate CDF of the random variable of interest. In the general setting we consider, with  $f$  being an arbitrary function, such as a sample path of the GP, the random variable  $f(x \diamond \xi)$  (induced by  $\xi \sim P(\xi)$ ) does not admit a known CDF. Thus, to compute a QMC estimate of MVAR, we first need to compute the empirical CDF corresponding to the QMC samples of  $f(x \diamond \xi)$ .

Computing the empirical CDF of a random variable is a conceptually simple operation. All we need to do is to count the number of samples that are dominated by a given point and divide that by the total number of samples. Keeping all other objectives constant, the empirical CDF can be seen as a step function over the domain of the given objective that changes its value only at the points that correspond to one of the sample values of that objective. Since there are  $n_\xi$  samples, ignoring the possibility that some samples may have equal value for some objectives, this defines an  $M$ -dimensional grid with  $n_\xi$  points on each dimension, on which the empirical CDF can change its value. Thus, to fully compute the empirical CDF, we need to compute the number of samples that dominate a grid of  $n_\xi^M$  points, at a cost of  $\mathcal{O}(M)$  comparison per point, leading to a total  $\mathcal{O}(n_\xi^M M)$  cost for computing the empirical CDF.

Once the empirical CDF is computed, the MVAR set can easily be computed by taking the Pareto set of the points on the grid with a CDF greater than or equal to  $\alpha$ . This part of the computation, fortunately, has a lower complexity than the CDF computations.

A careful reader might have noticed that the MVAR (in the minimization setting) is bounded from below by the independent VAR of each objective and bounded from above by the maximum value observed for that given objective. We can leverage this fact to lower the cost of computing MVAR significantly. Instead of considering the full grid of  $n_\xi^M$  points, we can only compute the empirical CDF for the grid formed by the objective values that exceed the independent VAR of each objective, of which there are  $(1 - \alpha)n_\xi$ , reducing the complexity of MVAR computations to  $\mathcal{O}((1 - \alpha)^M n_\xi^M M)$ .

Within our code base, we provide two implementations for computing MVAR. One implementation is geared towards batched calculations with small to moderate  $n_\xi$ , and the other is geared towards less memory intensive calculations with large  $n_\xi$ . Both utilize efficient vectorized computations to exploit modern computing hardware. However, even with these highly optimized implementations, computing MVAR can easily become a bottleneck in a BO method, since MVAR has to be computed many times during acquisition optimization. Thus, the methods for direct optimization of MVAR that we present can still be prohibitively expensive unless the objective evaluations are significantly expensive.

## K. Multivariate Extensions of CVAR

CVAR is another popular risk measure that is commonly used with univariate random variables (Rockafellar and Uryasev, 2002). Similar to VAR, it has also been extended to the multi-variate case by various authors (e.g., Cousin and Di Bernardino (2014) and Meraklı and Küçükyavuz (2018)). However, these extensions in general do not admit a natural interpretation, whereas MVAR provides interpretable objective specifications that the objectives (under input noise) for a given design will meet with high-probability.






Figure 10: The effect of different noise models on the GMM problem.

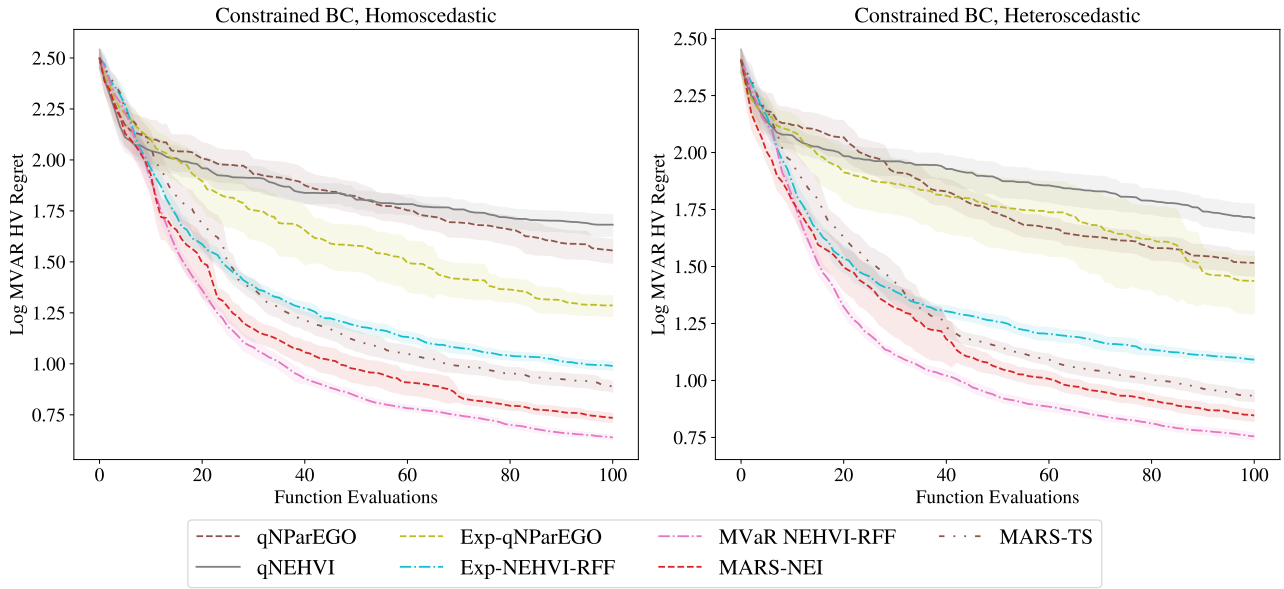


Figure 11: The effect of different noise models on the Constrained Branin Currin problem.




Figure 12: Final MVAR hypervolume regret obtained using different  $n_\xi$  on Constrained Branin-Currin and Disc Brake.