

---

# $\{\text{PF}\}^2\text{ES}$ : Parallel Feasible Pareto Frontier Entropy Search for Multi-Objective Bayesian Optimization

---

**Jixiang Qing**  
Ghent University -imec  
Ghent, Belgium  
Jixiang.Qing@Ugent.be

**Henry B. Moss**  
Secondmind.ai

**Tom Dhaene**  
Ghent University -imec  
Ghent, Belgium

**Ivo Couckuyt**  
Ghent University -imec  
Ghent, Belgium

## Abstract

We present Parallel Feasible Pareto Frontier Entropy Search ( $\{\text{PF}\}^2\text{ES}$ ) — a novel information-theoretic acquisition function for multi-objective Bayesian optimization supporting unknown constraints and batch queries. Due to the complexity of characterizing the mutual information between candidate evaluations and (feasible) Pareto frontiers, existing approaches must either employ crude approximations that significantly hamper their performance or rely on expensive inference schemes that substantially increase the optimization’s computational overhead. By instead using a variational lower bound,  $\{\text{PF}\}^2\text{ES}$  provides a low-cost and accurate estimate of the mutual information. We benchmark  $\{\text{PF}\}^2\text{ES}$  against other information-theoretic acquisition functions, demonstrating its competitive performance for optimization across synthetic and real-world design problems.

## 1 INTRODUCTION

The problem of optimizing multiple objectives is common across science, machine learning and industry and is known as *Multi-Objective Optimization* (MOO). Due to the conflicting nature of the multiple objectives, there is not one single solution and we instead seek the *Pareto Frontier* — a set of optimal solutions that provide a trade-off among different objectives. More precisely, the Pareto frontier consists of all solutions from which one cannot improve the performance of a specific objective without degrading another. For constrained MOO problems, the goal is to find the

*feasible Pareto frontier*, i.e., the Pareto frontier containing only points that satisfy the constraints.

*Multi-Objective Bayesian Optimization* (MOBO) (Yang et al., 2019; Daulton et al., 2021; Feliot et al., 2017; Qing et al., 2022a) is a well-established framework for solving expensive MOO problems with strict evaluation budgets. In order to achieve data-efficiency, BO uses cheap probabilistic *surrogate models* to predict the performance of not-yet-evaluated configurations. Heuristic search strategies, known as *acquisition functions*, then use the posterior belief of this model to direct subsequent objective function evaluations into promising areas of the search space. The choice of the acquisition function plays a crucial role on the performance of BO and so acquisition function design is an active research area. Existing common MOBO acquisition functions include (Yang et al., 2019; Daulton et al., 2020; Couckuyt et al., 2014; Abdolshah et al., 2018), to name a few.

A new class of acquisition functions for MOBO has recently arisen based on information theory. Motivated by state-of-the-art performance achieved by information-theoretic approaches for single-objective BO (Hernández-Lobato et al., 2014; Wang and Jegelka, 2017; Moss et al., 2021; Hvarfner et al., 2022), recent works (Hernández-Lobato et al., 2016; Garrido-Merchán and Hernández-Lobato, 2019; Belakaria and Deshwal, 2019; Suzuki et al., 2020; Belakaria et al., 2020; Fernández-Sánchez et al., 2020) have provided several information-theoretic approaches for MOO. These acquisition functions follow the intuitive goal of seeking to reduce the amount of uncertainty (as quantified by differential entropy (Hennig and Schuler, 2012)) held by our surrogate model about the Pareto frontier.

However, the full potential of multi-objective information-theoretic acquisition functions has not yet been realised, with existing work not demonstrating the state-of-the-art performance exhibited by their single-objective counterparts. This poor performance is due to the difficulty in providing a proper and efficient characterization of the mutual information between sampled observation and the Pareto

frontier. Existing approaches rely on either computationally expensive approximate inference schemes like expectation propagation (Hernández-Lobato et al., 2016; Garrido-Merchán and Hernández-Lobato, 2019) and density filtering (Fernández-Sánchez et al., 2020), or on making coarse assumptions about the structure of the Pareto front (Suzuki et al., 2020; Belakaria and Deshwal, 2019; Belakaria et al., 2020). Indeed, both types of approximation lead to poor empirical performance (Daulton et al., 2021).

In this work, we provide {PF}<sup>2</sup>ES, a new information-theoretic acquisition function for MOBO. Inspired by the recent work of (Poole et al., 2019; Takeno et al., 2022), we use a variational lower bound to propose a novel cheap, accurate and explainable approximation to the joint mutual information between (batches of) candidate evaluations and (feasible) Pareto fronts. Our primary contributions are summarised as follows:

- We introduce a new information-theoretic acquisition function for multi-objective optimization that **supports both constrained and unconstrained problems**.
- We propose an efficient parallelisation strategy  $q$ -{PF}<sup>2</sup>ES which provides effective **batch optimization** across batches of  $q$  points.
- We provide **theoretical links** between {PF}<sup>2</sup>ES and a common MOBO acquisition function: multi-objective Probability of Improvement (Yang et al., 2019; Hawe and Sykulski, 2007), the first such analysis across information-theoretic MOBO.
- We provide a new taxonomy of comparing different information theoretic acquisition functions focusing on the output space (see Section 2.3) through the **uncertainty calibration** of the Pareto frontier.
- We demonstrate {PF}<sup>2</sup>ES’s **competitive performance** against existing acquisition functions across a range of synthetic and real-life batch optimization problems.

## 2 PRELIMINARIES

### 2.1 Multi-Objective Optimization with Unknown Constraints

$$\begin{aligned} & \text{Maximize } \mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_M(\mathbf{x})) \\ & \text{s.t. } (g_1(\mathbf{x}) \geq 0, \dots, g_C(\mathbf{x}) \geq 0) := \mathbf{g}(\mathbf{x}) \geq \mathbf{0} \quad (1) \\ & \mathbf{x} \in \mathcal{X} \in \mathbb{R}^d \end{aligned}$$

We consider constrained multi-objective optimization (CMOO) problems, formally expressed as finding the **maximum** of a vector value function:  $\mathbf{f} = \{f_1, \dots, f_M\} : \mathbf{x} \rightarrow \mathbb{R}^M$  in a bounded design space  $\mathcal{X} \subset \mathbb{R}^d$  that need to consider a set of *unknown constraints*  $\mathbf{g} = \{g_1, \dots, g_C\}$  (i.e., the analytic formulation is unknown), where  $M$  and  $C$  represents

the number of objectives and constraints respectively. A candidate  $\mathbf{x}$  is *feasible* if  $\mathbf{g}(\mathbf{x}) \geq \mathbf{0}$ . The optimal comparison of different feasible candidates is determined through the following *ranking* mechanism: a feasible candidate  $\mathbf{x}$  is preferable to  $\mathbf{x}'$  in the sense that  $\forall j \in M : f_j(\mathbf{x}) \geq f_j(\mathbf{x}')$  and  $\exists j \in M : f_j(\mathbf{x}) > f_j(\mathbf{x}')$ . This specific ranking strategy is termed as *dominance* ( $\succ$ ) and described as  $\mathbf{f}(\mathbf{x})$  dominates  $\mathbf{f}(\mathbf{x}')$ :  $\mathbf{f}(\mathbf{x}) \succ \mathbf{f}(\mathbf{x}')$ . A feasible candidate input  $\mathbf{x}$  is called a *Pareto optimal candidate* if there do not exist any other feasible candidates in the design space that are able to dominate it; the set containing all the Pareto optimal candidates is called *Pareto set* and denoted as  $\mathbf{x}_{\mathcal{F}}$ . The goal of CMOO is to efficiently identify the *feasible Pareto frontier*<sup>1</sup>  $\mathcal{F} := \{\mathbf{f}(\mathbf{x}) | \forall \mathbf{x}' \in \mathcal{X} \setminus \mathbf{x}_{\mathcal{F}} \text{ s.t. } \mathbf{g}(\mathbf{x}') \geq \mathbf{0}, \exists \mathbf{x} \in \mathbf{x}_{\mathcal{F}} \text{ s.t. } \mathbf{f}(\mathbf{x}) \succ \mathbf{f}(\mathbf{x}')\}$ , which is constructed by the Pareto set.

### 2.2 Multi-Objective Bayesian Optimization

For many real-world problems, the exact form of  $\mathbf{f}$  and  $\mathbf{g}$  is unknown, and the evaluation of the objective functions and constraints functions  $\mathbf{h}_{\mathbf{x}} := \{f_{1_{\mathbf{x}}}, \dots, f_{M_{\mathbf{x}}}, g_{1_{\mathbf{x}}}, \dots, g_{C_{\mathbf{x}}}\}$  at an input location  $\mathbf{x}$  is computationally expensive. In these settings, it is crucial to restrict the total number of observations required to find the Pareto frontier  $\mathcal{F}$ .

In order to achieve data efficiency, BO (Frazier, 2018; Shahriari et al., 2015; Garnett, 2022) leverages a probabilistic *surrogate model* as a computationally efficient approximation of the original expensive objective function. Within this research we focus our discussion on the standard Gaussian Process (GP) (Rasmussen, 2003) framework. Given  $N_{tr}$  expensive observations  $D = \{\{\mathbf{x}\}_i, \{\mathbf{h}\}_i\}_{i=1}^{N_{tr}} := \{\mathbf{X}_{tr}, \mathbf{H}\}$ , a GP is able to provide a Gaussian posterior distribution of any not-yet-evaluated  $\mathbf{h}$ . Here we follow a common and most generic assumption that each objective (and constraint) are statistically independent. Consequently, the posterior distribution of  $i$ 'th outcome  $\mathbf{h}^i$  at unknown candidate(s)  $\mathbf{x}$  is a (multivariate) Gaussian with mean and (co)variance defined as:

$$\begin{aligned} \mathbf{m}^i(\mathbf{x}|D) &= \mathbf{k}(\mathbf{x})^T \mathbf{K}^{-1} \mathbf{H}^i \\ \mathbb{V}^i(\mathbf{x}|D) &= \mathbf{k}(\mathbf{x}, \mathbf{x}) - \mathbf{k}(\mathbf{x})^T \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}) \end{aligned} \quad (2)$$

where  $\mathbf{k}(\cdot, \cdot)$  represents kernel and  $\mathbf{K}$  is the Gram matrix building upon existing input  $\mathbf{X}_{tr}$ ,  $\mathbf{H}^i$  represents the  $i$ th output of training data. See Rasmussen (2003) for a comprehensive introduction to GPs.

To guide the search into promising areas of the search space and provide highly efficient optimization, BO relies on an *acquisition function* that uses the posterior belief to predict

<sup>1</sup>This formulation can be easily adapted to MOO problem without constraints  $\mathbf{g}$ . Since we aim to tackle both MOO and CMOO problems using our acquisition function, we overload  $\mathcal{F}$  (as well as the term *Pareto frontier*) to represent both feasible Pareto frontier in CMOO and Pareto frontier in the MOO problem.

the utility of making an evaluation at any candidate input. The original expensive objective functions and constraints are then evaluated at the input with the largest predicted utility and the resulting evaluation is used to update the surrogate model. This model updating, acquisition function building, and objective function evaluation pattern iterates until a predefined stopping criterion has been met.

### 2.3 Information-Theoretic Multi-Objective Bayesian Optimization

One increasingly popular class of acquisition functions are those based on the now well-established information-theoretic framework. Here, we seek evaluations that provide maximal information about a given *target quantity*. In the context of MOO, the target is to reduce our uncertainty about the set of optimal (feasible) trade-offs (i.e., Pareto set or Pareto frontier). However, this can be formulated in two distinct ways. First, we can use the *input-space formulation* (Hernández-Lobato et al., 2016; Garrido-Merchán and Hernández-Lobato, 2019) and calculate our uncertainty over where the optimal trade-offs lie in our search space, i.e., the Pareto set. More recently, *output-space* methods (Belakaria and Deshwal, 2019; Belakaria et al., 2020; Fernández-Sánchez et al., 2020) have been proposed that seek to reduce the uncertainty in the (feasible) Pareto frontier  $\mathcal{F}$  directly. As the discrete Pareto frontier is just an  $M \cdot |\mathcal{F}|$ -dimensional quantity in contrast to the  $d \cdot |\mathcal{F}|$ -dimensional quantities in the Pareto set, the output-space method enjoys simpler calculations and easier numerical approximations (at least when  $d > M$ ). Unfortunately, although existing output-based MOO methods are much cheaper than their input-based alternatives (Belakaria and Deshwal, 2019), they all employ coarse approximations that hamper their performance and hinder their interpretation. For example, existing approaches employ approximations like assuming factorized conditional probability distributions (Fernández-Sánchez et al., 2020), or using an overdone heuristic approximation to collapse  $\mathcal{F}$  to its outcome-wise max (Belakaria et al., 2021).

## 3 {PF}<sup>2</sup>ES FOR CONSTRAINED MULTI-OBJECTIVE OPTIMIZATION

The goal is to determine *where to sample  $\mathbf{x}$  in order to learn as much as possible about the Pareto frontier  $\mathcal{F}$* ? Under the information-theoretic framework, the act of learning is characterized by the Shannon mutual information  $I(\cdot, \cdot)$  (Cover, 1999) between our target quantity  $\mathcal{F}$ , and the possible distribution of the concatenated objective-constraint observations  $\mathbf{h}_x$  according to our GP surrogate models:

$$I(\mathcal{F}; \mathbf{h}_x) \quad (3)$$

Current output-space methods for information-theoretic MOO (Belakaria et al., 2020; Suzuki et al., 2020) (i.e.,

those that seek to learn the Pareto Frontier, as discussed in Section 2.3) recast the mutual information (Eq. 3's) as the reduction in differential entropy  $\mathbb{H}$  of  $\mathcal{F}$  provided by the candidate evaluation, i.e., using the expansion  $I(\mathcal{F}; \mathbf{h}_x) = \mathbb{H}(\mathcal{F}) - \mathbb{H}(\mathcal{F}|\mathbf{h}_x)$ . Following this formulation, the primary difficulty is providing a reliable and efficient approximation of the differential entropy  $\mathbb{H}(\mathcal{F}|\mathbf{h}_x)$ .

### 3.1 A Variational Lower Bound of the Mutual Information $I(\mathcal{F}; \mathbf{h}_x)$

For {PF}<sup>2</sup>ES we avoid the difficulties of approximating the differential entropy by directly approximating the mutual information itself (Eq. 3). In particular, we follow the ideas of (Poole et al., 2019; Takeno et al., 2022) and replace the mutual information with the following variational lower bound:

$$\begin{aligned} I(\mathcal{F}; \mathbf{h}_x) &= \int_{\mathcal{F}} \int_{\mathbf{h}_x} p(\mathbf{h}_x, \mathcal{F}) \log \frac{p(\mathbf{h}_x, \mathcal{F})}{p(\mathbf{h}_x)p(\mathcal{F})} d\mathcal{F} d\mathbf{h}_x \\ &= \mathbb{E}_{\mathcal{F}} \left[ \int_{\mathbf{h}_x} p(\mathbf{h}_x|\mathcal{F}) \log \frac{p(\mathbf{h}_x|\mathcal{F})q(\mathbf{h}_x|\mathcal{F})}{p(\mathbf{h}_x)q(\mathbf{h}_x|\mathcal{F})} d\mathbf{h}_x \right] \\ &= \mathbb{E}_{\mathcal{F}} \left[ \int_{\mathbf{h}_x} p(\mathbf{h}_x|\mathcal{F}) \log \frac{q(\mathbf{h}_x|\mathcal{F})}{p(\mathbf{h}_x)} d\mathbf{h}_x + \right. \\ &\quad \left. D_{KL}(p(\mathbf{h}_x|\mathcal{F})||q(\mathbf{h}_x|\mathcal{F})) \right] \\ &\geq \mathbb{E}_{\mathcal{F}} \left[ \int_{\mathbf{h}_x} p(\mathbf{h}_x|\mathcal{F}) \log \frac{q(\mathbf{h}_x|\mathcal{F})}{p(\mathbf{h}_x)} d\mathbf{h}_x \right] \end{aligned} \quad (4)$$

where  $D_{KL}$  represents the KL-divergence and the density  $q(\mathbf{h}_x|\mathcal{F})$  is a variational approximation of the ground truth conditional distribution  $p(\mathbf{h}_x|\mathcal{F})$ . The inequality in the final line of Eq. 4 is due to the non-negativity of the KL-divergence. As the gap between the true mutual information and our lower bound can be explicitly seen as the KL-divergence between the variational approximation and ground truth conditional distribution, the suitability of this lower bound depends on our ability to build a reasonable approximation  $q(\mathbf{h}_x|\mathcal{F}) \approx p(\mathbf{h}_x|\mathcal{F})$ .

We now explain how to build an effective variational approximation  $q(\mathbf{h}_x|\mathcal{F})$ . Given a realization of the Pareto frontier  $\mathcal{F}$  (illustrated in Fig. 1a for an unconstrained problem), the *output space*  $\mathbb{R}^{M+C}$  can be partitioned into two complementary regions  $A(\mathcal{F})$  and  $\bar{A}(\mathcal{F})$ , i.e.,  $A(\mathcal{F}) \cup \bar{A}(\mathcal{F}) = \mathbb{R}^{M+C}$ , where  $A(\mathcal{F}) \subset \mathbb{R}^{M+C}$  is the *feasible non-dominated region*  $\{\{\mathbf{f}_x, \mathbf{g}_x\} : \exists \mathbf{x}^* \in \mathbf{x}_{\mathcal{F}} \text{ s.t. } \mathbf{f}_{x^*} \succ \mathbf{f}_x, \mathbf{g}_x \geq \mathbf{0}\}$  in the output space. We then set  $q(\mathbf{h}_x|\mathcal{F})$  as (an extension of the

<sup>2</sup>We note the output space is actually defined in a bounded space formulated by two extreme points: *ideal point* and *anti-ideal point* as shown in Fig. 1a to allow a practical objective space partition, however, in {PF}<sup>2</sup>ES we set these two extreme points to a constant vector with large values (i.e., 1e20) and omit the notation of them afterwards.




Figure 1: Partitioning the output space (bounded by the ideal point and anti-ideal point) according to a Pareto frontier: a) Given  $\mathcal{F}$ : Partitioning into a dominated region  $\bar{A}$  and a non-dominated region  $A$ . b) Given a discrete Pareto frontier approximation  $\tilde{\mathcal{F}}$ : Partitioning the output space into an approximated dominated region  $\tilde{\bar{A}}$  and non-dominated region  $\tilde{A}$  by hypervolume based decomposition  $\mathcal{P}$ . c) A partition of the output space given a discretized feasible Pareto frontier  $\tilde{\mathcal{F}}$  for a constrained optimization problem.

concept of *Pareto frontier truncated normal distribution* (Suzuki et al., 2020):

$$q(\mathbf{h}_x | \mathcal{F}) = \begin{cases} \frac{p(\mathbf{h}_x)}{Z_{\bar{A}}} & \mathbf{h}_x \in \bar{A} \\ 0 & \text{else} \end{cases} \quad (5)$$

where  $Z_{\bar{A}} := \int_{\bar{A}} p(\mathbf{h}_x) d\mathbf{h}_x = 1 - Z_A$  is the probability of  $\mathbf{h}_x$  be in  $\bar{A}$ , which is constructed based on the frontier realization  $\bar{A}(\mathcal{F})$ . The choice of variational approximation  $q$  is motivated by that, given a fixed  $\mathcal{F}$ , for any  $\mathbf{x}, \mathbf{h}_x$ ; we assign  $q(\mathbf{h}_x | \mathcal{F}) = 0$  for  $\mathbf{h}_x$  to be feasible and non-dominated (i.e., there exists  $\mathbf{x}' \in \mathbf{x}_{\mathcal{F}}$  such that  $\mathbf{f}_{\mathbf{x}} \succ \mathbf{f}(\mathbf{x}')$  and  $\mathbf{g}_{\mathbf{x}} > \mathbf{0}$ ). This of course assumes that we are able to obtain the expensive observations without noise which is a common scenario (e.g., (Vakili et al., 2020)).

By substituting our approximate conditional distribution (Eq. 5) into Eq. 4, the variational information lower bound has the following simple expression<sup>3</sup>

$$\begin{aligned} I(\mathcal{F}; \mathbf{h}_x) &\geq \mathbb{E}_{\mathcal{F}} \left[ \int_{\mathbf{h}_x \in \bar{A}(\mathcal{F})} p(\mathbf{h}_x | \mathcal{F}) d\mathbf{h}_x \cdot \log \frac{1}{Z_{\bar{A}}(\mathcal{F})} \right] \\ &= -\mathbb{E}_{\mathcal{F}} [\log(1 - Z_A(\mathcal{F}))] = \alpha_{\{\text{PF}\}^2\text{ES}} \end{aligned} \quad (6)$$

which leads to the proposed {PF}<sup>2</sup>ES acquisition function.

**Explainability of {PF}<sup>2</sup>ES:** Given the {PF}<sup>2</sup>ES's expression in Eq. 6, we are able to link it with the common Multi-Objective Probability of Improvement (MOPI) (Yang et al., 2019; Hawe and Sykulski, 2007). MOPI is the multi-objective version of Probability of Improvement (Kushner, 1964) and measures utility by the probability that a new candidate will locate in the non-dominated region  $A(\mathcal{F})$ , which

<sup>3</sup>By convention (Cover, 1999)  $0 \log 0 = 0$  and so  $\int_{\mathbf{h}_x \in A} p(\mathbf{h}_x | \mathcal{F}) \log p(\mathbf{h}_x | \mathcal{F}) d\mathbf{h}_x = 0$ . See Appendix A of (Takeno et al., 2022) for a discussion.

can be constructed based on a Pareto frontier  $\mathcal{F}$ . MOPI can be adapted to the CMOO setting, where it can be multiplied with a Probability of Feasibility (PoF) term ( $\alpha_{\text{MOPI}} \times \alpha_{\text{PoF}}$ ) following the approach of (Hawe and Sykulski, 2008; Gardner et al., 2014). This, under the assumption that  $\mathbf{f}$  and  $\mathbf{g}$  is statistically independent, is equivalent to the probability that a new candidate is feasible and located in the non-dominated region.

The link between {PF}<sup>2</sup>ES and MOPI can be established through the following remark (see Appendix. A for proof):

**Remark 1.** *The following acquisition functions lead to the same maximal candidate  $\mathbf{x} = \arg \max_{\mathbf{x} \in \mathcal{X}} \alpha(\mathbf{x})$ :*

1. MOO

- 1.1.  $\alpha_{\text{MOPI}}$  using  $\mathcal{F}$  as Pareto frontier
- 1.2.  $\alpha_{\{\text{PF}\}^2\text{ES}}$  with  $|\tilde{\mathbf{F}}| = 1$  and use the same  $\mathcal{F}$  as 1.1.

2. Constrained MOO when

- 2.1.  $\alpha_{\text{MOPI}} * \alpha_{\text{PoF}}$  using  $\mathcal{F}$  as reference Pareto frontier
- 2.2.  $\alpha_{\{\text{PF}\}^2\text{ES}}$  with  $|\tilde{\mathbf{F}}| = 1$  and use the same  $\mathcal{F}$  as 2.1.

where  $\tilde{\mathbf{F}}$  is defined in Eq. 8. We note this linkage provides insights into the empirical performance of {PF}<sup>2</sup>ES as well.

## 4 PRACTICAL CALCULATION OF {PF}<sup>2</sup>ES AND ITS PARALLELIZATION

Practical calculation of the variational lower bound (Eq. 6) requires some additional steps. The primary difficulty is that the region  $A$  (for any given Pareto frontier  $\mathcal{F}$ ) does not have an analytical form. A popular approach (Belakaria and Deswal, 2019; Suzuki et al., 2020; Hernández-Lobato et al., 2016; Garrido-Merchán and Hernández-Lobato, 2019)



is to use a finite representation of the Pareto frontier  $\tilde{\mathcal{F}} := \{\mathbf{f} | \mathbf{f} \in \mathcal{F}\}$  s.t.  $|\tilde{\mathcal{F}}| < \infty$ . Based on  $\tilde{\mathcal{F}}$ , a hypervolume-based output space partitioning strategy (Lacour et al., 2017; Couckuyt et al., 2012), denoted as  $\mathcal{P}$ , can be used to partition the output space into  $N_p$  disjoint hypercubes, thus obtaining an approximation of  $A$  denoted by *hypervolume based approximated non-dominated region*:

$$A \approx \tilde{A}(\tilde{\mathcal{F}}) = \mathcal{P}(\tilde{\mathcal{F}}) = \{\tilde{A}_1, \tilde{A}_2, \dots, \tilde{A}_{N_p}\} \quad (7)$$

which is the complement of  $\tilde{\tilde{A}}(\tilde{\mathcal{F}}) := \{\mathbf{f}_x \in \mathbb{R}^M | \exists \mathbf{f} \in \tilde{\mathcal{F}}, \mathbf{f}_x \preceq \mathbf{f}\}$ . Such partitioning is illustrated in Fig. 1b for unconstrained MOO problems. For CMOO problem,  $\tilde{A}(\tilde{\mathcal{F}})$  can be obtained by first partitioning the non-dominated region in the *objective space*:  $\mathbb{R}^M$  and then concatenated with the constraint space where  $\mathbf{g} \geq \mathbf{0}$ , see Fig. 1c for an illustration.

Given  $\tilde{A}$ , we are able to reach the following analytically tractable approximation of Eq. 6:

$$-\mathbb{E}_{\tilde{\mathcal{F}}} [\log(1 - Z_{A(\tilde{\mathcal{F}})})] \approx -\frac{1}{|\tilde{\mathbf{F}}|} \sum_{\tilde{\mathcal{F}} \in \tilde{\mathbf{F}}} \left[ \log \left( 1 - \sum_{i=1}^{N_p} Z_{\tilde{A}_i} \right) \right] \quad (8)$$

where  $\tilde{\mathbf{F}}$  is a set of sampled Pareto frontiers providing a Monte Carlo (MC) approximation of the outer expectation in Eq. 6. We build this MC sample of  $\tilde{\mathcal{F}}$  by using a standard multi-objective optimization algorithm (e.g., NSGAII (Deb et al., 2002)) on samples from the GP’s spectral posterior (also known as sampling GP posterior via its Fourier features (Rahimi and Recht, 2007)), which is a common strategy utilized in acquisition functions for MOO (Belakaria and Deshwal, 2019; Suzuki et al., 2020; Daulton et al., 2022).  $Z_{\tilde{A}_i}$  is the probability that  $\mathbf{h}_x$  is in the  $i$ ’th hypercube  $\tilde{A}_i$  and can be calculated as (Suzuki et al., 2020):

$$Z_{\tilde{A}_i} = \prod_{k=1}^{M+C} \left( \Phi \left( \frac{\tilde{A}_{i_u}^k - \mathbf{m}_x^k}{\boldsymbol{\sigma}_x^k} \right) - \Phi \left( \frac{\tilde{A}_{i_l}^k - \mathbf{m}_x^k}{\boldsymbol{\sigma}_x^k} \right) \right). \quad (9)$$

where  $\Phi(\cdot)$  represents the cumulative density function of the standard normal distribution,  $\tilde{A}_{i_u}^k, \tilde{A}_{i_l}^k$  represents the  $k$ ’th dimensional upper and lower bound of the  $i$ ’th hypercube, respectively.  $\mathbf{m}_x^k, \boldsymbol{\sigma}_x^k$  represent the  $k$ ’th dimensional GP posterior mean and standard deviation, respectively.

Unfortunately, due to the subtlety of our employed partitioning strategy, we cannot simply use the r.h.s of Eq. 8 as  $\tilde{\alpha}_{\{\text{PF}\}^2\text{ES}}$ . In particular, the hyper-volume based partition  $\mathcal{P}$ , operating on the finite representation  $\tilde{\mathcal{F}}$ , means that  $\tilde{A}$  is an **overestimation** of the non-dominated region  $A$ . Therefore, we are no longer guaranteed to satisfy the lower bound property in Eq. 6 since  $Z_{A(\tilde{\mathcal{F}})} < Z_{\tilde{A}(\tilde{\mathcal{F}})}$ . More importantly, calculating  $\{\text{PF}\}^2\text{ES}$  based on the above  $\tilde{A}$  results in the same empirical *clustering issue* observed in MOPI, i.e. that the probability of improvement rewards tiny improvements of the Pareto frontier approximation  $\tilde{\mathcal{F}}$  if they are

achieved with high confidence (Emmerich et al., 2020) (see Appendix. B for a detailed elaboration).

In order to mitigate this clustering issue whilst ensuring that our acquisition function remains true to its motivation as a lower bound to the mutual information, we propose a small change to how we construct our approximation of  $A$ . Inspired by the common empirical strategy employed in MOPI and PI acquisition functions ((Wang et al., 2016; Emmerich et al., 2020; Kushner, 1964)), we introduce a small positive penalization vector  $\boldsymbol{\varepsilon} = [\varepsilon_1, \dots, \varepsilon_M]$  to shift the discrete Pareto frontier approximation:  $\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}} := \tilde{\mathcal{F}} + \boldsymbol{\varepsilon}$  to avoid any overestimation of  $A$  resulting from the hypervolume decomposition. In Appendix. D, we show that it is possible to construct an  $\boldsymbol{\varepsilon}$  that ensures that our acquisition function remains a lower bound. However, building such an epsilon is costly, so we also propose the practical solution of setting  $\boldsymbol{\varepsilon}$  through a simple heuristic that achieves very similar empirical performance (see Appendix. D). In particular, we set epsilon as  $\varepsilon_k = c \left( \max(\tilde{\mathcal{F}}^k) - \min(\tilde{\mathcal{F}}^k) \right) \forall k \in M, \tilde{\mathcal{F}} \in \tilde{\mathbf{F}}$ , i.e. a proportion of the maximal total variation observed over the Pareto front. We show empirically that  $\{\text{PF}\}^2\text{ES}$ ’s performance is robust to the choice of  $c$  through sensitivity analysis (Appendix. F). We set  $c = 0.04$  for all our experiments.

Combining all these steps, we can finally state our proposed acquisition function as

$$\tilde{\alpha}_{\{\text{PF}\}^2\text{ES}} = -\frac{1}{|\tilde{\mathbf{F}}|} \sum_{\tilde{\mathcal{F}} \in \tilde{\mathbf{F}}} \left[ \log \left( 1 - \sum_{i=1}^N Z_{\tilde{A}_i(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}})} \right) \right]. \quad (10)$$

See Algorithm 1 in Appendix. D for a high-level algorithmic summary of our calculation strategy.

#### 4.1 Parallelisation of $\{\text{PF}\}^2\text{ES}$

For many practical MOO problems, it is common to have parallel or distributed evaluation resources available, i.e. we recommend  $q > 1$  evaluations during each BO step. We now propose an extension of  $\{\text{PF}\}^2\text{ES}$  that allows it to be used for such batch optimizations. This time, the principled question is *at which  $q$  points should we sample simultaneously to learn as much as possible about  $\mathcal{F}$ ?*

Consider the random variable  $\mathbf{h}_X = \{\mathbf{h}_{x_1}, \dots, \mathbf{h}_{x_q}\}$ , i.e., the possible observations that could arise from a parallel evaluation of a batch of  $q$  points. In this scenario, we want to allocate a batch of  $q$  points  $\mathbf{X}$  that provide the most mutual info about the (feasible) Pareto frontier. Following the same derivation as Eq. 4 - 6, we are able to have our proposed extension of  $\{\text{PF}\}^2\text{ES}$  for batch design:

$$\tilde{\alpha}_{q-\{\text{PF}\}^2\text{ES}} \approx -\frac{1}{|\tilde{\mathbf{F}}|} \sum_{\tilde{\mathcal{F}} \in \tilde{\mathbf{F}}} \left[ \log \left( 1 - Z_{\tilde{A}_q(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}})} \right) \right], \quad (11)$$

where  $Z_{\tilde{A}_q} := \int_{\tilde{A}_q} p(\mathbf{h}_X) d\mathbf{h}_X$  (see Eq. 17). i.e., the probability that there exists at least one batch element  $\mathbf{h}_{x_i} \in \mathbf{h}_X$

such that  $\mathbf{h}_{x_i} \in \tilde{A}(\tilde{\mathcal{F}}_{\boldsymbol{\epsilon}}) = \{\tilde{A}_1(\tilde{\mathcal{F}}_{\boldsymbol{\epsilon}}), \dots, \tilde{A}_N(\tilde{\mathcal{F}}_{\boldsymbol{\epsilon}})\}$ . Calculating the probability of  $Z_{\tilde{A}}$  involves a complex interaction between joint batch points within the complex region  $\tilde{A}$ ; we hence propose to approximate this probability through an MC approximation (a detailed derivation is provided in Appendix. C.1):

$$\begin{aligned} Z_{\tilde{A}_q} &\approx \frac{1}{N_{MC}} \sum_{j=1}^{N_{MC}} \left( \bigcup_{i=1}^{N_p} \left( \bigcup_{l=1}^q \left( \prod_{k=1}^{M+C} \left( \mathbb{1}(A_{i_l}^k \leq \mathbf{h}_{x_{ij}}^k \leq A_{i_l}^k) \right) \right) \right) \right) \\ &\approx \frac{1}{N_{MC}} \sum_{j=1}^{N_{MC}} \left( \max_i \left( \max_l \left( \prod_{k=1}^{M+C} \left( \sigma\left(\frac{\mathbf{h}_{x_{ij}}^k - A_{i_l}^k}{\tau}\right) \right) \right) \right) \right) \\ &\quad \left( \sigma\left(\frac{A_{i_l}^k - \mathbf{h}_{x_{ij}}^k}{\tau}\right) \right) \right), \end{aligned} \tag{12}$$

where  $\mathbf{h}_{x_{ij}}^k$  represents the  $k$ 'th output dimensionality of  $j$ 'th MC sample of  $l$ 'th batch point  $\mathbf{h}_{x_i}$ , and  $\sigma(\cdot)$  is the sigmoid function,  $N_{MC}$  represents the MC sample size. To ensure the differentiability of our acquisition function, we follow a common strategy in BO (Wilson et al., 2018; Daulton et al., 2020) to relax the categorical event imposed by  $\mathbb{1}(\cdot)$  in the first line of Eq. 12 by replacing it with a sigmoid function and a small non-negative temperature parameter  $\tau$ . We note Eq. 12 can be explained as calculating the event that **whether any batch element is within  $\tilde{A}$** , averaged through  $N_{MC}$  sample numbers. For joint sampling of the batch outcome  $\mathbf{h}_{\mathbf{X}}$ , we use the reparameterization trick in combination with a sample average approximation (SAA) (Balandat et al., 2020) to perform a continuous acquisition function optimization process, where the base sample of SAA has been generated through a randomized quasi-Monte Carlo, details of the approximation and a demonstration of q-{PF}<sup>2</sup>ES are provided in Appendix. C.2.

We stress the existence of remark. 1 confirms that our MC-based parallelisation strategy (Eq. 12) can be directly applied to the MOPI acquisition function as well.

**Complexity Analysis** Lastly, we provide complexity analysis in Appendix G, A run-time comparison with other acquisition functions utilized in the next section is also conducted in Appendix H.

## 5 EXPERIMENTAL VALIDATION OF {PF}<sup>2</sup>ES

We now present the empirical performance of {PF}<sup>2</sup>ES across constrained and unconstrained synthetic benchmarks and real-life application problems. The primary focus of these experiments is to compare our methods with other output-based entropy search methods, we have also included the standard EHVI acquisition functions as a performance

reference <sup>4</sup>.

**Sequential MOO problem** We compare against PFES (Suzuki et al., 2020), MESMO (Belakaria et al., 2020), EHVI (Yang et al., 2019), PESMO (Hernández-Lobato et al., 2016), as well as providing random search as a common baseline.

**Parallel MOO problem** We compare against qEHVI (Daulton et al., 2020), {PF}<sup>2</sup>ES with a fantasizing method (i.e., the Kriging Believer (KB) method (Ginsbourger et al., 2010)) and random search.

**Sequential CMOO problem** We compare against EHVI-PoF (the multiplication of EHVI with PoF as a common strategy for CMOO (e.g., Martínez-Frutos and Herrero-Pérez (2016))), MESMOC (Belakaria et al., 2021), MESMOC+ (Fernández-Sánchez et al., 2020) and random search.

**Parallel CMOO problem** We compare against qEHVI (Daulton et al., 2020)<sup>5</sup>, PPESMOC (Garrido-Merchán and Hernández-Lobato, 2020), random Search and {PF}<sup>2</sup>ES using a fantasize method (Ginsbourger et al., 2010).

For PESMO, MESMOC+ and PPESMOC, we used the public implementation provided by the papers. The settings of all acquisition functions are detailed in Appendix. E.3. For the surrogate models, we build GPs with Matérn 5/2 kernels using maximum a posteriori estimates for the kernel parameters. For optimizing the acquisition functions, we use a multi-start L-BFGS-B optimizer starting from the  $\min(10 \times qd, 100)$  <sup>6</sup> best locations from 5000 random starting locations. For the hyperparameter of the acquisition function, we empirically choose  $\tau = 1e - 3$ . For sampling the (feasible) Pareto frontier  $\tilde{\mathcal{F}}$ , we use the open-source NS-GAII optimizer in PyMOO (Blank and Deb, 2020), where constraints are handled by the *parameter-less approach* (Deb and Agrawal, 1999). For all our information-theoretic acquisition functions, we use five sampled Pareto frontiers. All of the following experiments start with  $2d + 1$  random sampled initial points. For the recommendation of optimal candidates, we use the out-of-sample strategy (performing recommendation based on the GP model) as a common approach in the information-theoretic acquisition function (see Appendix. E.2 for details).

<sup>4</sup>Our code is available at [https://github.com/TsingQAQ/trieste/tree/PF2ES\\_preview\\_notebook](https://github.com/TsingQAQ/trieste/tree/PF2ES_preview_notebook).

<sup>5</sup>We note that the EHVI acquisition function typically requires a reference point to calculate, since the setting of reference point can introduce bias hence performance difference, in our numerical experiments we assume such a **reference point is unknown beforehand and need to be calculated dynamically in each iteration**, we refer to Appendix. E.1 on how such a reference point is specified.

<sup>6</sup>The upper limit is set to restrict the computation cost of multi-start L-BFGS-B grow continuously with  $q$  and  $d$ , especially for higher input dimensionality and large batch size, see Fig. 12 in Appendix H.

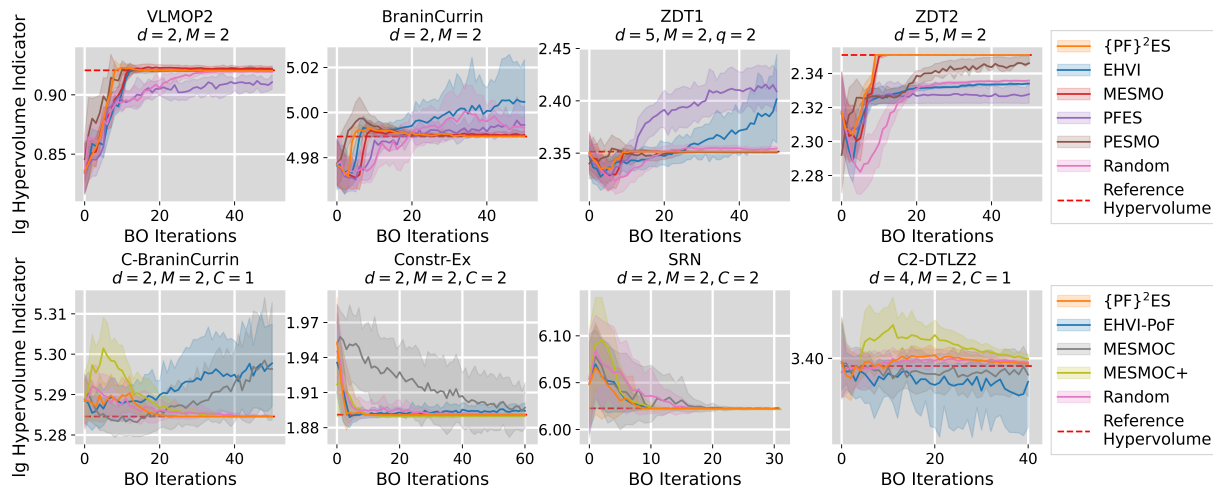


Figure 2: Uncertainty Calibration of inferred Pareto Frontier during BO iterations.

First, we consider a suite of popular synthetic benchmarks. Details of the benchmark functions, the reference point settings, and benchmark results on in-sample recommendations are provided in Appendix. E.1, I.1 respectively.

### 5.1 Uncertainty Calibration of Inferred Pareto Frontier

To help discern between  $\{PF\}^2ES$  and other information-theoretic acquisition functions that choose points with the overarching goal of reducing uncertainty in the Pareto frontier, it is natural to include a performance metric that directly measures this property.

Therefore, inspired by an approach in active learning (Qing et al., 2021), we propose an uncertainty calibration procedure on a *hypervolume based Pareto frontier indicator distribution* that reflects the uncertainty of GPs w.r.t the Pareto frontier. More specifically, the hypervolume indicator function  $I(\cdot)$  (Guerreiro et al., 2020) is applied to each sampled Pareto frontier  $\hat{\mathcal{F}}$  from the GPs. This way, a one-dimensional distribution of hypervolume indicators is obtained. Note that if the sampled Pareto frontiers have converged to the (finite approximation of the) Pareto frontier  $\mathcal{F}$ , the distribution will collapse to the hypervolume indicator of the (finite approximation of the) Pareto frontier  $\mathcal{F}$ . Hence, the expectation and uncertainty of the hypervolume based Pareto frontier indicator distribution can serve as a representation of the Pareto frontier estimation accuracy as well as the uncertainty based on GPs, respectively.

The proposed uncertainty calibration metric is illustrated in Fig. 2. For each BO iteration, the Pareto frontier indicator distribution  $I(\hat{\mathcal{F}}(GP))$  is approximated by 10 MC samples of  $\hat{\mathcal{F}}$  and we present the aggregate results across 30 repetitions of experiments by their median and 10-90 percentile. It can be seen that  $\{PF\}^2ES$  achieved the fastest

convergence to the reference hypervolume, implying low uncertainty about the Pareto frontier indicator distribution.

### 5.2 Multi-Objective Bayesian Optimization

#### Synthetic Problems

We now present our main experimental results. The log-hypervolume difference achieved across 30 repetitions for each algorithm is reported in Fig. 3. We see that  $\{PF\}^2ES$  leads to competitive performance in both MOO and CMOO problems in the case of sequential and batch sampling. Note that the performance of EHVI, often regarded as the gold standard, relies on setting a reasonable reference point. For problems where there is no prior information known about the reference point location, dynamic reference updating strategies must be used. For example, EHVI performs poorly on C2DTLZ2 which has disjoint Pareto frontiers (Deb, 2019). Additional experimental results on larger batch sizes and larger output dimensionalities are provided in Appendix. I.2 and I.3 respectively.

#### Four Bar Truss Design

We consider a real-life mechanical design problem (MOO) (Tanabe and Ishibuchi, 2020; Dauert, 1993), where we seek a four elastic truss structure system with small structure volume whilst having small deformation under external forces. The cross-section areas of the four bars in the truss system are chosen as the design variables for the unconstrained multi-objective optimization problem. Note the scale of the two objectives is significantly different for this problem.

The results are depicted in Fig. 4. For the sequential case,  $\{PF\}^2ES$  demonstrates the overall fastest convergence speed. While all acquisition functions tend to have similar performance in the batch scenario.

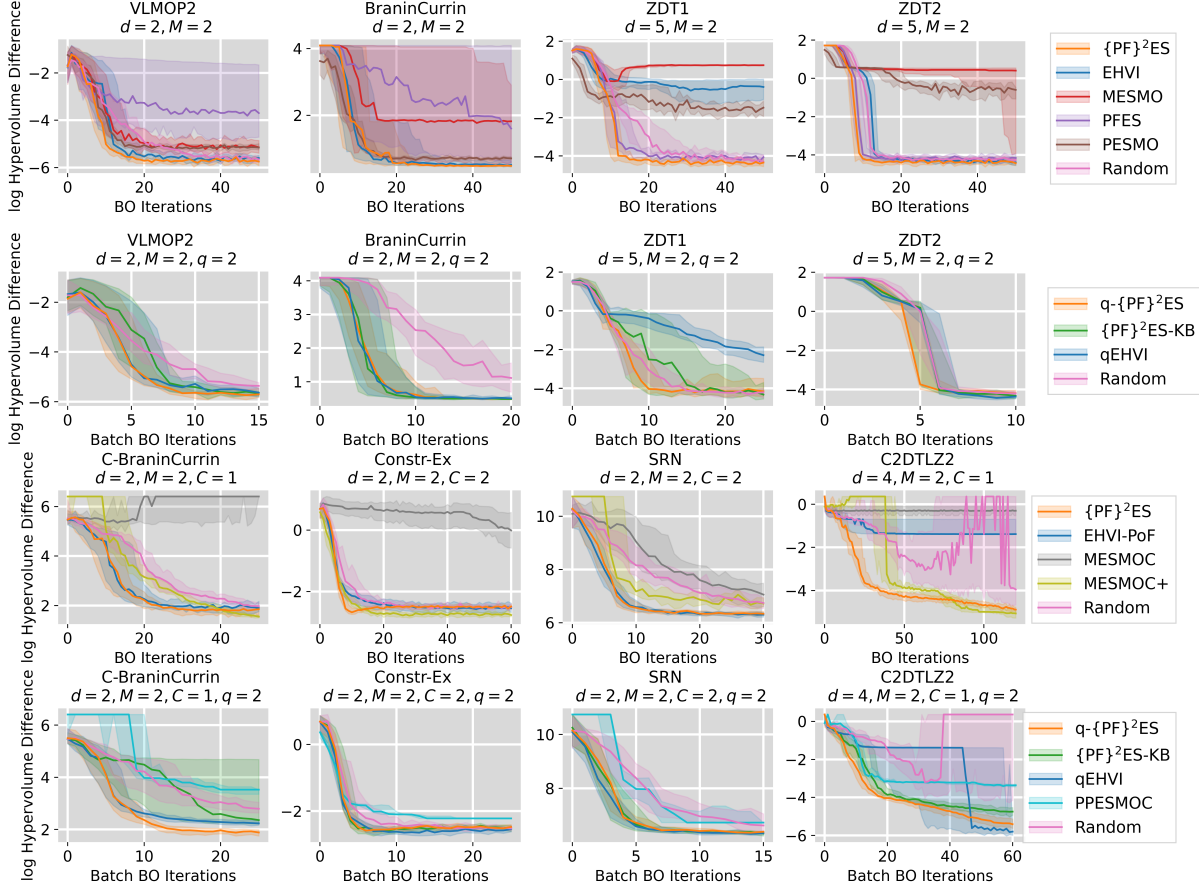


Figure 3: Comparison of different acquisition functions for multi-objective Bayesian optimization. Results are shown through median performances and interquartile ranges.




Figure 4: Four Bar Truss experimental results (with the same legend as Fig. 3).




Figure 5: Disc Brake Design experimental results (with the same legend as Fig. 3).

**Disc Brake Design** Finally, we consider a real-life disc brake design problem (CMOO) (Tanabe and Ishibuchi, 2020; Ray and Liew, 2002). Here, the objectives are the mass of the brake and the achieved stopping time. The design variables are the inner radius and outer radius of the discs, the engaging force, and the number of friction surfaces. Four constraints are presented in this problem: the minimum distance between the radii, maximum length of the brake, pressure, temperature, and torque limitations.

The results are depicted in Fig. 5. {PF}<sup>2</sup>ES demonstrate the overall fast convergence speed, while {PF}<sup>2</sup>ES-KB and qEHVI are slightly better than q-{PF}<sup>2</sup>ES in terms of stable

performance.

## 6 CONCLUSION

We have presented a new information-theoretic acquisition function {PF}<sup>2</sup>ES. By using a variational lower bound to the mutual information, {PF}<sup>2</sup>ES provides efficient (batch) multi-objective optimization for problems with unknown constraints. Extensive experimentation demonstrates the competitive performance of {PF}<sup>2</sup>ES over other entropy-based methods. We also advocate the advantages of {PF}<sup>2</sup>ES to other greedy acquisition functions (e.g., EHVI),

i.e., it is free from the requirement of a proper configuration of a reference point, which arguably is often unknown for real-life applications and can affect the performance of acquisition functions to a large extent.

**Limitations and Future Work** The current out-of-sample recommendation is based on the posterior mean of the GP. Hence, its performance depends on the modeling accuracy of the GP itself which means it can suffer from the curse of input dimensionality. Furthermore, the current out-of-sample recommendation strategy can be extended to a Bayesian version. Future research will focus on tackling these aspects and making the parallelization scalable to much larger batch sizes.

## 7 SOCIAL IMPACT

This paper presents a fundamental approach for multi-objective Bayesian Optimization and has no direct societal impact or ethical consequences. It can be used in generic real-life applications from, e.g., machine learning to engineering. Of course, the applications themselves can have an impact. As a Bayesian Optimization technique,  $\{PF\}^2ES$  aims for identifying the Pareto frontier hopefully with less expensive black-box function evaluations (e.g., expensive biomedical experiments, power-consuming simulations), hence contributing to less energy consumption.

## Acknowledgments

This research receives funding from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme and the “Fonds Wetenschappelijk Onderzoek (FWO)” programme, and Chinese Scholarship Council (grant number 201906290032). We sincerely thank Victor Picheny for the insightful comment on the necessity of considering the effect of spacing between discrete Pareto points, which has inspired proposition 1 and Algorithm 1. We also gratefully thank reviewers for providing extensive comments for improving the paper.

## References

- Abdolshah, M., Shilton, A., Rana, S., Gupta, S., and Venkatesh, S. (2018). Expected hypervolume improvement with constraints. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 3238–3243. IEEE.
- Balandat, M., Karrer, B., Jiang, D., Daulton, S., Letham, B., Wilson, A. G., and Bakshy, E. (2020). Botorch: a framework for efficient monte-carlo bayesian optimization. *Advances in neural information processing systems*, 33:21524–21538.
- Belakaria, S. and Deshwal, A. (2019). Max-value entropy search for multi-objective bayesian optimization. In *International Conference on Neural Information Processing Systems (NeurIPS)*.
- Belakaria, S., Deshwal, A., and Doppa, J. R. (2020). Max-value entropy search for multi-objective bayesian optimization with constraints. *arXiv preprint arXiv:2009.01721*.
- Belakaria, S., Deshwal, A., and Doppa, J. R. (2021). Output space entropy search framework for multi-objective bayesian optimization. *Journal of Artificial Intelligence Research*, 72:667–715.
- Berkeley, J., Moss, H. B., Artemev, A., Pascual-Diaz, S., Granta, U., Stojic, H., Couckuyt, I., Qing, J., Loka, N., Paleyes, A., Ober, S. W., Goodall, A., Ghani, K., and Picheny, V. (2022). Trieste.
- Blank, J. and Deb, K. (2020). pymoo: Multi-objective optimization in python. *IEEE Access*, 8:89497–89509.
- Couckuyt, I., Deschrijver, D., and Dhaene, T. (2012). Towards efficient multiobjective optimization: Multiobjective statistical criterions. *2012 IEEE Congress on Evolutionary Computation, CEC 2012*, pages 10–15.
- Couckuyt, I., Deschrijver, D., and Dhaene, T. (2014). Fast calculation of multiobjective probability of improvement and expected improvement criteria for pareto optimization. *Journal of Global Optimization*, 60(3):575–594.
- Cover, T. M. (1999). *Elements of information theory*. John Wiley & Sons.
- Dauert, J. (1993). Multicriteria optimization in engineering: a tutorial and survey. *Progress In Astronautics and Aeronautics: Structural Optimization: Status and Promise*, 150:209.
- Daulton, S., Balandat, M., and Bakshy, E. (2020). Differentiable expected hypervolume improvement for parallel multi-objective bayesian optimization. *arXiv preprint arXiv:2006.05078*.
- Daulton, S., Balandat, M., and Bakshy, E. (2021). Parallel bayesian optimization of multiple noisy objectives with expected hypervolume improvement. *Advances in Neural Information Processing Systems*, 34.
- Daulton, S., Eriksson, D., Balandat, M., and Bakshy, E. (2022). Multi-objective bayesian optimization over high-dimensional search spaces. In *Uncertainty in Artificial Intelligence*, pages 507–517. PMLR.
- Deb, K. (2019). Constrained multi-objective evolutionary algorithm. In *Evolutionary and Swarm Intelligence Algorithms*, pages 85–118. Springer.
- Deb, K. and Agrawal, S. (1999). A niched-penalty approach for constraint handling in genetic algorithms. In *Artificial neural nets and genetic algorithms*, pages 235–243. Springer.
- Deb, K. and Kalyanmoy, D. (2001). *Multi-Objective Optimization Using Evolutionary Algorithms*. John Wiley & Sons, Inc., USA.

- Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6(2):182–197.
- Deb, K., Thiele, L., Laumanns, M., and Zitzler, E. (2005). Scalable test problems for evolutionary multiobjective optimization. In *Evolutionary multiobjective optimization*, pages 105–145. Springer.
- Emmerich, M. and Deutz, A. H. (2018). A tutorial on multiobjective optimization: fundamentals and evolutionary methods. *Natural computing*, 17(3):585–609.
- Emmerich, M., Yang, K., and Deutz, A. H. (2020). Infill criteria for multiobjective bayesian optimization. In *High-Performance Simulation-Based Optimization*, pages 3–16. Springer.
- Feliot, P., Bect, J., and Vazquez, E. (2017). A bayesian approach to constrained single-and multi-objective optimization. *Journal of Global Optimization*, 67(1):97–133.
- Fernández-Sánchez, D., Garrido-Merchán, E. C., and Hernández-Lobato, D. (2020). Max-value entropy search for multi-objective bayesian optimization with constraints. *arXiv e-prints*, pages arXiv–2011.
- Frazier, P. I. (2018). A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*.
- Gardner, J. R., Kusner, M. J., Xu, Z. E., Weinberger, K. Q., and Cunningham, J. P. (2014). Bayesian optimization with inequality constraints. In *ICML*, volume 2014, pages 937–945.
- Garnett, R. (2022). *Bayesian Optimization*. Cambridge University Press. in preparation.
- Garrido-Merchán, E. C. and Hernández-Lobato, D. (2019). Predictive entropy search for multi-objective bayesian optimization with constraints. *Neurocomputing*, 361:50–68.
- Garrido-Merchán, E. C. and Hernández-Lobato, D. (2020). Parallel predictive entropy search for multi-objective bayesian optimization with constraints. *arXiv preprint arXiv:2004.00601*.
- Ginsbourger, D., Riche, R. L., and Carraro, L. (2010). Kriging is well-suited to parallelize optimization. In *Computational intelligence in expensive optimization problems*, pages 131–162. Springer.
- Guerreiro, A. P., Fonseca, C. M., and Paquete, L. (2020). The hypervolume indicator: Problems and algorithms. *arXiv preprint arXiv:2005.00515*.
- Hawe, G. and Sykulski, J. (2008). Probability of improvement methods for constrained multi-objective optimization. In *2008 IET 7th International Conference on Computation in Electromagnetics*, pages 50–51. IET.
- Hawe, G. I. and Sykulski, J. K. (2007). An enhanced probability of improvement utility function for locating pareto optimal solutions.
- Hennig, P. and Schuler, C. J. (2012). Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13(6).
- Hernández-Lobato, D., Hernández-Lobato, J., Shah, A., and Adams, R. (2016). Predictive entropy search for multi-objective bayesian optimization. In *International Conference on Machine Learning*, pages 1492–1501. PMLR.
- Hernández-Lobato, J. M., Hoffman, M. W., and Ghahramani, Z. (2014). Predictive entropy search for efficient global optimization of black-box functions. *Advances in neural information processing systems*, 27.
- Hvarfner, C., Hutter, F., and Nardi, L. (2022). Joint entropy search for maximally-informed bayesian optimization. *arXiv preprint arXiv:2206.04771*.
- Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Knudde, N., van der Herten, J., Dhaene, T., and Couckuyt, I. (2017). Gpflowopt: A bayesian optimization library using tensorflow. *arXiv preprint arXiv:1711.03845*.
- Kushner, H. J. (1964). A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise.
- Lacour, R., Klamroth, K., and Fonseca, C. M. (2017). A box decomposition algorithm to compute the hypervolume indicator. *Computers & Operations Research*, 79:347–360.
- Martínez-Frutos, J. and Herrero-Pérez, D. (2016). Kriging-based infill sampling criterion for constraint handling in multi-objective optimization. *Journal of Global Optimization*, 64(1):97–115.
- Moss, H. B., Leslie, D. S., Gonzalez, J., and Rayson, P. (2021). Gibbon: General-purpose information-based bayesian optimisation. *arXiv preprint arXiv:2102.03324*.
- Osyczka, A. and Kundu, S. (1995). A new method to solve generalized multicriteria optimization problems using the simple genetic algorithm. *Structural optimization*, 10(2):94–99.
- Parsons, M. G. and Scott, R. L. (2004). Formulation of multicriterion design optimization problems for solution with scalar numerical optimization methods. *Journal of Ship Research*, 48(01):61–76.
- Poole, B., Ozair, S., Van Den Oord, A., Alemi, A., and Tucker, G. (2019). On variational bounds of mutual information. In *International Conference on Machine Learning*, pages 5171–5180. PMLR.
- Qing, J., Couckuyt, I., and Dhaene, T. (2022a). A robust multi-objective bayesian optimization framework considering input uncertainty. *arXiv preprint arXiv:2202.12848*.
- Qing, J., Dhaene, T., and Couckuyt, I. (2022b). Spectral representation of robustness measures for optimization under input uncertainty. In *ICML2022, the 39th International Conference on Machine Learning*, pages 1–26.

- Qing, J., Knudde, N., Garbuglia, F., Spina, D., Couckuyt, I., and Dhaene, T. (2021). Adaptive sampling with automatic stopping for feasible region identification in engineering design. *Engineering with Computers*, pages 1–18.
- Rahimi, A. and Recht, B. (2007). Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20.
- Rasmussen, C. E. (2003). Gaussian processes in machine learning. In *Summer school on machine learning*, pages 63–71. Springer.
- Ray, T. and Liew, K. (2002). A swarm metaphor for multi-objective design optimization. *Engineering optimization*, 34(2):141–153.
- Schubert, E., Sander, J., Ester, M., Kriegel, H. P., and Xu, X. (2017). Dbscan revisited, revisited: why and how you should (still) use dbscan. *ACM Transactions on Database Systems (TODS)*, 42(3):1–21.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and De Freitas, N. (2015). Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175.
- Srinivas, N. and Deb, K. (1994). Multiobjective optimization using nondominated sorting in genetic algorithms. *Evol. Comput.*, 2(3):221–248.
- Suzuki, S., Takeno, S., Tamura, T., Shitara, K., and Karasuyama, M. (2020). Multi-objective bayesian optimization using pareto-frontier entropy. In *International Conference on Machine Learning*, pages 9279–9288. PMLR.
- Takeno, S., Tamura, T., Shitara, K., and Karasuyama, M. (2022). Sequential and parallel constrained max-value entropy search via information lower bound. In *International Conference on Machine Learning*, pages 20960–20986. PMLR.
- Tanabe, R. and Ishibuchi, H. (2020). An easy-to-use real-world multi-objective optimization problem suite. *Applied Soft Computing*, 89:106078.
- Ungredda, J. and Branke, J. (2021). Bayesian optimisation for constrained problems. *arXiv preprint arXiv:2105.13245*.
- Vakili, S., Picheny, V., and Durrande, N. (2020). Regret bounds for noise-free bayesian optimization. *arXiv preprint arXiv:2002.05096*.
- Van Veldhuizen, D. A. and Lamont, G. B. (1999). Multi-objective evolutionary algorithm test suites. In *Proceedings of the 1999 ACM symposium on Applied computing*, pages 351–357.
- Wang, Z. and Jegelka, S. (2017). Max-value entropy search for efficient Bayesian optimization. In Precup, D. and Teh, Y. W., editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3627–3635. PMLR.
- Wang, Z., Zhou, B., and Jegelka, S. (2016). Optimization as estimation with gaussian processes in bandit settings. In *Artificial Intelligence and Statistics*, pages 1022–1031. PMLR.
- Wilson, J., Hutter, F., and Deisenroth, M. (2018). Maximizing acquisition functions for bayesian optimization. *Advances in neural information processing systems*, 31.
- Yang, K., Emmerich, M., Deutz, A., and Bäck, T. (2019). Efficient computation of expected hypervolume improvement using box decomposition algorithms. *Journal of Global Optimization*, 75(1):3–34.
- Zitzler, E., Deb, K., and Thiele, L. (2000). Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary computation*, 8(2):173–195.



## A RELATIONSHIP WITH THE MULTI-OBJECTIVE PROBABILITY OF IMPROVEMENT

We provide additional insights for {PF}<sup>2</sup>ES by linking Eq. 6 to the Multi-objective Probability of Improvement (MOPI) (Yang et al., 2019). First, we define the following generic formulation of MOO acquisition functions:

$$\alpha_{MOO} = \int_E g \left( \int_F \ell(F, E) p(F) dF \right) p(E) dE, \quad (15)$$

where  $\ell(\cdot)$  and  $g(\cdot)$  are arbitrary functions, and  $E$  and  $F$  are random events. Evidently, {PF}<sup>2</sup>ES is a special case obtained by setting  $E := \mathcal{F}$ ,  $F := \mathbf{h}_x$ ,  $\ell(\mathbf{h}_x, \mathcal{F}) := \mathbb{1}(\mathbf{h}_x \in A(\mathcal{F}))$ , and  $g(\cdot) := -\log(1 - \cdot)$ . Here  $\mathbb{1}(\cdot)$  is the indicator function.

Subsequently, by setting  $\ell(E, F) := \mathbb{1}(\mathbf{h}_x \in \tilde{A}(\mathcal{F}))$  (i.e.,  $F := \mathbf{h}_x$  and  $E := \mathcal{F}$ ),  $g(\cdot)$  as the identity function, and specifying a Dirac delta distribution over the Pareto frontier  $\mathcal{F}$ , we see that the inner integration  $\int_F \ell(F, E) p(F) dF$  for both {PF}<sup>2</sup>ES and MOPI evaluates to  $p(\mathbf{h}_x \in \tilde{A}(\mathcal{F}))$ . Therefore, as  $g_{\text{{PF}}^2\text{ES}}(\cdot)$  and  $g_{\text{MOPI}}(\cdot)$  are both monotonic increasing over  $(0, 1)$ , the acquisition functions have the same maximizer, i.e., where  $p(\mathbf{h}_x \in \tilde{A}(\mathcal{F}))$  is largest. The constraint case (i.e., remark 1.2) is obtained by assuming independency of  $\mathbf{f}$  and  $\mathbf{g}$ .

## B MITIGATING THE CLUSTERING ISSUE

We first demonstrate the *clustering issue* on the inverted VLMOP2 problem (Van Veldhuizen and Lamont, 1999), i.e., we adopt VLMOP2 for maximization by taking the negative of each objective function. For illustrative purposes, the parameters for extracting the Pareto frontier in  $\tilde{\alpha}_{\text{{PF}}^2\text{ES}}(\tilde{A}(\mathcal{F}))$  are set to  $|\tilde{\mathbf{F}}| = 5$  and  $|\tilde{\mathcal{F}}| = 5 \forall \mathcal{F} \in |\tilde{\mathbf{F}}|$ .

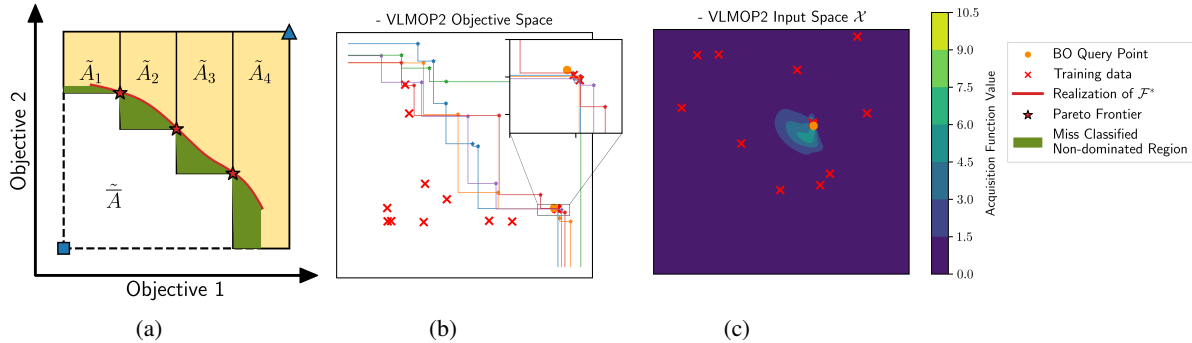


Figure 5: The hypervolume partition strategy for a discrete Pareto frontier  $\mathcal{F}$  is an overestimation of the actual non-dominated space  $A$ . a) Part of the dominated region has been misclassified as non-dominated (green area). b) This overestimation will make the clustering more severe when calculating {PF}<sup>2</sup>ES based on  $\tilde{A}(\mathcal{F})$ : For a certain iteration, given the MC samples of  $\mathcal{F}$ , {PF}<sup>2</sup>ES favors the 'false positive' location in the output space ( $\sigma_x = [0.00158, 0.00112]$ ) that is certain (high confidancy) to (marginally) improve. c) This causes a *clustering issue* as candidates very close to existing training data points are preferred.

A contour of {PF}<sup>2</sup>ES based on the approximation of Eq. 8 is depicted in Fig. 5, which shows that the new candidate point is believed to make an improvement with high confidence (due to the small GP posterior variance) in 4 of the 5 Pareto frontier samples. Unfortunately, the improvement is marginal as the predicted observation is located in the (relatively small) *false positive* region introduced by the hypervolume decomposition  $\mathcal{P}$ , see Fig. 5a. In practice, {PF}<sup>2</sup>ES will tend to keep sampling in a region that it is confident to make tiny (or no) improvements, causing the sampled points to be densely clustered in a small region of the input domain  $\mathcal{X}$ .

In order to mitigate this issue and encourage the acquisition function to focus more on where the Pareto frontier is uncertain, we utilize a parameter  $\epsilon$  to shift the Pareto frontier approximation  $\mathcal{F}$  and shrink the false positive region. The effect is demonstrated in Fig. 6. It can be seen that the modified acquisition function  $\tilde{\alpha}_{\text{{PF}}^2\text{ES}}(\tilde{A}(\mathcal{F}_\epsilon))$  tend to be more explorative than the naive approach  $\tilde{\alpha}_{\text{{PF}}^2\text{ES}}(\tilde{A}(\mathcal{F}))$  (i.e., based on Eq. 8).

Finally, this significant performance improvement is validated by the sensitivity analysis of  $\epsilon$  in Fig. 10, where the performance of the original approach  $\tilde{\alpha}_{\text{{PF}}^2\text{ES}}(\tilde{A}(\mathcal{F}))$  is denoted as  $\epsilon = 0$ .



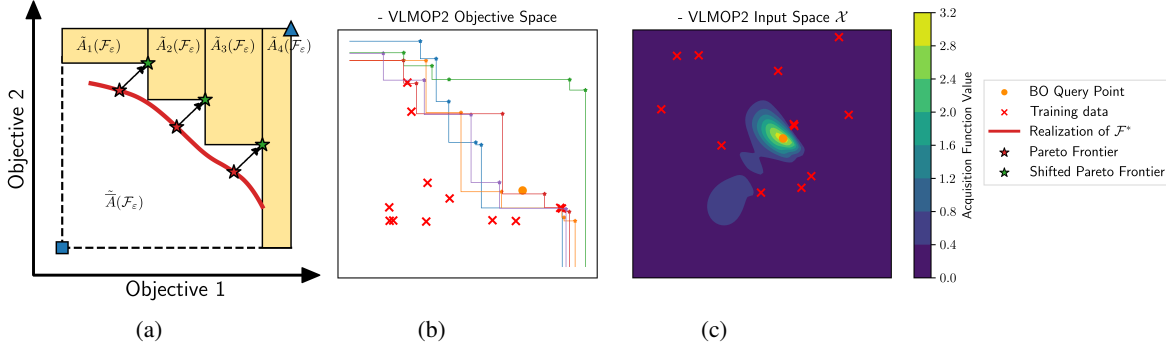


Figure 6: The  $\varepsilon$  approach for mitigating clustering issues. a) The non-dominated region is extracted from the partitioning  $\mathcal{P}$  based on the shifted Pareto frontier  $\tilde{\mathcal{F}}_\varepsilon$ . b) Relative to Fig. 5, the conservative estimation of  $\tilde{A}$  has encouraged more exploration, i.e., candidate  $(\sigma_{\mathbf{x}} = [0.053400, 0.032628])$ . c) The corresponding exploration around the Pareto set in the input space.

## C PARALLEL VERSION q- $\{\text{PF}\}^2\text{ES}$

### C.1 Derivation of q- $\{\text{PF}\}^2\text{ES}$

Consider the random variable  $\mathbf{h}_{\mathbf{X}} = \{\mathbf{h}_{\mathbf{x}_1}, \dots, \mathbf{h}_{\mathbf{x}_q}\}$ , i.e., the objective-constraint observations that could arise from a parallel evaluation of a batch of  $q$  points. In this scenario, we want to allocate a batch of  $q$  points  $\mathbf{X}$  that provide the most mutual information about the (feasible) Pareto frontier. With Eq. 4, we have:

$$I(\mathcal{F}; \mathbf{h}_{\mathbf{X}}) \geq \mathbb{E}_{\mathcal{F}} \left[ \int_{\mathbf{h}_{\mathbf{X}}} p(\mathbf{h}_{\mathbf{X}} | \mathcal{F}) \log \frac{q(\mathbf{h}_{\mathbf{X}} | \mathcal{F})}{p(\mathbf{h}_{\mathbf{X}})} d\mathbf{h}_{\mathbf{X}} \right] \quad (16)$$

Given the Pareto frontier  $\mathcal{F}$ , we set the variational approximation of the ground-truth conditional distribution  $p(\mathbf{h}_{\mathbf{X}} | \mathcal{F})$  as:

$$q(\mathbf{h}_{\mathbf{X}} | \mathcal{F}) = \begin{cases} \frac{p(\mathbf{h}_{\mathbf{X}})}{Z_{\bar{A}_q}} & \mathbf{h}_{\mathbf{X}} \in \bar{A}_q \\ 0 & \text{otherwise,} \end{cases} \quad (17)$$

where  $A_q \cup \bar{A}_q = \mathbb{R}^{q(M+C)}$ ,  $A_q \in \mathbb{R}^{q(M+C)}$  is the union of the non-dominated feasible regions:  $\{\mathbf{h}_{\mathbf{X}} \in \mathbb{R}^{q(M+C)} | \exists i \in \{1, \dots, q\}, \mathbf{h}_{\mathbf{x}_i} \in A\}$  and  $\bar{A}_q \in \mathbb{R}^{q(M+C)}$  is its complement  $\{\mathbf{h}_{\mathbf{X}} \in \mathbb{R}^{q(M+C)} | \forall i \in \{1, \dots, q\}, \mathbf{h}_{\mathbf{x}_i} \in \bar{A}\}$ . An intuitive illustration of  $A_q$  and  $q(\mathbf{h}_{\mathbf{X}} | \mathcal{F})$  is provided in Fig. 7. Note, the practical interpretation of the conditional probability Eq. 17 is that, given  $\mathcal{F}$ , the outcome of the objective-constraint observation is expected to **have zero probability of being in the feasible non-dominated region  $A$  for any candidate point within the batch**, as otherwise  $\mathcal{F}$  is not a valid Pareto frontier.

Following the same steps to motivate  $\{\text{PF}\}^2\text{ES}$ , we arrive at the following MC estimate of the resulting variational bound and the proposed batch extension of  $\{\text{PF}\}^2\text{ES}$ :

$$\alpha_{q-\{\text{PF}\}^2\text{ES}} \approx -\frac{1}{|\tilde{\mathbf{F}}|} \sum_{\mathcal{F} \in \tilde{\mathbf{F}}} \left[ \log \left( 1 - Z_{\bar{A}_q(\mathcal{F}_\varepsilon)} \right) \right], \quad (18)$$

where  $Z_{\bar{A}_q} := \int_{\bar{A}_q} p(\mathbf{h}_{\mathbf{X}}) d\mathbf{h}_{\mathbf{X}}$ . i.e., the probability that **there exists at least one batch element  $\mathbf{h}_{\mathbf{x}_i} \in \mathbf{h}_{\mathbf{X}}$  such that  $\mathbf{h}_{\mathbf{x}_i} \in \tilde{A} = \{\tilde{A}_1, \dots, \tilde{A}_{N_p}\}$** . We omit the notion  $\tilde{\mathcal{F}}_\varepsilon$  here as well as in Eq. 19,12 for notational simplicity.

To demonstrate the calculation of  $Z_{\bar{A}_q}$ , it is convenient to utilize the notion of events. Define  $\omega(\tilde{A}, \mathbf{h}_{\mathbf{X}})$  representing the event that there exists at least one batch element  $\mathbf{h}_{\mathbf{x}_l} \in \mathbf{h}_{\mathbf{X}}$  such that  $\mathbf{h}_{\mathbf{x}_l} \in \tilde{A}$ , it can be seen that  $Z_{\bar{A}_q} = p(\omega(\tilde{A}, \mathbf{h}_{\mathbf{X}}))$ . Now, since  $\tilde{A} = \{\tilde{A}_1, \dots, \tilde{A}_{N_p}\}$ , we know that  $\omega(\tilde{A}, \mathbf{h}_{\mathbf{X}}) = \bigcup_{i=1}^{N_p} \omega(\tilde{A}_i, \mathbf{h}_{\mathbf{X}})$ , where  $\omega(\tilde{A}_i, \mathbf{h}_{\mathbf{X}})$  represents there exists at least one batch element  $\mathbf{h}_{\mathbf{x}_l} \in \mathbf{h}_{\mathbf{X}}$  such that  $\mathbf{h}_{\mathbf{x}_l} \in \tilde{A}_i$ . It can be also seen that  $\omega(\tilde{A}_i, \mathbf{h}_{\mathbf{X}}) = \bigcup_{l=1}^q \omega(\tilde{A}_i, \mathbf{h}_{\mathbf{x}_l})$  where  $\omega(\tilde{A}_i, \mathbf{h}_{\mathbf{x}_l})$  is the  $l$ th batch

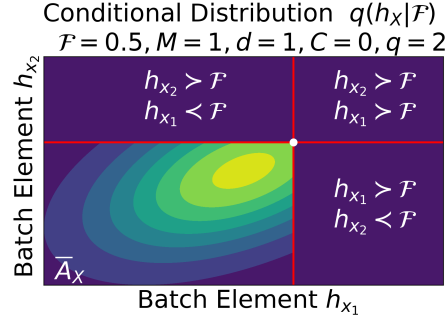


Figure 7: The batch conditional probability density  $q(\mathbf{h}_{\mathbf{X}}|\mathcal{F})$  for  $M = 1$ ,  $q = 2$  and  $C = 0$ . The problem degenerates to a batch single-objective optimization problem, where  $\mathcal{F}$  collapses to a scalar value (the maximum) and the dominance operation is reduced to  $>$ . In this scenario, by conditioning on  $\mathcal{F}$  (white dot), the objective space can be partitioned into 2 different subspaces: the union of the non-dominated feasible regions  $A_q$  and its complement  $\bar{A}_q$ , see bottom-left plot. The condition of being in the region of  $\bar{A}_q$  can be interpreted as: *no point inside the batch is larger than  $\mathcal{F}$* . The concept can be generalized for arbitrary  $M$ ,  $q$  and  $C$ .

element  $h_{x_l}$  is in  $\tilde{A}_i$ . Combine the two unions, and the fact that  $\omega(\tilde{A}_i, h_{\mathbf{X}}) = \prod_{k=1}^{M+C} \left( \mathbb{1}(A_{i_l}^k \leq \mathbf{h}_{\mathbf{x}_{l_j}} \leq A_{i_u}^k) \right)$  representing  $\mathbf{h}_{\mathbf{x}_l}$  is within the decomposed hypercube  $\tilde{A}_i$ , the event of  $\omega(\tilde{A}, h_{\mathbf{X}})$  can be written as:

$$\omega(\tilde{A}, h_{\mathbf{X}}) = \bigcup_{i=1}^{N_p} \left( \bigcup_{l=1}^q \omega(\tilde{A}_i, h_{\mathbf{x}_l}) \right) = \bigcup_{i=1}^{N_p} \left( \bigcup_{l=1}^q \left( \prod_{k=1}^{M+C} \left( \mathbb{1}(\tilde{A}_{i_l}^k \leq \mathbf{h}_{\mathbf{x}_{l_j}} \leq \tilde{A}_{i_u}^k) \right) \right) \right) \quad (19)$$

Finally, we are able to MC approximate the probability of this event:

$$\begin{aligned} Z_{\tilde{A}_q} &= p(\omega(\tilde{A}, h_{\mathbf{X}})) \\ &\approx \frac{1}{N_{MC}} \sum_{j=1}^{N_{MC}} \left( \bigcup_{i=1}^{N_p} \left( \bigcup_{l=1}^q \left( \prod_{k=1}^{M+C} \left( \mathbb{1}(\tilde{A}_{i_l}^k \leq \mathbf{h}_{\mathbf{x}_{l_j}} \leq \tilde{A}_{i_u}^k) \right) \right) \right) \right) \\ &\approx \frac{1}{N_{MC}} \sum_{j=1}^{N_{MC}} \left( \max_i \left( \max_l \left( \prod_{k=1}^{M+C} \left( \sigma \left( \frac{\mathbf{h}_{\mathbf{x}_{l_j}} - \tilde{A}_{i_l}^k}{\tau} \right) \cdot \sigma \left( \frac{\tilde{A}_{i_u}^k - \mathbf{h}_{\mathbf{x}_{l_j}}}{\tau} \right) \right) \right) \right) \right), \end{aligned} \quad (20)$$

Note that the additional approximation introduced in Eq. 20 is caused by the relaxation of the categorical event (imposed by the indicator function) to a continuous approximation:  $\mathbb{1}(\tilde{A}_{i_l}^k \leq \mathbf{h}_{\mathbf{x}_{l_j}} \leq \tilde{A}_{i_u}^k) \approx \sigma \left( \frac{\mathbf{h}_{\mathbf{x}_{l_j}} - \tilde{A}_{i_l}^k}{\tau} \right) \cdot \sigma \left( \frac{\tilde{A}_{i_u}^k - \mathbf{h}_{\mathbf{x}_{l_j}}}{\tau} \right)$ . This makes the MC approximation suitable for gradient-based optimizers.

## C.2 Monte Carlo Approximation of q-{PF}<sup>2</sup>ES and its Demonstration

We elaborate on the implementation of the MC approximation. Let  $\ell : \mathbb{R}^{q(M+C)} \times A \rightarrow \mathbb{R}$  represent the *utility function* that calculates  $Z_{\tilde{A}_q}$  as in Eq. 12, we can write the MC approximated q-{PF}<sup>2</sup>ES in terms of  $\ell$  as:

$$\alpha_{q\text{-}\{\text{PF}\}^2\text{ES}} \approx \tilde{\alpha}_{q\text{-}\{\text{PF}\}^2\text{ES}} = -\frac{1}{|\tilde{\mathbf{F}}|} \sum_{\mathcal{F} \in \tilde{\mathbf{F}}} \left[ \log \left( 1 - \frac{1}{N_{MC}} \sum_{j=1}^{N_{MC}} \ell(\mathbf{h}_{\mathbf{X}_j}, \tilde{A}_q(\mathcal{F}_{\boldsymbol{\epsilon}})) \right) \right], \quad (21)$$

where  $\tilde{\alpha}_{q\text{-}\{\text{PF}\}^2\text{ES}}$  represents the MC approximated acquisition function. Note that  $\tilde{A}_q(\mathcal{F}_{\boldsymbol{\epsilon}})$  is independent of  $\mathbf{X}$ . With the aim of a consistent MC acquisition function gradient, we leverage the reparameterization trick (Kingma and Welling, 2013) to generate samples of  $\mathbf{h}_{\mathbf{X}}$ : since  $\mathbf{h}_{\mathbf{X}} \sim \mathcal{N}(\mathbf{m}_{\mathbf{X}}, L_{\mathbf{X}} L_{\mathbf{X}}^T)$  where  $\mathbf{m}_{\mathbf{X}} \in \mathbb{R}^{q(M+C)}$ , Cholesky factor  $L_{\mathbf{X}} \in \mathbb{R}^{q(M+C) \times q(M+C)}$ , it can be generated via  $\mathbf{h}_{\mathbf{X}} = \mathbf{m}_{\mathbf{X}} + L_{\mathbf{X}} \boldsymbol{\lambda}$ , where  $\boldsymbol{\lambda} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  and the *base sample* sets:  $\{\boldsymbol{\lambda}_j\}_{j=1}^{N_{MC}}, \boldsymbol{\lambda}_j \in \mathbb{R}^{q(M+C)}$  is holding through the whole process of acquisition function maximization, the largely preserved consistency can be shown via the chain

rule of derivative as  $\nabla_{\mathbf{x}} \tilde{\alpha}_{q\text{-}\{\text{PF}\}^2\text{ES}} = -\frac{1}{|\tilde{\mathcal{F}}|} \sum_{\tilde{\mathcal{F}} \in \tilde{\mathcal{F}}} \left[ \frac{-\frac{1}{N_{\text{MC}}} \sum_{j=1}^{N_{\text{MC}}} \nabla_{\mathbf{x}} \ell(\mathbf{h}_{\mathbf{x}_j, \tilde{A}_q(\tilde{\mathcal{F}}_{\mathbf{e}})})}{\left(1 - \frac{1}{N_{\text{MC}}} \sum_{j=1}^{N_{\text{MC}}} \ell(\mathbf{h}_{\mathbf{x}_j, \tilde{A}_q(\tilde{\mathcal{F}}_{\mathbf{e}})})\right)} \right]$ , where each element within the *max* function  $\left( \prod_{k=1}^{M+C} \left( \sigma \left( \frac{\mathbf{h}_{x_{l_j}^k} - \tilde{A}_{l_j}^k}{\tau} \right) \cdot \sigma \left( \frac{\tilde{A}_{l_j}^k - \mathbf{h}_{x_{l_j}^k}}{\tau} \right) \right) \right)$  is differentiable. We use quasi-Monte Carlos to generate the base samples and refer Balandat et al. (2020); Daulton et al. (2020) for details of such an approximation.

As a proof of concept, we illustrate the effectiveness of  $q\text{-}\{\text{PF}\}^2\text{ES}$  (with batch size  $q = 2$ , and qMC sample size 128) on a one-dimensional Sinlinear-Forrester function (Table 3 of Qing et al. (2022a)). The progress plot of batch queries is provided in Fig. 8. It can be seen that the out-of-sample strategy is able to accurately recommend a Pareto frontier without false positives after five batches BO iterations.

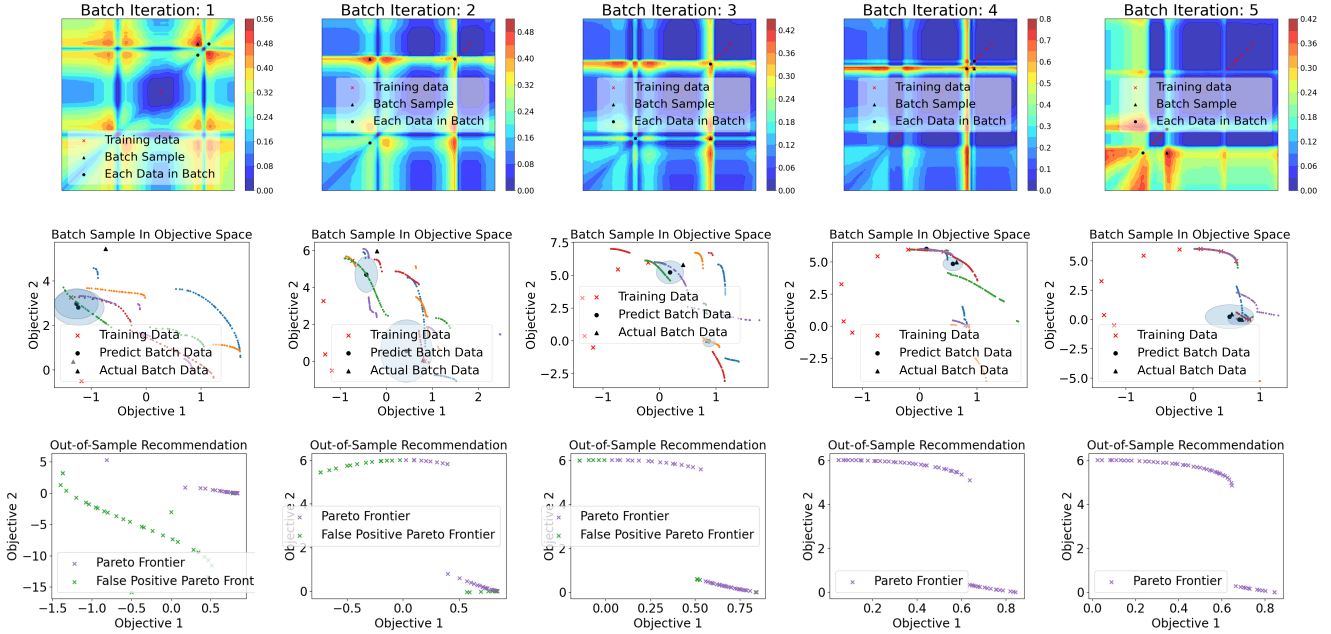


Figure 8:  $q\text{-}\{\text{PF}\}^2\text{ES}$  on the inverted Sinlinear-Forrester problem ( $d = 1, M = 2, C = 0$ ), i.e., for maximization we take the negative of the objective functions, with  $q = 2$  starting from three training data points. The first row is contour of  $q\text{-}\{\text{PF}\}^2\text{ES}$  in the concatenated batch input space. The second row represents the sampled Pareto frontiers  $\tilde{\mathcal{F}}$  and the queried batches in the objective space (the shaded oval area is the  $2\sigma$  of the GP uncertainty). The third row shows the out-of-sample recommended Pareto frontiers after the batch query, with false positives denoted in green.

## D ALTERNATIVE STRATEGY FOR SETTING $\varepsilon$

In the case of common bi-objective optimization ( $M = 2$ ) scenarios,  $\varepsilon$  can also be specified using an alternative strategy. The motivation for this approach is that, under mild additional assumptions,  $\varepsilon$  is closer to the infimum of the *mutual information lower bound set* (i.e., an ordered set of  $\varepsilon$  with the lower bound property of the mutual information).

**Proposition 1.** For a continuous Pareto Frontier  $\mathcal{F}$  with  $M = 2$ , given its finite approximation  $\tilde{\mathcal{F}} := \{\mathbf{f}^* | \mathbf{f}^* \in \mathcal{F}\}$  that satisfy  $\forall k \in M: \exists \mathbf{f}^* \in \tilde{\mathcal{F}} \text{ s.t. } \mathbf{f}_k^* = \max_{\mathbf{x} \in \mathcal{X}}(\mathbf{f}_k(\mathbf{x}))$  and  $|\tilde{\mathcal{F}}| \geq 2$ , define its corresponding set for  $k$ th output  $\tilde{\mathcal{F}}_{\text{ord}}^k$  ordered with  $\leq$ . Define  $\boldsymbol{\varepsilon}_{\mathbf{t}} := [\varepsilon_1, \dots, \varepsilon_M]$ , where  $\forall k \in M, \varepsilon_k = \max_{1 \leq i \leq |\tilde{\mathcal{F}}| - 1} (\tilde{\mathcal{F}}_{\text{ord}_{i+1}}^k - \tilde{\mathcal{F}}_{\text{ord}_i}^k)$ .  $\forall \boldsymbol{\varepsilon} \geq \boldsymbol{\varepsilon}_{\mathbf{t}}$ , with the hypervolume partition  $\mathcal{P}(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}})$  based on  $\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}} := \{\mathbf{f} + \boldsymbol{\varepsilon} | \mathbf{f} \in \tilde{\mathcal{F}}\}$ ,  $I(\mathcal{F}; \mathbf{h}_{\mathbf{x}}) \geq \lim_{|\tilde{\mathcal{F}}| \rightarrow \infty} \tilde{\alpha}_{\{\text{PF}\}^2\text{ES}}(\tilde{A}(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}}))$ .

Proposition 1 states that under the *extrema assumption* (i.e., the discrete Pareto frontier preserves all extrema points in  $\mathcal{F}$ ), specifying  $\varepsilon$  as the maximum inner spacing between Pareto points within  $\tilde{\mathcal{F}}$  per output dimension. One can subsequently define the non-dominated region  $\tilde{A}$  based on the shifted Pareto frontier  $\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}}$ , which guarantees the acquisition function to be a lower bound of mutual information estimation (Eq. 8) when the Pareto frontier samples  $|\tilde{\mathcal{F}}| \rightarrow \infty$ . Here we elaborate on

the proof of Proposition 1. We start with assuming an unconstrained multi-objective optimization problem. Note that the following proof also utilizes the well-established definition of *strict* and *weak dominance* (i.e.,  $\succ$  and  $\succeq$  respectively) and the property of transitivity of dominance (i.e.,  $a \succeq b \succeq c \Rightarrow a \succeq c$ ), for which we refer to standard multi-objective optimization textbooks, e.g., Emmerich and Deutz (2018).

**Lemma 1.** *For any discrete Pareto frontier  $\tilde{\mathcal{F}}^*$  when  $M = 2$ , with  $\boldsymbol{\varepsilon}_t$  defined as in proposition 1,  $\forall \mathbf{f} \in \tilde{A}(\tilde{\mathcal{F}}^*) \cap \{\mathbf{f} | \forall k \in M : \mathbf{f}_k \leq \max(\tilde{\mathcal{F}}_k^*)\}$ ,  $\exists \mathbf{f}^* \in \tilde{\mathcal{F}}^*$  s.t.  $\mathbf{f}^* \preceq \mathbf{f} + \boldsymbol{\varepsilon}_t$  and  $\mathbf{f}^*$  is not unique.*

*Proof.* For any  $k$ th output, given that  $\mathbf{f}_k \leq \max(\tilde{\mathcal{F}}_k^*)$ , we know that  $\exists k_i \in [1, |\tilde{\mathcal{F}}| - 1]$ :  $\tilde{\mathcal{F}}_{ord_{k_i}}^k \leq \mathbf{f}_k \leq \tilde{\mathcal{F}}_{ord_{k_{i+1}}}^k$ . 1°: if for  $l \neq k$ ,  $\tilde{\mathcal{F}}_{ord_{l_{|\tilde{\mathcal{F}}|-i}}^l} \leq \mathbf{f}_l \leq \tilde{\mathcal{F}}_{ord_{l_{|\tilde{\mathcal{F}}|-i+1}}^l}$ , then we know that  $\mathbf{f}_k + \boldsymbol{\varepsilon}_k \geq \tilde{\mathcal{F}}_{ord_{k_{i+1}}}^k$  and  $\mathbf{f}_l + \boldsymbol{\varepsilon}_l \geq \tilde{\mathcal{F}}_{ord_{l_{|\tilde{\mathcal{F}}|-i+1}}^l}$ , hence  $\mathbf{f} + \boldsymbol{\varepsilon} \succeq (\tilde{\mathcal{F}}_{ord_{k_{i+1}}}^k, \tilde{\mathcal{F}}_{ord_{l_{|\tilde{\mathcal{F}}|-i+1}}^l})$ . Since  $(\tilde{\mathcal{F}}_{ord_{k_{i+1}}}^k, \tilde{\mathcal{F}}_{ord_{l_{|\tilde{\mathcal{F}}|-i+1}}^l}) \succeq (\tilde{\mathcal{F}}_{ord_{k_i}}^k, \tilde{\mathcal{F}}_{ord_{l_{|\tilde{\mathcal{F}}|-i+1}}^l})$  and  $(\tilde{\mathcal{F}}_{ord_{k_{i+1}}}^k, \tilde{\mathcal{F}}_{ord_{l_{|\tilde{\mathcal{F}}|-i+1}}^l}) \succeq (\tilde{\mathcal{F}}_{ord_{k_{i+1}}}^k, \tilde{\mathcal{F}}_{ord_{l_{|\tilde{\mathcal{F}}|-i}}^l})$ , the lemma holds; 2°: if for  $l \neq k$ ,  $\mathbf{f}_l > \tilde{\mathcal{F}}_{ord_{l_{|\tilde{\mathcal{F}}|-i+1}}^l}$ . Finally it can be shown that  $(\tilde{\mathcal{F}}_{ord_{k_{i+1}}}^k, \tilde{\mathcal{F}}_{ord_{l_{|\tilde{\mathcal{F}}|-i+1}}^l}) \succeq (\tilde{\mathcal{F}}_{ord_{k_i}}^k, \tilde{\mathcal{F}}_{ord_{l_{|\tilde{\mathcal{F}}|-i+1}}^l})$  and  $(\tilde{\mathcal{F}}_{ord_{k_{i+1}}}^k, \tilde{\mathcal{F}}_{ord_{l_{|\tilde{\mathcal{F}}|-i+1}}^l}) \succeq (\tilde{\mathcal{F}}_{ord_{k_{i+1}}}^k, \tilde{\mathcal{F}}_{ord_{l_{|\tilde{\mathcal{F}}|-i}}^l})$ .  $\square$

**Lemma 2.** *For a given discrete Pareto frontier  $\tilde{\mathcal{F}}$  with  $\boldsymbol{\varepsilon}_t$  as defined in proposition 1,  $\tilde{A}(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}_t}) \subseteq A(\mathcal{F})$ .*

*Proof.* Note the necessary condition for  $\tilde{A}(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}_t}) \subseteq A(\mathcal{F})$  is  $\forall \mathbf{f} \in \tilde{A}(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}_t}) : \mathbf{f} \in A(\mathcal{F})$ . Proof by contradiction: suppose  $\exists \mathbf{f} \in \tilde{A}(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}_t})$  s.t.  $\mathbf{f} \notin A(\mathcal{F})$ , hence  $\mathbf{f} \in \bar{A}(\mathcal{F})$ . This implies  $\exists \mathbf{f}' \in \mathcal{F} : \mathbf{f}' \succeq \mathbf{f} \Rightarrow \mathbf{f}' + \boldsymbol{\varepsilon}_t \succeq \mathbf{f} + \boldsymbol{\varepsilon}_t$ . 1°: Suppose  $\forall k \in M : \mathbf{f}_k \leq \max(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}_t, k})$ , with lemma 1, by setting  $\tilde{\mathcal{F}}^* = \tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}_t}$  we know  $\exists \mathbf{f}'' \in \tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}_t}$  s.t.  $\mathbf{f}'' \preceq \mathbf{f} + \boldsymbol{\varepsilon}_t$ , then  $\mathbf{f}' + \boldsymbol{\varepsilon}_t \succeq \mathbf{f} + \boldsymbol{\varepsilon}_t \succeq \mathbf{f}'' \Rightarrow \mathbf{f}' \succeq \mathbf{f}'' - \boldsymbol{\varepsilon}_t$ , since  $\mathbf{f}'' - \boldsymbol{\varepsilon}_t \in \mathcal{F}$ , due to the definition of a Pareto frontier this can only hold when  $\mathbf{f}'' - \boldsymbol{\varepsilon}_t = \mathbf{f}'$ , which contradicts with its non-unique property, hence lemma holds; 2°: Suppose  $\exists k \in M$  s.t.  $\mathbf{f}_k > \max(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}_t, k})$ , since according to the definition of  $\tilde{\mathcal{F}}$ ,  $\max(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}_t, k}) = \max(\mathcal{F}_k) + \boldsymbol{\varepsilon}_{t, k}$ , this results  $\mathbf{f}_k > \max(\mathcal{F}_k)$  and hence  $\mathbf{f} \in A(\mathcal{F})$ , contradicting with  $\mathbf{f} \in \bar{A}(\mathcal{F})$ .  $\square$

Eventually, the proof of Proposition 1 becomes straightforward: with lemma 2 and Eq. 6 we have  $I(\mathcal{F}; \mathbf{h}_x) \geq -\mathbb{E}_{\mathcal{F}} [\log(1 - Z_{A(\mathcal{F})})] = \lim_{|\mathbf{F}| \rightarrow \infty} -\frac{1}{|\mathbf{F}|} \sum_{\mathcal{F} \in \mathbf{F}} [\log(1 - Z_{A(\mathcal{F})})] \geq \lim_{|\mathbf{F}| \rightarrow \infty} \tilde{\alpha}_{\{\text{PF}\}^2\text{ES}}(\tilde{A}(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}}))$  since  $Z_{A(\mathcal{F})} \geq Z_{\tilde{A}(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}})}$ .

Remark that the proposition 1 is also valid for CMOO when  $\mathcal{F} \neq \emptyset$  (the extrema condition in CMOO needs to be generalized for constraints). In reality,  $\mathcal{F}$  is not continuous for all problems, unfortunately. Clustering techniques (e.g., Schubert et al. (2017)) can be utilized as a pre-processing step to satisfy the condition of the proposition. Nevertheless, this does increase the complexity of the method. For  $M > 2$ , proposition 1 does not hold since  $\boldsymbol{\varepsilon}$ , as defined using the maximum inner spacing, can not guarantee that  $\tilde{A}(\tilde{\mathcal{F}}_{\boldsymbol{\varepsilon}}) \subseteq A(\mathcal{F})$ , and the choice of  $\boldsymbol{\varepsilon}$  can be correlated with the choice of partitioning strategy  $\mathcal{P}$ .

The practical implementation of the original acquisition function {PF}<sup>2</sup>ES-c, as well as using the proper  $\boldsymbol{\varepsilon}$  approach based on proposition 1 ({PF}<sup>2</sup>ES-lb) is detailed in Algorithm 1 and 2, respectively. Note that lines 2-9 in both algorithms only need to be calculated once per BO iteration.

## D.1 Comparing strategies for setting $\boldsymbol{\varepsilon}$


We investigate the two different approaches for setting  $\boldsymbol{\varepsilon}$  for two bi-objective optimization problems. Namely, VLMOP2 and C-BrarinCurrin as provided in Table 1, both with continuous Pareto frontiers. The results are shown in Fig. 9 where we refer to the heuristical approach as {PF}<sup>2</sup>ES-4% (i.e.,  $c = 0.04$ ) and the approach above as {PF}<sup>2</sup>ES-lb. From the in-sample regret, it can be seen that {PF}<sup>2</sup>ES-lb can be more greedy than {PF}<sup>2</sup>ES-4%, although it is less competitive in terms of out-of-sample recommendations for C-BrarinCurrin. As conclusion, both strategies lead to comparable performance. Hence, we propose to simply utilize the heuristic approach which does not depend on extra assumptions and generalizes better with respect to  $M$ .

**Algorithm 1** Practical implementation of  $\{\text{PF}\}^2\text{ES-c}$ 

- 1: **Input:** a candidate  $\mathbf{x}$ ; GP models for objectives (and constraints):  $\mathcal{M}$ , heuristic shifting percentage  $c$ .
- 2: **for**  $i = 1, \dots, |\mathcal{F}|$  **do**
- 3:   generate  $i$ th GP parametric trajectory approximations (see e.g., Eq. 3, 4 of Qing et al. (2022b)).
- 4:   Apply NSGA2 to GP trajectories to obtain Pareto frontier samples  $\tilde{\mathcal{F}}$
- 5:   **for**  $k = 1, \dots, M$  **do:**
- 6:      $\varepsilon_k = c \cdot (\max(\tilde{\mathcal{F}}^k) - \min(\tilde{\mathcal{F}}^k))$
- 7:   **end for**
- 8:    $\tilde{\mathcal{F}}_\varepsilon = \tilde{\mathcal{F}} + \varepsilon$
- 9:   Partition  $\tilde{A} : \tilde{A} = \mathcal{P}(\tilde{\mathcal{F}}_\varepsilon)$
- 10: **end for**
- 11: **Return**  $\{\text{PF}\}^2\text{ES-c}$  (Eq. 10).

**Algorithm 2** Practical implementation of  $\{\text{PF}\}^2\text{ES-lb}$ 

- 1: **Input:** a candidate  $\mathbf{x}$ ; GP models for objectives (and constraints):  $\mathcal{M}$ .
- 2: **for**  $i = 1, \dots, |\mathcal{F}|$  **do**
- 3:   generate  $i$ th GP parametric trajectory approximations (see, e.g., Eq. 3, 4 of Qing et al. (2022b)).
- 4:   Apply NSGA2 to GP trajectories to obtain Pareto frontier samples  $\tilde{\mathcal{F}}$
- 5:   **for**  $k = 1, \dots, M$  **do:**
- 6:      $\varepsilon_k = \max_{1 \leq i \leq |\mathcal{F}|-1} (\tilde{\mathcal{F}}^k_{\text{ord}_{i+1}} - \tilde{\mathcal{F}}^k_{\text{ord}_i})$  as in proposition 1
- 7:   **end for**
- 8:    $\tilde{\mathcal{F}}_\varepsilon = \tilde{\mathcal{F}} + \varepsilon$
- 9:   Partition  $\tilde{A} : \tilde{A} = \mathcal{P}(\tilde{\mathcal{F}}_\varepsilon)$
- 10: **end for**
- 11: **Return**  $\{\text{PF}\}^2\text{ES-lb}$  (Eq. 10).


 Figure 9: Comparing strategies for setting  $\varepsilon$  for two synthetic benchmark functions.

## E EXPERIMENTAL DETAILS

### E.1 Synthetic problems and setting the reference point

We benchmark on synthetic problems regularly used within the literature. Details are given in Table 1, together with the reference point used in the uncertainty calibration<sup>7</sup> (i.e., Fig. 3) and for performance metric (i.e., Fig. 4, 5). Note that the ideal hypervolume utilized to calculate the log hypervolume difference performance metric is set optimistically to guarantee a well-behaved performance metric.

### E.2 Out-of-sample recommendation strategy

For both MOO and CMOO problems, the feasible Pareto optimal inputs  $\mathbf{X}_r$  is recommended using a conservative model-based approach similar to Garrido-Merchán and Hernández-Lobato (2020); Ungredda and Branke (2021). The recommendation task can be formulated as a MOO problem:

$$\begin{aligned}
 & \underset{\mathbf{x} \in \mathcal{X}}{\text{maximize}} \quad \mathbf{m}^1, \mathbf{m}^2, \dots, \mathbf{m}^M \\
 & \text{s.t.}, 1 - \Phi\left(\frac{\mathbf{m}^{M+1} - \eta^1}{\sigma^{M+1}}\right) \geq C_{Fea}, \dots, 1 - \Phi\left(\frac{\mathbf{m}^{M+C} - \eta^C}{\sigma^{M+C}}\right) \geq C_{Fea}
 \end{aligned} \tag{22}$$

where  $m$  and  $\sigma$  represent the GP poster mean and standard deviation, respectively. We expect all the recommendations to be feasible in our empirical experiments, hence  $C_{Fea}$  is set to 0.95 and is decreased with 0.05 in case no feasible

<sup>7</sup>Note that since the Pareto frontier samples from the GP posterior, in case of an inaccurate model, it can significantly deviate from the real Pareto frontier (e.g., at the initial stages of the BO process). Hence, the reference point utilized in the hypervolume indicator calculation is set more conservatively than in the regret calculation.

Table 1 : Details of the synthetic problems and reference point settings

Problem	Formulation	$d$	$M$	$C$	Reference point (uncertainty calibration)	Reference point (regret)	Ideal Hypervolume (regret)
VLMOP2	Table 1 of Van Veldhuizen and Lamont (1999)	2	2	0	[3, 3]	[1.2, 1.2]	0.77157
BraninCurri	Appendix E.2 of (Daulton et al., 2020)	2	2	0	[2000, 50]	[18.0, 6.0]	60.0000
ZDT1	Eq. 7 of Zitzler et al. (2000)	5	2	0	[15, 15]	[2.5, 2.5]	5.90453
ZDT2	Eq. 8 of Zitzler et al. (2000)	5	2	0	[15, 15]	[2.5, 2.5]	5.57236
Four Bar Truss	Eq. 21 of Dautert (1993)	4	2	0	NA	[3400, 0.05]	81.9590
DTLZ4 (App 1.2)	Eq. 24 of Deb et al. (2005)	8	2	0	NA	[2.5, 2.5]	5.45327
DTLZ4 (App 1.3)	Eq. 24 of Deb et al. (2005)	5	3	0	NA	[2.5, 2.5, 2.5]	15.0080
C-BraninCurri	Appendix E.2 of Daulton et al. (2020)	2	2	1	[2000, 100]	[80.0, 12.0]	606.752
Constr-EX	Eq. 7.2 of Deb and Kalyanmoy (2001)	2	2	2	[2, 50]	[1.1, 10.0]	5.30186
SRN	Eq. 10 of Srinivas and Deb (1994)	2	2	2	[1500, 500]	[250, 50]	46205.0
C2-DTLZ2	Section 4.6 of Deb (2019)	4	2	1	[10, 250]	[2.5, 2.5]	5.43594
Disc Brake Design	Section 3.4 of Ray and Liaw (2002)	4	2	5	NA	[8.0, 4.0]	17.6882
OSY (App 1.2)	Sec 4.3 of Osyczka and Kundu (1995)	6	2	6	NA	[50, 100]	25344.9
Marine Design (App 1.3)	CRE32 of Tanabe and Ishibuchi (2020)	6	3	9	NA	[-700, 13000, 3500]	6991669

solution exists (i.e.,  $\mathbf{X}_r = \phi$ ),  $\eta^C = 0.5\% \cdot (\max(\mathbf{H}^C) - \min(\mathbf{H}^C))$ . In practice, we empirically find that while an aggressive recommendation strategy (e.g.,  $C_{Fea} = 0.5, \eta = 0$ ) can possibly contain infeasible candidates, the recommended candidates are able to cover the real Pareto frontier faster and the constraints are often only violated with a tiny amount (e.g.,  $1e-3$  level), which is tolerable for some applications. Hence, the practitioner is recommended to choose  $C_{Fea}$  and  $\eta$  according to the tolerance to constraint violations.

### E.3 Setup details of different acquisition functions

We provide configuration details for different information-theoretic acquisition functions used in the comparison.

**EHVI** is using the Trieste (Berkeley et al., 2022) implementation. The dynamic reference point specification strategy (Knudde et al., 2017) is utilized:

$$r_{ref} = \max(\tilde{\mathcal{F}}) + 2 \cdot \frac{\max(\tilde{\mathcal{F}}) - \min(\tilde{\mathcal{F}})}{|\tilde{\mathcal{F}}|} \quad (23)$$

**qEHVI** is based on the Trieste (Berkeley et al., 2022) implementation and extended with a random quasi-Monte Carlo (qMC) method for batch reparameterization sampling (see Eq. 15 of Daulton et al. (2020)). 128 qMC samples are used to approximate the hypervolume improvement.

The qEHVI paper does not provide a dynamic reference point strategy for CMOO. Denoting the outcome of objective function on training data as  $\mathbf{F}$ , the following strategy is used:

$$r_{ref} = \begin{cases} \max(\tilde{\mathcal{F}}) + 2 \cdot \frac{\max(\tilde{\mathcal{F}}) - \min(\tilde{\mathcal{F}})}{|\tilde{\mathcal{F}}|} & \tilde{\mathcal{F}} \neq \phi \\ \max(\mathbf{F}) + 2 \cdot \frac{\max(\mathbf{F}) - \min(\mathbf{F})}{|\mathbf{F}|} & \tilde{\mathcal{F}} = \phi \end{cases} \quad (24)$$

**EHVI-PoF** is a common strategy for handling CMOO. In case there are no predicted feasible observations on the training data, the probability of feasibility function will be used for locating a feasible observation first. For EHVI-PoF, the same reference point setting as EHVI is utilized, where  $\tilde{\mathcal{F}}$  is the **feasible** Pareto frontier.

**Information-theoretic acquisition functions** For PESMO, MESMOC+ and PPESMOC, the implementations of the accompanying papers are used<sup>8</sup> and implement the remaining acquisition functions ourselves. While we aim to use similar settings for all algorithms (e.g., the Pareto frontier number, the population sizes for NSGA2), the provided implementation of PESMO, MESMOC+, and PPESMOC still have different settings, see Table 2<sup>9</sup>.

Table 2: Configurations of different information-theoretic acquisition functions.

	Pareto Frontier Sample Strategy	Pareto Frontier Samples	Out-of-sample Recommendation Strategy
{PF} <sup>2</sup> ES (for MOO)	NSGA2	5	NSGA2
PFES	NSGA2	5	NSGA2
MESMO	NSGA2	5	NSGA2
PESMO	Monte Carlo	1	NSGA2
q-{PF} <sup>2</sup> ES (for MOO)	NSGA2	5	NSGA2
{PF} <sup>2</sup> ES (for CMOO)	NSGA2*	5	NSGA2*
MESMOC	NSGA2*	5	NSGA2*
MESMOC+	Monte Carlo	1	Monte Carlo
q-{PF} <sup>2</sup> ES (for CMOO)	NSGA2*	5	NSGA2*
PPESMOC	Monte Carlo	1	Monte Carlo

<sup>8</sup>We are not able to use a continuous optimizer for out-of-sample recommendations for MESMOC+ and PPESMOC.

<sup>9</sup>NSGA2\* is a parameter-less NSGA2 approach (Deb and Agrawal, 1999) to perform CMOO. We use the available training data to initialize NSGA2, which leads to improved performance.

## F SENSITIVITY ANALYSIS OF THE HYPERPARAMETERS OF {PF}<sup>2</sup>ES AND q-{PF}<sup>2</sup>ES

### F.1 Parameter $\epsilon$

We conduct a sensitivity analysis for setting  $\epsilon$  to determine its effect on the performance of the acquisition function. The results are depicted in Fig. 10, with results of both in-sample, and out-of-sample regret, as well as the uncertainty calibration of the Pareto frontier.

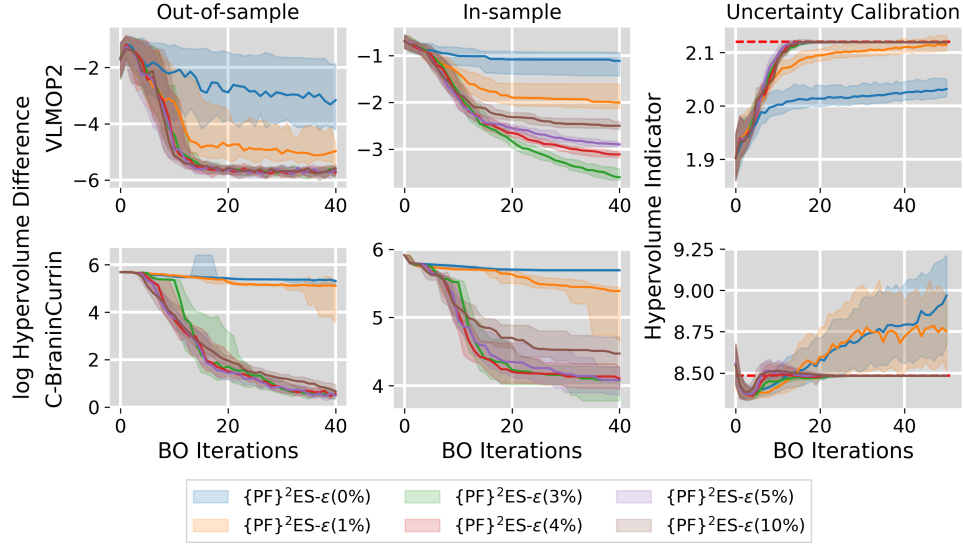


Figure 10: Sensitivity analysis of  $\epsilon$  on the performance of {PF}<sup>2</sup>ES.

It can be seen that  $\epsilon$  is less sensitive in terms of out-of-sample regret within the range of 3% – 5%. A too-small value will not mitigate the clustering issue, while a large value of  $\epsilon$  can cause too much exploration resulting in a deterioration of the in-sample regret. Importantly, using  $\epsilon$  is helpful for a faster reduction of the uncertainty of the Pareto frontier.

### F.2 Number of Monte Carlo samples for q-{PF}<sup>2</sup>ES

We conduct an empirical sensitivity analysis of the MC sample size used in approximating the q-{PF}<sup>2</sup>ES (i.e., Eq. 12) and its effect on the performance. The results on two different benchmark functions are provided in Fig. 11. As expected, it can be seen that with an out-of-sample recommendation strategy, a larger MC sample size can positively affect the performance of the acquisition function. The benefit is more obvious for the in-sample recommendation strategy.




Figure 11: The performance difference of q-{PF}<sup>2</sup>ES with respect to different MC sample sizes.



## G COMPUTATIONAL COMPLEXITY OF $\{\text{PF}\}^2\text{ES}$ AND $\mathbf{q}\text{-}\{\text{PF}\}^2\text{ES}$

For the information-theoretic acquisition function, generating the Pareto frontier approximation (i.e., line 3-9 of Algorithm 1,2) is treated separately as the *initialization cost*, while the cost of evaluating the acquisition functions is termed as *evaluation cost*.

**Initialization cost** The initialization of  $\{\text{PF}\}^2\text{ES}$  consists of two parts: generation of  $\tilde{\mathcal{F}}$  samples and calculate  $\tilde{A}(\tilde{\mathcal{F}})$  by utilizing a hypervolume decomposition  $\mathcal{P}$  on  $\tilde{\mathcal{F}}$ . The cost of the former, ignoring the non-dominated sorting complexity of NSGA2 which is  $\mathcal{O}(M|\tilde{\mathcal{F}}|^2)$  (Deb et al., 2002), breaks down to the cost of evaluating the GP spectral samples, which is linear with the number of random Feature features. The hypervolume decomposition strategy (Lacour et al., 2017) utilized in this research has a complexity of  $\mathcal{O}(|\tilde{\mathcal{F}}|^{\lfloor \frac{M}{2} \rfloor + 1})$ . However, both  $\{\text{PF}\}^2\text{ES}$  and  $\mathbf{q}\text{-}\{\text{PF}\}^2\text{ES}$  are agnostic to the choice of  $\mathcal{P}$ .

### Evaluation cost

- $\{\text{PF}\}^2\text{ES}$  By denoting the maximum number of cells across the set of Pareto frontier samples  $\tilde{\mathbf{F}}$  as  $\max(N_p)$ , the computational complexity of evaluate  $\{\text{PF}\}^2\text{ES}$  is:  $\mathcal{O}(\max(N_p)M)^{10}$ .
- $\mathbf{q}\text{-}\{\text{PF}\}^2\text{ES}$  We omit the cost of the batch reparameterization sampling (which is  $\mathcal{O}((M+C)q)$ ) and hence, the evaluation cost of  $\mathbf{q}\text{-}\{\text{PF}\}^2\text{ES}$  itself is:  $\mathcal{O}(\max(N_p)qN_{MC}M)$ .

## H COMPARING RUNNING TIMES

The following experiments are conducted in parallel per batch of ten on a Linux server with 256 GB RAM. The results are depicted in Table 3. In general, the joint batch acquisition function  $\mathbf{q}\text{-}\{\text{PF}\}^2\text{ES}$  requires a longer time to query a batch of samples than  $\{\text{PF}\}^2\text{ES}$ . Besides the additional computational complexity introduced by the MC approximation, cfr. Section G, this can be attributed to an increase in difficulty for the multi-start L-BFGS-B optimizer as illustrated in Fig. 12. We also note that, though the  $\{\text{PF}\}^2\text{ES}\text{-KB}$  consumes similar time as other batch strategies, its computationally cost grows much faster w.r.t.  $q$  as it requires additional sampling of the Pareto frontier  $\tilde{\mathcal{F}}$ . For PPESMOC, we generally observe a longer average query time than the original paper which reports the median (Table 1 of Garrido-Merchán and Hernández-Lobato (2020)), and the query time grows drastically with the number of iterations (e.g., the batch query time for C2DTLZ2 grows from around 20 minutes per batch sample to almost 3 hours when approaching the maximum number of batch iterations).




Figure 12: Comparison of  $\{\text{PF}\}^2\text{ES}$  and the Monte Carlo approximation of  $\mathbf{q}\text{-}\{\text{PF}\}^2\text{ES}$  (with  $q = 1$ ) on the SRN benchmark function. The effect of the MC approximation on the gradient-based acquisition function optimizer can be seen in the last sub-figure: the multi-start gradient-based optimization generally requires larger function evaluations to converge (i.e., the norm of the derivatives reaches a lower value).

## I ADDITIONAL EXPERIMENTAL RESULTS

### I.1 Experimental results for the in-sample recommendation strategy

We report the experimental results of the in-sample recommendation strategy in the main paper, see Fig. 13, the in-sampled Pareto frontiers are extracted from all of the data obtained so far. However, the in-sample recommendation strategy is not the intrinsic strategy for non-myopic information-theoretic acquisition functions.

<sup>10</sup>Note that for both  $\{\text{PF}\}^2\text{ES}$  and  $\mathbf{q}\text{-}\{\text{PF}\}^2\text{ES}$ , the complexity involved with constraint number  $C$  is independent with the decomposed grid size  $\max(N_p)$  given independent assumption across each outcome.

Table 3: Wall time for optimizing the acquisition function in seconds on a CPU (2x Silver 4210 CPU @ 2.20GHz). The mean and 2 times the standard deviations are reported across 30 runs.

	VLMOP2	BraninCurrin	ZDT1	ZDT2
{PF} <sup>2</sup> ES	47.38 ± 4.28	47.01 ± 2.58	46.35 ± 2.08	45.86 ± 2.23
PFES	63.68 ± 32.95	53.90 ± 26.29	84.02 ± 63.6	63.32 ± 44.74
MESMO	34.70 ± 1.56	34.85 ± 1.50	34.21 ± 1.79	34.00 ± 1.64
PESMO	319.91 ± 133.79	321.83 ± 135.60	335.36 ± 396.14	432.59 ± 334.89
EHVI	0.29 ± 1.19	1.08 ± 1.48	0.43 ± 1.39	0.3226 ± 1.35
q-{PF} <sup>2</sup> ES (q=2)	113.96 ± 85.54	128.08 ± 115.16	132.23 ± 215.71	117.43 ± 202.28
{PF} <sup>2</sup> ES-KB (q=2)	89.96 ± 79.00	90.43 ± 65.71	94.80 ± 80.17	94.80 ± 80.17
qEHVI (q=2)	6.49 ± 3.90	10.44 ± 9.90	4.98 ± 9.83	4.20 ± 8.69
	C-BraninCurrin	Constr-Ex	SRN	C2DTLZ2
{PF} <sup>2</sup> ES	63.05 ± 4.58	75.52 ± 5.92	75.81 ± 6.39	65.43 ± 4.54
EHVI-PoF	7.46 ± 2.9	10.44 ± 3.90	8.47 ± 3.32	6.81 ± 1.85
MESMOC	46.41 ± 2.59	54.47 ± 2.81	52.22 ± 2.93	48.89 ± 5.06
MESMOC+	58.52 ± 38.62	70.02 ± 54.44	70.78 ± 72.39	83.42 ± 72.88
PPESMOC (q = 2)	2152.51 ± 2981.10	2735.67 ± 3953.59	2372.81 ± 3791.25	4276.52 ± 7232.86
q-{PF} <sup>2</sup> ES (q=2)	186.72 ± 177.68	259.27 ± 167.52	98.13 ± 50.22	217.71 ± 217.48
{PF} <sup>2</sup> ES-KB (q=2)	170.56 ± 126.00	169.40 ± 156.65	196.04 ± 160.60	169.11 ± 94.08
qEHVI (q=2)	33.07 ± 14.43	30.30 ± 22.13	35.57 ± 14.25	32.99 ± 27.26

## I.2 Experimental results for larger batch sizes

Additional experimental results for larger batch sizes  $q$  are given in Fig. 14. Besides the ZDT1 and C2DTLZ2 synthetic functions reported in the main paper, we also add the results of two new synthetic experiments, namely the DTLZ4 problem and the Osy problem, see Table 1 for detailed settings. In general, with the same black box function evaluation budget, q-{PF}<sup>2</sup>ES demonstrates similar performance. It is also expected that with an increase in batch size, the acquisition function performance can be degraded slightly due to the increasing complexity of the acquisition function optimization process.

## I.3 Experimental results for $M > 2$ objectives

Two additional results for a larger number of objectives are seen in Fig. 15 based on two new problems. For the unconstrained version, we use DTLZ4 with 3 objective functions. For the constraint case, we use a real-life conceptual Marine design problem (Parsons and Scott, 2004; Tanabe and Ishibuchi, 2020). Details of the problem are also provided in Table 1. It can be observed that {PF}<sup>2</sup>ES and q-{PF}<sup>2</sup>ES provide competitive performance. For the conceptual Marine design problem, we generally observe that {PF}<sup>2</sup>ES-KB provides the fastest converge speed. We reason this is because given the relatively high constraint numbers, the approximated non-dominated feasible region  $\tilde{A}$  is relatively small. For Monte Carlo-based acquisition functions like q-{PF}<sup>2</sup>ES, this can impose difficulties for approximation and optimization since  $Z_{\tilde{A}_q}$  (Eq. 20) is rather small.

Note that, as discussed in Appendix. G, the computational complexity for providing the exact approximation of  $A$  will increase exponentially with the number of objectives  $M$ , meaning that it will become the dominant factor in the calculation cost of {PF}<sup>2</sup>ES as  $M$  increases. This is a common problem for hypervolume partition strategy-based acquisition functions (e.g., EHVI, MOPI).

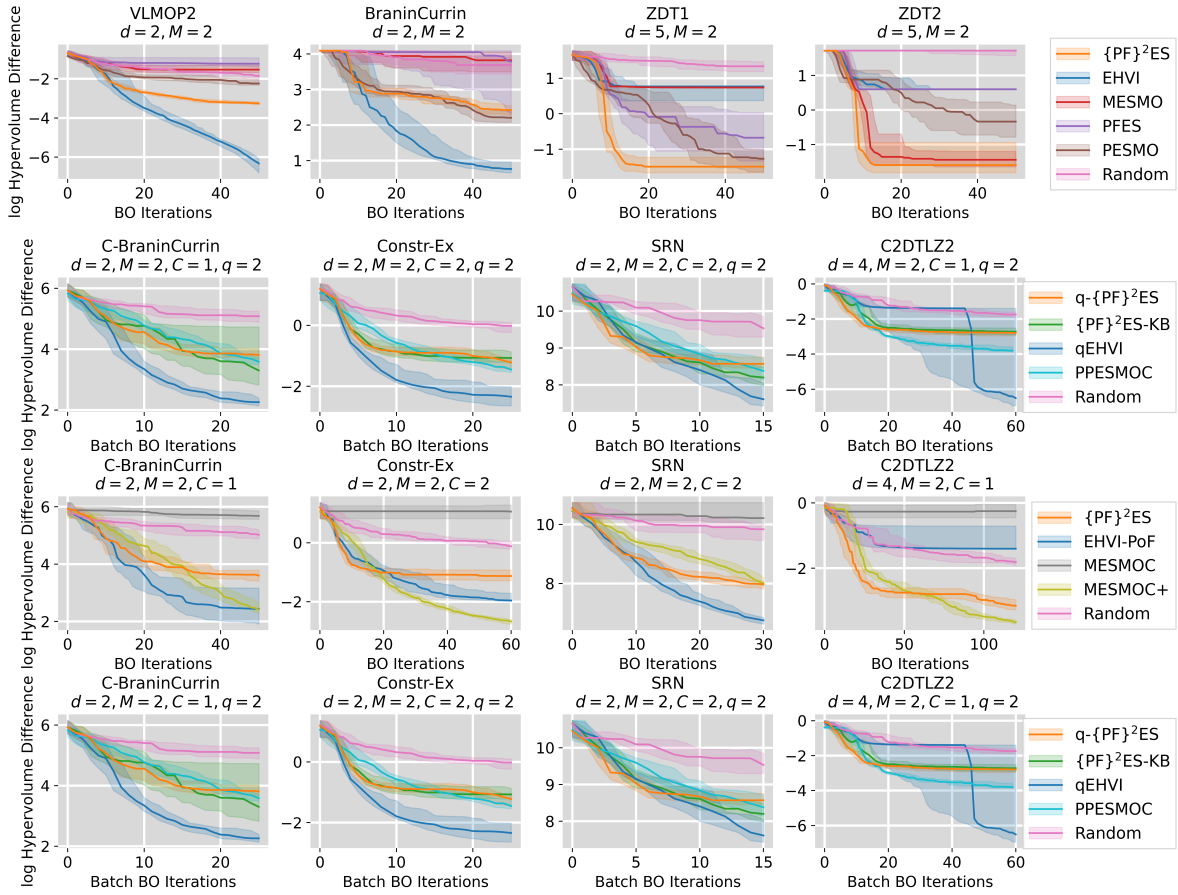


Figure 13: Comparison of different acquisition functions for multi-objective Bayesian optimization with respect to an in-sample recommendation strategy. In case of encountering a numerical issue (i.e., Cholesky decomposition issue) a different initial sample set is utilized.

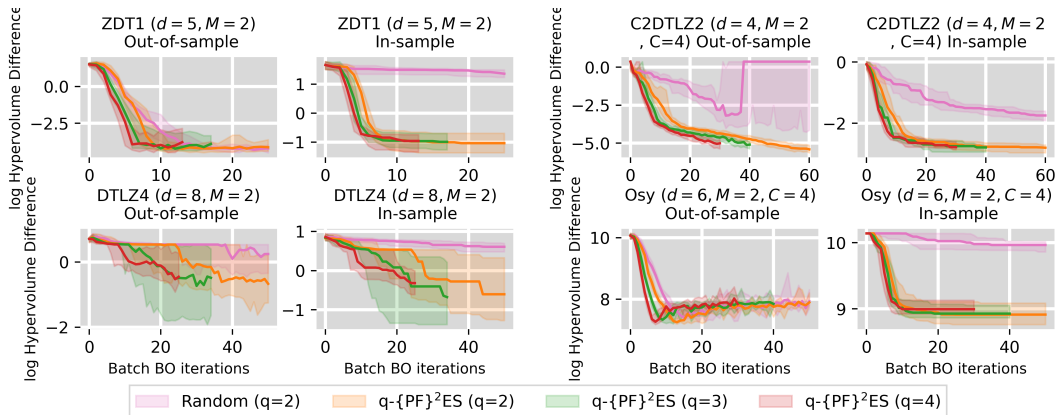


Figure 14: Numerical experiments of  $q\text{-}\{PF\}^2ES$  on larger batch sizes. Note all the batch methods terminate on roughly the same number of function evaluations.

**{PF}<sup>2</sup>ES: Parallel Feasible Pareto Frontier Entropy Search for Multi-Objective Bayesian Optimization**

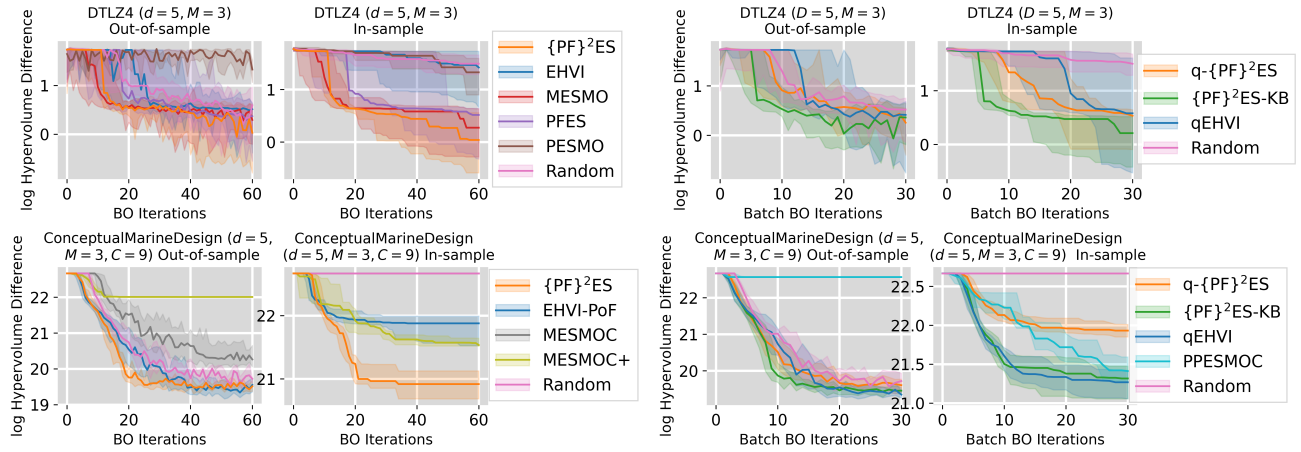


Figure 15: Experimental comparison for larger output dimensionality.