

---

# Information-Theoretic Multi-Objective Bayesian Optimization with Continuous Approximations

---

Syrine Belakaria, Aryan Deshwal, Janardhan Rao Doppa

School of EECS, Washington State University

{syrine.belakaria, aryan.deshwal, jana.doppa}@wsu.edu

## Abstract

Many real-world applications involve black-box optimization of multiple objectives using continuous function approximations that trade-off accuracy and resource cost of evaluation. For example, in rocket launching research, we need to find designs that trade-off return-time and angular distance using continuous-fidelity simulators (e.g., varying tolerance parameter to trade-off simulation time and accuracy) for design evaluations. The goal is to approximate the optimal Pareto set by minimizing the cost for evaluations. In this paper, we propose a novel approach referred to as *information-Theoretic Multi-Objective Bayesian Optimization with Continuous Approximations (iMOCA)* to solve this problem. The key idea is to select the sequence of input and function approximations for multiple objectives which maximize the information gain per unit cost for the optimal Pareto front. Our experiments on diverse synthetic and real-world benchmarks show that iMOCA significantly improves over existing single-fidelity methods.

## 1 Introduction

A large number of real-world engineering and scientific design problems involve optimizing *multiple* expensive black-box functions with *continuous approximations* (also called as fidelities) which vary in accuracy and resource cost. For example, in rocket launching research, the goal is to find designs that trade-off return time and angular distance using continuous-fidelity simulators which trade-off accuracy of evaluation and simulation time by varying the tolerance parameter. Solving these problems require finding the optimal Pareto set of design inputs (as all objectives cannot be optimized simultaneously) while minimizing the total resource cost of function evaluations.

For solving such expensive blackbox optimization problems, Bayesian optimization (BO) [39, 4, 17, 15, 16] is an effective framework. BO methods intelligently select the sequence of inputs for evaluation using the following three key elements: **1**) Surrogate statistical model learned from past function evaluations data, e.g., Gaussian process (GP); **2**) Acquisition function parametrized by the statistical model to score usefulness of evaluating different inputs, e.g., expected improvement (EI); and **3**) Acquisition function optimization to select the best input for function evaluation in each BO iteration. There are two broad categories of acquisition functions in the BO literature: *myopic* family accounts for local utility of improvement, e.g., EI; and *non-myopic* family relying on global utility for solving the optimization problem when evaluating a candidate input, e.g., predictive entropy search [23]. Recent work has shown the theoretical and practical benefits of non-myopic acquisition functions over myopic ones [23, 24, 46].

In this paper, we propose a novel approach referred as *information-Theoretic Multi-Objective Bayesian Optimization with Continuous Approximations (iMOCA)* to solve multi-objective optimization problems using continuous-fidelity function evaluations. *To the best of our knowledge, this is the first work to study this continuous fidelity in the multi-objective setting within the ML literature.* The key idea behind iMOCA is to select the sequence of candidate input and fidelity-vector pairs for

evaluation which maximize the information gained per unit cost about the optimal Pareto front (i.e., non-myopic acquisition function). To efficiently compute entropy, which is an important step for iMOCA, we develop two qualitatively different approximations that make different trade-offs in terms of computation-time and accuracy. iMOCA extends the single-fidelity multi-objective algorithm MESMO ([5]) and its discrete-fidelity version ([7]) to the more general continuous-fidelity setting.

**Contributions.** We make the following specific contributions.

- Development of a principled information-theoretic approach referred as iMOCA to solve multi-objective blackbox optimization problems using continuous function approximations. Providing two qualitatively different approximations for iMOCA.
- Experimental evaluation on diverse synthetic and real-world benchmark problems to demonstrate the effectiveness of iMOCA over existing single-fidelity MO algorithms and a naive continuous-fidelity MO approach.

## 2 Related work

**Multi-fidelity single-objective optimization.** acquisition function for single-fidelity and single-objective BO has been extensively studied ([39]). Canonical examples of myopic acquisition function include expected improvement (EI) and upper-confidence bound (UCB). EI was extended to multi-fidelity setting ([26, 36, 32]). The popular GP-UCB method ([42]) was also extended to multi-fidelity setting with discrete fidelities ([27]) and continuous fidelities ([28]).

Entropy based methods fall under the category of *non-myopic* acquisition function. Some examples include entropy search (ES) ([21]) and predictive entropy search (PES) ([23]). Their multi-fidelity extensions include MT-ES ([44, 29]) and MF-PES ([48, 33]). Unfortunately, they inherit the computational difficulties of the original ES and PES. Max-value entropy search (MES) and output space predictive entropy search ([46, 24]) are recent approaches that rely on the principle of output space entropy (OSE) search. Prior work ([46]) has shown advantages of OSE search in terms of compute-time, robustness, and accuracy over input space entropy search methods. Recent work ([41]) proposed a general approach based on mutual information. [45] extended MES to multi-fidelity setting and showed its effectiveness over MF-PES. [34] extended MES to the continuous fidelity and multi-task setting

**Single-fidelity multi-objective optimization.** Multi-objective algorithms can be classified into three families. *Scalarization methods* are model-based algorithms that reduce the problem to single-objective optimization. ParEGO method ([30]) employs random scalarization for this purpose. ParEGO is simple and fast, but more advanced approaches often outperform it. *Pareto hypervolume optimization methods* optimize the Pareto hypervolume (PHV) metric ([18]) that captures the quality of a candidate Pareto set. This is done by extending the standard acquisition functions to PHV objective, e.g., expected improvement in PHV ([18]) and probability of improvement in PHV ([35]). Unfortunately, algorithms to optimize PHV based acquisition functions scale very poorly and are not feasible for more than three objectives. To improve scalability, methods to reduce the search space are also explored ([37]). A common drawback of this family is that reduction to single-objective optimization can potentially lead to more exploitative behavior.

*Uncertainty reduction methods* like PAL ([51]), PESMO ([22]) and the concurrent works USeMO ([9]) and MESMO ([5]) are principled algorithms based on information theory. In each iteration, PAL selects the candidate input for evaluation towards the goal of minimizing the size of uncertain set. PAL provides theoretical guarantees, but it is only applicable for input space  $\mathcal{X}$  with finite set of discrete points. USeMO is a general framework that iteratively generates a cheap Pareto front using the surrogate models and then selects the point with highest uncertainty as the next query. USeMOC [8] is an extension of USeMO for handling constraints and was applied to design analog circuits [11]. PESMO relies on input space entropy search and iteratively selects the input that maximizes the information gained about the optimal Pareto set  $\mathcal{X}^*$ . Unfortunately, optimizing this acquisition function poses significant challenges: a) requires a series of approximations, which can be potentially sub-optimal; and b) optimization, even after approximations, is expensive c) performance is strongly dependent on the number of Monte-Carlo samples. MESMO ([5]) relies on output space entropy search and its advantages over PESMO were demonstrated in a recent work [5]. MESMOC [6] is an

extension of MESMO for handling constraints and was applied to design electrified aviation power systems [10].

**Multi-fidelity multi-objective optimization.** Prior work outside ML literature has considered domain-specific methods that employ single-fidelity multi-objective approaches in the context of multi-fidelity setting by using the lower fidelities *only as an initialization* ([31, 2]). Specifically, ([2]) employs the single-fidelity algorithm based on expected hypervolume improvement acquisition function and ([31]) employs an algorithm that is very similar to SMSego. Additionally, both these methods model all fidelities with the same GP and assume that higher fidelity evaluation is a sum of lower-fidelity evaluation and offset error. These are strong assumptions and may not hold in general multi-fidelity settings including the problems we considered in our experimental evaluation. A recent work [7] proposed a generic approach by generalizing the MESMO algorithm [5] for the discrete-fidelity setting. However, this approach assumes the existence of different functions for each fidelity which makes it hard to extend/apply to the continuous-fidelity setting. We are not aware of any continuous-fidelity algorithms for MO problems.

### 3 Problem Setup

**Multi-objective optimization with continuous function approximations.** Suppose  $\mathcal{X} \subseteq \mathbb{R}^d$  is a continuous input space. In multi-objective optimization (MO) problems, our goal is to minimize  $K \geq 2$  expensive objective functions  $f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_K(\mathbf{x})$ . The evaluation of each input  $\mathbf{x} \in \mathcal{X}$  results in a vector of  $K$  function values  $\mathbf{y} = (y_1, y_2, \dots, y_K)$ , where  $y_i = f_i(x)$  for all  $i \in \{1, 2, \dots, K\}$ . An input  $\mathbf{x} \in \mathcal{X}$  is said to *Pareto-dominate* another input  $\mathbf{x}' \in \mathcal{X}$  if  $f_i(\mathbf{x}) \leq f_i(\mathbf{x}') \forall i$  and there exists some  $j \in \{1, 2, \dots, K\}$  such that  $f_j(\mathbf{x}) < f_j(\mathbf{x}')$ . The solution for this MO problem referred as optimal *Pareto set*  $\mathcal{X}^* \subset \mathcal{X}$  is a set of non-dominated inputs and optimal *Pareto front* is the corresponding set of function value vectors.

In continuous-fidelity MO problems, we have access to  $g_i(\mathbf{x}, z_i)$  where  $g_i$  is an alternative function through which we can evaluate cheaper approximations of  $f_i$  by varying the fidelity variable  $z_i \in \mathcal{Z}$ . Without loss of generality, let  $\mathcal{Z}=[0, 1]$  be the fidelity space. Fidelities for each function  $f_i$  vary in the amount of resources consumed and the accuracy of evaluation, where  $z_i=0$  and  $z_i^*=1$  refer to the lowest and highest fidelity respectively. At the highest fidelity  $z_i^*$ ,  $g_i(\mathbf{x}, z_i^*) = f_i(\mathbf{x})$ . Let  $\mathcal{C}_i(\mathbf{x}, z_i)$  be the cost of evaluating  $g_i(\mathbf{x}, z_i)$ . Evaluation of an input  $\mathbf{x} \in \mathcal{X}$  with fidelity vector  $\mathbf{z} = [z_1, z_2, \dots, z_K]$  produces an evaluation vector of  $K$  values denoted by  $\mathbf{y} \equiv [y_1, y_2, \dots, y_K]$ , where  $y_i = g_i(\mathbf{x}, z_i)$  for all  $i \in \{1, 2, \dots, K\}$ , and the normalized cost of evaluation is  $\mathcal{C}(\mathbf{x}, \mathbf{z}) = \sum_{i=1}^K (\mathcal{C}_i(\mathbf{x}, z_i)/\mathcal{C}_i(\mathbf{x}, z_i^*))$ . We normalize the cost of each function by the cost of its highest fidelity because the cost units of different objectives can be different. If the cost is known, it can be directly injected in the latter expression. However, in some real world setting, the cost of a function evaluation can be only known after the function evaluation. For example, in hyper-parameter tuning of a neural network, the cost of the experiment is defined by the training and inference time. However, we cannot know the exact needed time until after the experiment is finalised. In this case, the cost can be modeled by an independent Gaussian process. The predictive mean can be used during the optimization. The final goal is to recover  $\mathcal{X}^*$  while minimizing the total cost of function evaluations.

**Continuous-fidelity GPs as surrogate models.** Let  $D = \{(\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i)\}_{i=1}^{t-1}$  be the training data from past  $t-1$  function evaluations, where  $\mathbf{x}_i \in \mathcal{X}$  is an input and  $\mathbf{y}_i = [y_1, y_2, \dots, y_K]$  is the output vector resulting from evaluating functions  $g_1, g_2, \dots, g_K$  for  $\mathbf{x}_i$  at fidelities  $z_1, z_2, \dots, z_K$  respectively. We learn  $K$  surrogate statistical models  $\mathcal{GP}_1, \dots, \mathcal{GP}_K$  from  $D$ , where each model  $\mathcal{GP}_j$  corresponds to the  $j$ th function  $g_j$ . Continuous fidelity GPs (CF-GPs) are capable of modeling functions with continuous fidelities within a single model. Hence, we employ CF-GPs to build surrogate statistical models for each function. Specifically, we use the CF-GP model proposed in [28]. W.l.o.g, we assume that our functions  $g_j$  are defined over the spaces  $\mathcal{X} = [0, 1]^d$  and  $\mathcal{Z} = [0, 1]$ . Let  $g_j \sim \mathcal{GP}_j(0, \kappa_j)$  such that  $y_j = g_j(z_j, \mathbf{x}) + \epsilon$ , where  $\epsilon \sim \mathcal{N}(0, \eta^2)$  and  $\kappa : (\mathcal{Z} \times \mathcal{X})^2 \rightarrow \mathbb{R}$  is the prior covariance matrix defined on the product of input and fidelity spaces.

$$\kappa_j([z_j, \mathbf{x}], [z'_j, \mathbf{x}']) = \kappa_{j\mathcal{X}}(\mathbf{x}, \mathbf{x}') \cdot \kappa_{j\mathcal{Z}}(z_j, z'_j)$$

where  $\kappa_{j\mathcal{X}}, \kappa_{j\mathcal{Z}}$  are radial kernels over  $\mathcal{X}$  and  $\mathcal{Z}$  spaces respectively.  $\mathcal{Z}$  controls the smoothness of  $g_j$  over the fidelity space to be able to share information across different fidelities. A key advantage of this model is that it integrates all fidelities into one single GP for information sharing. We denote

the posterior mean and standard deviation of  $g_j$  conditioned on  $D$  by  $\mu_{g_j}(\mathbf{x}, z_j)$  and  $\sigma_{g_j}(\mathbf{x}, z_j)$ . We denote the posterior mean and standard deviation of the highest fidelity functions  $f_j(\mathbf{x}) = g_j(\mathbf{x}, z_j^*)$  by  $\mu_{f_j}(\mathbf{x}) = \mu_{g_j}(\mathbf{x}, z_j^*)$  and  $\sigma_{f_j}(\mathbf{x}) = \sigma_{g_j}(\mathbf{x}, z_j^*)$  respectively. We define  $\sigma_{g_j, f_j}^2(x)$  as the predictive co-variance between a lower fidelity  $z_j$  and the highest fidelity  $z_j^*$  at the same  $\mathbf{x}$ .

**Table of Notations.** For the sake of reader, Table 1 provides all the mathematical notations used in this paper.

Notation	Definition
$\mathbf{x}, \mathbf{y}, \mathbf{z}$	Bold notation represents vectors
$f_1, f_2, \dots, f_K$	Highest fidelity objective functions
$g_1, g_2, \dots, g_K$	General objective functions with low and high fidelities
$\tilde{f}_j$	Function sampled from the highest fidelity of the $j$ th Gaussian process model
$\tilde{g}_j$	Function sampled from the $j$ th Gaussian process model at fidelity $z_j$
$\mathbf{x}$	Input vector
$z_1, z_2, \dots, z_K$	The fidelity variables for each function
$\mathbf{z}$	Fidelity vector
$\mathbf{z}^* = [z_1^*, z_2^*, \dots, z_K^*]$	Fidelity vector with all fidelities at their highest value
$y_j$	$j$ th function $g_j$ evaluated at fidelity $z_j$
$\mathbf{y} = [y_1, y_2, \dots, y_K]$	Output vector resulting from evaluating $g_1, g_2, \dots, g_K$ for $\mathbf{x}_i$ at fidelities $z_1, z_2, \dots, z_K$ respectively
$\mathbf{f} = [f_1, f_2, \dots, f_K]$	Output vector resulting from evaluating functions $f_1, f_2, \dots, f_K$ or equivalently $g_1, g_2, \dots, g_K$ for $\mathbf{x}_i$ at fidelities $z_1^*, z_2^*, \dots, z_K^*$ respectively
$\mathcal{F}^*$	true Pareto front of the highest fidelity functions $[f_1, f_2, \dots, f_K]$
$\mathcal{F}_s^*$	Pareto front of the sampled highest fidelity functions $[f_1, f_2, \dots, f_K]$
$\mathcal{C}_j(\mathbf{x}, z_j)$	cost of evaluating $j$ th function $g_j$ at fidelity $z_j$
$\mathcal{C}(\mathbf{x}, \mathbf{z})$	total normalized cost $\mathcal{C}(\mathbf{x}, \mathbf{z}) = \sum_{i=1}^K (C_i(\mathbf{x}, z_i)/C_i(\mathbf{x}, z_i^*))$
$\mathcal{X}$	Input space
$\mathcal{Z}$	Fidelity space
$\mathcal{Z}_t^{(j)}$	Reduced fidelity space for function $g_j$ at iteration $t$
$\mathcal{Z}_r$	Reduced fidelity space over all the functions
$\xi$	Information gap
$\beta_t^{(j)}$	Exploration/exploitation parameter for function $g_j$ at iteration $t$
$I$	Information gain

Table 1: Mathematical notations and their associated definition.

## 4 iMOCA Algorithm with Two Approximations

We first describe the key idea behind our proposed iMOCA algorithm including the main challenges. Next, we present our algorithmic solution to address those challenges.

**Key Idea of iMOCA:** The acquisition function behind iMOCA employs principle of output space entropy search to select the sequence of input and fidelity-vector (one for each objective) pairs. iMOCA is applicable for solving MO problems in both continuous and discrete fidelity settings. *The key idea is to find the pair  $\{\mathbf{x}_t, \mathbf{z}_t\}$  that maximizes the information gain  $I$  per unit cost about the Pareto front of the highest fidelities (denoted by  $\mathcal{F}^*$ ), where  $\{\mathbf{x}_t, \mathbf{z}_t\}$  represents a candidate input  $\mathbf{x}_t$  evaluated at a vector of fidelities  $\mathbf{z}_t = [z_1, z_2, \dots, z_K]$  at iteration  $t$ . Importantly, iMOCA performs joint search over input space  $\mathcal{X}$  and reduced fidelity space  $\mathcal{Z}_r$  over fidelity vectors for this selection.*

$$(\mathbf{x}_t, \mathbf{z}_t) \leftarrow \arg \max_{\mathbf{x} \in \mathcal{X}, \mathbf{z} \in \mathcal{Z}_r} \alpha_t(\mathbf{x}, \mathbf{z}), \text{ where } \alpha_t(\mathbf{x}, \mathbf{z}) = I(\{\mathbf{x}, \mathbf{y}, \mathbf{z}\}, \mathcal{F}^* | D) / \mathcal{C}(\mathbf{x}, \mathbf{z}) \quad (4.1)$$

In the following sections, we describe the details and steps of our proposed algorithm iMOCA. We start by explaining the bottlenecks of continuous fidelity optimization due to the infinite size of the fidelity space followed by describing a principled approach to reduce the fidelity space. Subsequently, we present the computational steps of our proposed acquisition function: Information gain per unit cost for each candidate input and fidelity-vector pair.

## 4.1 Approach to Reduce Fidelity Search Space

In this work, we focus primarily on MO problems with continuous fidelity space. The continuity of this space results in infinite number of fidelity choices. Thus, selecting an informative and meaningful fidelity becomes a major bottleneck. Therefore, we reduce the search space over fidelity-vector variables in a principled manner guided by the learned statistical models [28]. Our fidelity space reduction method is inspired from BOCA for single-objective optimization [28]. We apply the method in BOCA to each of the objective functions to be optimized in MO setting.

A favourable setting for continuous-fidelity methods would be for the lower fidelities  $g_j$  to be informative about the highest fidelity  $f_j$ . Let  $h_j$  be the bandwidth parameter of the fidelity kernel  $\kappa_{j\mathcal{Z}}$  and let  $\xi : \mathcal{Z} \rightarrow [0, 1]$  be a measure of the gap in information about  $g_j(\cdot, z_j^*)$  when queried at  $z_j \neq z_j^*$  with  $\xi(z_j) \approx \frac{\|z_j - z_j^*\|}{h_j}$  for the squared exponential kernels [28]. A larger  $h_j$  will result in  $g_j$  being smoother across  $\mathcal{Z}$ . Consequently, lower fidelities will be more informative about  $f_j$  and the information gap  $\xi(z_j)$  will be smaller.

To determine an informative fidelity for each function in iteration  $t$ , we reduce the space  $\mathcal{Z}$  and select  $z_j$  from the subset  $\mathcal{Z}_t^{(j)}$  defined as follows:

$$\mathcal{Z}_t^{(j)}(\mathbf{x}) = \{\{z_j \in \mathcal{Z} \setminus \{z_j^*\}, \sigma_{g_j}(\mathbf{x}, z_j) > \gamma(z_j), \xi(z_j) > \beta_t^{(j)}\|\xi\|_\infty\} \cup \{z_j^*\}\} \quad (4.2)$$

where  $\gamma(z_j) = \xi(z_j) \left(\frac{C_j(\mathbf{x}, z_j)}{C_j(\mathbf{x}, z_j^*)}\right)^q$  and  $q = \frac{1}{p_j + d + 2}$  with  $p_j, d$  the dimensions of  $\mathcal{Z}$  and  $\mathcal{X}$  respectively. Without loss of generality, we assume that  $p_j = 1$ .  $\beta_t^{(j)} = \sqrt{0.5 \ln\left(\frac{2t+1}{h_j}\right)}$  is the exploration/exploitation parameter [42]. We denote by  $\mathcal{Z}_r = \{\mathcal{Z}_t^{(j)}, j \in \{1 \dots K\}\}$ , the reduced fidelity space for all  $K$  functions.

We filter out the fidelities for each objective function at BO iteration  $t$  using the above-mentioned two conditions. We provide intuitive explanation of these conditions below.

**The first condition**  $\sigma_{g_j}(\mathbf{x}, z_j) > \gamma(z_j)$ : A reasonable multi-fidelity strategy would query the cheaper fidelities in the beginning to learn about the function  $g_j$  by consuming the least possible cost budget and later query from higher fidelities in order to gain more accurate information. Since the final goal is to optimize  $f_j$ , the algorithm should also query from the highest-fidelity. However, the algorithm might never query from higher fidelities due to their high cost. This condition will make sure that lower fidelities are likely to be queried, but not excessively and the algorithm will move toward querying higher fidelities as iterations progress. Since  $\gamma(z_j)$  is monotonically increasing in  $C_j$ , this condition can be easily satisfied by cheap fidelities. However, if a fidelity is very far from  $z_j^*$ , then the information gap  $\xi$  will increase and hence, uninformative fidelities would be discarded. Therefore,  $\gamma(z_j)$  will guarantee achieving a good trade-off between resource cost and information.

**The second condition**  $\xi(z_j) > \beta_t^{(j)}\|\xi\|_\infty$ : We recall that if the first subset of  $\mathcal{Z}_t^{(j)}$  is empty, the algorithm will automatically evaluate the highest-fidelity  $z_j^*$ . However, if it is not empty, and since the fidelity space is continuous (infinite number of choices for  $z_j$ ), the algorithm might query fidelities that are very close to  $z_j^*$  and would cost nearly the same as  $z_j^*$  without being as informative as  $z_j^*$ . The goal of this condition is to prevent such situations by excluding fidelities in the small neighborhood of  $z_j^*$  and querying  $z_j^*$  instead. Since  $\beta_t^{(j)}$  increases with  $t$  and  $\xi$  is increasing as we move away from  $z_j^*$ , this neighborhood is shrinking and the algorithm will eventually query  $z_j^*$ .

### 4.1.1 Naive-CFMO: A Simple Continuous-Fidelity MO Baseline

In this section, we first describe a simple baseline approach referred as *Naive-CFMO* to solve continuous-fidelity MO problems by combining the above-mentioned fidelity space reduction approach with existing multi-objective BO methods. Next, we explain the key drawbacks of Naive-CFMO and how our proposed iMOCA algorithm overcomes them.

A straightforward way to construct a continuous-fidelity MO method is to perform a two step selection process similar to the continuous-fidelity single-objective BO algorithm proposed in [28]:

**Step 1)** Select the input  $\mathbf{x}$  that maximizes the acquisition function at the *highest fidelity*. This can be done using any existing multi-objective BO algorithm.

**Step 2)** Evaluate  $\mathbf{x}$  at the *cheapest valid fidelity for each function* in the reduced fidelity space  $\mathcal{Z}_t^{(j)}(\mathbf{x})$  computed using the reduction approach mentioned in the previous section. Since we are studying information gain based methods in this work, we instantiate Naive-CFMO using the state-of-the-art information-theoretic MESMO algorithm [5] for Step 1. Algorithm 2 shows the complete pseudo-code of Naive-CFMO.

**Drawbacks of Naive-CFMO:** Unfortunately, Naive-CFMO has two major drawbacks.

- The acquisition function solely relies on the highest-fidelity  $f_j$ . Therefore, it does not capture and leverage the statistical relation between different fidelities and full-information provided by the global function  $g_j$ .
- Generally, there is a dependency between the fidelity space and the input space in continuous-fidelity problems. Therefore, selecting an input that maximizes the highest-fidelity and then evaluating it at a different fidelity can result in a mismatch in the evaluation process leading to poor performance and slower convergence.

**iMOCA vs. Naive-CFMO:** Our proposed iMOCA algorithm overcomes the drawbacks of Naive-CFMO as follows.

- iMOCA’s acquisition function maximizes the information gain per unit cost across all fidelities by capturing the relation between fidelities and the impact of resource cost on information gain.
- iMOCA performs joint search over input and fidelity space to select the input variable  $\mathbf{x} \in \mathcal{X}$  and fidelity variables  $\mathbf{z} \in \mathcal{Z}_r$  while maximizing the proposed acquisition function. Indeed, our experimental results demonstrate the advantages of iMOCA over Naive-CFMO.

## 4.2 Information-Theoretic Continuous-Fidelity Acquisition Function

In this section, we explain the technical details of the acquisition function behind iMOCA. We propose two approximations for the computation of information gain per unit cost.

The information gain in equation 4.1 is defined as the expected reduction in entropy  $H(\cdot)$  of the posterior distribution  $P(\mathcal{F}^*|D)$  due to evaluating  $\mathbf{x}$  at fidelity vector  $\mathbf{z}$ . Based on the symmetric property of information gain, the latter can be rewritten as follows:

$$I(\{\mathbf{x}, \mathbf{y}, \mathbf{z}\}, \mathcal{F}^*|D) = H(\mathbf{y}|D, \mathbf{x}, \mathbf{z}) - \mathbb{E}_{\mathcal{F}^*}[H(\mathbf{y}|D, \mathbf{x}, \mathbf{z}, \mathcal{F}^*)] \quad (4.3)$$

In equation 4.3, the first term is the entropy of a  $K$ -dimensional Gaussian distribution that can be computed in closed form as follows:

$$H(\mathbf{y}|D, \mathbf{x}, \mathbf{z}) = \sum_{j=1}^K \ln(\sqrt{2\pi e} \sigma_{g_j}(\mathbf{x}, z_j)) \quad (4.4)$$

In equation 4.3, the second term is an expectation over the Pareto front of the highest fidelities  $\mathcal{F}^*$ . This term can be approximated using Monte-Carlo sampling:

$$\mathbb{E}_{\mathcal{F}^*}[H(\mathbf{y}|D, \mathbf{x}, \mathbf{z}, \mathcal{F}^*)] \simeq \frac{1}{S} \sum_{s=1}^S [H(\mathbf{y}|D, \mathbf{x}, \mathbf{z}, \mathcal{F}_s^*)] \quad (4.5)$$

where  $S$  is the number of samples and  $\mathcal{F}_s^*$  denote a sample Pareto front obtained over the highest fidelity functions sampled from  $K$  surrogate models. To compute Equation 4.5, we provide algorithmic solutions to construct Pareto front samples  $\mathcal{F}_s^*$  and to compute the entropy with respect to a given Pareto front sample  $\mathcal{F}_s^*$ .

**Computing Pareto front samples:** We first sample highest fidelity functions  $\tilde{f}_1, \dots, \tilde{f}_K$  from the posterior CF-GP models via random fourier features [23, 38]. This is done similar to prior work [22, 46]. We solve a cheap MO optimization problem over the  $K$  sampled functions  $\tilde{f}_1, \dots, \tilde{f}_K$  using the popular NSGA-II algorithm [14] to compute the sample Pareto front  $\mathcal{F}_s^*$ .

**Entropy computation for a given Pareto front sample:** Let  $\mathcal{F}_s^* = \{\mathbf{v}^1, \dots, \mathbf{v}^l\}$  be the sample Pareto front, where  $l$  is the size of the Pareto front and each  $\mathbf{v}^i = \{v_1^i, \dots, v_K^i\}$  is a  $K$ -vector

evaluated at the  $K$  sampled highest-fidelity functions. The following inequality holds for each component  $y_j$  of  $\mathbf{y} = (y_1, \dots, y_K)$  in the entropy term  $H(\mathbf{y}|D, \mathbf{x}, \mathbf{z}, \mathcal{F}_s^*)$ :

$$y_j \leq f_s^{j*} \quad \forall j \in \{1, \dots, K\} \quad (4.6)$$

where  $f_s^{j*} = \max\{v_j^1, \dots, v_j^l\}$ . Essentially, this inequality says that the  $j^{\text{th}}$  component of  $\mathbf{y}$  (i.e.,  $y_j$ ) is upper-bounded by a value, which is the maximum of  $j^{\text{th}}$  components of all  $l$  vectors  $\{\mathbf{v}^1, \dots, \mathbf{v}^l\}$  in the Pareto front  $\mathcal{F}_s^*$ . For the ease of notation, we drop the dependency on  $\mathbf{x}$  and  $\mathbf{z}$ . We use  $f_j$  to denote  $f_j(x) = g_j(x, z_j^*)$ , the evaluation of the highest fidelity at  $x$  and  $y_j$  to denote  $g_j(x, z_j)$  the evaluation of  $g_j$  at a lower fidelity  $z_j \neq z_j^*$ .

The proof of 4.6 can be divided into two cases: **a)** If  $y_j$  is evaluated at its highest fidelity (i.e.,  $z_j = z_j^*$  and  $y_j = f_j$ ), we provide a proof by contradiction for inequality 4.6. Suppose there exists some component  $f_j$  of  $\mathbf{f}$  such that  $f_j > f_s^{j*}$ . However, by definition,  $\mathbf{f}$  is a non-dominated point because no point dominates it in the  $j^{\text{th}}$  dimension. This results in  $\mathbf{f} \in \mathcal{F}_s^*$  which is a contradiction. Therefore, inequality 4.6 holds. **b)** If  $y_j$  is evaluated at one of its lower fidelities (i.e.,  $z_j \neq z_j^*$ ), the proof follows from the assumption that the value of lower fidelity of an objective is usually smaller than the corresponding higher fidelity, i.e.,  $y_j \leq f_j \leq f_s^{j*}$ . This is especially true for real-world optimization problems. For example, in optimizing the accuracy of a neural network with respect to its hyper-parameters, a commonly employed fidelity is the number of data samples used for training. It is reasonable to believe that the accuracy is always higher for the higher fidelity.

By combining the inequality 4.6 and the fact that each function is modeled as an independent CF-GP, a common property of entropy measure allows us to decompose the entropy of a set of independent variables into a sum over entropies of individual variables [13]:

$$H(\mathbf{y}|D, \mathbf{x}, \mathbf{z}, \mathcal{F}_s^*) \simeq \sum_{j=1}^K H(y_j|D, \mathbf{x}, z_j, f_s^{j*}) \quad (4.7)$$

The computation of 4.7 requires the computation of the entropy of  $p(y_j|D, \mathbf{x}, z_j, f_s^{j*})$ . This is a conditional distribution that depends on the value of  $z_j$  and can be expressed as  $H(y_j|D, \mathbf{x}, z_j, y_j \leq f_s^{j*})$ . This entropy can be computed using two different approximations as described below:

**Truncated Gaussian approximation (iMOCA-T):** As a consequence of 4.6, which states that  $y_j \leq f_s^{j*}$  also holds for all fidelities, the entropy of  $p(y_j|D, \mathbf{x}, z_j, f_s^{j*})$  can also be approximated by the entropy of a truncated Gaussian distribution and expressed as follows:

$$H(y_j|D, \mathbf{x}, z_j, y_j \leq f_s^{j*}) = \ln(\sqrt{2\pi e} \sigma_{g_j}) + \ln \Phi(\gamma_s^{(g_j)}) - \frac{\gamma_s^{(g_j)} \phi(\gamma_s^{(g_j)})}{2\Phi(\gamma_s^{(g_j)})} \quad \text{where } \gamma_s^{(g_j)} = \frac{f_s^{j*} - \mu_{g_j}}{\sigma_{g_j}} \quad (4.8)$$

From equations 4.5, 4.4, and 4.8, we get the final expression of iMOCA-T as shown below:

$$\alpha_t(\mathbf{x}, \mathbf{z}, \mathcal{F}^*) = \frac{1}{\mathcal{C}(\mathbf{x}, \mathbf{z})S} \sum_{j=1}^K \sum_{s=1}^S \frac{\gamma_s^{(g_j)} \phi(\gamma_s^{(g_j)})}{2\Phi(\gamma_s^{(g_j)})} - \ln(\Phi(\gamma_s^{(g_j)})) \quad (4.9)$$

**Extended-skew Gaussian approximation (iMOCA-E):** Although equation 4.9 is sufficient for computing the entropy, this entropy can be mathematically interpreted and computed with a different approximation. The condition  $y_j \leq f_s^{j*}$ , is originally expressed as  $f_j \leq f_s^{j*}$ . Substituting this condition with it's original equivalent, the entropy becomes  $H(y_j|D, \mathbf{x}, z_j, f_j \leq f_s^{j*})$ . Since  $y_j$  is an evaluation of the function  $g_j$  while  $f_j$  is an evaluation of the function  $f_j$ , we observe that  $y_j|f_j \leq f_s^{j*}$  follows an extended-skew Gaussian (ESG) distribution [34, 3]. It has been shown that the differential entropy of an ESG does not have a closed form expression [1]. Therefore, we derive a simplified expression where most of the terms are analytical by manipulating the components of the entropy. We use the derivation of the entropy based on ESG formulation, proposed in ([34]), for the multi-objective setting.

In order to simplify the calculation  $H(y_j|D, \mathbf{x}, z_j, f_j \leq f_s^{j*})$ , let us define the normalized variable  $\Gamma_{f_s^{j*}}$  as  $\Gamma_{f_s^{j*}} = \frac{y_j - \mu_{g_j}}{\gamma_{g_j}} | f_j \leq f_s^{j*}$ .  $\Gamma_{f_s^{j*}}$  is an ESG with p.d.f whose mean  $\mu_{\Gamma_{f_s^{j*}}}$  and variance  $\sigma_{\Gamma_{f_s^{j*}}}$  are defined in Appendix A. We define the predictive correlation between  $y_j$  and  $f_j$  as  $\tau = \frac{\sigma_{g_j}^2 \cdot f_j}{\sigma_{g_j} \sigma_{f_j}}$ . The entropy can be computed using the following expression. Due to lack of space, we only provide

the final expression. Complete derivation for equations 4.10 and 4.11 are provided in Appendix A.

$$H(y_j|D, \mathbf{x}, z_j, f_j \leq f_s^{j*}) = \ln(\sqrt{2\pi e} \sigma_{g_j}) + \ln(\Phi(\gamma_s^{(f_j)})) - \tau^2 \frac{\phi(\gamma_s^{(f_j)}) \gamma_s^{(f_j)}}{2\Phi(\gamma_s^{(f_j)})} - \mathbb{E}_{u \sim \Gamma_{f_s^{j*}}} \left[ \ln\left(\Phi\left(\frac{\gamma_s^{(f_j)} - \tau u}{\sqrt{1 - \tau^2}}\right)\right) \right] \quad (4.10)$$

From equations 4.5, 4.4 and 4.10, the final expression of iMOCA-E can be expressed as follow:

$$\alpha_t(\mathbf{x}, \mathbf{z}, \mathcal{F}^*) = \frac{1}{\mathcal{C}(\mathbf{x}, \mathbf{z})S} \sum_{j=1}^K \sum_{s=1}^S \tau^2 \frac{\gamma_s^{(f_j)} \phi(\gamma_s^{(f_j)})}{2\Phi(\gamma_s^{(f_j)})} - \ln(\Phi(\gamma_s^{(f_j)})) + \mathbb{E}_{u \sim \Gamma_{f_s^{j*}}} \left[ \ln\left(\Phi\left(\frac{\gamma_s^{(f_j)} - \tau u}{\sqrt{1 - \tau^2}}\right)\right) \right] \quad (4.11)$$

The expression given by equation 4.11 is mostly analytical except for the last term. We perform numerical integration via Simpson's rule using  $\mu_{\Gamma_{f_s^{j*}}} \mp \gamma \sqrt{\sigma(\Gamma_{f_s^{j*}})}$  as the integral limits. Since this integral is over one-dimension variable, numerical integration can result in a tight approximation with low computational cost. Complete pseudo-code of iMOCA is shown in Algorithm 1.

**Generality of the two approximations:** We observe that for any fixed value of  $\mathbf{x}$ , when we choose the highest-fidelity for each function  $\mathbf{z}=\mathbf{z}^*$ : **a)** For iMOCA-T, we will have  $g_i = f_j$ ; and **b)** For iMOCA-E, we will have  $\tau = 1$ . Consequently, both equation 4.9 and equation 4.11 will degenerate to the acquisition function of MESMO optimizing only highest-fidelity functions:

$$\alpha_t(\mathbf{x}, \mathcal{F}^*) = \frac{1}{S} \sum_{j=1}^K \sum_{s=1}^S \frac{\gamma_s^{(f_j)} \phi(\gamma_s^{(f_j)})}{2\Phi(\gamma_s^{(f_j)})} - \ln(\Phi(\gamma_s^{(f_j)})) \quad (4.12)$$

The main advantages of our proposed acquisition function are: cost-efficiency, computational-efficiency, and robustness to the number of Monte-Carlo samples. Indeed, our experiments demonstrate these advantages over state-of-the-art single-fidelity MO algorithms.

---

#### Algorithm 1 iMOCA Algorithm

---

**Input:** input space  $\mathcal{X}$ ;  $K$  blackbox functions  $f_j$  and their continuous approximations  $g_j$ ; total budget  $\mathcal{C}_{total}$

- 1: Initialize continuous fidelity gaussian process  $\mathcal{GP}_1, \dots, \mathcal{GP}_K$  by initial points  $D$
- 2: **While**  $\mathcal{C}_t \leq \mathcal{C}_{total}$  **do**
- 3:   for each sample  $s \in 1, \dots, S$ :
- 4:     Sample highest-fidelity functions  $\tilde{f}_j \sim \mathcal{GP}_j(\cdot, z_j^*)$
- 5:      $\mathcal{F}_s^* \leftarrow$  Solve *cheap* MOO over  $(\tilde{f}_1, \dots, \tilde{f}_K)$
- 6:     Find the query based on  $\mathcal{F}^* = \{\mathcal{F}_s^*, s \in \{1 \dots S\}\}$
- 7:     // Choose one of the two approximations
- 8:     **If** approx = T // Use eq 4.9 for  $\alpha_t$  (iMOCA-T)
- 9:       select  $(\mathbf{x}_t, \mathbf{z}_t) \leftarrow \arg \max_{\mathbf{x} \in \mathcal{X}, \mathbf{z} \in \mathcal{Z}_r} \alpha_t(\mathbf{x}, \mathbf{z}, \mathcal{F}^*)$
- 10:     **If** approx = E // Use eq 4.11 for  $\alpha_t$  (iMOCA-E)
- 11:       select  $(\mathbf{x}_t, \mathbf{z}_t) \leftarrow \arg \max_{\mathbf{x} \in \mathcal{X}, \mathbf{z} \in \mathcal{Z}_r} \alpha_t(\mathbf{x}, \mathbf{z}, \mathcal{F}^*)$
- 12:     Update the total cost:  $\mathcal{C}_t \leftarrow \mathcal{C}_t + \mathcal{C}(\mathbf{x}_t, \mathbf{z}_t)$
- 13:     Aggregate data:  $\mathcal{D} \leftarrow \mathcal{D} \cup \{(\mathbf{x}_t, \mathbf{y}_t, \mathbf{z}_t)\}$
- 14:     Update models  $\mathcal{GP}_1, \dots, \mathcal{GP}_K$
- 15:      $t \leftarrow t + 1$
- 16: **end while**
- 17: **return** Pareto front and Pareto set of black-box functions  $f_1(x), \dots, f_K(x)$

---



---

#### Algorithm 2 Naive-CFMO Algorithm

---

**Input:** input space  $\mathcal{X}$ ;  $K$  blackbox functions  $f_j$  and their continuous approximations  $g_j$ ; total budget  $\mathcal{C}_{total}$

- 1: Initialize continuous fidelity gaussian process  $\mathcal{GP}_1, \dots, \mathcal{GP}_K$  by evaluating at initial points  $D$
- 2: **While**  $\mathcal{C}_t \leq \mathcal{C}_{total}$  **do**
- 3:   for each sample  $s \in 1, \dots, S$ :
- 4:     Sample highest-fidelity functions  $\tilde{f}_j \sim \mathcal{GP}_j(\cdot, z_j^*)$
- 5:      $\mathcal{F}_s^* \leftarrow$  Solve *cheap* MOO over  $(\tilde{f}_1, \dots, \tilde{f}_K)$
- 6:     Find the query based on  $\mathcal{F}^* = \{\mathcal{F}_s^*, s \in \{1 \dots S\}\}$ :
- 7:     // Use eq 4.12 for  $\alpha_t$  (MESMO)
- 8:     select  $\mathbf{x}_t \leftarrow \arg \max_{\mathbf{x} \in \mathcal{X}} \alpha_t(\mathbf{x}, \mathcal{F}^*)$
- 9:     **for**  $j \in 1 \dots K$  **do**
- 10:       select  $z_j \leftarrow \arg \min_{\mathbf{z}_j \in \mathcal{Z}_t^{(j)}(\mathbf{x}_t) \cup \{z_j^*\}} \mathcal{C}_i(x_t, z_j)$
- 11:     Fidelity vector  $\mathbf{z}_t \leftarrow [z_1 \dots z_K]$
- 12:     Update the total cost:  $\mathcal{C}_t \leftarrow \mathcal{C}_t + \mathcal{C}(\mathbf{x}_t, \mathbf{z}_t)$
- 13:     Aggregate data:  $\mathcal{D} \leftarrow \mathcal{D} \cup \{(\mathbf{x}_t, \mathbf{y}_t, \mathbf{z}_t)\}$
- 14:     Update models  $\mathcal{GP}_1, \dots, \mathcal{GP}_K$
- 15:      $t \leftarrow t + 1$
- 16: **end while**
- 17: **return** Pareto front and Pareto set of black-box functions  $f_1(x), \dots, f_K(x)$

---

## 5 Experiments and Results

In our experiments, we employed CF-GP models as described in section 3 with squared exponential kernels. We initialize the surrogate models of all functions with same number of points selected randomly from both lower and higher fidelities.

We compare iMOCA with several baselines: six state-of-the-art single-fidelity MO algorithms (ParEGO, SMSego, EHI, SUR, PESMO, and MESMO), one naive continuous-fidelity baseline that



we proposed in Section 4.1.1. For experiments in discrete fidelity setting, the number of fidelities is very limited. Thus, the fidelity space reduction method deem meaningless in this case. Therefore, we employ iMOCA without the fidelity space reduction. Additionally, we compare to the state-of-the-art discrete fidelity method MF-OSEMO. MF-OSEMO has two variants: MF-OSEMO-TG and MF-OSEMO-NI. Since MF-OSEMO-TG has the same formulation as iMOCA-T and provide similar results, we compare only to MF-OSEMO-NI.

### 5.1 Synthetic Benchmarks

We evaluate our algorithm iMOCA and baselines on four different synthetic benchmarks. We construct two problems using a combination of benchmark functions for continuous-fidelity and single-objective optimization ([43]): *Branin, Currin* (with  $K=2, d=2$ ) and *Ackley, Rosen, Sphere* (with  $K=3, d=5$ ). To show the effectiveness of iMOCA on settings with discrete fidelities, we employ two of the known general MO benchmarks: *QV* (with  $K=2, d=8$ ) and *DTLZ1* (with  $K=6, d=5$ ) ([19, 40]). Due to lack of space, we provide their complete details in Appendix B.1. The titles of plots in Fig. 1, Fig. 2, and Fig. 3 denote the corresponding experiments.

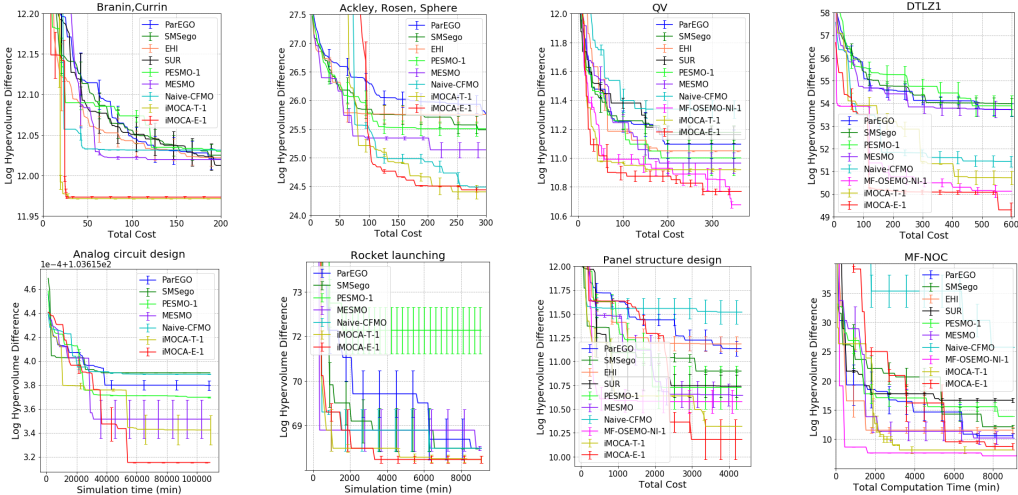


Figure 1: Results of iMOCA and the baseline algorithms on synthetic benchmarks and real-world problems. The PHV metric is presented against the total resource cost of function evaluations.

### 5.2 Real-world Engineering Design Optimization Problems

We evaluate iMOCA and baselines on four real-world design optimization problems from diverse engineering domains. We provide the details of these problems below.

**Analog circuit design optimization.** Design of a voltage regulator via Cadence circuit simulator that imitate the real hardware ([25]). Each candidate circuit design is defined by 33 input variables ( $d=33$ ). We optimize nine objectives: efficiency, four output voltages, and four output ripples. This problem has a continuous-fidelity space with cost varying from 10 mins to 120 mins.

**Panel structure design for large vessels.** The deck structure in large vessels commonly require the design of panels resisting uni-axial compression in the direction of the stiffeners ([49]). We consider optimizing the trade-off between two objective functions: weight and strength of the panel. These functions depend on six input variables ( $d=6$ ): one of them is the number of stiffeners used and five others relating to the plate thickness and stiffener dimensions. This problem has a discrete fidelity setting: two fidelities with computational costs 1 min and 21 mins respectively.

**Rocket launching simulation.** Rocket launching studies ([20]) require several long and computationally-expensive simulations to reach an optimal design. In this problem, we have three input variables ( $d = 3$ ): mass of fuel, launch height, and launch angle. The three objective functions

are return time, angular distance, and difference between the launch angle and the radius at the point of launch.

The simulator has a parameter that can be adjusted to perform continuous fidelity simulations. We employ the parameter range to vary the cost from 0.05 to 30 mins.

**Network-on-chip.** Communication infrastructure is critical for efficient data movement in hardware chips and they are designed using cycle-accurate simulators. We consider a dataset of 1024 configurations of a network-on-chip with ten input variables ( $d=10$ ) ([12]). We optimize two objectives: latency and energy. This problem has two discrete fidelities with costs 3 mins and 45 mins respectively.

### 5.3 Results and Analysis

We compare iMOCA with both approximations (iMOCA-T and iMOCA-E) to all baselines. We employ two known metrics for evaluating the quality of a given Pareto front: *Pareto hypervolume* ( $PHV$ ) metric and  $R_2$  indicator.  $PHV$  ([50]) is defined as the volume between a reference point and the given Pareto front; and  $R_2$  is a distance-based metric defined as the average distance between two Pareto-fronts. We report both the difference in the hyper-volume, and the average distance between an optimal Pareto front ( $\mathcal{F}^*$ ) and the best recovered Pareto front estimated by optimizing the posterior mean of the models at the highest fidelities ([22]). The mean and variance of  $PHV$  and  $R_2$  metrics across 10 different runs are reported as a function of the total cost.

Fig. 1 shows the  $PHV$  results of all the baselines and iMOCA for synthetic and real-world experiments (Fig. 3 in Appendix B.2 shows the corresponding  $R_2$  results). We observe that: 1) iMOCA consistently outperforms all baselines. Both iMOCA-T and iMOCA-E have significantly lower converge cost. 2) iMOCA-E shows a better convergence rate than iMOCA-T. This result can be explained by its tighter approximation. Nevertheless, iMOCA-T displays very close or sometimes better results than iMOCA-E. This demonstrates that even with loose approximation, using the iMOCA-T approximation can provide consistently competitive results using less computation time. 3) For experiments with the discrete fidelity setting, iMOCA most of the times outperformed MF-OSEMO or produced very close results. It is important to note that MF-OSEMO is an algorithm designed specifically for the discrete-fidelity setting. Therefore, the competitive performance of iMOCA shows its effectiveness and generalisability.

Figure 2 in appendix B.2 shows the results of evaluating iMOCA and PESMO with varying number of Monte-Carlo samples  $S \in \{1, 10, 100\}$ . For ease of comparison and readability, we present these results in two different figures side by side. We observe that the convergence rate of PESMO is dramatically affected by the number of MC samples  $S$ . However, iMOCA-T and iMOCA-E maintain a better performance consistently even with a single sample. These results strongly demonstrate that our method *iMOCA is much more robust to the number of Monte-Carlo samples*.

Table 2: *Best* convergence cost from all baselines  $\mathcal{C}_B$ , *Worst* convergence cost for iMOCA  $\mathcal{C}$ , and cost reduction factor  $\mathcal{G}$ .

Name	BC	ARS	Circuit	Rocket
$\mathcal{C}_B$	200	300	115000	9500
$\mathcal{C}$	30	100	55000	2000
$\mathcal{G}$	85%	66.6%	52.1%	78.9%

**Cost reduction factor.** We also provide the *cost reduction factor* for experiments with continuous fidelities, which is the percentage of gain in the convergence cost when compared to the best performing baseline (the earliest cost for which any of the single-fidelity baselines converge). Although this metric gives advantage to baselines, the results in Table 2 show a consistently high gain ranging from 52.1% to 85%.

## 6 Conclusions

We introduced a novel approach referred as iMOCA to solve multi-objective Bayesian optimization problems with continuous function approximations. The key idea is to select inputs and function

approximations for evaluation which maximizes the information gained per unit cost about the optimal Pareto front. Our experimental results on diverse benchmarks showed that iMOCA consistently outperforms state-of-the-art single-fidelity methods and a naive continuous-fidelity MO algorithm.

**Acknowledgements.** The authors gratefully acknowledge the support from National Science Foundation (NSF) grants IIS-1845922 and OAC-1910213. The views expressed are those of the authors and do not reflect the official policy or position of the NSF.

## References

- [1] Reinaldo B Arellano-Valle, JAVIER E CONTRERAS-REYES, and Marc G Genton. Shannon entropy and mutual information for multivariate skew-elliptical distributions. *Scandinavian Journal of Statistics*, 40(1), 2013.
- [2] Atthaphon Ariyarat and et al. Multi-fidelity multi-objective efficient global optimization applied to airfoil design problems. *Applied Sciences*, 7(12), 2017.
- [3] Adelchi Azzalini. A class of distributions which includes the normal ones. *Scandinavian Journal of Statistics*, 1985.
- [4] Ricardo Baptista and Matthias Poloczek. Bayesian optimization of combinatorial structures. In *Proceedings of the 35th International Conference on Machine Learning*, pages 462–471, 2018.
- [5] Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Max-value entropy search for multi-objective bayesian optimization. In *Conference on Neural Information Processing Systems*, pages 7823–7833, 2019.
- [6] Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Max-value entropy search for multi-objective bayesian optimization with constraints. *CoRR*, abs/2009.01721, 2020.
- [7] Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Multi-fidelity multi-objective bayesian optimization: An output space entropy search approach. In *AAAI*, pages 10035–10043, 2020.
- [8] Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Uncertainty aware search framework for multi-objective bayesian optimization with constraints. *CoRR*, abs/2008.07029, 2020.
- [9] Syrine Belakaria, Aryan Deshwal, Nitthilan Kannappan Jayakodi, and Janardhan Rao Doppa. Uncertainty-aware search framework for multi-objective bayesian optimization. In *AAAI conference on artificial intelligence*, 2020.
- [10] Syrine Belakaria, Derek Jackson, Yue Cao, Janardhan Rao Doppa, and Xiaonan Lu. Machine learning enabled fast multi-objective optimization for electrified aviation power system design. In *IEEE Energy Conversion Congress and Exposition (ECCE)*, 2020.
- [11] Syrine Belakaria, Zhiyuan Zhou, Aryan Deshwal, Janardhan Rao Doppa, Partha Pande, and Deuk Heo. Design of multi-output switched-capacitor voltage regulator via machine learning. In *Proceedings of the Twenty-Third IEEE/ACM Design Automation and Test in Europe Conference (DATE)*, 2020.
- [12] Shuai Che, Michael Boyer, Jiayuan Meng, David Tarjan, Jeremy W. Sheaffer, Sang-Ha Lee, and et al. Rodinia: A benchmark suite for heterogeneous computing. In *2009 IEEE international symposium on workload characterization (IISWC)*, 2009.
- [13] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley and Sons, 2012.
- [14] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, T Meyarivan, and A Fast. Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 2002.
- [15] Aryan Deshwal, Syrine Belakaria, and Janardhan Rao Doppa. Scalable combinatorial Bayesian optimization with tractable statistical models. *CoRR*, abs/2008.08177, 2020.

- [16] Aryan Deshwal, Syrine Belakaria, Janardhan Rao Doppa, and Alan Fern. Optimizing discrete spaces via expensive evaluations: A learning to search framework. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 3773–3780, 2020.
- [17] Aryan Deshwal, Nitthilan Kannappan Jayakodi, Biresh Kumar Joardar, Janardhan Rao Doppa, and Partha Pratim Pande. MOOS: A multi-objective design space exploration and optimization framework for noc enabled manycore systems. *ACM Transactions on Embedded Computing Systems*, 18(5s):77:1–77:23, 2019.
- [18] Michael Emmerich and Jan-willem Klinkenberg. The computation of the expected improvement in dominated hypervolume of pareto front approximations. *Technical Report, LU*, 34, 2008.
- [19] Ahsanul Habib, Hemant K Singh, and et al. A multiple surrogate assisted multi/many-objective multi-fidelity evolutionary algorithm. *Information Sciences*, 2019.
- [20] Javier E Hasbun. *Classical mechanics with MATLAB applications*. Jones & Bartlett Publishers, 2012.
- [21] Philipp Hennig and Christian J Schuler. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 2012.
- [22] Daniel Hernández-Lobato, Jose Hernandez-Lobato, Amar Shah, and et al. Predictive entropy search for multi-objective bayesian optimization. In *International Conference on Machine Learning*, 2016.
- [23] José Miguel Hernández-Lobato and et al. Predictive entropy search for efficient global optimization of black-box functions. In *Conference on Neural Information Processing Systems*, 2014.
- [24] Matthew W Hoffman and Zoubin Ghahramani. Output-space predictive entropy search for flexible global optimization. In *Conference on Neural Information Processing Systems workshop on Bayesian optimization*, 2015.
- [25] Wookpyo Hong and et al. A dual-output step-down switched-capacitor voltage regulator with a flying capacitor crossing technique for enhanced power efficiency. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 27(12), 2019.
- [26] Deng Huang, Theodore T Allen, William I Notz, and R Allen Miller. Sequential kriging optimization using multiple-fidelity evaluations. *Structural and Multidisciplinary Optimization*, 2006.
- [27] Kirthevasan Kandasamy, Gautam Dasarathy, Junier B Oliva, and et al. Gaussian process bandit optimisation with multi-fidelity evaluations. In *Conference on Neural Information Processing Systems*, 2016.
- [28] Kirthevasan Kandasamy, Gautam Dasarathy, Jeff Schneider, and et al. Multi-fidelity bayesian optimisation with continuous approximations. *International Conference on Machine Learning*, 2017.
- [29] Aaron Klein, Stefan Falkner, Simon Bartels, Philipp Hennig, and Frank Hutter. Fast bayesian optimization of machine learning hyperparameters on large datasets. In *International Conference on Artificial Intelligence and Statistics*, 2017.
- [30] Joshua Knowles. Parego: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 2006.
- [31] Spyridon G Kontogiannis, Jean Demange, T Kipouros, and et al. A comparison study of two multifidelity methods for aerodynamic optimization. In *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2018.
- [32] Rémi Lam, Douglas L Allaire, and et al. Multifidelity optimization using statistical surrogate modeling for non-hierarchical information sources. In *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2015.

- [33] Mark McLeod, Michael A Osborne, and Stephen J Roberts. Practical bayesian optimization for variable cost objectives. *arXiv preprint arXiv:1703.04335*, 2017.
- [34] Henry B Moss, David S Leslie, and Paul Rayson. Mumbo: Multi-task max-value bayesian optimization. *The European Conference on Machine Learning*, 2020.
- [35] Victor Picheny. Multi-objective optimization using gaussian process emulators via stepwise uncertainty reduction. *Statistics and Computing*, 25(6):1265–1280, 2015.
- [36] Victor Picheny, David Ginsbourger, and et al. Quantile-based optimization of noisy computer experiments with tunable precision. *Technometrics*, 2013.
- [37] Wolfgang Ponweiser and et al. Multiobjective optimization on a limited budget of evaluations using model-assisted s-metric selection. In *PPSN*. Springer, 2008.
- [38] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Conference on Neural Information Processing Systems*, 2008.
- [39] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 2016.
- [40] Leshi Shu, Ping Jiang, Qi Zhou, Xinyu Shao, Jiexiang Hu, and Xiangzheng Meng. An on-line variable fidelity metamodel assisted multi-objective genetic algorithm for engineering design optimization. *Applied Soft Computing*, 66, 2018.
- [41] Jialin Song, Yuxin Chen, and Yisong Yue. A general framework for multi-fidelity bayesian optimization with gaussian processes. *International Conference on Artificial Intelligence and Statistics*, 2019.
- [42] Niranjn Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv:0912.3995*, 2009.
- [43] S. Surjanovic and D. Bingham. Virtual library of simulation experiments: Test functions and datasets. Retrieved January 21, 2020, from <http://www.sfu.ca/~ssurjano>.
- [44] Kevin Swersky, Jasper Snoek, and Ryan P Adams. Multi-task bayesian optimization. In *Conference on Neural Information Processing Systems*, 2013.
- [45] Shion Takeno, Hitoshi Fukuoka, Yuhki Tsukada, Toshiyuki Koyama, Motoki Shiga, Ichiro Takeuchi, and Masayuki Karasuyama. Multi-fidelity bayesian optimization with max-value entropy search. *arXiv:1901.08275*, 2019.
- [46] Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient bayesian optimization. In *International Conference on Machine Learning*, 2017.
- [47] Jian Wu and Peter I Frazier. Continuous-fidelity bayesian optimization with knowledge gradient. 2018.
- [48] Yehong Zhang, Trong Nghia Hoang, and et al. Information-based multi-fidelity bayesian optimization. In *Conference on Neural Information Processing Systems Workshop on Bayesian Optimization*, 2017.
- [49] Jiandao Zhu, Yi-Jen Wang, and Matthew Collette. A multi-objective variable-fidelity optimization method for genetic algorithms. *Engineering Optimization*, 46(4):521–542, 2014.
- [50] Eckart Zitzler. *Evolutionary algorithms for multiobjective optimization: Methods and applications*. Citeseer, 1999.
- [51] Marcela Zuluaga, Guillaume Sergent, Andreas Krause, and Markus Püschel. Active learning for multi-objective optimization. In *International Conference on Machine Learning*, pages 462–470, 2013.

## A Full derivation of acquisition function

Our goal is to derive a full approximation for iMOCA algorithm. In this appendix, we provide the technical details of the extended-skew Gaussian approximation (iMOCA-E) for the computation of the information gain per unit cost.

The information gain in equation 4.1 is defined as the expected reduction in entropy  $H(\cdot)$  of the posterior distribution  $P(\mathcal{F}^*|D)$  due to evaluating  $\mathbf{x}$  at fidelity vector  $\mathbf{z}$ . Based on the symmetric property of information gain, we can rewrite it as shown below:

$$I(\{\mathbf{x}, \mathbf{y}, \mathbf{z}\}, \mathcal{F}^*|D) = H(\mathbf{y}|D, \mathbf{x}, \mathbf{z}) - \mathbb{E}_{\mathcal{F}^*}[H(\mathbf{y}|D, \mathbf{x}, \mathbf{z}, \mathcal{F}^*)] \quad (\text{A.1})$$

In equation A.1, the first term is the entropy of a  $K$ -dimensional Gaussian distribution that can be computed in closed form as follows:

$$H(\mathbf{y}|D, \mathbf{x}, \mathbf{z}) = \sum_{j=1}^K \ln(\sqrt{2\pi e} \sigma_{g_j}(\mathbf{x}, z_j)) \quad (\text{A.2})$$

The second term of equation A.1 is an expectation over the Pareto front of the highest fidelities  $\mathcal{F}^*$ . This term can be approximated using Monte-Carlo sampling:

$$\mathbb{E}_{\mathcal{F}^*}[H(\mathbf{y}|D, \mathbf{x}, \mathbf{z}, \mathcal{F}^*)] \simeq \frac{1}{S} \sum_{s=1}^S [H(\mathbf{y}|D, \mathbf{x}, \mathbf{z}, \mathcal{F}_s^*)] \quad (\text{A.3})$$

In the main paper, we showed that :

$$y_j \leq f_s^{j*} \quad \forall j \in \{1, \dots, K\} \quad (\text{A.4})$$

By combining the inequality A.4 and the fact that each function is modeled as an independent CF-GP, a common property of entropy measure allows us to decompose the entropy of a set of independent variables into a sum over entropies of individual variables [13]:

$$H(\mathbf{y}|D, \mathbf{x}, \mathbf{z}, \mathcal{F}_s^*) \simeq \sum_{j=1}^K H(y_j|D, \mathbf{x}, z_j, f_s^{j*}) \quad (\text{A.5})$$

In what follows, we provide details of iMOCA-E approximation to compute  $H(y_j|D, \mathbf{x}, z_j, f_s^{j*})$ .

The condition  $y_j \leq f_s^{j*}$ , is originally expressed as  $f_j \leq f_s^{j*}$ . Substituting this condition with its original equivalent, the entropy becomes  $H(y_j|D, \mathbf{x}, z_j, f_j \leq f_s^{j*})$ . Since  $y_j$  is an evaluation of the function  $g_j$  and  $f_j$  is an evaluation of the function  $f_j$ , we make the observation that  $y_j|f_j \leq f_s^{j*}$  follows an extended-skew Gaussian (ESG) distribution [3]. It had been shown that the differential entropy of an ESG does not have a closed-form expression [1]. Therefore, we derive a simplified expression where most of the terms are analytical by manipulating the components of the entropy as shown below.

In order to simplify the calculation  $H(y_j|D, \mathbf{x}, z_j, f_j \leq f_s^{j*})$ , we start by deriving an expression for its probability distribution. Based on the definition of the conditional distribution of a bi-variate normal,  $f_j|y_j$  is normally distributed with mean  $\mu_{f_j} + \frac{\sigma_{f_j}}{\sigma_{g_j}} \tau (y_j - \mu_{g_j})$  and variance  $\sigma_{f_j}^2 (1 - \tau)^2$ , where  $\tau = \frac{\sigma_{g_j, f_j}^2}{\sigma_{g_j} \sigma_{f_j}}$  is the predictive correlation between  $y_j$  and  $f_j$ . We can now write the cumulative distribution function for  $y_j|f_j \leq f_s^{j*}$  as shown below:

$$P(y_j \leq u | f_j \leq f_s^{j*}) = \frac{P(y_j \leq u, f_j \leq f_s^{j*})}{P(f_j \leq f_s^{j*})} = \frac{\int_{-\infty}^u \phi\left(\frac{\theta - \mu_{g_j}}{\sigma_{g_j}}\right) \Phi\left(\frac{f_s^{j*} - \mu_{f_j} - \frac{\sigma_{f_j}}{\sigma_{g_j}} \tau (\theta - \mu_{g_j})}{\sqrt{\sigma_{f_j}^2 (1 - \tau)^2}}\right) d\theta}{\sigma_{g_j} \Phi\left(\frac{f_s^{j*} - \mu_{f_j}}{\sigma_{f_j}}\right)}$$

Let us define the normalized variable  $\Gamma_{f_s^{j*}}$  as  $\Gamma_{f_s^{j*}} = \frac{y_j - \mu_{g_j}}{\gamma_{g_j}} | f_j \leq f_s^{j*}$ . After differentiating with respect to  $u$ , we can express the probability density function for  $\Gamma_{f_s^{j*}}$  as:

$$P(u) = \frac{\phi(u)}{\Phi(\gamma_s^{(f_j)})} \Phi\left(\frac{\gamma_s^{(f_j)} - \tau u}{\sqrt{1 - \tau^2}}\right)$$

which is the density of an ESG with mean and variance defined as follows:

$$\mu_{\Gamma_{f_s^{j*}}} = \tau \frac{\phi(\gamma_s^{(f_j)})}{\Phi(\gamma_s^{(f_j)})}, \sigma_{\Gamma_{f_s^{j*}}} = 1 - \tau^2 \frac{\phi(\gamma_s^{(f_j)})}{\Phi(\gamma_s^{(f_j)})} \left[ \gamma_s^{(f_j)} + \frac{\phi(\gamma_s^{(f_j)})}{\Phi(\gamma_s^{(f_j)})} \right] \quad (\text{A.6})$$

Therefore, we can express the entropy of the ESG as shown below:

$$H(\Gamma_{f_s^{j*}}) = - \int P(u) \ln(P(u)) du \quad (\text{A.7})$$

We also derive a more simplified expression of the iMOCA-E acquisition function based on ESG. For a fixed sample  $f_s^{j*}$ ,  $H(\Gamma_{f_s^{j*}})$  can be decomposed as follows:

$$H(\Gamma_{f_s^{j*}}) = \mathbb{E}_{u \sim \Gamma_{f_s^{j*}}} \left[ -\ln(\phi(u)) + \ln(\Phi(\gamma_s^{(f_j)})) - \ln\left(\Phi\left(\frac{\gamma_s^{(f_j)} - \tau u}{\sqrt{1 - \tau^2}}\right)\right) \right] \quad (\text{A.8})$$

We expand the first term as shown below:

$$\mathbb{E}_{u \sim \Gamma_{f_s^{j*}}} [-\ln(\phi(u))] = \frac{1}{2} \ln(2\pi) + \frac{1}{2} \mathbb{E}_{u \sim \Gamma_{f_s^{j*}}} [u^2] \quad (\text{A.9})$$

From the mean and variance of  $\Gamma_{f_s^{j*}}$  in equation A.6, we get:

$$\mathbb{E}_{u \sim \Gamma_{f_s^{j*}}} [u^2] = \mu_{\Gamma_{f_s^{j*}}}^2 + \sigma_{\Gamma_{f_s^{j*}}} = 1 - \tau^2 \frac{\phi(\gamma_s^{(f_j)}) \gamma_s^{(f_j)}}{\Phi(\gamma_s^{(f_j)})} \quad (\text{A.10})$$

We note that the final entropy can be computed using the following expression.

$$H(y_j | D, \mathbf{x}, z_j, y_j \leq f_s^{j*}) = H(\Gamma_{f_s^{j*}}) + \ln(\sigma_{g_j}) \quad (\text{A.11})$$

By combining equations A.8 and A.11, we get:

$$\begin{aligned} H(y_j | D, \mathbf{x}, z_j, f_j \leq f_s^{j*}) &= \ln(\sqrt{2\pi} e \sigma_{g_j}) + \ln(\Phi(\gamma_s^{(f_j)})) - \tau^2 \frac{\phi(\gamma_s^{(f_j)}) \gamma_s^{(f_j)}}{2\Phi(\gamma_s^{(f_j)})} \\ &\quad - \mathbb{E}_{u \sim \Gamma_{f_s^{j*}}} \left[ \ln\left(\Phi\left(\frac{\gamma_s^{(f_j)} - \tau u}{\sqrt{1 - \tau^2}}\right)\right) \right] \end{aligned} \quad (\text{A.12})$$

From equations A.3, A.2, and A.12, the final expression of iMOCA-E can be expressed as follows:

$$\alpha_t(\mathbf{x}, \mathbf{z}, \mathcal{F}^*) = \frac{1}{\mathcal{C}(\mathbf{x}, \mathbf{z}) S} \sum_{j=1}^K \sum_{s=1}^S \tau^2 \frac{\gamma_s^{(f_j)} \phi(\gamma_s^{(f_j)})}{2\Phi(\gamma_s^{(f_j)})} - \ln(\Phi(\gamma_s^{(f_j)})) + \mathbb{E}_{u \sim \Gamma_{f_s^{j*}}} [\ln(\Phi(\frac{\gamma_s^{(f_j)} - \tau u}{\sqrt{1 - \tau^2}}))]$$

Since the differential entropy of an ESG cannot be computed analytically, we perform numerical integration via Simpson's rule using  $\mu_{\Gamma_{f_s^{j*}}} \mp \gamma \sqrt{\sigma_{\Gamma_{f_s^{j*}}}}$  as the integral limits. In practice, we set  $\gamma$  to 5. Since this integral is over one-dimension variable, numerical integration can result in a tight approximation with small amount of computation.

## B Additional Experiments and Results

### B.1 Description of Synthetic Benchmarks

In what follows, we provide complete details of the synthetic benchmarks employed in this paper. Since our algorithm is designed for maximization settings, we provide the benchmarks in their maximization form.

### 1) Branin, Currin experiment

In this experiment, we construct a multi-objective problem using a combination of existing single-objective optimization benchmarks [28]. It has two functions with two dimensions ( $K=2$  and  $d=2$ ).

**Branin function:** We use the following function where  $\mathcal{C}(z) = 0.05 + z^{6.5}$

$$g(\mathbf{x}, z) = - (a(x_2 - b(z)x_1^2 + c(z)x_1 - r)^2 + s(1 - t(z))\cos(x_1) + s)$$

where  $a = 1$ ,  $b(z) = 5.1/(4\pi^2) - 0.01(1 - z)$ ,  $c(z) = 5/\pi - 0.1(1 - z)$ ,  $r = 6$ ,  $s = 10$  and  $t(z) = 1/(8\pi) + 0.05(1 - z)$ .

**Currin exponential function:** We use  $\mathcal{C}(z) = 0.1 + z^2$

$$g(\mathbf{x}, z) = - \left( 1 - 0.1(1 - z) \exp\left(\frac{-1}{2x_2}\right) \right) \left( \frac{2300x_1^3 + 1900x_1^2 + 2092x_1 + 60}{100x_1^3 + 500x_1^2 + 4x_1 + 20} \right).$$

### 2) Ackley, Rosen, Sphere experiment

In this experiment, we construct a multi-objective problem using a combination of existing single-objective optimization benchmarks [47]. It has three functions with five dimensions ( $K=3$  and  $d=5$ ). For all functions, we employed  $\mathcal{C}(z) = 0.05 + z^{6.5}$

**Ackley function**

$$g(\mathbf{x}, z) = - \left( -20 \exp \left[ -0.2 \sqrt{\frac{1}{d} \sum_{i=1}^d x_i^2} - \exp \left[ \frac{1}{d} \sum_{i=1}^d \cos(2\pi x_i) \right] + e + 20 \right) - 0.01(1 - z)$$

**Rosenbrock function:**

$$g(\mathbf{x}, z) = - \sum_{i=1}^{d-1} \left[ 100 (x_{i+1} - x_i^2 + 0.01(1 - z))^2 + (1 - x_i)^2 \right]$$

**Sphere function:**

$$g(\mathbf{x}, z) = - \sum_{i=1}^d x_i^2 - 0.01(1 - z)$$

### 3) DTLZ1 experiment

In this experiment, we solve a problem from the general multi-objective optimization benchmarks [19]. We have six functions with five dimensions ( $K=6$  and  $d=5$ ) with a discrete fidelity setting. Each function has three fidelities in which  $z$  takes three values from  $\{0.2, 0.6, 1\}$  with  $z^*=1$ . The cost of evaluating each fidelity function is  $\mathcal{C}(z)=\{0.01, 0.1, 1\}$

$$g_j(\mathbf{x}, z) = f_j(\mathbf{x}) - e(\mathbf{x}, z)$$

$$f_1(\mathbf{x}) = -(1 + r)0.5 \prod_{i=1}^5 x_i$$

$$f_j(\mathbf{x}) = -(1 + r)0.5(1 - x_{6-j+1}) \prod_{i=1}^{6-j} x_i \text{ with } j = 2 \dots 5$$

$$f_6(\mathbf{x}) = -(1 + r)0.5(1 - x_1)$$

$$r = 100[d + \sum_{i=1}^d ((x_i - 0.5)^2) - \cos(10\pi(x_i - 0.5))]$$

$$e(\mathbf{x}, z) = \sum_{i=1}^d \alpha(z) \cos(10\pi\alpha(z)x_i + 0.5\pi\alpha(z) + \pi) \text{ with } \alpha(z) = 1 - z$$



#### 4) QV experiment

In this experiment, we solve a problem from the general multi-objective optimization benchmarks [40]. We have two functions with eight dimensions ( $K=2$  and  $d=8$ ) with a discrete fidelity setting.

**Function 1** has only one fidelity which is the highest fidelity

$$f_1(\mathbf{x}) = -\left(\frac{1}{d} \sum_{i=1}^d (x_i^2 - 20\pi x_i + 10)\right)^{\frac{1}{4}}$$

**Function 2** has two fidelities with cost  $\{0.1, 1\}$  respectively and the following expressions:

$$\text{High fidelity: } f_2(\mathbf{x}, \text{High}) = -\left(\frac{1}{d} \sum_{i=1}^d ((x_i - 1.5)^2 - 20\pi(x_i - 1.5) + 10)\right)^{\frac{1}{4}}$$

$$\text{Low fidelity: } f_2(\mathbf{x}, \text{Low}) = -\left(\frac{1}{d} \left(\sum_{i=1}^d (\alpha[\mathbf{i}](x_i - 1.5)^2 - 20\pi(x_i - 1.5) + 10)\right)^{\frac{1}{4}}\right)$$

with  $\alpha=[0.9, 1.1, 0.9, 1.1, 0.9, 1.1, 0.9, 1.1]$

#### B.2 Additional results




Figure 2: Results of synthetic benchmarks showing the effect of varying the number of Monte-Carlo samples for iMOCA, MESMO, and PESMO. The hypervolume difference is shown against the total resource cost of function evaluations.

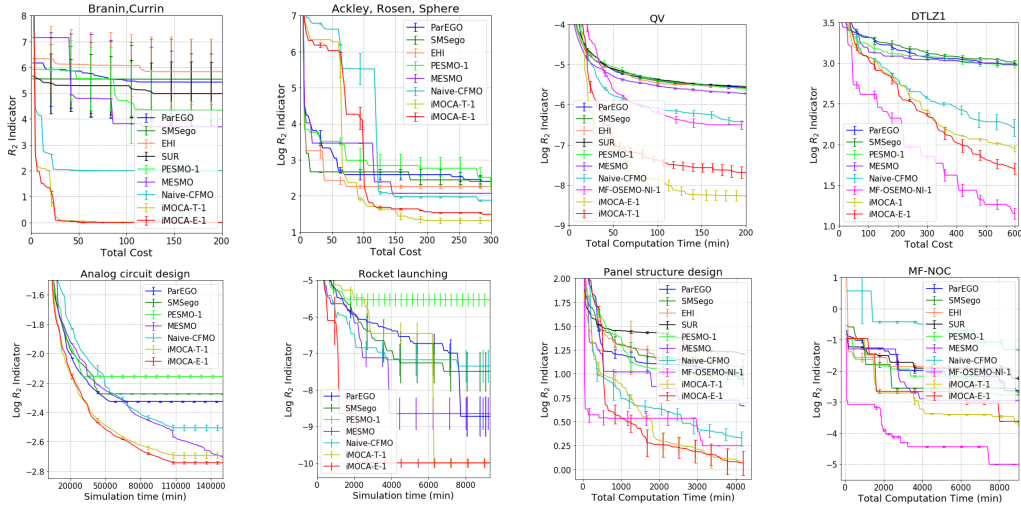


Figure 3: Results of iMOCA and the baselines algorithms on synthetic benchmarks and real-world problems. The  $R_2$  metric is presented against the total resource cost of function evaluations.