

Pareto Front-Diverse Batch Multi-Objective Bayesian Optimization

Alaleh Ahmadianshalchi^{*1}, Syrine Belakaria^{*2}, Janardhan Rao Doppa¹

¹ School of EECS, Washington State University

² Computer Science Department, Stanford University

a.ahmadianshalchi@wsu.edu, syrineb@stanford.edu, jana.doppa@wsu.edu

Abstract

We consider the problem of multi-objective optimization (MOO) of expensive black-box functions with the goal of discovering high-quality and diverse Pareto fronts where we are allowed to evaluate a batch of inputs. This problem arises in many real-world applications including penicillin production where diversity of solutions is critical. We solve this problem in the framework of Bayesian optimization (BO) and propose a novel approach referred to as **Pareto front-Diverse Batch Multi-Objective BO (PDBO)**. PDBO tackles two important challenges: 1) How to automatically select the best acquisition function in each BO iteration, and 2) How to select a diverse batch of inputs by considering multiple objectives. We propose principled solutions to address these two challenges. First, PDBO employs a multi-armed bandit approach to select one acquisition function from a given library. We solve a cheap MOO problem by assigning the selected acquisition function for each expensive objective function to obtain a candidate set of inputs for evaluation. Second, it utilizes Determinantal Point Processes (DPPs) to choose a Pareto-front-diverse batch of inputs for evaluation from the candidate set obtained from the first step. The key parameters for the methods behind these two steps are updated after each round of function evaluations. Experiments on multiple MOO benchmarks demonstrate that PDBO outperforms prior methods in terms of both the quality and diversity of Pareto solutions.

1 Introduction

A wide range of science and engineering applications, including materials design [Ashby 2000], biological sequence design [Taneda 2015], and drug/vaccine design [Nicolaou and Brown 2013] involves optimizing multiple *expensive-to-evaluate* objective functions. For example, in nanoporous materials design [Deshwal et al. 2021], the goal is to optimize the adsorption property and cost of synthesis guided by physical lab experiments. Since the experiments are expensive in terms of the consumed resources, our goal is to approximate the *optimal Pareto set* of solutions. In many of the aforementioned applications, there are two important considerations. First, we can perform multiple parallel experiments which should be leveraged to accelerate the discovery of high-quality solutions. Second, practitioners care about

diversity in solutions and their outcomes. For example, in penicillin production application, diverse solutions for objectives including penicillin production, the time to ferment, and the CO_2 byproduct [Birol et al. 2002].

We consider the problem of multi-objective optimization (MOO) over expensive-to-evaluate functions to find high-quality and diverse Pareto fronts when we are allowed to perform a batch of experiments. We solve this problem using Bayesian optimization (BO) [Shahriari et al. 2015] which has been shown to be highly effective for such problems. The key idea behind BO is to learn a surrogate model (e.g., Gaussian process) from past experiments and use it to intelligently select the sequence of experiments guided by an acquisition function (e.g., expected improvement). There is prior BO work on selecting a batch of experiments to find diverse high-performing solutions for single-objective optimization. However, there is very limited work on batch BO for MOO problems and to produce diverse MOO solutions. A key drawback of the existing BO methods for MOO is that they evaluate diversity in terms of input space, which is not appropriate for MOO (diversity in input space $\not\Rightarrow$ diversity in output space). To overcome this drawback and to measure the diversity of MOO solution in the output space, we define a new metric referred to as *Diversity of the Pareto Front (DPF)*.

We propose a novel approach referred to as **Pareto front-Diverse Batch Multi-Objective BO (PDBO)**. PDBO selects a batch B of inputs for evaluations in each iteration using two main steps. First, it employs a principled multi-arm bandit strategy to dynamically select one acquisition function (AF) from a given library of AFs within the single-objective BO literature. A cheap MOO problem is solved by assigning the selected AF for each expensive objective function to obtain a Pareto set. Second, a principled configuration of determinantal point processes (DPPs) [Borodin 2009, Kulesza et al. 2012] for multiple objectives is used to select B inputs for evaluation from the Pareto set of the first step to improve the diversity of the uncovered Pareto front. PDBO updates the key parameters of the algorithms for these two steps (dynamic selection of AF and selecting B Pareto-front diverse inputs from a candidate Pareto set) after obtaining the objective function evaluations. Our experiments on multiple benchmarks with varying input dimensions and number of objective functions demonstrate that PDBO outperforms prior methods in finding high-quality and diverse Pareto fronts.

^{*}These authors contributed equally.

Contributions. The key contribution of this paper is developing and evaluating the PDBO approach for solving MOO problems to find high-quality and diverse Pareto fronts.

Specific contributions are as follows:

- A multi-arm bandit strategy with a novel reward function to dynamically select one acquisition function from a given library for MOO problems.
- A novel DPP method for selecting a batch of inputs from a given Pareto set to maximize Pareto front diversity using a new mechanism to generate the DPP kernel for MOO.
- To demonstrate the effectiveness of PDBO and study the Pareto front diversity, we propose a new metric to measure the diversity of *Pareto fronts*. To the best of our knowledge, this is the first attempt at extensive experimental evaluation of Pareto front diversity within BO.
- Theoretical analysis of our PDBO algorithm in terms of asymptotic regret bounds.
- Experimental evaluation of PDBO and baselines on multiple benchmark MOO problems. The code for PDBO is publicly available at <https://github.com/Alaleh/PDBO>.

2 Problem Setup and Background

We first formally define the batch MOO problem along with the metrics to evaluate the quality and diversity of solutions. Next, we provide an overview of the BO framework.

Batch Multi-Objective Optimization. We consider a MOO problem where the goal is to optimize multiple conflicting functions. Let $\mathcal{X} \subset \mathbb{R}^d$ be the input space of d design variables, where each candidate input $\mathbf{x} \in \mathcal{X}$ is a d -dimensional input vector. And let $\{f_1, \dots, f_K\}$ with $K \geq 2$ be the objective functions defined over the input space \mathcal{X} where $f_1(\mathbf{x}), \dots, f_K(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}$. We denote the function evaluation at an input \mathbf{x} as $\mathbf{y} = [y_1, \dots, y_K]$, where $y_i = f_i(\mathbf{x})$ for all $i \in \{1, \dots, K\}$. Without loss of generality, we assume minimization for all K objective functions. The optimal solution of the MOO problem is a set of inputs $\mathcal{X}^* \subset \mathcal{X}$ such that no input $\mathbf{x}' \in \mathcal{X} \setminus \mathcal{X}^*$ Pareto-dominates another input $\mathbf{x} \in \mathcal{X}^*$. An input \mathbf{x} Pareto-dominates another point \mathbf{x}' if and only if $\forall j : f_j(\mathbf{x}) \leq f_j(\mathbf{x}')$ and $\exists j : f_j(\mathbf{x}) < f_j(\mathbf{x}')$. The set of input solutions \mathcal{X}^* is called the optimal *Pareto set* and the corresponding set of function values \mathcal{Y}^* is called the optimal *Pareto front*. We can select B inputs for parallel evaluation in each iteration, and our goal is to uncover a high-quality and diverse Pareto front while minimizing the total number of expensive function evaluations.

Metrics for Quality and Diversity of Pareto Front. Our goal is to find high-quality and diverse Pareto fronts. The diversity of the Pareto front has not been formally *evaluated* in any previous work. We introduce an appropriate evaluation metric to measure the diversity of the Pareto front and discuss an existing measure of Pareto front quality below.

Diversity of Pareto Front. Diversity is an important criterion for many optimization problems. Prior work on batch BO, both in the single-objective and MO settings focused on evaluating diversity with respect to the input space [Jain et al. 2022]. However, in most real-world MOO problems, diversity in the input space does not necessarily reflect diversity

in the output space. In MOO, practitioners might care more about the diversity of the Pareto front rather than the Pareto set. Yet, little work has gone into understanding, formalizing, and measuring Pareto front diversity in MOO. In most cases, finding a more diverse set of points in the output space leads to a higher hypervolume [Zitzler and Thiele 1999]. However, a higher hypervolume does not necessarily correspond to a more diverse Pareto front. [Konakovic Lukovic et al. 2020] is the only prior work that proposed a diversity-guided approach for batch MOO. However, the diversity of the produced Pareto front *was not evaluated*. To fill this gap, we propose an evaluation metric to fill this gap to assess the *Diversity of the Pareto Front (DPF)*. Given a Pareto front \mathcal{Y}_t , $DPF(\mathcal{Y}_t)$ is the average pairwise distance between points (i.e., output vectors) in Pareto front \mathcal{Y}_t . It is important to clarify that the pairwise distances are computed in the output space between different vector pairs $(\mathbf{y}, \mathbf{y}') \in \mathcal{Y}_t$, unlike previously used metrics to assess input space diversity in the single-objective setting [Angermueller et al. 2020].

$$DPF(\mathcal{Y}_t) = \frac{\sum_{(i,j) \in \mathcal{I}} \|\mathbf{y}_i - \mathbf{y}_j\|}{|\mathcal{I}|}$$

$$\text{with } \mathcal{I} = \{(i, j) \forall i, j \in \{1 \dots t\}, i < j\}$$

We provide a more detailed discussion of existing metrics and some illustrative results on previous metrics and their utility in evaluating diversity in the Appendix.

Hypervolume Indicator. The hypervolume indicator [Zitzler and Thiele 1999] is the most commonly employed measure to evaluate the quality of a given Pareto front. Given a set of functions evaluations $\mathcal{Y}_t = \{\mathbf{y}_0, \dots, \mathbf{y}_t\}$, the Pareto hypervolume (PHV) indicator is the volume between a predefined reference point \mathbf{r} and the given Pareto front.

2.1 Bayesian Optimization Framework

BO [Shahriari et al. 2015] is a general framework for solving expensive black-box optimization problems in a sample-efficient manner. BO algorithms iterate through three steps.

1) Build a probabilistic *surrogate model* of the true expensive objective function. Gaussian process (GPs) [Williams and Rasmussen 2006] are widely employed in BO.

2) Define an *acquisition function (AF)* to score the utility of unevaluated inputs. It uses the surrogate model’s predictions as a cheap substitute for expensive function evaluations and strategically trades off exploitation and exploration.

Some examples of widely used acquisition functions in single-objective optimization include expected improvement (EI) [Mockus et al. 1978], upper confidence bound (UCB) [Auer 2002], Thompson sampling (TS) [Thompson 1933] and Identity (ID).

$$EI(\mathbf{x}) = \sigma(\mathbf{x})(\alpha\Phi(\alpha) + \phi(\alpha)), \alpha = \left(\tau - \frac{\mu(\mathbf{x})}{\sigma(\mathbf{x})}\right) \quad (1)$$

$$UCB(\mathbf{x}) = \mu(\mathbf{x}) + \sqrt{\beta}\sigma(\mathbf{x}) \quad (2)$$

$$TS(\mathbf{x}) = \hat{f}(\mathbf{x}) \text{ with } \hat{f} \sim \text{GP} \quad (3)$$

$$ID(\mathbf{x}) = \mu(\mathbf{x}) \quad (4)$$

Here, β is a parameter to balance exploration and exploitation in the UCB acquisition function [Srinivas et al. 2009],

τ is the best-uncovered function value; and Φ and ϕ are the CDF and PDF of a standard normal distribution respectively.

3) Select the input with the highest utility score for function evaluation by optimizing the acquisition function.

3 Related Work

Multi-Objective BO. There is relatively less work on MOO in comparison with single-objective BO. Prior work builds on the insights from single-objective methods [Shahriari et al. 2015, Hernández-Lobato et al. 2014, Wang and Jegelka 2017, Hvarfner et al. 2022] for MOO. Recent work on MOBO includes Predictive Entropy Search [Hernández-Lobato et al. 2016], Max-value Entropy Search [Belakaria et al. 2019], Multi-Objective Regionalized BO [Daulton et al. 2022], Uncertainty-aware Search [Belakaria et al. 2020a], Pareto-Frontier Entropy Search [Suzuki et al. 2020], and Expected Hypervolume Improvement [Daulton et al. 2020, Emmerich and Klinkenberg 2008]. These methods are shown to perform well on a variety of MOO problems.

Batch Multi-objective BO. The Batch BO problem in the multi-objective setting is even much less studied. Diversity-Guided Efficient Multi-Objective Optimization (DGEMO) [Konakovic Lukovic et al. 2020] approximates and analyzes a piecewise-continuous Pareto set representation which allows the algorithm to introduce a batch selection strategy that optimizes for both hypervolume improvement and diversity of selected samples. However, DGEMO work did not study or evaluate the Pareto-front diversity of the produced solutions. qEHVI [Daulton et al. 2020] is an exact computation of the joint EHVI of q new candidate points (up to Monte-Carlo integration error).

qPAREGO is a novel extension of ParEGO [Knowles 2006, Daulton et al. 2020] that supports parallel evaluation and constraints. More recent work [Lin et al. 2022] proposed an approach to address MOO problems with continuous/infinite Pareto fronts by approximating the whole Pareto set via a continuous manifold. This approach enables a better preference-based exploration strategy for practitioners compared to prior work [Abdolshah et al. 2019, Paria et al. 2020, Astudillo and Frazier 2020]. However, it is typically unknown to the user if the Pareto front is dense/continuous, especially in expensive function settings where the data is limited. *It is not known whether any of these batch methods produce diverse Pareto fronts or not, as they were not evaluated on diversity metrics.* We perform an experimental evaluation to answer this question.

DPPs for Batch Single-Objective BO. DPPs are elegant probabilistic models [Borodin and Olshanski 2005, Borodin 2009] that characterize the property of repulsion in a set of vectors and are well-suited for the selection of a diverse subset of inputs from a predefined set. Prior work used DPPs for selecting batches for evaluation in the single-objective BO literature [Kathuria et al. 2016, Nava et al. 2022, Wang et al. 2017]. In the context of MOO for cheap objective functions using evolutionary algorithms, DPP was deployed using non-learning-based kernels such as cosine function applied to points in the Pareto front while disregarding the input space [Wang et al. 2022, Zhang et al. 2020, Okoth et al. 2022].

However, *to the best of our knowledge, there is no work on using DPPs for multi-objective BO to uncover diverse Pareto fronts using a learned kernel that captures the trade-off between multiple objectives, similarity in the input space, and diversity of the Pareto front.*

Adaptive Acquisition Function Selection. There has been a plethora of research on finding efficient and reliable acquisition functions (AFs). However, prior work has shown that no single acquisition function is universally efficient and consistently outperforms all others. GP-Hedge [Hoffman et al. 2011] proposed to use a portfolio of acquisition functions. The optimization of each AF will nominate an input, and the algorithm will select one of them for evaluation using the selection probabilities. The GP-Hedge method uses the Hedge strategy [Freund and Schapire 1997], a multi-arm bandit method designed to choose one action amongst a set of different possibilities using selection probabilities calculated based on the reward (performance given by function values) collected from previous evaluations. Vasconcelos et al. extended Hoffman et al. by proposing to use discounted cumulative reward and Vasconcelos et al. suggested using Thompson sampling to automatically set the hedge hyperparameter η .

4 Proposed PDBO Algorithm

We start by providing an overview of the proposed PDBO algorithm illustrated in Figure 1. Next, we explain our algorithms for two key components of PDBO, namely, adaptive acquisition function selection via a multi-arm bandit strategy and diverse batch selection via determinantal point processes for multi-objective output space diversity.

Overview of PDBO. PDBO is an iterative algorithm. It introduces novel methods for selecting varying acquisition functions and for bringing the diversity of inputs into the multi-objective BO setting. The method builds K independent Gaussian processes GP_1, \dots, GP_K as surrogates for each of the objective functions. Its three key steps at each iteration t to select B inputs for evaluation are:

1. *Solving multiple cheap MOO problems:* PDBO takes as input a portfolio of acquisition functions, $\mathcal{P} = \{AF_1, \dots, AF_M\}$ for single-objective BO. It constructs M cheap MOO problems, each corresponding to one AF. The multiple objectives defining the cheap MOO problems are acquisition functions respectively corresponding to the K objective functions. Solving cheap MOO problems will generate M cheap Pareto-sets of solutions $\mathcal{X}_c^1 \dots \mathcal{X}_c^M$.

2. *Diverse batch selection:* From each cheap Pareto set \mathcal{X}_c^j , a batch $X_t^{Bj} \subset \mathcal{X}_c^j$ of B inputs is selected using a diversity-aware approach based on DPPs. Importantly, the adapted DPP is configured to favor the diversity in the output space and to handle multiple objective settings by using a principally fitted convex combination of the kernels of the K Gaussian processes. The convex combination scalars are strategically set to maximize the likelihood of selecting a diverse subset of inputs with respect to the Pareto front.

3. *Acquisition function selection:* From $\{X_t^{Bj}; \forall j \in [1, \dots, M]\}$, only one nominated subset would be selected using a multi-arm bandit strategy. The keys to this selection

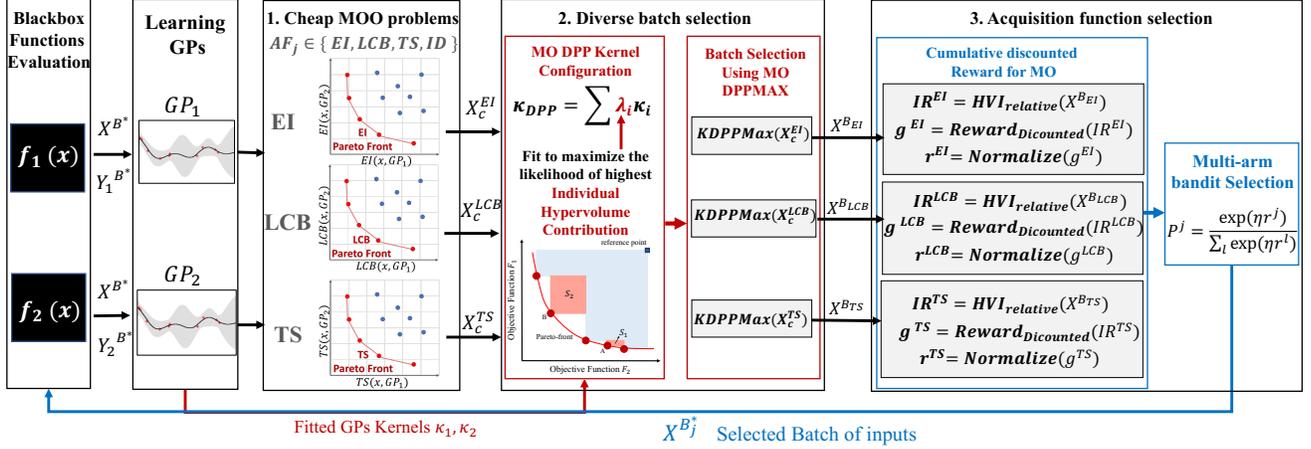


Figure 1: High-level overview of PDBO algorithm illustrating its three key components.

are probabilities p^1, \dots, p^j , one for each acquisition function to capture their performance based on past iterations. p^j is the probability of selecting the batch generated by the acquisition function AF_j , defined in equation 9. These probabilities are updated based on the discounted cumulative reward r^j of each of the respective acquisition functions. The reward values r^j are updated based on the quality of the batches nominated by the respective acquisition functions.

Algorithm 1 provides a pseudocode with high-level steps of the PDBO approach. The DPP-SELECT and ADAPTIVE-AF-SELECT represent the second and third key steps. The details of these methods and their corresponding pseudocodes are provided in Sections 4.1 and 4.2, respectively.

4.1 Multi-arm Bandit Strategy for Adaptive Acquisition Function Selection

In this section, we propose a multi-arm bandit (MAB) approach to *adaptively* select one acquisition function (AF) from a given library of AFs in each iteration of PDBO.

Multi-arm Bandit Formulation. We are given a portfolio of M acquisition functions $\mathcal{P} = \{AF_1, \dots, AF_M\}$ and our goal is to adaptively select one AF in each iteration of PDBO. Each acquisition function in \mathcal{P} corresponds to one arm, and we need to select an arm based on the performance of past selections for solving the MOO problem. Inspired by the previous work on AF selection and algorithm selection in the single objective setting [Hoffman et al. 2011, Vasconcelos et al. 2019, 2022], we propose an adaptive acquisition function selection approach for the MOO setting (see Algorithm 2). We explain the two main steps of this approach below.

Nominating Promising Candidates via Cheap MOO. In each PDBO iteration, we employ the updated statistical models $\{GP_1 \dots GP_K\}$ and the portfolio \mathcal{P} to generate M sets of candidate points. For each AF_j , the algorithm constructs a cheap MOO problem with the objectives defined as $AF_j(GP_1, \mathbf{x}) \dots AF_j(GP_K, \mathbf{x})$. Assuming minimization, the cheap MOO generate M Pareto-sets of solutions $\mathcal{X}_c^1 \dots \mathcal{X}_c^M$

Algorithm 1: Pareto front-Diverse Batch Multi-Objective BO

Input: \mathfrak{X} input space; $\{f_1, \dots, f_K\}$, K black-box objective functions; $\mathcal{P} = \{AF_1, \dots, AF_M\}$ portfolio of acquisition functions; B batch size; and T_{max} number of iterations

- 1: Initialize data $\mathcal{D}_0 = \{\mathcal{X}_0, \mathcal{Y}_0\}$ with N_0 initial points
- 2: **for** each iteration $t \in [1, T_{max}]$ **do**
- 3: Fit statistical models GP_1, \dots, GP_K using \mathcal{D}_{t-1}
- 4: **for** each acquisition function $AF_j \in \mathcal{P}$ **do**
- 5: $\mathcal{X}_c^j \leftarrow \arg \min_{\mathbf{x} \in \mathfrak{X}} (AF_j(GP_1, \mathbf{x}), \dots, AF_j(GP_K, \mathbf{x}))$
// Solve cheap MOO problem
- 6: $X_t^{Bj} \leftarrow \text{DPP-SELECT}(\mathcal{X}_c^j, \{GP_1, \dots, GP_K\}, \mathcal{D}_{t-1})$
// Select a batch of inputs from \mathcal{X}_c^j using DPPs
- 7: **end for**
- 8: $AF_{j^*} \leftarrow \text{ADAPTIVE-AF-SELECT}(\mathcal{D}_{t-1}, \{X_{t-1}^{Bj}\}_{j=1}^M)$
// Select an AF using previously aggregated data
- 9: $X_t^B = X_t^{Bj^*}$ // Choose the batch nominated by AF_{j^*}
- 10: $Y_t^B \leftarrow \{[f_1(\mathbf{x}), \dots, f_K(\mathbf{x})]; \forall \mathbf{x} \in X_t^B\}$
// Evaluate objective functions for batch of inputs X_t^B
- 11: $\mathcal{D}_t = \{\mathcal{X}_t, \mathcal{Y}_t\} \leftarrow \{\mathcal{X}_{t-1}, \mathcal{Y}_{t-1}\} \cup \{(X_t^B, Y_t^B)\}$
- 12: **end for**
- 13: **return** Pareto set $\mathcal{X}_{T_{max}}$ and Pareto front $\mathcal{Y}_{T_{max}}$

(one for each acquisition function) defined as:

$$\mathcal{X}_c^j \leftarrow \arg \min_{\mathbf{x} \in \mathfrak{X}} \{AF_j(GP_1, \mathbf{x}), \dots, AF_j(GP_K, \mathbf{x})\} \quad (5)$$

We employ the algorithm proposed by Deb et al. to solve cheap MOO problems defined in 5. From each \mathcal{X}_c^j , a batch $X_t^{Bj} \subset \mathcal{X}_c^j$ of B points is selected in iteration t using a diversity-aware approach described in Section 4.2. We denote the function evaluations of inputs in X_t^{Bj} by Y_t^{Bj} .

Multi-Objective Reward Update. We employ the relative hypervolume improvement as the quality metric to define our reward. The Pareto hypervolume captures the quality of nominated batches from a Pareto-dominance perspective and carries information about the represented trade-off between the multiple objectives. Defining the immediate reward as the raw Pareto hypervolume of the nominated batch

$IR_t^j = HV(Y_t^j)$ can lead to an undesirable assessment of the suitable acquisition function since a batch of points can have a large hypervolume value at iteration t but does not provide a significant improvement over the previous Pareto front \mathcal{Y}_{t-1} while another batch nominated in iteration $t-1$ may have a smaller hypervolume $HV(Y_{t-1}^j)$ yet provides a higher improvement over \mathcal{Y}_{t-2} . Additionally, initial iterations might provide drastic hypervolume improvements even if the selected points are not optimal. To mitigate these issues, we use the relative hypervolume improvement as the immediate reward instead of the hypervolume. In each BO iteration t , the immediate reward IR_t^j for each acquisition function AF_j ; $\forall j \in \{1 \dots M\}$ is defined as follows.

$$IR_t^j = \frac{HV(\tilde{\mathcal{Y}}_{t-1} \cup \tilde{Y}_t^j) - HV(\tilde{\mathcal{Y}}_{t-1})}{HV(\tilde{\mathcal{Y}}_{t-1})} \quad (6)$$

where $\tilde{\mathcal{Y}}_{t-1}$ is the Pareto front at iteration $t-1$ and \tilde{Y}_t^j is the evaluation of the batch of points X_t^j nominated by AF_j computed using the *predictive mean of the updated GP based statistical models*. As optimization progresses, statistical models provide a better representation of the objective functions, and the batches nominated by each AF become more informative about the quality of its selections. Hence, the impact of the early iterations may become irrelevant later. Consequently, we employ a *discounted cumulative reward* for each acquisition function AF_j at iteration t (denoted g_t^j).

$$g_t^j = \gamma g_{t-1}^j + IR_t^j = \sum_{t' \leq t} \gamma^{t-t'} IR_{t'}^j \quad (7)$$

Where γ is a decay rate that trades off past and recent improvements. The use of the decay rate can lead to equal or comparable rewards in advanced iterations, causing the algorithm to select the acquisition function randomly. To address this problem, the discounted cumulative reward g_t^j should be normalized [Vasconcelos et al. 2019]. The rewards at the first iteration are all initialized to zero and then updated at each iteration t using the following expression:

$$r_t^j = \frac{g_t^j - g_{max}^j}{g_{max}^j - g_{min}^j} \quad (8)$$

where $g_{max}^j = \max(\{g_{t'}^j, t' \in [1, t]\})$ and $g_{min}^j = \min(\{g_{t'}^j, t' \in [1, t]\})$. Finally, the probability of selection of each AF at each iteration t is calculated using equation 9.

$$p_t^j = \frac{\exp(\eta r_t^j)}{\sum_{l=1}^M \exp(\eta r_t^l)} \text{ for } j = 1 \dots M \quad (9)$$

The proposed MAB approach is an extension of the Hedge algorithm but is fundamentally distinct in its methodology and applicability. While it does generalize certain aspects of Hedge, it introduces critical variations that make it unique within the context of this study. Our approach diverges from the original Hedge algorithm by incorporating two key modifications: **1)** The use of discounted rewards and the application of normalization. Unlike the conventional Hedge, where rewards are typically used without any discounting or normalization, our strategy accounts for these factors, enhancing

Algorithm 2: ADAPTIVE-AF-SELECT

Input: training data \mathcal{D}_t ; Batches nominated by different AFs $\{X_t^{Bj}, \forall j \in \{1 \dots M\}\}$

- 1: **if** $t == 0$ **then**
- 2: $r_t^j = 0$ for $j = 1 \dots M$
- 3: **else**
- 4: Compute rewards: r_t^j for $j = 1 \dots M$ using Equation 8
- 5: Update probabilities: $p_t^j = \frac{\exp(\eta r_t^j)}{\sum_{l=1}^M \exp(\eta r_t^l)}$ for $j = 1 \dots M$
- 6: **end if**
- 7: Select AF_{j^*} according to the probabilities $\{p_t^j\}_{j=1}^M$
- 8: **return** acquisition function AF_{j^*}

its adaptability to the specific problem domain. **2)** Another significant departure lies in the problem setting itself. Hedge was initially designed for single-objective optimization while our proposed approach solves the more challenging problem of multi-objective optimization. This shift in focus has substantial implications, as it requires an entirely different set of considerations and techniques to address the complexities introduced by multiple conflicting objectives.

Remark. It is important to note that we are using a full information multi-arm bandit strategy that requires the reward to be updated for all possible actions (i.e., for all acquisition functions) at each iteration. Since we evaluate only the batch nominated by the selected AF, we achieve this by computing the reward using the predictive mean functions of the *updated surrogate models*. For this reason, we solve a cheap MOO problem for each $AF_j \in \mathcal{P}$ even though the acquisition function is selected based on the data from the previous iterations. Algorithm 2 provides the pseudocode of the adaptive AF selection based on the estimated rewards and probabilities.

4.2 DPPs for Batch Selection

We explain our approach to select a batch of diverse inputs by configuring DPPs to promote *output space diversity*.

Determinantal Point Processes. (DPPs) [Kulesza et al. 2012] are well-suited to model samples of a diverse subset of k points from a predefined set of n of points. Given a similarity function over a pair of points, DPPs assign a high probability of selection to the most diverse subsets according to the similarity function. The similarity function is typically defined as a kernel. Formally, given a DPP kernel defined over a set S of n elements, the k -DPP distribution is defined as selecting a subset S' of size k with $S' \subset S$ with probability proportional to the determinant of the kernel:

$$Pr(S') = \frac{\det(\kappa(S'))}{\sum_{|s|=k} \det(\kappa(s))} \quad (10)$$

[Kathuria et al. 2016] introduced the use of DPP for batch BO in the single-objective setting. Given the surrogate GP of the objective function, the covariance of the GP is used as the similarity function for the DPP. The approach selects the first point in the batch by maximizing the UCB acquisition function. Next, it creates a set of points referred to as relevance region by bounding the search space with the maximizer of the LCB acquisition function and manually discretizing the

bounded space into a grid of n points. DPP selects the remaining $(k - 1)$ points out of the n points in the relevance region. [Wang et al. 2017, Oh et al. 2021, Nava et al. 2022] used similar techniques to apply DPPs to high-dimensional and discrete spaces.

There exist two approaches to selecting a diverse subset with a fixed size via DPP: 1) Choosing the subset that maximizes the determinant referred to as *DPP-max*; and 2) Sampling with the determinantal probability measure referred to as *DPP-sample*. In this paper, we will focus on DPP-max. Although selecting the subset that maximizes the determinant is an NP-Hard problem, several approximations were proposed [Nikolov 2015]. A greedy strategy [Kathuria et al. 2016] provides an approximate solution and was adopted in several BO papers [Wang et al. 2017]

Limitations of Prior Work and Challenges for MOO. We list the key limitations of prior methods for DPP-based batch selection in the single-objective setting as they are applicable to the multi-objective setting too. **L1)** How can we overcome the limitation of selecting the first point separately regardless of the DPP diversity? **L2)** How can we prevent the potential limitation of under-explored search space caused by the discretization of the space to create the relevance region set?

The key challenges to employing DPPs for batch selection in the MOO setting include **C1)** How to define a kernel that captures the diversity for multiple objectives given that we have K separate surrogate models and their corresponding kernel? **C2)** How can the DPP kernel capture the Pareto front diversity and the trade-off between the objectives without compromising the Pareto quality of selected points?

DPPs for Multi-objective BO. We propose principled methods to address the limitations of prior work on DPPs for batch BO (**L1** and **L2**) and the challenges **C1** and **C2** for MOO.

Multi-objective Relevance Region. Our proposed algorithm naturally mitigates the two limitations of the single-objective DPP approach. Recall that the first step of PDBO algorithm (Section 4.1) proposes to generate cheap approximate Pareto-sets which capture the trade-offs between the objectives in the utility space and might include, with high probability, optimal points [Belakaria et al. 2020a, Konakovic Lukovic et al. 2020]. We consider the cheap Pareto sets as the multi-objective relevance region. Our approach allows for generating the relevance region without manually discretizing the search space. Also, the full batch is selected from the multi-objective relevance region leading to a better diversity among all the points in the batch.

Multi-objective DPP Kernel Fitting. To overcome the challenges of using DPPs in the MOO setting, we build a new kernel κ_{dpp} that is defined as a convex combination of the K kernels of the statistical models (GPs) representing each of the black-box objective functions. Let $\Lambda = [\lambda_1, \dots, \lambda_K]$ be a vector of size K where each λ_i corresponds to the convex combination scalar associated with kernel κ_i of the objective function f_i . The DPP kernel κ_{DPP} is defined as:

$$\kappa_{DPP} = \sum_{i=1}^K \lambda_i \cdot \kappa_i \quad \text{s.t.} \quad \sum_{i=1}^K \lambda_i = 1 \quad (11)$$

Algorithm 3: DPP-SELECT

Input: cheap Pareto set \mathcal{X}_c ; surrogate models GP_1, \dots, GP_K ; data \mathcal{D}_t .

- 1: $\mathbf{C}_t = [HVC(\mathbf{y}), \forall \mathbf{y} \in \mathcal{Y}_t]$ // Calculate the individual hypervolume contribution for each input $\in \mathcal{D}_t$
- 2: $\kappa_{DPP} = \sum_{i=1}^K \lambda_i \cdot \kappa_i$ s.t. $\sum_{i=1}^K \lambda_i = 1$
// Construct κ_{DPP} as a convex combination of function kernels
- 3: $\Lambda^* = \arg \min_{\Lambda \in [0,1]^K} \log p(\mathbf{C}_t | \mathcal{X}_t)$ s.t. $\sum_{i=1}^K \lambda_i = 1$
// Select λ_i values by maximizing the LML
- 4: Use the fitted κ_{DPP} kernel to select the most diverse points X_t^B from the cheap Pareto set \mathcal{X}_c via DPP-max
- 5: **return** the selected B inputs, X_t^B

The hyperparameters of the kernels $\kappa_i \forall i \in \{1, 2, \dots, K\}$ are fixed during the fitting of the GPs. To set the convex combination scalars Λ in a principled manner that promotes diverse batch selection, we propose to set the Λ by maximizing the log marginal likelihood of selecting points with the highest individual hypervolume contribution.

The individual hypervolume contribution (HVC) of each point in the evaluated Pareto front \mathcal{Y}_t (via evaluated training data \mathcal{D}_t) is the reduction in hypervolume if the point is removed from the Pareto front. HVC is considered a Pareto front (PF) diversity indicator [Daulton et al. 2022]. Points in crowded regions of the Pareto front have smaller HVC values. Therefore, more Pareto front points with high HVC indicate more output space coverage and consequently, higher PF diversity.

$$HVC(\mathbf{y}) = HV(\mathcal{Y}_t) - HV(\mathcal{Y}_t \setminus \{\mathbf{y}\}) \quad \forall \mathbf{y} \in \mathcal{Y}_t \quad (12)$$

Given equation 12, we can construct the training set for the fitting of Λ . Let $\mathbf{C}_t = [HVC(\mathbf{y}); \forall \mathbf{y} \in \mathcal{Y}_t]$ be a vector of the individual HV contributions of evaluated points $\mathcal{X}_t \in \mathcal{D}_t$.

$$\Lambda^* = \operatorname{argmin}_{\Lambda \in [0,1]^K} \log p(\mathbf{C}_t | \mathcal{X}_t) \quad \text{s.t.} \quad \sum_{i=1}^K \lambda_i = 1 \quad (13)$$

$$\begin{aligned} \text{where } \log p(\mathbf{C}_t | \mathcal{X}_t) = & -\frac{1}{2} \mathbf{C}_t^T \mathcal{K}_{DPP}^{-1} \mathbf{C}_t \\ & -\frac{1}{2} \log |\mathcal{K}_{DPP}| - \frac{n}{2} \log 2\pi \end{aligned}$$

Algorithm 3 provides the pseudo-code for selecting the diverse batch from a given candidate/cheap Pareto set \mathcal{X}_c .

5 Theoretical Analysis

Prior work developed regret bound for different single objective acquisition functions including UCB [Srinivas et al. 2009, Belakaria et al. 2020a]. However, the theoretical analysis of our proposed MAB algorithm is more challenging as it involves input selection based on different acquisition functions at each iteration t . The choices made at any given iteration t influence the state and subsequent rewards of all future iterations [Hoffman et al. 2011] and therefore there is a need to adapt prior theoretical analysis. Additionally, regret bounds for Hedge MAB strategies have been developed independently outside the context of acquisition function selection [Cesa-Bianchi and Lugosi 2006]. We follow similar

steps suggested by Hoffman et al. [2011] to derive a suitable regret bound for our MOO setting.

We assume maximization of objectives and further assume that the UCB acquisition function is in the portfolio of acquisition functions used by PDBO. We make this choice for the sake of clarity and ease of readability as we build our theoretical analysis on prior seminal work [Srinivas et al. 2009, Hoffman et al. 2011]. Notably, this is not a restrictive assumption, and with minimal mathematical manipulations, the same derived regret bound holds for the case of minimization with the LCB acquisition function being in the portfolio instead.

To simplify the proof and solely for the sake of theoretical regret bound, we consider the instant reward at iteration t to be the sum of predictive means of the Gaussian processes

$$IR_t = \sum_{i=1}^k \mu_{i,t-1}(\mathbf{x}_t) \quad (14)$$

where $\mu_{i,t-1}$ is the posterior mean of function i . The cumulative reward over T_{max} iterations that would have been obtained using acquisition function AF_j is defined as:

$$r_{T_{max}}^j = \sum_{t=1}^{T_{max}} IR_t = \sum_{t=1}^{T_{max}} \sum_{i=1}^k \mu_{i,t-1}(\mathbf{x}_t^j). \quad (15)$$

It is important to note that in our proposed PDBO algorithm, we use a different and better instant reward IR_t and cumulative reward $r_{T_{max}}^i$. The rewards in equation 14 is a design choice to achieve the following regret bound. In Section 5.1, we provide a discussion accompanied by an experimental ablation study comparing the reward function used in theory to the reward function used in PDBO.

Theorem 5.1. *Let \mathbf{x}^* be a point in the optimal Pareto set \mathcal{X}^* . Let \mathbf{x} be a point in the Pareto set \mathcal{X}_t estimated by PDBO via solving cheap MOO problem at the t^{th} iteration. Let the cumulative regret for the multiple objectives be defined as $R_{T_{max}}(\mathbf{x}^*) = \sum_{t=1}^{T_{max}} \sum_{i=1}^k f_i(\mathbf{x}^*) - f_i(\mathbf{x}_t)$*

Assuming maximization of objectives and that UCB is in the portfolio of acquisition functions, let β_t be the UCB parameter and $\gamma_{T_{max}}^i$ be the bound on the information gained for function i at points selected by PDBO after T_{max} iterations, then with probability at least $1 - \delta$ the cumulative regret is bounded by

$$R_{T_{max}}(\mathbf{x}^*) \leq \mathcal{O}(\sqrt{T_{max}}) + \sum_{t=1}^{T_{max}} \sum_{i=1}^k \sqrt{\beta_t \sigma_{i,t-1}(\mathbf{x}_t^{UCB})} + \sqrt{C_i T_{max} \beta_{T_{max}} \gamma_{T_{max}}^i}.$$

We provide complete proof in the Appendix. The theorem suggests that our regret is bounded by two sublinear terms and another term that might include points suggested by UCB but not necessarily selected by the Hedge strategy. Additionally, the theoretical proof accounts only for sequential input selection. Extension to batch selection using DPP is possible by carefully accounting for results introduced by Kathuria et al. [2016].

5.1 Analysis and Ablation Study

To simplify the proof for regret bound, we defined a new instant reward and cumulative reward in equations 14 and 15 that are different from the rewards we use for PDBO in equations 6, 7, and 8. The goal of this section was to define a reward that allows tractable theoretical analysis. However, this reward is not intuitive and has several practical issues: **1)** It is defined as a summation over predictive means of the functions and does not provide any insight on the quality of the selected points; **2)** It does not account for improvement with respect to previous iterations which is uninformative in terms of the quality of points selected by different acquisition functions at different iterations; **3)** It is non-discounted and does not account for the importance of the iterative progress of the input selection; and **4)** It is not normalized and therefore can lead to random selection in the advanced iterations. All stated issues have been addressed by our proposed reward function which we carefully designed to be intuitive and informative about the different acquisition strategies and to mitigate potential numerical issues. We performed an ablation study to compare the performance of PDBO when using our proposed reward function and the reward function in the theoretical proof. Our results in the Appendix show superior performance for our proposed reward strategy.

6 Experiments and Results

We provide experimental details and compare PDBO to baseline methods on multiple MOO benchmarks and varying batch sizes. We evaluate all methods using the hypervolume indicator and diversity of Pareto front (DPF) measure.

Benchmarks. We conduct experiments on benchmarks with varying numbers of input and output dimensions to show the versatility and flexibility of our method. We use several synthetic problems: ZDT-1, ZDT-2, ZDT-3 [Zitzler et al. 2000], DTLZ-1, DTLZ-3, DTLZ-5 [Deb et al. 2005] and three real world problems: the gear train design problem [Deb and Srinivasan 2006, Konakovic Lukovic et al. 2020], SWL-LVM [Siegmond et al. 2012] and Unmanned aerial vehicle power system design [Belakaria et al. 2020b]. More details and descriptions of problem settings are included in the Appendix. In the ablation studies, we provide additional experiments with synthetic benchmarks where we vary the input and output dimensions.

Baselines. We compare our PDBO method to state-of-the-art batch MOO methods: DGEMO, qEHVI, qPAREGO, and USEMO-EI. We also include NSGA-II as the evolutionary algorithm baseline and random input selection. We set the hyperparameters of PDBO to $\gamma = 0.7$ and $\tau = 4$ as recommended by [Hoffman et al. 2011, Vasconcelos et al. 2019]. We define the AF portfolio as $\mathcal{P} = \{EI, TS, UCB, ID\}$.

Experimental Setup. All experiments are initialized with five random inputs/evaluations and run for at least 250 function evaluations. We conduct experiments with four different batch sizes $B \in \{2, 4, 8, 16\}$ and adjust the number of iterations accordingly. For instance, when using a batch size of two, we run the algorithm for 125 iterations. Each experiment is repeated 25 times, and we report the average and standard deviation of the hypervolume indicator and the

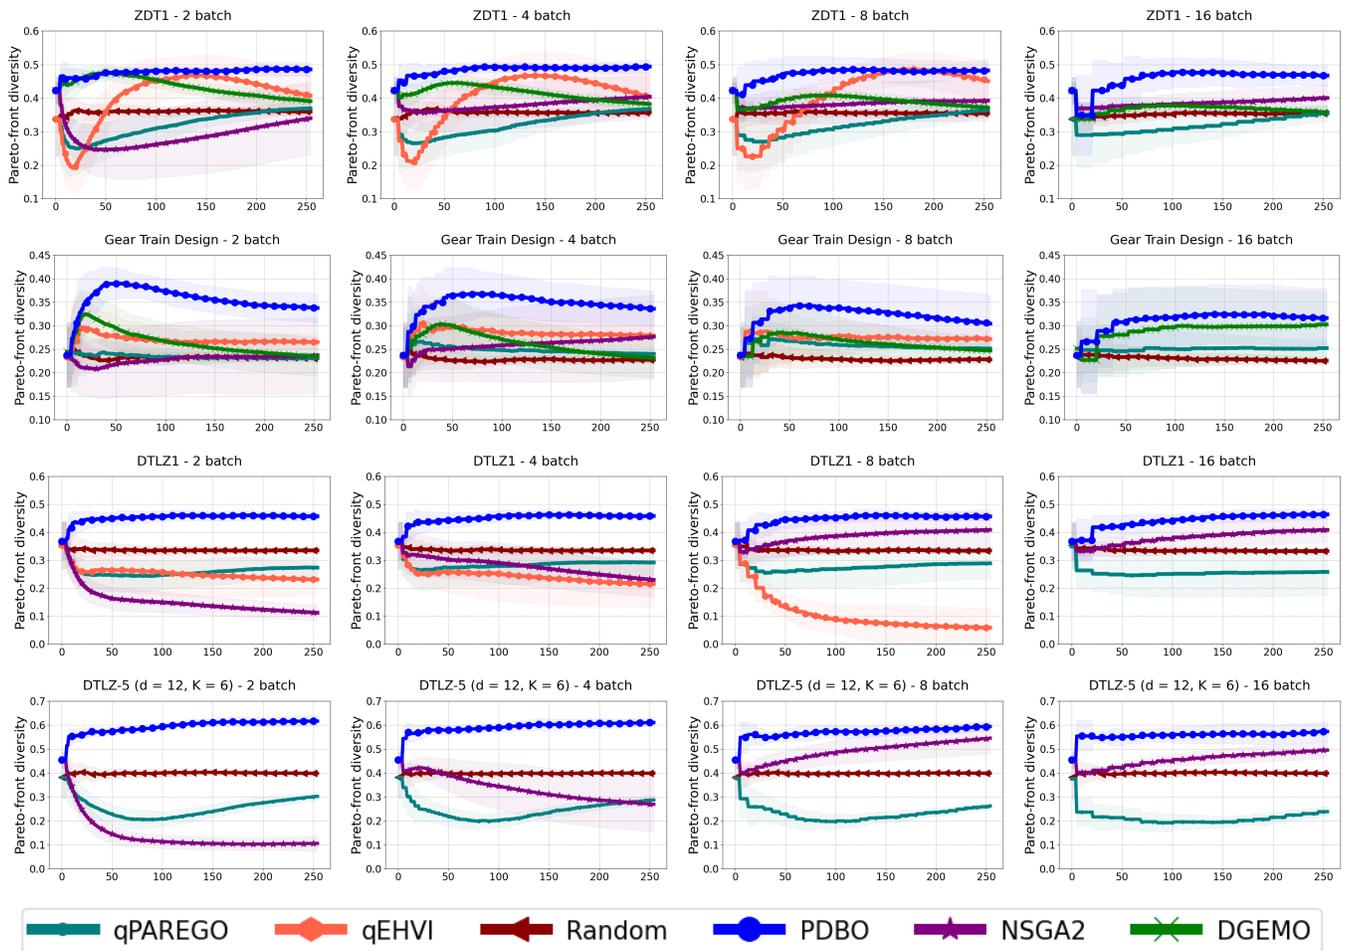


Figure 2: Diverse Pareto front (DPF) results evaluated on multiple benchmarks and batch sizes.

DPF metric. To solve the constrained optimization problem in the DPP algorithm, we utilize an implementation of the SQP method [Lalee et al. 1998, Nocedal and Wright 2006] from the Python SciPy library [Virtanen et al. 2020]. For baselines, we use the codes and hyperparameters provided in the open source repositories of DGEMO¹ and Botorch². We provide the details for the NSGA-II baseline and cheap MO solver, and more details about the setup for fitting the hyperparameters of GP models in the Appendix.

Results and Discussion. Figure 2 demonstrates that PDBO outperforms other baselines with respect to the Pareto-front diversity metric. Additionally, Figure 3 demonstrates that PDBO outperforms all baseline methods in most experiments with respect to the Hypervolume indicator and provides a competitive performance on the others.

In the Appendix, we present a comprehensive set of additional results and analyses. This includes the evaluation of hypervolume and DPF on various benchmarks. We also introduce results using other metrics, notably the Inverted

Generational Distance (IGD) and a modified version of DPF, accompanied by a relevant discussion. Additionally, we compare the run-time of all baseline methods. For visual insight into the diversity of solutions, we include scatter plots representing the Pareto front for problems with two objectives. Lastly, we provide statistics on the selection of AF.

PDBO Advantages. PDBO is fast and effective in producing high-quality and diverse Pareto fronts. While outperforming the baseline methods, it can also be used with any number of input and output dimensions as well as being flexible to run with any batch size. The two state-of-the-art methods are DGEMO and qEHVI. The DGEMO method fails to run for experiments with more than three objective functions as the graph cut algorithm consistently crashes (same observation was made by [Daulton et al. 2021]). qEHVI fails to run with batch sizes higher than eight as the method becomes extremely memory-consuming even with GPUs. We provide a more detailed discussion about these limitations in the Appendix. Therefore, PDBO’s ability to easily run with any input and output dimensions as well as any batch size is an advantage for practitioners. PDBO is capable of proactively creating a diverse Pareto front while improving

¹<https://github.com/yunshengtian/DGEMO>

²<https://github.com/pytorch/botorch>

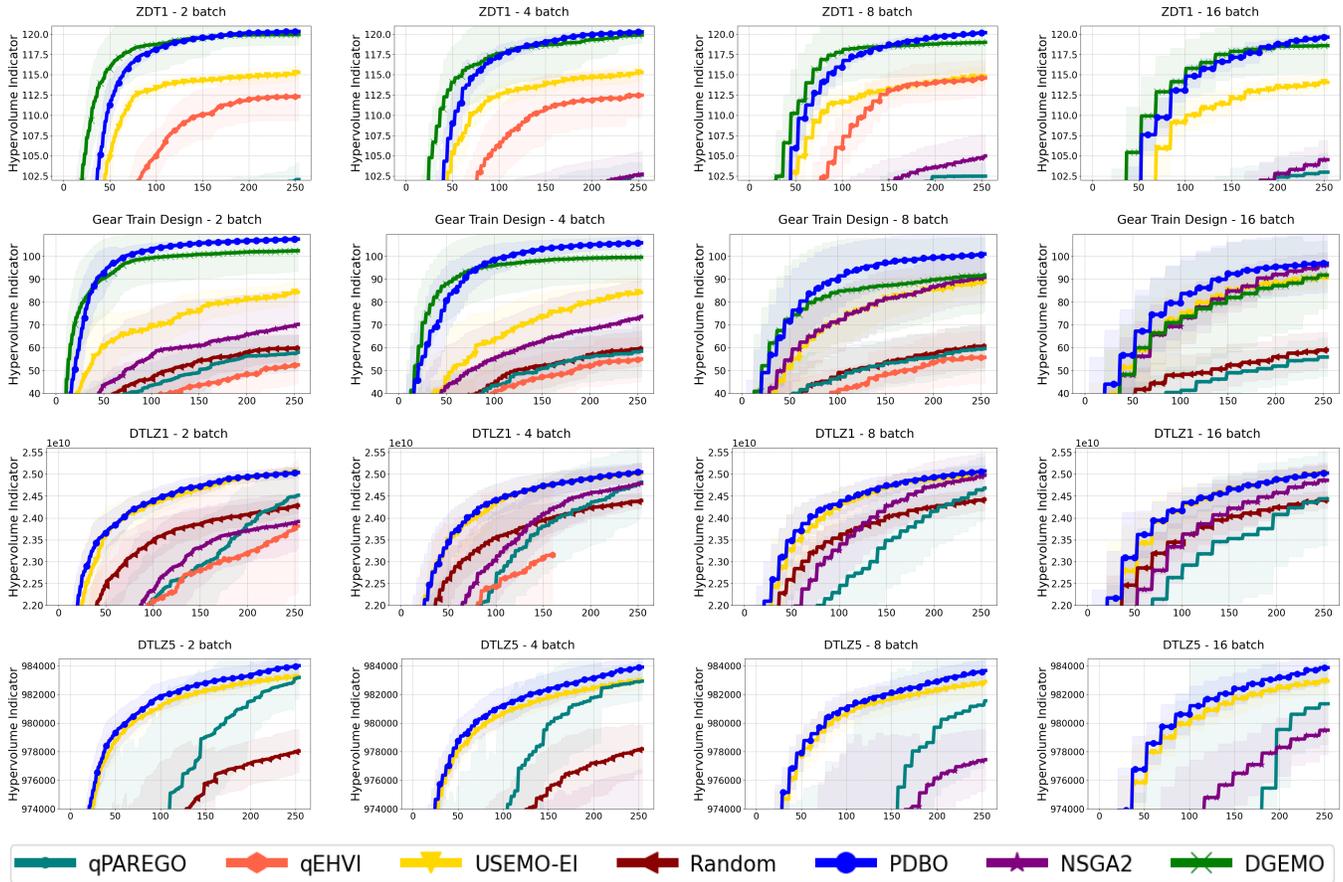


Figure 3: Hypervolume results evaluated on multiple benchmarks and batch sizes.

or maintaining the quality of the Pareto front.

Given that PDBO incorporates two key contributions, namely adaptive acquisition function selection and multi-objective batch selection using DPPs, we examine the individual contributions of each component to the overall performance by conducting ablation experiments.

Merits of Adaptive AF Selection. We demonstrate the superiority of the adaptive AF selection method, as outlined in Section 4.1, compared to using a static AF from the portfolio. To isolate the impact of this component from the batch selection process, we conduct an ablation study using the USEMO baseline. With a batch size of one, we consider USEMO with UCB, TS, ID, and EI as baselines. We then evaluate the efficacy of our MAB method by incorporating the adaptive AF selection approach into USEMO. Results shown in the Appendix consistently demonstrate the superior performance of the MAB strategy over using a static AF.

Merits of DPP-Based Batch Selection for MOO. Following a similar ablation approach, we employ the USEMO-EI baseline to examine the impact of the DPP-based batch selection. USEMO-EI selects the next input for evaluation from the cheap Pareto set based on an uncertainty metric. To perform this ablation, we replace the input selection mechanism utilized in USEMO with our proposed DPP selection strategy

and compare their performance. The ablation is conducted across different batch sizes $B \in \{2, 4, 8, 16\}$. The results presented in the Appendix reveal that the proposed DPP selection strategy, referred to as DPP-EI, surpasses the USEMO selection strategy in terms of diversity while simultaneously improving the quality of hypervolume.

7 Summary

We studied the Pareto front-Diverse Batch Multi-Objective BO (PDBO) method based on the BO framework. It employs a full information multi-arm bandit algorithm with discounted reward to adaptively select the most suitable acquisition function in each iteration. We also proposed an appropriate reward based on the relative hypervolume contribution of each acquisition function and a multi-objective DPP approach configured to select a batch of Pareto-diverse inputs for evaluation. Experiments on multiple benchmarks demonstrate that PDBO outperforms prior methods in terms of both diversity and quality of Pareto-front solutions.

Acknowledgements The authors gratefully acknowledge the in part support from National Science Foundation (NSF) grants IIS-1845922, SII-2030159, and CNS-2308530. The views expressed are those of the authors and do not reflect the official policy or position of the NSF.

References

- Majid Abdolshah, Alistair Shilton, Santu Rana, Sunil Gupta, and Svetha Venkatesh. Multi-objective bayesian optimisation with preferences over objectives. *Advances in neural information processing systems*, 32, 2019.
- Christof Angermueller, David Belanger, Andreea Gane, Zeldia Mariet, David Dohan, Kevin Murphy, Lucy Colwell, and D Sculley. Population-based black-box optimization for biological sequence design. In *International Conference on Machine Learning*, pages 324–334. PMLR, 2020.
- MF Ashby. Multi-objective optimization in material design and selection. *Acta materialia*, 2000.
- Raul Astudillo and Peter Frazier. Multi-attribute bayesian optimization with interactive preference learning. In *International Conference on Artificial Intelligence and Statistics*, pages 4496–4507. PMLR, 2020.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *JMLR*, 2002.
- Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Max-value entropy search for multi-objective Bayesian optimization. In *Conference on Neural Information Processing Systems*, 2019.
- Syrine Belakaria, Aryan Deshwal, Nitthilan Kannappan Jayakodi, and Janardhan Rao Doppa. Uncertainty-aware search framework for multi-objective bayesian optimization. In *AAAI*, 2020a.
- 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)*, 2020b.
- Gulnur Birol, Cenk Undey, and Ali Cinar. A modular simulation package for fed-batch fermentation: penicillin production. *Computers and chemical engineering*, 2002.
- Alexei Borodin. Determinantal point processes. *arXiv preprint arXiv:0911.1153*, 2009.
- Alexei Borodin and Grigori Olshanski. Harmonic analysis on the infinite-dimensional unitary group and determinantal point processes. *Annals of mathematics*, pages 1319–1422, 2005.
- Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- Carlos A Coello Coello and Margarita Reyes Sierra. A study of the parallelization of a coevolutionary multi-objective evolutionary algorithm. In *MICAI 2004: Third Mexican International Conference on Artificial Intelligence, Mexico City, Mexico, April 26-30, 2004. Proceedings 3*. Springer, 2004.
- Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Differentiable expected hypervolume improvement for parallel multi-objective bayesian optimization. *Advances in Neural Information Processing Systems*, 33:9851–9864, 2020.
- Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Parallel bayesian optimization of multiple noisy objectives with expected hypervolume improvement. *NeurIPS*, 34, 2021.
- Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Multi-objective bayesian optimization over high-dimensional search spaces. In *The 38th Conference on Uncertainty in Artificial Intelligence*, 2022. URL <https://openreview.net/forum?id=r5IEvvIs9xq>.
- Kalyanmoy Deb and Aravind Srinivasan. Innovization: Innovating design principles through optimization. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 1629–1636, 2006.
- Kalyanmoy Deb, Ram Bhushan Agrawal, et al. Simulated binary crossover for continuous search space. *Complex systems*, 1995.
- Kalyanmoy Deb, Mayank Goyal, et al. A combined genetic adaptive search (geneas) for engineering design. *Computer Science and informatics*, 26:30–45, 1996.
- Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, T Meyarivan, and A Fast. Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002a.
- Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6, 2002b.
- Kalyanmoy Deb, Lothar Thiele, Marco Laumanns, and Eckart Zitzler. Scalable test problems for evolutionary multi-objective optimization. In *Evolutionary multiobjective optimization*. Springer, 2005.
- Aryan Deshwal, Cory M Simon, and Janardhan Rao Doppa. Bayesian optimization of nanoporous materials. *Molecular Systems Design and Engineering*, 2021.
- Michael Emmerich and Jan-willem Klinkenberg. The computation of the expected improvement in dominated hypervolume of pareto front approximations. *Leiden University*, 2008.
- David Eriksson, Michael Pearce, Jacob Gardner, Ryan D Turner, and Matthias Poloczek. Scalable global optimization via local bayesian optimization. *NeurIPS*, 2019.
- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1), 1997.
- Daniel Hernández-Lobato, Jose Hernandez-Lobato, Amar Shah, and Ryan Adams. Predictive entropy search for multi-objective Bayesian optimization. In *Proceedings of International Conference on Machine Learning (ICML)*, pages 1492–1501, 2016.
- José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In *NeurIPS*, 2014.
- Matthew Hoffman, Eric Brochu, Nando De Freitas, et al. Portfolio allocation for bayesian optimization. In *UAI*, pages 327–336, 2011.
- Carl Hvarfner, Frank Hutter, and Luigi Nardi. Joint entropy search for maximally-informed bayesian optimization. *arXiv preprint arXiv:2206.04771*, 2022.
- Moksh Jain, Sharath Chandra Rapparthi, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Yoshua Bengio, Santiago Miret,

- and Emmanuel Bengio. Multi-objective gflownets. *arXiv preprint arXiv:2210.12765*, 2022.
- Tarun Kathuria, Amit Deshpande, and Pushmeet Kohli. Batched gaussian process bandit optimization via determinantal point processes. *NeurIPS*, 29, 2016.
- Joshua Knowles. Parego: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1):50–66, 2006.
- Mina Konakovic Lukovic, Yunsheng Tian, and Wojciech Matusik. Diversity-guided multi-objective bayesian optimization with batch evaluations. *Advances in Neural Information Processing Systems*, 2020.
- Alex Kulesza, Ben Taskar, et al. Determinantal point processes for machine learning. *Foundations and Trends in Machine Learning*, 2012.
- Marucha Lalee, Jorge Nocedal, and Todd Plantenga. On the implementation of an algorithm for large-scale equality constrained optimization. *SIAM Journal on Optimization*, 8(3):682–706, 1998.
- Xi Lin, Zhiyuan Yang, Xiaoyuan Zhang, and Qingfu Zhang. Pareto set learning for expensive multi-objective optimization. *arXiv preprint arXiv:2210.08495*, 2022.
- Jonas Mockus, Vytautas Tiesis, and Antanas Zilinskas. The application of bayesian methods for seeking the extremum. *Towards global optimization*, 1978.
- Elvis Nava, Mojmir Mutny, and Andreas Krause. Diversified sampling for batched bayesian optimization with determinantal point processes. In *International Conference on Artificial Intelligence and Statistics*, pages 7031–7054. PMLR, 2022.
- Christos A Nicolaou and Nathan Brown. Multi-objective optimization methods in drug design. *Drug Discovery Today: Technologies*, 2013.
- Aleksandar Nikolov. Randomized rounding for the largest simplex problem. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 861–870, 2015.
- J Nocedal and SJ Wright. Numerical optimization (springer, new york, 1999)., 2006.
- Changyong Oh, Roberto Bondesan, Efstratios Gavves, and Max Welling. Batch bayesian optimization on permutations using acquisition weighted kernels. *arXiv preprint arXiv:2102.13382*, 2021.
- Michael Aggrey Okoth, Ronghua Shang, Licheng Jiao, Jehangir Arshad, Ateeq Ur Rehman, and Habib Hamam. A large scale evolutionary algorithm based on determinantal point processes for large scale multi-objective optimization problems. *Electronics*, 11(20):3317, 2022.
- Biswajit Paria, Kirthevasan Kandasamy, and Barnabás Póczos. A flexible framework for multi-objective bayesian optimization using random scalarizations. In *Uncertainty in Artificial Intelligence*, pages 766–776. PMLR, 2020.
- Thomas Pierrot, Guillaume Richard, Karim Beguir, and Antoine Cully. Multi-objective quality diversity optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 2022.
- 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*, 2015.
- Norbert Siegmund, Sergiy S Kolesnikov, Christian Kästner, Sven Apel, Don Batory, Marko Rosenmüller, and Gunter Saake. Predicting performance via automated feature-interaction detection. In *Proceedings of the 34th International Conference on Software Engineering (ICSE)*, pages 167–177, 2012.
- Niranjana Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.
- Shinya Suzuki, Shion Takeno, Tomoyuki Tamura, Kazuki Shitara, and Masayuki Karasuyama. Multi-objective bayesian optimization using pareto-frontier entropy. In *International Conference on Machine Learning*, pages 9279–9288. PMLR, 2020.
- Akito Taneda. Multi-objective optimization for rna design with multiple target secondary structures. *BMC bioinformatics*, 2015.
- William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25:285–294, 1933.
- Thiago de P Vasconcelos, Daniel ARMA de Souza, César LC Mattos, and João PP Gomes. No-past-bo: Normalized portfolio allocation strategy for bayesian optimization. In *International Conference on Tools with Artificial Intelligence (ICTAI)*. IEEE, 2019.
- Thiago de P Vasconcelos, Daniel Augusto RMA de Souza, Gustavo C de M Virgolino, César LC Mattos, and João PP Gomes. Self-tuning portfolio-based bayesian optimization. *Expert Systems with Applications*, 188:115847, 2022.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.
- Mengzhen Wang, Fangzhen Ge, Debao Chen, and Huaiyu Liu. A many-objective evolutionary algorithm using determinantal point process in potential region. In *Proceedings of the 6th International Conference on Control Engineering and Artificial Intelligence*, pages 83–91, 2022.
- Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient Bayesian optimization. In *ICML*, 2017.
- Zi Wang, Chengtao Li, Stefanie Jegelka, and Pushmeet Kohli. Batched high-dimensional bayesian optimization via struc-

tural kernel learning. In *International Conference on Machine Learning*, pages 3656–3664. PMLR, 2017.

Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006.

Peng Zhang, Jinlong Li, Tengfei Li, and Huanhuan Chen. A new many-objective evolutionary algorithm based on determinantal point processes. *IEEE Transactions on Evolutionary Computation*, 2020.

Eckart Zitzler and Lothar Thiele. Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach. *IEEE transactions on Evolutionary Computation*, 1999.

Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary computation*, 8(2):173–195, 2000.

8 Appendix

8.1 Theoretical Analysis

In this section, we assume maximization and assume that the UCB acquisition function is in the portfolio of acquisition functions used during the multi-objective Bayesian optimization process. We make this choice for the sake of clarity and ease of readability as we build our theoretical analysis on prior seminal work [Srinivas et al. 2009, Hoffman et al. 2011]. It is important to clarify that this is not a restrictive assumption and that with minimal mathematical transformations, the same derived regret bound holds for the case of minimization with the LCB acquisition function being in the portfolio instead.

In order to simplify the proof and solely for the sake of theoretical regret bound, we consider the instant reward of iteration t to be the sum of predictive means of the Gaussian processes.

$$IR_t = \sum_{i=1}^k \mu_{i,t-1}(\mathbf{x}_t) \quad (16)$$

with $\mu_{i,t-1}$ is the posterior mean of function i .

The cumulative reward over T_{max} iterations that would have been obtained using acquisition function j is defined as:

$$r_{T_{max}}^j = \sum_{t=1}^{T_{max}} IR_t = \sum_{t=1}^{T_{max}} \sum_{i=1}^k \mu_{i,t-1}(\mathbf{x}_t^j). \quad (17)$$

It is important to note that in our proposed algorithm, we use different and better-designed instant reward IR_t and cumulative reward $r_{T_{max}}^j$ functions. The rewards in equation 16 is a design choice to achieve the following regret bound. In section 5.1, we provide a discussion accompanied with an ablation study comparing the reward function used in theory to the reward function used in our proposed approach.

To bound the regret of Hedge with respect to the gain, we define the maximum strategy as: $r_{T_{max}}^{\max} = \max_j r_{T_{max}}^j$

Lemma 8.1. *With probability at least $1 - \delta_1$ and $\eta = \sqrt{8 \ln k / T_{max}}$, the regret is bounded by*

$$r_{T_{max}}^{\max} - r_{T_{max}}^{\text{Hedge}} \leq \mathcal{O}(\sqrt{T_{max}}). \quad (18)$$

This result follows directly from [Cesa-Bianchi and Lugosi 2006, Section 4.2] for rewards in the range $[0, 1]$. This result is also used in the proof of regret bound in [Hoffman et al. 2011]. [Hoffman et al. 2011] discussed possible generalizations that come at the cost of worsening the bound with a multiplicative or additive constant. These relaxations include a non-restrictive reward bound and a time variant η term. These generalizations also hold in the context of our proof. We refer the reader to [Hoffman et al. 2011] for more details. It is important to note that this lemma holds for any choice of IR_t where the reward is a result of the action taken by the Hedge algorithm.

The next lemmas are defined in [Srinivas et al. 2009] and [Hoffman et al. 2011]. We will refer the reader to [Srinivas et al. 2009, Lemma 5.1 and 5.3] and [Hoffman et al. 2011, Lemma 4 and 5] for proofs. We provide them here for the sake

of completeness. It is important to clarify that these lemmas only depend on the surrogate Gaussian process models and can be used regardless of the UCB acquisition function.

Lemma 8.2. *Assume $\delta_2 \in (0, 1)$, a finite sample space $|\mathcal{X}| < \infty$, and $\beta_t = 2 \log(|\mathcal{X}| \pi_t / \delta_2)$ where $\sum_t \pi_t^{-1} = 1$ and $\pi_t > 0$. Then with probability at least $1 - \delta_2$, the absolute deviation of the mean is bounded by*

$$|f_i(\mathbf{x}) - \mu_{i,t-1}(\mathbf{x})| \leq \sqrt{\beta_t} \sigma_{i,t-1}(\mathbf{x}) \quad \forall \mathbf{x} \in \mathcal{X}, \forall t \geq 1.$$

Lemma 8.3. *Given points \mathbf{x}_t selected by the algorithm, the following bound holds for the sum of variances:*

$$\sum_{t=1}^{T_{max}} \beta_t \sigma_{i,t}^2(\mathbf{x}_t) \leq C_i \beta_{T_{max}} \gamma_{T_{max}}^i,$$

where $C_i = 2 / \log(1 + \sigma_i^{-2})$.

The next lemma is defined in [Hoffman et al. 2011] and follows directly from [Srinivas et al. 2009, Lemma 5.2]. This lemma depends only on the definition of the UCB acquisition function, and does not require that points at any previous iteration were selected via UCB acquisition function.

Lemma 8.4. *The cheap multi-objective optimization (MOO) problem when the UCB acquisition function is selected is defined as follows:*

$$\max_{\mathbf{x} \in \mathcal{X}} (UCB_{1,t}(\mathbf{x}), \dots, UCB_{k,t}(\mathbf{x})) \quad (19)$$

Assuming that the cheap MOO solver achieves optimality, leading to the optimal Pareto set for the above defined problem, with the Pareto set defined as \mathcal{X}_t , either there exists a $\mathbf{x}_t^{\text{UCB}} \in \mathcal{X}_t$ such that

$$UCB_{i,t}(\mathbf{x}^*) \leq UCB_{i,t}(\mathbf{x}_t^{\text{UCB}}), \forall i \in \{1, \dots, k\} \quad (20)$$

or \mathbf{x}^* is in the optimal Pareto set \mathcal{X}_t generated by cheap MOO solver (i.e., $\mathbf{x}_t^{\text{UCB}} = \mathbf{x}^*$).

Now, if the bound from Lemma 8.2 holds, then for a point $\mathbf{x}_t^{\text{UCB}}$ proposed by UCB with parameters β_t , the following bound holds for any function f_i ,

$$f_i(\mathbf{x}^*) \leq UCB_{i,t}(\mathbf{x}^*) \leq UCB_{i,t}(\mathbf{x}_t^{\text{UCB}}) \quad (21)$$

Leading to the following bound

$$f_i(\mathbf{x}^*) - \mu_{i,t-1}(\mathbf{x}_t^{\text{UCB}}) \leq \sqrt{\beta_t} \sigma_{i,t-1}(\mathbf{x}_t^{\text{UCB}}). \quad (22)$$

We can now combine these results to construct the proof of Theorem 5.1.

Proof of Theorem 5.1. With probability at least $1 - \delta_1$, the result of Lemma 8.1 holds. If we assume that UCB acquisition function is included in the portfolio of acquisition functions, we have

$$-r_{T_{max}}^{\text{Hedge}} \leq \mathcal{O}(\sqrt{T_{max}}) - r_{T_{max}}^{\text{UCB}}$$

and by adding $\sum_{t=1}^{T_{max}} \sum_{i=1}^k f_i(\mathbf{x}^*)$ to both sides of the inequality, we have:

$$\begin{aligned} & \sum_{t=1}^{T_{max}} \sum_{i=1}^k f_i(\mathbf{x}^*) - \mu_{i,t-1}(\mathbf{x}_t) \\ & \leq \mathcal{O}(\sqrt{T_{max}}) + \sum_{t=1}^{T_{max}} \sum_{i=1}^k f_i(\mathbf{x}^*) - \mu_{i,t-1}(\mathbf{x}_t^{\text{UCB}}) \end{aligned}$$

With probability at least $1 - \delta_2$ the bound from Lemma 8.2 can be applied to the left-hand-side and the result of Lemma 8.4 can be applied to the right side of the inequality leading to the following inequality

$$\begin{aligned} & \sum_{t=1}^{T_{max}} \sum_{i=1}^k f_i(\mathbf{x}^*) - f_i(\mathbf{x}_t) - \sqrt{\beta_t} \sigma_{i,t-1}(\mathbf{x}_t) \\ & \leq \mathcal{O}(\sqrt{T_{max}}) + \sum_{t=1}^{T_{max}} \sum_{i=1}^k \sqrt{\beta_t} \sigma_{i,t-1}(\mathbf{x}_t^{\text{UCB}}) \end{aligned}$$

which means that the regret is bounded by

$$\begin{aligned} R_{T_{max}}(\mathbf{x}^*) &= \sum_{t=1}^{T_{max}} \sum_{i=1}^k f_i(\mathbf{x}^*) - f_i(\mathbf{x}_t) \\ &\leq \mathcal{O}(\sqrt{T_{max}}) + \sum_{t=1}^{T_{max}} \sum_{i=1}^k \sqrt{\beta_t} \sigma_{i,t-1}(i, \mathbf{x}_t^{\text{UCB}}) \\ &\quad + \sum_{t=1}^{T_{max}} \sum_{i=1}^k \sqrt{\beta_t} \sigma_{i,t-1}(\mathbf{x}_t) \\ &\leq \mathcal{O}(\sqrt{T_{max}}) + \sum_{t=1}^{T_{max}} \sum_{i=1}^k \sqrt{\beta_t} \sigma_{i,t-1}(\mathbf{x}_t^{\text{UCB}}) \\ &\quad + \sqrt{C_i T_{max} \beta_{T_{max}} \gamma_{T_{max}}^i}. \end{aligned}$$

We should note that we cannot use Lemma 8.3 to further simplify the terms involving $\mathbf{x}_t^{\text{UCB}}$. This is because the lemma only holds for points that are sampled by the algorithm, which may not include those proposed by UCB acquisition function.

Finally, this derivation requires Lemmas 8.1 and 8.4 to hold. Using a union bound argument, we can see that both lemmas hold with probability at least $1 - \delta_1 - \delta_2$. By setting $\delta_1 = \delta_2 = \delta/2$, we recover our result [Hoffman et al. 2011]. \square

8.2 PDBO Hypervolume Experiments

In this section, we present supplementary experiments that focus on the comparison of hypervolume between our proposed method, PDBO, and other existing baselines. We refer the interested reader to Figure 11 for the additional results.

8.3 PDBO Diversity Experiments

In terms of Pareto front diversity, we conduct a comprehensive comparison between PDBO and state-of-the-art methods introduced in Section 6. It is worth noting that, apart from DGEMO, none of the other baseline methods explicitly address Pareto front diversity. The results depicted in Figure 8 demonstrate the superior performance of PDBO, outperforming all existing methods in terms of Pareto front diversity measure.

8.4 Visualization of Diversity of Pareto front

In Figure 4, we present illustrative examples of Pareto front figures achieved by each baseline. In each experiment, we deliberately vary the batch size to showcase the consistent diversity of results, irrespective of the batch size chosen.

It’s worth noting that several of benchmarks featured in this paper involve more than two objective functions, rendering it impractical to effectively visualize the Pareto front points in a coherent and legible manner.

8.5 Acquisition Function Selection Statistics

Table 1 contains statistics regarding the selection of each acquisition function by the MAB algorithm. It is important to clarify that the goal of the acquisition function selection multi-armed bandit algorithm is not solely to converge to a single acquisition function. Instead, its purpose is to dynamically choose the most appropriate acquisition function at each iteration to maximize overall performance improvement. Therefore, we believe that these statistics are interesting but do not necessarily lead to the conclusion that the most selected acquisition function is necessarily the best for the experiment.

8.6 Additional Real-World Experiments

This section offers insights into two supplementary real-world experiments carried out to further demonstrate the effectiveness of PDBO when compared to existing baselines in terms of the hypervolume metric and diversity of the Pareto front.

Unmanned Aerial Vehicle (UAV) power system design:

This real-world benchmark [Belakaria et al. 2020b] has two objective functions: minimizing the energy consumption and minimizing the UAV’s total mass. The search space has five input variables defined as the battery cells in series (ranging between 10 and 18), the battery cells in parallel (ranging between 16 and 70), the motor quantity (ranging between 6 and 10), the height of the stator structure (ranging between 80 and 260), and the motor stator winding turns (ranging between 100 and 550). We use a high-fidelity simulator to evaluate the two objective functions for any candidate input configuration.

SW-LLVM compiler settings optimization: SW-LLVM

[Siegmond et al. 2012] is a compiler settings problem determined by d=10 compiler configurations. The goal of this experiment is to find a setting of the LLVM compiler that optimizes the memory footprint and performance on a given set of software programs.

In Figure 15 and Figure 16, we provide hypervolume and diverse Pareto front (DPF) results for these two additional real-world problems.

8.7 Ablation Studies

Merits of DPP-Based Batch Selection compared to Greedy HVC selection.

As detailed in Section 4.2 of our paper, we have established that Hypervolume alone may not effectively capture the diversity of individual data points. For a more nuanced assessment of each new point in relation to previously evaluated designs, we turn to Individual Hypervolume Contributions (HVCs). However, selecting points solely based on the highest HVC values presents several potential issues: 1. It may lead to exploitative behavior, as it relies solely on predictive mean information; 2. It may undermine diversity considerations, especially when multiple points exhibit closely ranked HVC values; 3. It may overlook diversity in the input domain, as DPP’s utilization of a kernel allows for a

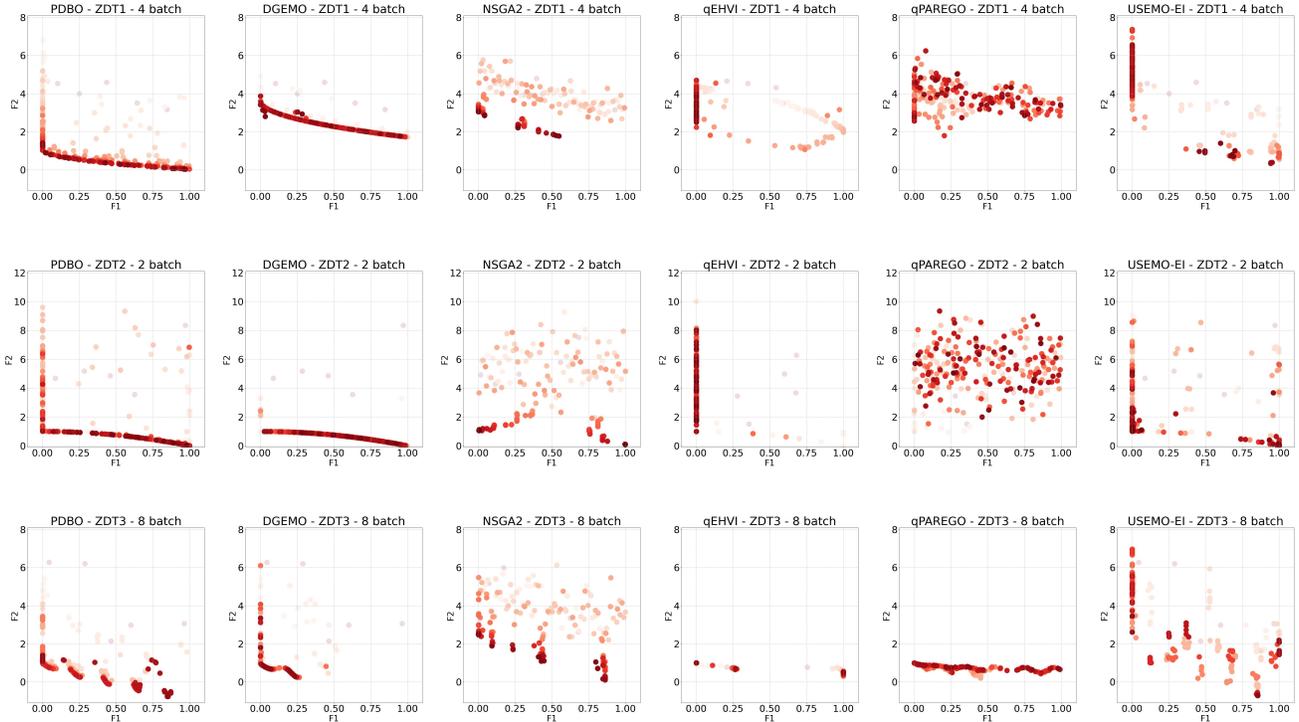


Figure 4: The Pareto front of the inputs selected by each algorithm. Points from later batches are depicted in darker colors.

Problem Name	2-Batch				4-Batch				8-Batch				16-Batch			
	TS	EI	UCB	ID	TS	EI	UCB	ID	TS	EI	UCB	ID	TS	EI	UCB	ID
ZDT1	12±12	9±7	1±1	75±15	13±18	14±12	4±3	67±17	16±28	10±12	5±5	67±30	10±21	5±7	11±10	72±21
ZDT2	23±17	22±11	19±10	33±18	17±18	25±18	16±17	40±30	9±19	31±27	21±24	39±27	7±19	40±38	24±29	28±34
ZDT3	1±5	8±6	4±5	85±10	1±2	12±11	2±3	83±12	1±2	16±14	7±8	74±17	1±2	28±33	7±3	62±35
DTLZ1	13±10	13±9	7±5	65±16	11±12	13±11	7±6	67±17	10±13	12±17	10±10	66±27	14±24	20±26	16±14	48±35
DTLZ3	7±6	9±11	3±4	79±15	16±19	9±10	5±6	69±23	13±20	17±24	9±10	60±32	26±36	15±32	21±24	37±38
DTLZ5	4±3	85±8	3±3	5±4	5±6	82±8	3±2	8±6	6±8	68±20	10±10	14±19	7±17	61±29	14±9	17±22
Gear Train Design	1±4	2±3	1±1	94±5	1±4	2±3	2±2	93±6	2±7	5±6	3±1	88±8	2±6	10±15	8±8	79±18

Table 1: The percentage of each acquisition function used for each problem and batch size, averaged among 25 separate runs

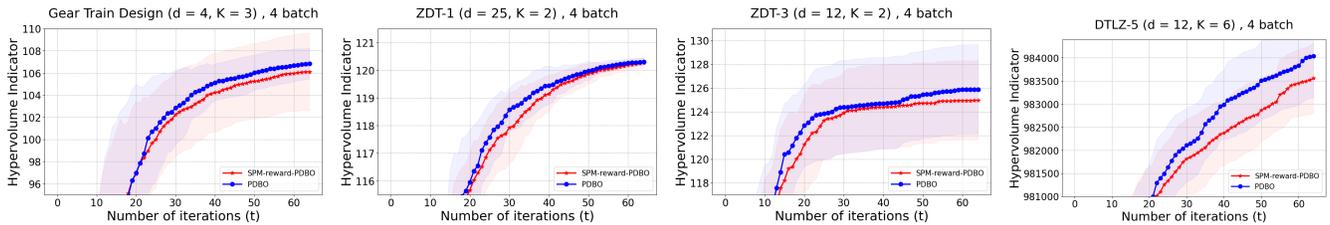


Figure 5: Comparing PDBO against PDBO with the sum of posterior means (SPM) as the reward.

more comprehensive representation of diversity across the input space. Figure 17 presents the results of our ablation study on both a real-world problem and a synthetic problem offering a comparative evaluation of batch selection methods: one based on the highest hypervolume contribution (HVC) and the other based on our proposed multi-objective DPP-based

batch selection approach.

8.8 Additional Diversity Metrics

We present examples of our results evaluated using the Inverted Generational Distance (IGD) metric [Coello Coello and Reyes Sierra 2004]. However, we contend that IGD may

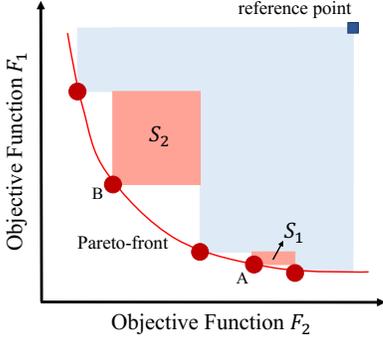


Figure 6: Sections S_1 and S_2 exhibit the specific hypervolume contributions of points A and B, respectively. This figure indicates a positive correlation between higher individual hypervolume contributions and increased diversity among the points.

not be a suitable metric for assessing diversity. This is because IGD is heavily influenced by the number of points discovered, which can lead to misleading results. As an example, if one baseline uncovers only two points very close to the ideal Pareto front, it may achieve a heavily more favorable IGD value than a method finding 50 diverse points slightly further away. Clearly, the set of two points is not inherently more diverse. Furthermore, the IGD metric relies on a predefined ground truth Pareto front, which is often unavailable in black-box settings, rendering its use less meaningful. Moreover, the choice of the ideal Pareto front shape can impact whether IGD primarily reflects diversity or performance. Recent work [Pierrot et al. 2022] proposed an approach to evaluate diversity for cheap multi-objective problems. However, the proposed approach requires the user to manually define space descriptors. Depending on the descriptors’ definition, the results provided by the metric will change leading to potentially inconsistent results, especially in the black-box expensive setting where it is harder to define the descriptors. To our understanding, there is no clear/known strategy to define the descriptors.

As additional metrics, we employ the Pareto Hypervolume (PHV) to evaluate the quality of the uncovered optimal Pareto front and use the Diversity of Pareto Front (DPF) metric to assess the algorithm’s ability to select diverse designs. We believe that incorporating DPF alongside PHV provides a more comprehensive evaluation of the algorithm since these metrics offer complementary insights into different aspects of performance.

Here, we include two new metrics. First, we calculate the DPF value exclusively on the optimal discovered Pareto frontier as shown in Figure 9. Additionally, we include experiments where the IGD metric is computed as shown in Figure 10.

8.9 Experimental Details

This section includes supplementary implementation elements and experimental details.

Table 2 provides information regarding the dimensions,

objectives, and reference points utilized for each benchmark problem in the primary hypervolume and DPF results. Additionally, Table 3 presents the hyperparameters of the GP surrogate model. In all experiments, a consistent definition of the GP surrogate model is employed, featuring a zero mean function and an anisotropic Matern 5/2 kernel.

Problem Name	d	K	reference point
ZDT-1	25	2	[11.0, 11.0]
ZDT-2	4	2	[11.0, 11.0]
ZDT-3	12	2	[11.0, 11.0]
SW-LLVM	10	2	[1000.0, 500.0]
UAV	6	2	[1e9, 1000.0]
Gear Train Design	4	3	[6.6764, 59.0, 0.4633]
DTLZ-1	10	4	[400.0, ..., 400.0]
DTLZ-3	9	4	[10000.0, ..., 10000.0]
DTLZ-5	12	6	[10.0, ..., 10.0]

Table 2: The problem details for hypervolume and diversity experiments

Hyperparameter name	Hyperparameter value
initial l	$(1, \dots, 1) \in \mathcal{R}^d$
l range	$(\sqrt{10^{-3}}, \sqrt{10^3})$
initial σ_f	1
σ_f range	$(\sqrt{10^{-3}}, \sqrt{10^3})$
initial σ_n	10^{-2}
σ_n range	10^{-2}

Table 3: Hyperparameters used in the GP

The NSGA-II multi-objective optimization (MOO) solver adopts simulated binary crossover [Deb et al. 1995] with a parameter of $\mu_c = 15$, and polynomial mutation [Deb et al. 1996] with a parameter of $\mu_m = 20$. It utilizes a population size of 100 and runs for a total of 200 generations to explore the Pareto front of acquisition function values. The initial population is derived from the best current samples identified using non-dominated sorting [Deb et al. 2002b]. As for the baseline NSGA-II algorithm, it employs the same crossover and mutation parameters (μ_c and μ_m), but the population size is adjusted to match the batch size, and the number of generations aligns with the number of algorithm iterations. These hyperparameters are as mentioned in [Konakovic Lukovic et al. 2020]

Table 4 presents a comprehensive comparison of the runtime performance for all the methods discussed. The qEHVI and qPAREGO techniques have the capability to leverage GPUs for their execution. However, considering the limited

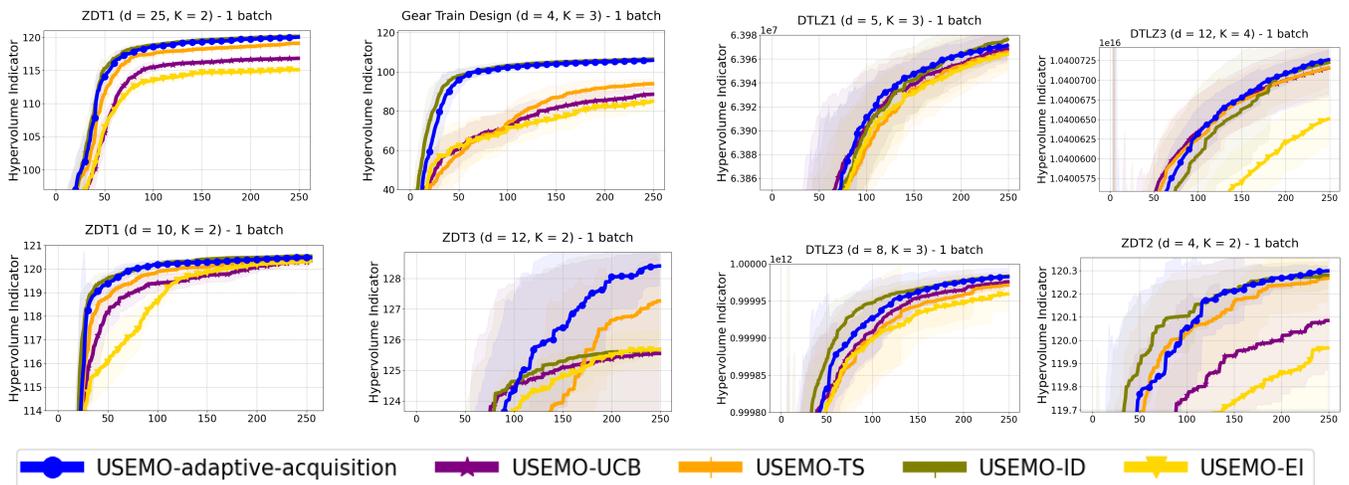


Figure 7: Acquisition functions ablation study results

accessibility of GPUs for many BO problems, we also provide an analysis of these methods when executed on CPUs. It should be noted that the runtime of qEHVI-CPU and qPAREGO-CPU is notably longer when compared to their counterparts. Therefore, for these two methods, we report the average runtime based on 4 runs, whereas for all other experiments, we report the runtime based on 25 runs. The experiments utilizing GPUs were conducted using four NVIDIA Quadro RTX 6000 GPUs, while the CPU experiments were performed on an AMD EPYC 7451 24-Core Processor.

Practical Issues and Limitations of Existing Baselines

The two state-of-the-art methods are DGEMO and qEHVI. The both have practical issues and limitations leading to their inapplicability in some experiments. Below we provide a detailed explanation of these limitations:

- **DGEMO:** One algorithmic step of the DGEMO algorithm is to divide the points in the Pareto front into different regions (line 9, Algorithm 1 in [Konakovic Lukovic et al. 2020]). Depending on the number of objective functions, the algorithm used to divide the Pareto front into regions would be different. In the implementation of DGEMO, only the cases of 2 and 3 functions were implemented, any case beyond 3 objectives is not handled. So this limitation of the DGEMO approach prevents it from being applicable to any problem. The inability to handle more than 3 objectives was not explicitly discussed in the paper. Therefore, no solution was suggested for it and no implementation was provided for this case. Finding the appropriate algorithm for the splitting approach when the number of objectives exceeds 3 and implementing it is not straightforward, and a research investigation in itself.
- **qEHVI:** qEHVI algorithm has the capability to use auto differentiation and batch selection parallelization. However, this leads to high memory consumption when the batch sizes increase or when the input dimension or the number of objective functions increases, leading to an out-of-memory issue even on (GPU) machines with 40 GB of memory. Unfortunately, we do not have access to higher

memory machines. We believe that this is a limitation of the approach since most scientists and practitioners using BO algorithms might not necessarily have access to very high-memory GPU machines.

8.10 Limitations

Our approach overcomes the adaptive acquisition function selection problem via a MAB approach. However, our algorithm requires the proper selection of the MAB algorithm hyperparameters, namely, the decay factor and the probability of selection hyperparameter. We set these hyperparameters manually based on a previous study of their suitable value [Hoffman et al. 2011]. We leave the study of the adaptive selection of hyperparameters to future work. Another limitation of our work is its inability to handle high-dimensional search space. However, it is important to note that this work is not proposed in the context of high-dimensional BO, and it can be synergistically combined with effective high-dimensional BO approaches such as [Eriksson et al. 2019] to handle high-dimensional search spaces.

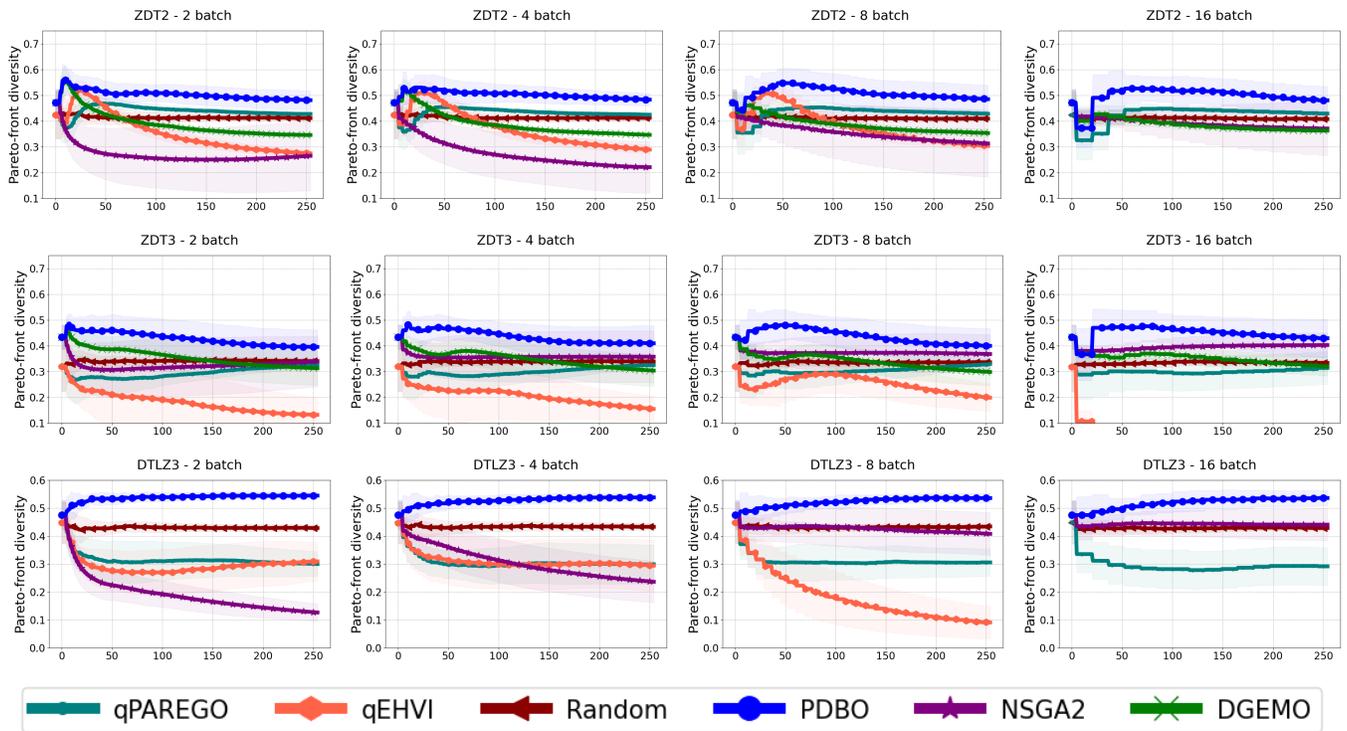


Figure 8: DPF results for all benchmarks

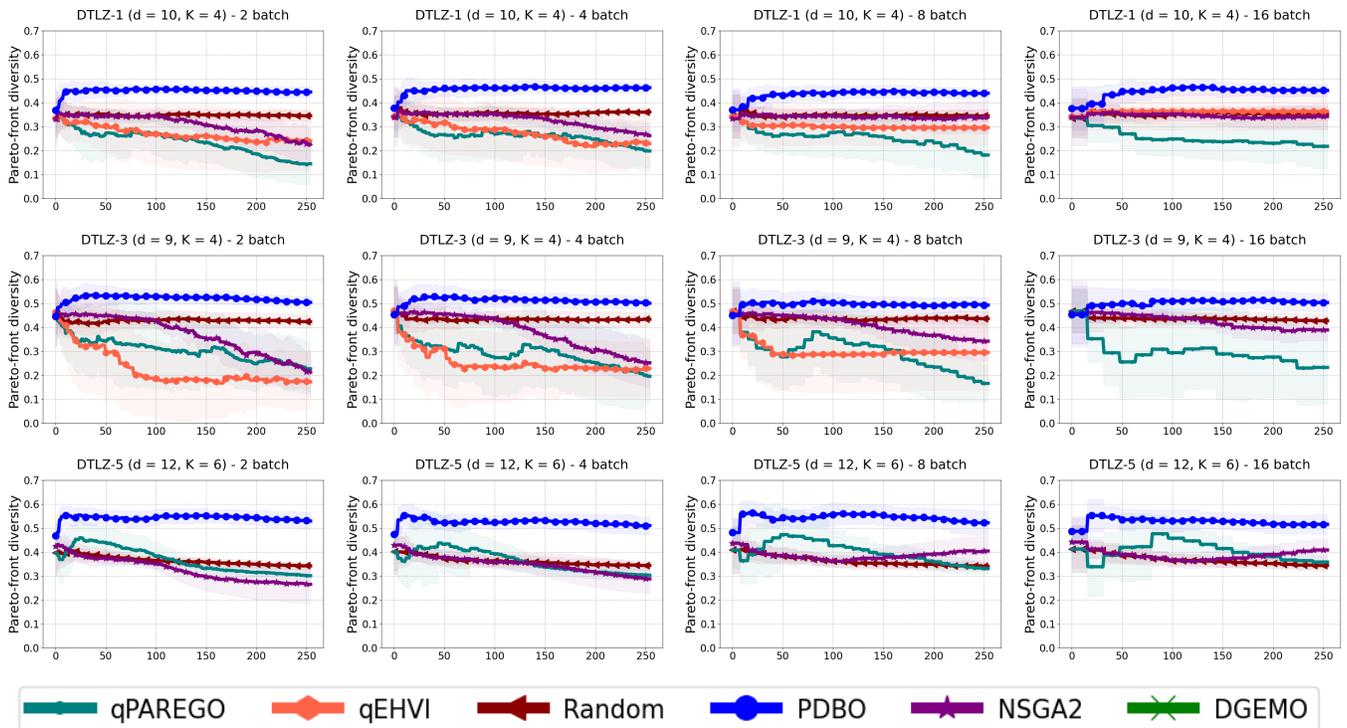


Figure 9: DPF results calculated using the optimal Pareto front per algorithm iteration

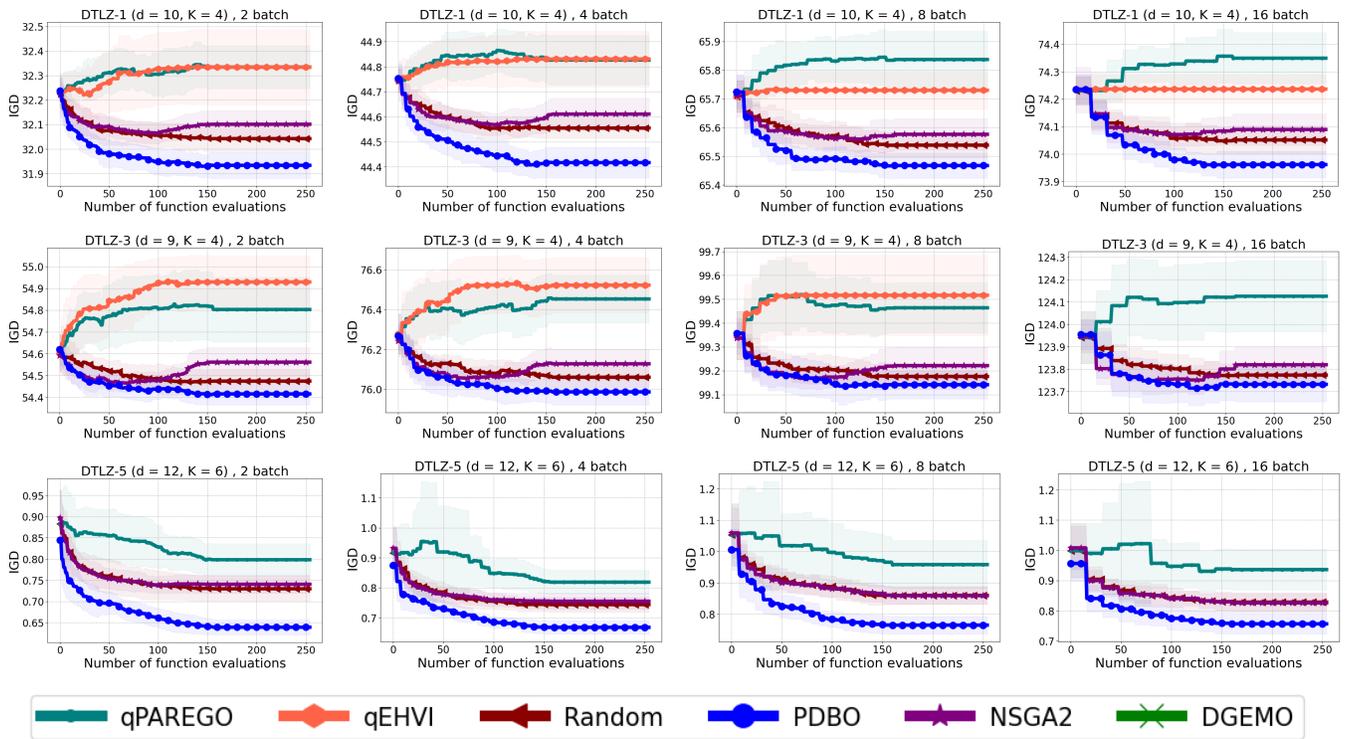


Figure 10: IGD results on several benchmarks; A lower IGD value indicates a solution of higher quality.

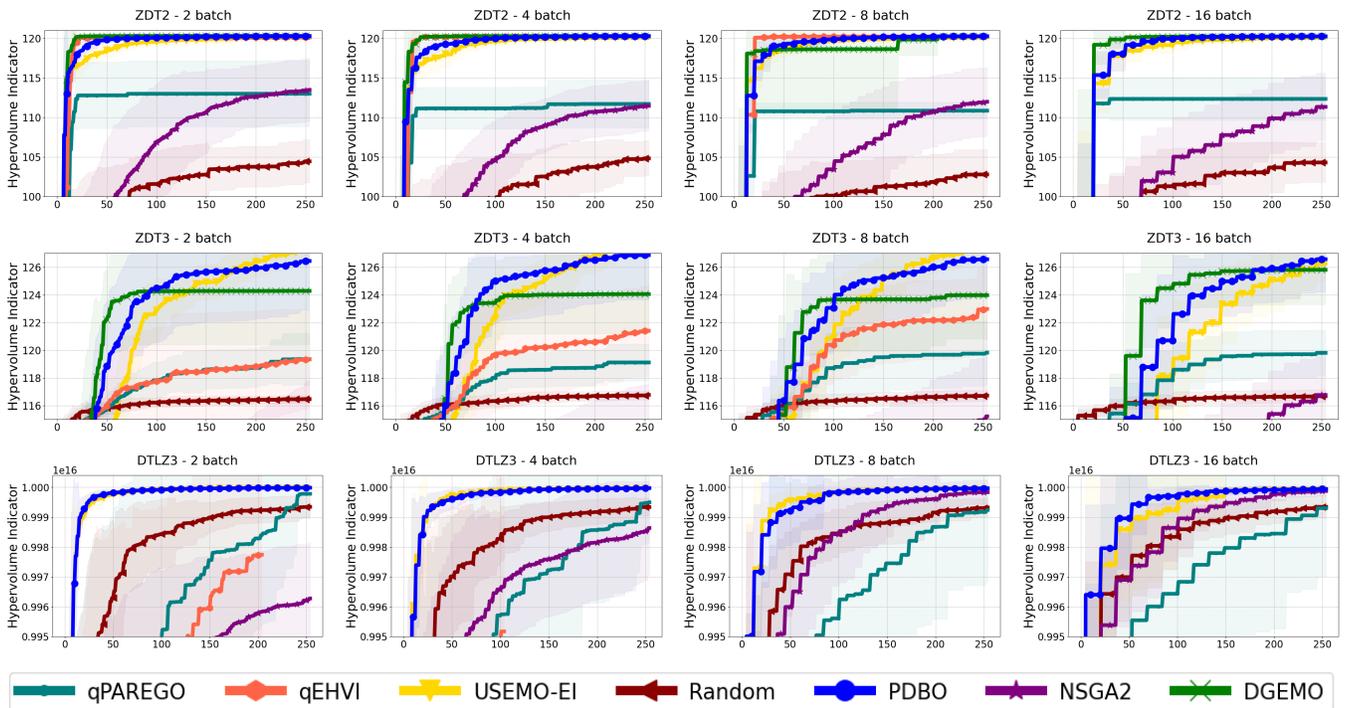


Figure 11: Average hypervolume results on additional benchmarks

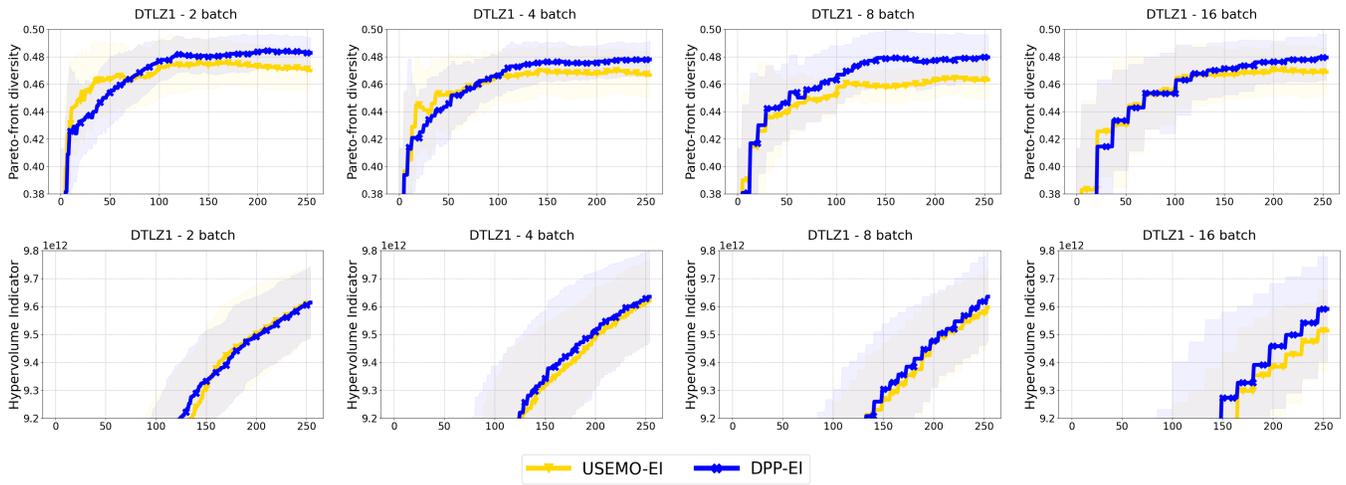


Figure 12: DPF and hypervolume results for DTLZ1 problem, $d = 13$, $K = 5$

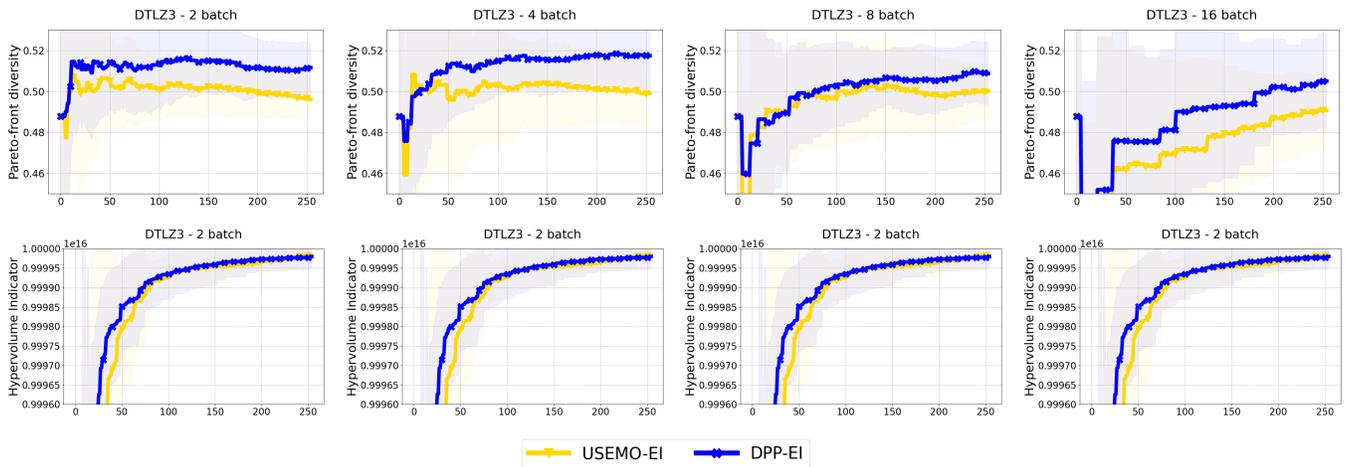


Figure 13: DPF and hypervolume results - DTLZ3 problem, $d = 8$, $K = 4$

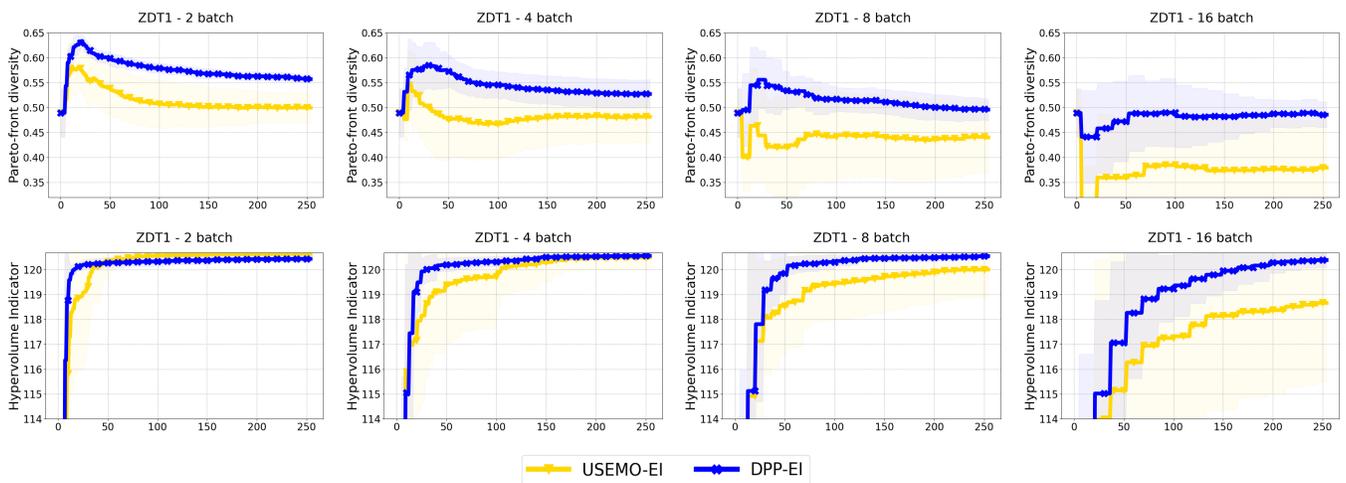


Figure 14: DPF and hypervolume results - ZDT1 problem, $d = 3$, $K = 2$

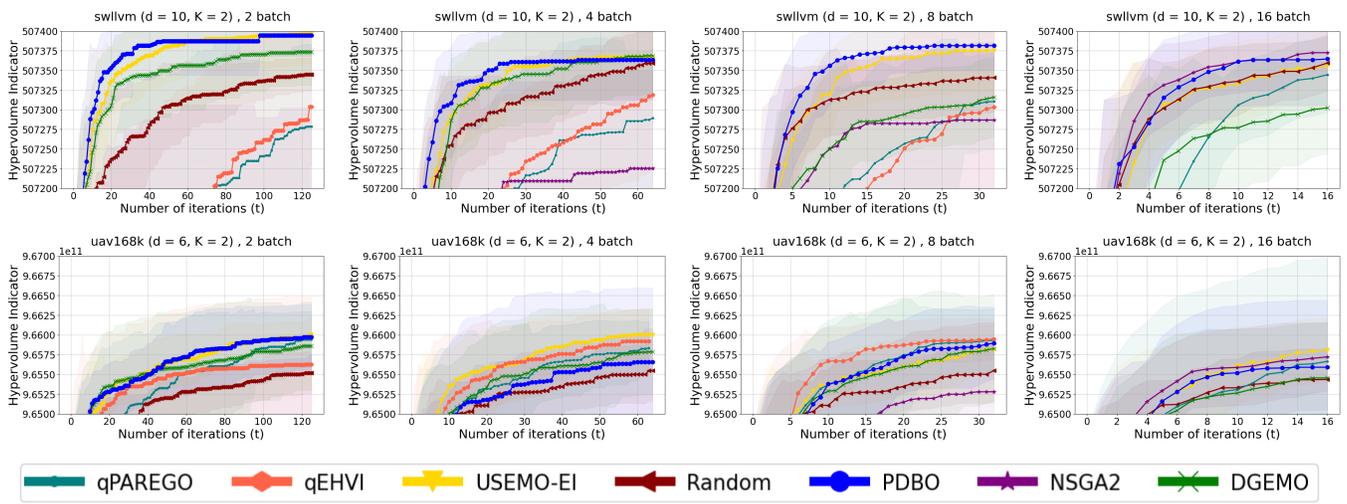


Figure 15: Additional real-world experiments hypervolume results

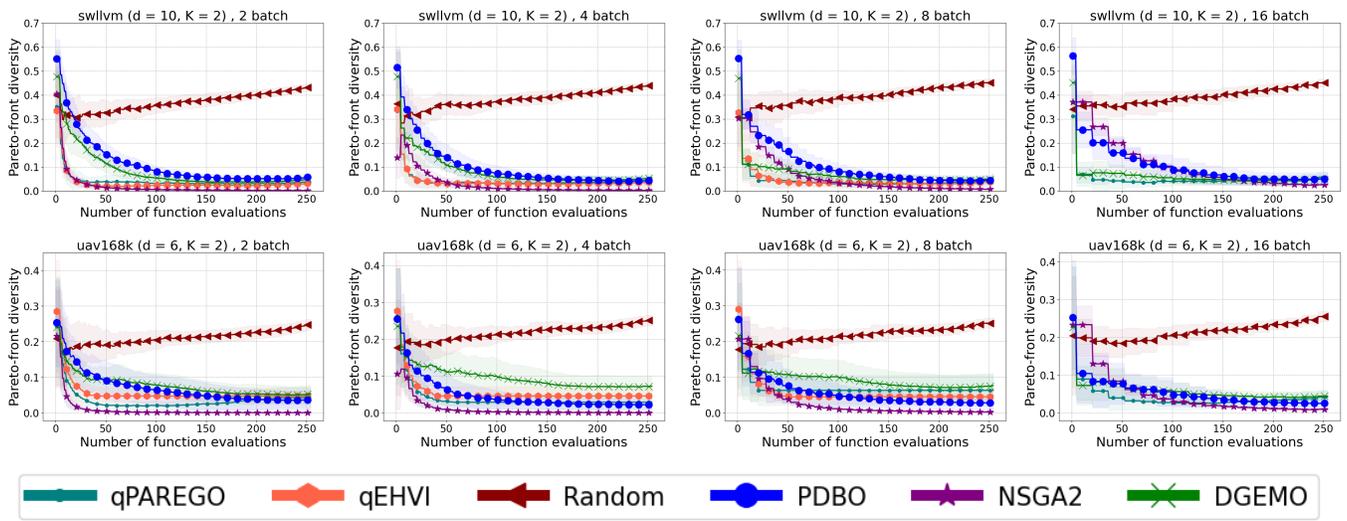


Figure 16: Additional real-world experiments DPF results

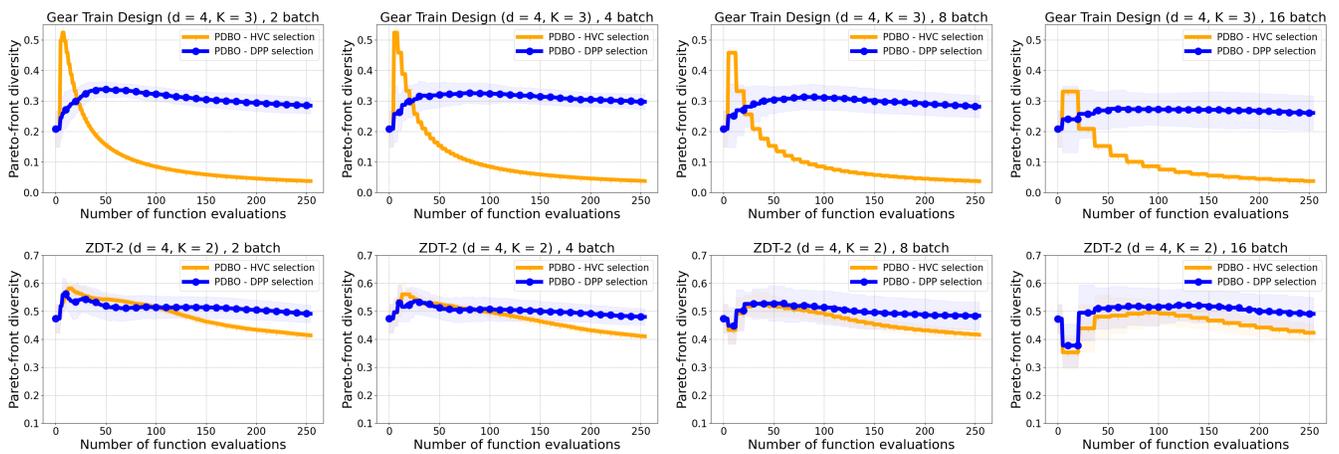


Figure 17: Comparing DPF batch selection to greedy selection based on the highest individual hypervolume contribution

Problem Name	Batch Size	PDBO	DGEMO	qEHVI		NSGA-II	qPAREGO		USEMO-EI
				CPU	GPU		CPU	GPU	
ZDT-1	2	52.66	58.22	192.64	7.31	0.01	66.57	3.64	1.34
	4	57.73	57.93	510.11	16.17	0.01	113.00	5.94	1.47
	8	62.57	60.24	1781.94	53.67	0.01	226.20	11.93	1.55
	16	70.40	65.28	NA	NA	0.01	667.65	44.92	1.50
ZDT-2	2	52.37	51.47	48.22	2.95	0.01	16.58	1.37	0.89
	4	56.70	51.96	127.10	5.25	0.01	26.40	2.04	0.97
	8	60.25	52.14	626.71	15.67	0.01	56.98	3.43	0.95
	16	64.91	51.94	NA	NA	0.01	163.08	9.05	0.91
ZDT-3	2	49.52	61.81	201.73	8.95	0.01	69.55	4.23	1.18
	4	54.46	61.42	530.65	22.54	0.01	114.35	9.48	1.33
	8	57.74	60.34	1871.59	68.70	0.01	247.07	18.39	1.33
	16	65.46	60.20	NA	NA	0.01	745.80	43.90	1.24
Gear Train Design	2	30.42	77.65	60.01	1.54	0.01	17.82	0.92	1.09
	4	34.51	77.15	316.89	5.62	0.01	43.32	2.19	1.08
	8	40.50	76.83	1534.78	37.66	0.01	200.46	7.05	1.08
	16	51.12	74.88	NA	NA	0.01	780.58	17.70	0.96
DTLZ-1	2	55.50	NA	267.98	25.99	0.02	112.70	7.43	2.28
	4	60.86	NA	NA	NA	0.02	240.43	13.41	2.25
	8	64.68	NA	NA	NA	0.03	545.08	21.82	2.23
	16	75.46	NA	NA	NA	0.05	1492.85	45.54	2.16
DTLZ-3	2	54.29	NA	NA	NA	0.02	111.90	9.20	2.07
	4	58.44	NA	NA	NA	0.02	265.97	16.58	2.10
	8	63.90	NA	NA	NA	0.03	603.90	25.97	2.05
	16	74.11	NA	NA	NA	0.04	1676.98	50.30	2.02
DTLZ-5	2	242.78	NA	NA	NA	1.15	146.71	6.89	4.86
	4	244.08	NA	NA	NA	4.64	201.05	9.08	4.81
	8	268.75	NA	NA	NA	1.59	368.43	12.69	4.41
	16	303.58	NA	NA	NA	2.06	935.33	29.46	4.20

Table 4: Average runtime (seconds per iteration) of each baseline