

Improving Pareto Set Learning for Expensive Multi-objective Optimization via Stein Variational Hypernetworks

Minh-Duc Nguyen¹, Phuong Mai Dinh^{1, 2}, Quang-Huy Nguyen^{1, 3},
Long P. Hoang⁴, Dung D. Le¹

¹ College of Engineering and Computer Science, VinUniversity, Vietnam

² VinUni-Illinois Smart Health Center, VinUniversity, Vietnam

³ Department of Computer Science and Engineering, The Ohio State University, USA

⁴ Information Systems Technology and Design, Singapore University of Technology and Design, Singapore
duc.nm2@vinuni.edu.vn, phuong.dm2@vinuni.edu.vn, nguyen.2959@osu.edu,
long_hoang@mymail.sutd.edu.sg, dung.ld@vinuni.edu.vn

Abstract

Expensive multi-objective optimization problems (EMOPs) are common in real-world scenarios where evaluating objective functions is costly and involves extensive computations or physical experiments. Current Pareto set learning methods for such problems often rely on surrogate models like Gaussian processes to approximate the objective functions. These surrogate models can become fragmented, resulting in numerous small uncertain regions between explored solutions. When using acquisition functions such as the Lower Confidence Bound (LCB), these uncertain regions can turn into pseudo-local optima, complicating the search for globally optimal solutions. To address these challenges, we propose a novel approach called SVH-PSL, which integrates Stein Variational Gradient Descent (SVGD) with Hypernetworks for efficient Pareto set learning. Our method addresses the issues of fragmented surrogate models and pseudo-local optima by collectively moving particles in a manner that smooths out the solution space. The particles interact with each other through a kernel function, which helps maintain diversity and encourages the exploration of underexplored regions. This kernel-based interaction prevents particles from clustering around pseudo-local optima and promotes convergence towards globally optimal solutions. Our approach aims to establish robust relationships between trade-off reference vectors and their corresponding true Pareto solutions, overcoming the limitations of existing methods. Through extensive experiments across both synthetic and real-world MOO benchmarks, we demonstrate that SVH-PSL significantly improves the quality of the learned Pareto set, offering a promising solution for expensive multi-objective optimization problems.

Introduction

Multi-objective optimization (MOO) has numerous essential and practical applications across various fields, from text-to-image generation (Lee et al. 2024) to ejector design for fuel cell systems (Hou, Chen, and Pei 2024). However, real-world MOO problems often involve multiple conflicting and expensive-to-evaluate objectives. For example, recommendation systems must achieve a balance between precision and efficiency (Le and Lauw 2017) or between precision and

revenue (Milojkovic et al. 2019), battery usage optimization requires trade-offs between performance and lifetime (Attia et al. 2020), robotic radiosurgery involves coordination of internal and external motion (Schweikard et al. 2000). Traditional MOO methods, such as hyperparameter tuning, seek a finite set of Pareto-optimal solutions representing trade-offs between objectives. However, as the number of objectives increases, approximating the Pareto front becomes exponentially costly, leading to a trade-off between generalization and computational time (Swersky, Snoek, and Adams 2013).

Pareto Set Learning (PSL) (Navon et al. 2021; Hoang et al. 2023; Tuan et al. 2024) is a promising approach that allows users to explore the entire Pareto front for MOO problems by learning a parametric mapping to align trade-off preference weights assigned to objectives with their corresponding Pareto optimal solutions. The optimized mapping model enables real-time adjustments between objectives.

Expensive objective problems refer to a class of real-world problems in which evaluating each objective is expensive. For example, evaluating a battery's lifetime costs a newly developed battery or evaluating robotic radiosurgery costs a lot of one-time-use medical equipment. To address this challenge, researchers have developed surrogate model techniques (He et al. 2023) that estimate the actual objective functions, thus minimizing the necessity for costly evaluations. Lin et al. (Lin et al. 2022) is one of the pioneer teams in the study of Pareto set learning for the EMOPs by acquiring knowledge of the entire Pareto front, referred to as PSL-MOBO. This work employs the Pareto set learning with Multi-objective Bayesian Optimization (MOBO) (Laumanns and Ocenasek 2002) to effectively solve black-box expensive optimization problems by minimizing the number of function evaluations.

Challenge. We observed that PSL-MOBO often struggles with instability due to its tendency to become trapped in pseudo-local optima. This issue arises because the optimization process can mistakenly identify these false optima as true optimal solutions, causing the algorithm to converge prematurely and fail to explore the solution space thoroughly. As a result, PSL-MOBO might not effectively approximate the true Pareto front, especially in complex or high-dimensional spaces where numerous local optima are present. Consequently, learning the Pareto front requires an advanced algo-




Figure 1: Approximate Pareto front comparison after the first 6 function evaluations using PSL-MOBO and our method, SVH-PSL, on the ZDT1 problem (2 objectives, 20 dimensions).

rithm to avoid premature convergence while simultaneously utilizing Hypernetwork and sampling multiple samples in parallel during the learning process.

Approach. In this paper, we propose a novel Pareto set learning model that simultaneously samples multiple solutions during the learning process. By leveraging the mutual interactions among these samples, our approach helps them escape local optima. These interactions also promote mutual repulsion, encouraging the samples to explore a diverse solution space. This results in a more diverse Pareto set, enhancing exploration while ensuring convergence to the optimal front. As shown in Figure 1, the Pareto front approximated by PSL-MOBO becomes stuck during the initial function evaluation steps, whereas our method, SVH-PSL, effectively overcomes this issue and successfully captures the entire Pareto front in the ZDT1 problem.

Building upon the previously established framework, we propose a novel methodology for multi-objective Bayesian optimization (MOBO) that employs Stein Variational Gradient Descent (SVGD) to enhance the learning of the Pareto set. This approach tackles challenges in expensive multi-objective optimization problems (EMOPs), where the evaluation of objective functions incurs high computational costs. Our contributions are as follows:

- First, we present the mathematical formulation for controllable Pareto set learning in the context of expensive multi-objective optimization tasks.
- Second, we present SVH-PSL, an innovative framework for Pareto set learning in EMOPs, integrating Stein Variational Gradient Descent with a Hypernetwork and introducing a novel kernel for greater efficiency.
- Third, we perform extensive experiments on synthetic and real-world multi-objective optimization problems to assess the effectiveness of our proposed SVH-PSL method compared to baseline approaches.

Preliminary

Expensive Multi-Objective Optimization

We consider the following expensive continuous multi-objective optimization problem:

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} \mathbf{f}(\mathbf{x}), \quad (1)$$

$$\mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_m(\mathbf{x}))$$

where $\mathcal{X} \subset \mathbb{R}^n$ is the decision space, $f_i : \mathcal{X} \rightarrow \mathbb{R}$ is a black-box objective function. For a nontrivial problem, no single solution can optimize all objectives simultaneously, and there will always be a trade-off among them. We have the following definitions for multi-objective optimization:

Definition 1 (Dominance) A solution \mathbf{x}^a is said to dominate another solution \mathbf{x}^b if and only if $f_i(\mathbf{x}^a) \leq f_i(\mathbf{x}^b)$, $\forall i \in \{1, \dots, m\}$ and $\mathbf{f}(\mathbf{x}^a) \neq \mathbf{f}(\mathbf{x}^b)$. We denote this relationship as $\mathbf{x}^a \prec \mathbf{x}^b$.

Definition 2 (Pareto Optimality) A solution \mathbf{x}^* is called Pareto optimal solution if $\nexists \mathbf{x}^b \in \mathcal{X} : \mathbf{x}^b \prec \mathbf{x}^*$.

Definition 3 (Pareto Set/Front) The set of Pareto optimal is Pareto set, denoted by $\mathcal{P} = \{\mathbf{x}^*\} \subseteq \mathcal{X}$ and the corresponding images in objectives space are Pareto front $\mathcal{P}_f = \{\mathbf{f}(\mathbf{x}) \mid \mathbf{x} \in \mathcal{P}\}$.

Definition 4 (Hypervolume) Hypervolume (Zitzler and Thiele 1999) is the area dominated by the Pareto front. Therefore, the quality of a Pareto front is proportional to its hypervolume. Given a set of n points $\mathbf{y} = \{y^{(i)} \mid y^{(i)} \in \mathbb{R}^m; i = 1, \dots, n\}$ and a reference point $\rho \in \mathbb{R}^m$, the Hypervolume of \mathbf{y} is measured by the region of non-dominated points bounded above by $y^{(i)} \in \mathbf{y}$, then the hypervolume metric is defined as follows:

$$HV(\mathbf{y}) = VOL \left(\bigcup_{y^{(i)} \in \mathbf{y}, y^{(i)} \prec \rho} \Pi_{i=1}^n [y^{(i)}, \rho_i] \right) \quad (2)$$

where ρ_i is the i^{th} coordinate of the reference point ρ and $\Pi_{i=1}^n [y^{(i)}, \rho_i]$ is the operator creating the n -dimensional hypercube from the ranges $[y^{(i)}, \rho_i]$.

Gaussian Process and Bayesian Optimization

A Gaussian Process with a single objective is characterized by a prior distribution defined over the function space as:

$$f(\mathbf{x}) \sim GP(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x})), \quad (3)$$

where $\mu : \mathcal{X} \rightarrow \mathbb{R}$ represents the mean function, and $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^2$ is the covariance kernel function. Given n evaluated solutions $\mathbf{D} = \{\mathbf{X}, \mathbf{y}\} = \{(\mathbf{x}^{(i)}, f(\mathbf{x}^{(i)})) \mid i = 1, \dots, n\}$, the posterior distribution can be updated by maximizing the marginal likelihood based on the available data. For a new solution \mathbf{x}^{n+1} , the posterior mean and variance are given by:

$$\begin{aligned} \hat{\mu}(\mathbf{x}^{(n+1)}) &= \mathbf{k}^T \mathbf{K}^{-1} \mathbf{y}, \\ \hat{\sigma}^2(\mathbf{x}^{(n+1)}) &= k(\mathbf{x}^{(n+1)}, \mathbf{x}^{(n+1)}) - \mathbf{k}^T \mathbf{K}^{-1} \mathbf{k}, \end{aligned}$$

where $\mathbf{k} = k(\mathbf{x}^{(n+1)}, \mathbf{X})$ is the kernel vector and $\mathbf{K} = k(\mathbf{X}, \mathbf{X})$ is the kernel matrix.

Bayesian optimization involves searching for the global optimum of a black-box function $f(\cdot)$ by wisely choosing the next evaluation via the current Gaussian process and acquisition functions. A new evaluation \mathbf{x}^{n+1} is determined through an acquisition function α , which guides the search for the optimal solution. More specifically, the next evaluation \mathbf{x}^{n+1} is selected as the optimal solution of the acquisition in order to make the best improvement:

$$\mathbf{x}^{n+1} = \arg \max_{\mathbf{x}} \alpha(\mathbf{x}; \mathbf{D}), \quad (4)$$

When a Gaussian process is used to approximate the unknown objective function, the acquisition function helps balance the trade-off between exploration (sampling areas with high uncertainty) and exploitation (sampling areas with promising results based on the model's predictions). There are many options available in the field of surrogate models, such as Expected Improvement (EI), Upper Confidence Bound (UCB), and Lower Confidence Bound (LCB). In the scope of this study, we choose to use LCB as our acquisition function, which is:

$$\hat{f}(\mathbf{x}) = \hat{\mu}(\mathbf{x}) - \lambda \hat{\sigma}(\mathbf{x}) \quad (5)$$

Additionally, in the setting of multi-objective optimization, hypervolume is a common criterion for determining the quality of the Pareto front. For multi-objective Bayesian optimization, we use the concept of Hypervolume Improvement (HVI) is a valuable metric in multi-objective optimization that quantifies the increase in hypervolume achieved by adding the set of new solutions $\{\mathbf{X}_+ = \{\mathbf{x}^{(i)}\}_{i=1}^b, \mathbf{Y}_+ = \hat{f}(\mathbf{X}_+)\}$ to a current set of solutions $\{\mathbf{X}, \mathbf{Y}\}$:

$$HVI(\mathbf{Y}_+, \mathbf{Y}) = HV(\mathbf{Y}_+ \cup \mathbf{Y}) - HV(\mathbf{Y}) \quad (6)$$

Here b represents the number of solutions. To select the optimal \mathbf{X}_+ for the next evaluation, we choose those that maximize the HVI value. In other words, we use HVI as the additional acquisition function for selecting new solutions:

$$\mathbf{X}_+ = \arg \max_{\mathbf{X}_+ \in \mathcal{X}} HVI(\mathbf{Y}_+, \mathbf{Y}) \quad (7)$$

We select the set \mathbf{X}_+ in a sequential greedy manner from \mathbf{X} where $|\mathbf{X}| = 1000$ is approximated by the Pareto set model, which we present in the section below.

Pareto Set Learning

Pareto Set Learning approximates the entire Pareto Front of Problem (1) by directly approximating the mapping between an arbitrary preference vector r and a corresponding Pareto optimal solution computed by surrogate models:

$$\theta^* = \arg \min_{\theta} \mathbb{E}_{r \sim \text{Dir}(\alpha)} g(\hat{f}(\mathbf{x}_r)|r) \quad (8)$$

$$\text{s.t } \mathbf{x}_r = h(r|\theta) \in \mathcal{P}, h(\mathbb{S}^m|\theta^*) = \mathcal{P}$$

where $\text{Dir}(\alpha)$ is the flat Dirichlet distribution with $\alpha = (\frac{1}{m}, \dots, \frac{1}{m}) \in \mathbb{R}^m$, $\mathbb{S}^m = \{r \in \mathbb{R}_{>0}^m : \sum_i r_i = 1\}$ is the feasible space of preference vectors r , $\hat{f}_i(\cdot) : \mathcal{X} \rightarrow$

$\mathbb{R}, \forall i \in \{1, \dots, m\}$ are surrogate models, $\hat{f}(\cdot) = \left[\hat{f}_i(\cdot) \right]_{i=1}^m$, the scalarization function $g : \mathbb{R}^m \times \mathbb{S}^m \rightarrow \mathbb{R}$ helps us map a given preference vector with a Pareto solution, and $h(\cdot, \cdot) : \mathbb{S}^m \times \Theta \rightarrow \mathcal{X}$ is called the Pareto set model, approximating the mentioned mapping.

The landscape provides various choices for scalarization functions like Linear Scalarization, Chebyshev, and Inverse Utility. However, we decided to use the Chebyshev function,

$$g(\hat{f}(\mathbf{x})|r) = \max_i \{r_i |\hat{f}_i(\mathbf{x}) - z_i^*|\} \quad \forall i \in \{1, \dots, m\} \quad (9)$$

where z^* is an ideal objective vector. In black-box optimization, where the form of the objectives is unknown (and potentially non-convex), the Chebyshev scalarization is useful because it allows us to work with the problem in a simpler, scalarized form. Unlike traditional linear scalarization, which can struggle with non-convex Pareto fronts (as they may miss points on non-convex regions), Chebyshev scalarization can handle non-convex Pareto fronts more effectively by finding solutions across the entire front.

In truth, obtaining a complete and accurate approximation of the entire variable space using surrogate models is impossible. However, the essence of Pareto set learning lies in the precise approximation of the feasible optimal space. This space, existing as a continuous manifold, represents a distinct subset within the broader variable space. Through a well-considered strategy, Pareto Set Learning demonstrates its ability to establish a highly accurate mapping between the preference vector space and the Pareto continuous manifold. However, optimizing the Pareto Set Model can be challenging and unstable if Gaussian processes do not approximate well black-box functions.

Stein Variational Gradient Descent (SVGD)

We present the Stein Variational Gradient Descent method, as proposed by (Liu and Wang 2016). SVGD serves as an effective approach for approximating intricate distributions by utilizing a collection of particles. These particles are iteratively updated to match a target distribution more closely.

To approximate a given target distribution $p(x)$ by a set of particles $\{x_i\}_{i=1}^n$, we first draw a set of initial particles $\{x_i^0\}_{i=1}^n$ from the initial distribution $q(x)$, and then iteratively updating them with a deterministic transformation of form:

$$\mathbf{x}_i \leftarrow \mathbf{x}_i + \epsilon \phi_k^*(\mathbf{x}_i), \forall i = 1, \dots, n, \quad (10)$$

$$\phi_k^* = \arg \max_{\phi \in \mathcal{B}_k} \left\{ -\frac{d}{d\epsilon} \mathbf{KL}(q_{[\epsilon\phi]} || p) \Big|_{\epsilon=0} \right\} \quad (11)$$

here ϵ is a step size, ϕ^* is an optimal transform chosen to maximize the decreasing rate of the KL divergence between the distribution of particles the target p , and $q_{[\epsilon\phi]}$ is defined the distribution of the updated particles, and \mathcal{B}_k is a unit ball of a reproducing kernel Hilbert space (RKHS) $\mathcal{H}_k^d := \mathcal{H}_k \times \dots \times \mathcal{H}_k$, \mathcal{H}_k is a Hilbert space associated with a positive definite kernel $k(\mathbf{x}, \mathbf{x}')$.

(Liu and Wang 2016) refers to \mathcal{P} as the Stein operator and shows that the optimal transform ϕ^* expressed:

$$\begin{aligned} \phi_k^* &\propto \mathbb{E}_{\mathbf{x} \sim q} [\mathcal{P}k(\mathbf{x}, \cdot)] \\ &= \mathbb{E}_{\mathbf{x} \sim q} [\nabla_{\mathbf{x}} \log p(\mathbf{x}) k(\mathbf{x}, \cdot) + \nabla_{\mathbf{x}} k(\mathbf{x}, \cdot)] \end{aligned} \quad (12)$$

Pareto Set Learning with SVGD

This section outlines our primary framework. We present the fundamental concept of the alignment between Hypernetwork for sampling and the SVGD algorithm. This study introduces a novel methodology for acquiring the Pareto front by employing SVGD for gradient updates while incorporating a control factor through scalarization techniques. Next, we develop our local kernel to improve the efficiency of the SVGD optimization technique.

The key idea is to observe the utilization of hypernetwork to generate a set of initial arbitrary particles. In the context of SVGD, choosing an optimal transform ϕ^* corresponds to the learning gradient to update the parameter of the hypernetwork. Our objective is to advance the particles toward the Pareto set.

Stein Variational Hypernetworks

The proposal involves employing Stein Variational Gradient Descent to adjust the gradient for the update direction in a method we refer to as Stein Variational Hypernetworks (SVH). This approach leverages the concept of Multi-Sample Hypernetwork (Hoang et al. 2023), which involves sampling a set of preference vectors as input to the hypernetwork and generating a corresponding set of \mathbf{x}_r as output.

We sample K preference vectors r_i and fit them to the Pareto set model $h(r|\theta)$ to generate solutions $\{\mathbf{x}_{r_i}\}_{i=1}^K$. For a problem with m objectives, the objective functions are represented as $\mathcal{F}_i(\mathbf{x}_{r_i}) = [\hat{f}_1(\mathbf{x}_{r_i}), \dots, \hat{f}_m(\mathbf{x}_{r_i})]$, where \mathcal{F}_i is considered as a particle. The set of particles $\mathbf{F} = \{\mathcal{F}_i\}_{i=1}^K$ serves as the initialization set. We iteratively move the particles towards the Pareto front using the update rule:

$$\theta_{t+1} = \theta_t - \xi \nabla_{\theta} g(\mathcal{F}(\mathbf{x}_r) | r) \quad (13)$$

where ξ is the learning rate.

Stein Variational Gradient Descent plays a role in adjusting the direction of the gradient update and helps push the points away from each other, thereby enhancing the diversity of the Pareto front. Consequently, the following modifications will be applied to Formula 13:

$$\theta_{t+1} = \theta_t - \xi \sum_{i=1}^K \sum_{j=1}^K \nabla_{\theta} g(\mathcal{F}_i | r_i) k(\mathcal{F}_i, \mathcal{F}_j) - \alpha \nabla_{\theta} k(\mathcal{F}_i, \mathcal{F}_j) \quad (14)$$

here, $k(\mathcal{F}_i, \mathcal{F}_j)$ is kernel matrix. Gaussian kernel is used: $k(\mathcal{F}_i, \mathcal{F}_j) = \exp(-\frac{1}{2c^2} \|\mathcal{F}_i - \mathcal{F}_j\|^2)$, c is bandwidth, α is a positive coefficient that controls the importance of the divergence term.

We can see in this formula that there are two main ideas:

- When considering $\nabla_{\theta} g(\mathcal{F}_i | r_i)$, it is our priority to guarantee that the gradient update direction aligns with our intended direction and is maintainable under control.
- The second term $\nabla_{\theta} k(\mathcal{F}_i, \mathcal{F}_j)$ plays a crucial role in the optimization process by facilitating the separation of particles. This term is primarily responsible for generating a repulsive force that pushes particles apart from each other (Liu, Tong, and Liu 2021).

Notably, our study proposes an innovative approach which is different from MOO-SVGD by (Liu, Tong, and Liu 2021). We employ the gradient of the scalarization function $\nabla_{\theta} g(\mathcal{F}_i | r_i)$ as a substitute for Multi-Gradient Descent Algorithm (MGDA) (Désidéri 2012) in their work, which guarantees the convergence of the Pareto set learning method with preference vector, while simultaneously enhancing the diversity of the Pareto front in accordance with SVGD theory.

Design of SVGD Local Kernel

The selection of the kernel significantly influences the efficacy of SVGD. The kernel employed in SVGD plays a pivotal role in determining the direction in which the particles are transformed, as it assigns weights to the contributions of each particle. Consequently, selecting an appropriate kernel is of utmost importance; while the Radial Basis Function (RBF) kernel utilizing a median heuristic is frequently adopted, it tends to be suboptimal for more challenging tasks. The concept of multiple kernel learning is presented by (Ai et al. 2023) under the designation MK-SVGD. This approach uses a composite kernel for approximation, with each kernel weighted by its importance.

In the context of expensive multi-objective optimization problems, surrogate models such as Gaussian Processes are often employed to approximate objective functions due to the high cost of direct evaluations. However, these models can lead to fragmented representations of the objective landscape, resulting in numerous pseudo-local optima—apparent optimal points that do not represent true global optima. Within our SVH-PSL framework, we generate a diverse set of particles using a hypernetwork, which can initially distribute across these various pseudo-local optima, which can lead to difficulties in particle movement and gradient updates. Traditional kernel computation considers the overall distance between points, which may not capture each dimension's influence in complex spaces.

To address this challenge, we introduce the concept of the local kernel for adapting to complex landscapes by calculating the kernel for each dimension separately, capturing its individual influence on particle movement.

$$\mathbf{k}(\mathcal{F}_i, \mathcal{F}_j) = \sum_{n=1}^m \mathbf{k}_n(\hat{f}_n(\mathbf{x}_{r_i}), \hat{f}_n(\mathbf{x}_{r_j})) \cdot id_n \quad (15)$$

Here $\mathbf{k}_n(\cdot)$ is the kernel for n th dimension, id_n is 1 in case n is the index of the function \hat{f}_n that achieves the maximum formula (9), 0 otherwise.

This methodology facilitates a more detailed evaluation of particle interactions, which can help in overcoming the challenges posed by pseudo-local optima. By considering the influence of each dimension separately, we can achieve a more accurate modification of particle locations. With a more refined kernel computation, the gradient updates in SVGD can become more accurate, leading to improved convergence properties. Particles are less likely to get stuck in pseudo-local optima, as the dimensional influence can guide them more effectively towards globally optimal regions. Our method might enhance the robustness of SVGD by allowing




Figure 2: Comparison of the impact of neighboring particles on a particle with and without the local kernel. Here we are considering $K = 10$

it to adapt better to the geometry of the objective space, particularly when dealing with complex optimization problems.

Discussion. Intuitively, in SVH-PSL executes direct operations via the kernel to modify the gradient. It is well established that gradients are highly sensitive and can easily experience issues such as gradient explosion or vanishing. Figure 2 illustrates that when a particle is significantly affected by the presence of other particles, this interaction can facilitate a more rapid convergence. However, it may also lead to instability, potentially diverting the particle from reaching the global optimum. The implementation of our local kernel facilitates precise adjustments aimed at minimizing the potential for instabilities.

Algorithm 1: SVH-PSL main algorithm

Input: Black-box multi-objective objectives $\mathbf{f}(\mathbf{x}) = \{f_j(\mathbf{x}), j \in 1, \dots, m\}$ and initial evaluation samples $\{\mathbf{x}_0, \mathbf{f}(\mathbf{x}_0)\}$

- 1: $\mathbf{D} \leftarrow \{\mathbf{x}_0, \mathbf{f}(\mathbf{x}_0)\}$
- 2: Initialize Pareto set model $h(r|\theta_0)$
- 3: **for** $i \leftarrow 0$ to N **do**
- 4: Training GP $\hat{f}_j(\mathbf{x})$ for each $f_j(\mathbf{x})$ on \mathbf{D}
- 5: **for** $t \leftarrow 0$ to T **do**
- 6: Randomly sample K vectors $\{r_i\}_{i=1}^K \sim \mathbb{S}^m$
- 7: Update θ_i with gradient descent by Formula (14)
- 8: **end for**
- 9: Randomly sample B vectors $\{r_i\}_{i=1}^B \sim \mathbb{S}^m$, then compute $\mathbf{x}_{r_i} = h(r_i|\theta)$, $\mathbf{X} = \{\mathbf{x}_{r_i}\}_{i=1}^B$
- 10: Selecting subset $\{\mathbf{x}\}_b \in \mathbf{X}$ that has highest HVI
- 11: $\mathbf{D} \leftarrow \mathbf{D} \cup \{\mathbf{x}, \mathbf{f}(\mathbf{x})\}_b$
- 12: **end for**

Output: Total evaluated solutions $\mathbf{D} = \{\mathbf{x}, \mathbf{f}(\mathbf{x})\}$ and the final parameterized Pareto set model $h(r|\theta)$

For each iteration, we trained the Pareto Set Model $h(r|\theta)$ with T training step, whereas we randomly sampled K preference $\{r_i\}_{i=1}^K \sim \text{Dir}(\alpha)$, $\alpha \in \mathbb{R}^m$ for optimizing θ . Then, we sampled B preference $\{r_i\}_{i=1}^B \sim \text{Dir}(\alpha)$ to compute the corresponding evaluation \mathbf{x}_{r_i} . Finally, we selected a subset $\{\mathbf{x}\}_b \in \{\mathbf{x}_{r_i}\}_{i=1}^B$ including b elements that achieve the high-

est HVI for the next batch of expensive evaluations:

$$\begin{aligned} \{\mathbf{x}\}_b &= \arg \max_{\mathbf{x}_+} HVI(\hat{\mathbf{f}}(\mathbf{x}_+), D_y) \\ \text{s.t. } |\mathbf{x}_+| &= b, \mathbf{x}_+ \in \{\mathbf{x}_{r_i}\}_{i=1}^B \end{aligned} \quad (16)$$

Experiments

Evaluation Metric. We use Log Hypervolume Difference (LHD) to evaluate the quality of a learned Pareto front, denoted as \mathcal{P}_f , compared to the true approximate Pareto front for the synthetic/real-world problems, represented as $\hat{\mathcal{P}}_f$. The calculation of LHD involves taking the logarithm of the difference in hypervolumes between $\hat{\mathcal{P}}_f$ and \mathcal{P}_f as follows:

$$LHD(\mathcal{P}_f, \hat{\mathcal{P}}_f) = \log \left(\text{HV}(\mathcal{P}_f) - \text{HV}(\hat{\mathcal{P}}_f) \right)$$

Baseline Methods. We compare SVH-PSL with the current state-of-the-art MOBO PFL methods, including TS-TCH (Paria, Kandasamy, and Póczos 2020), USeMO-EI (Belakaria et al. 2020), MOEA/D-EGO (Zhang et al. 2009), TSEMO (Bradford, Schweidtmann, and Lapkin 2018), DGEMO (Konakovic Lukovic, Tian, and Matusik 2020), qParEGO (Knowles 2006), qEHVI (Daulton, Balandat, and Bakshy 2020), qNEHVI (Daulton, Balandat, and Bakshy 2021), DAPSL (Lu, Li, and Zhou 2024) and PSL-MOBO (Lin et al. 2022).

Synthetic and Real-World Benchmarks. To demonstrate the effectiveness of our proposed method, we perform experiments in synthetic benchmarks and real-world application datasets; each dataset contains 2 or 3 constraint objectives. For the synthetic test, we choose 5 problems in the class of problem ZDT (Deb and Srinivasan 2006) and F2, VLMOP2 (Van Veldhuizen and Lamont 1999) with 2 objectives. For the real-world issues, we test our hypothesis in 7 over 16 datasets summarized by Tanabe and Ishibuchi (Tanabe and Ishibuchi 2020). In our experiments, RE problems are denoted as RE x , where x is the number of objectives and y is the id of the problems, ie, the RE21 means the real-world problems with 2 objectives and has id 1, associated with the issues RE2-4-1. The same definition applies to RE32, RE33, RE36, RE37, RE41, and RE42.

Experiment Settings

For all experiments¹, we randomly generate 20 initial solutions for expensive evaluations. The model is trained over 20 iterations with a batch size of $b = 5$. In Formula 14, the hyperparameters are set as $c = med$ (the median of the pairwise distances between samples) and $\alpha = 0.1$. Each problem is trained 5 times to eliminate randomness.

Experimental Results and Analysis

MOBO Performance. Figure 3 compares the log hypervolume difference (LHD) of baseline methods, with solid lines showing the mean and shaded regions indicating standard deviation. Our proposed SVH-PSL demonstrates superior performance, converging rapidly in synthetic experiments

¹Code is available at <https://github.com/nguyenduc810/SVH-PSL>

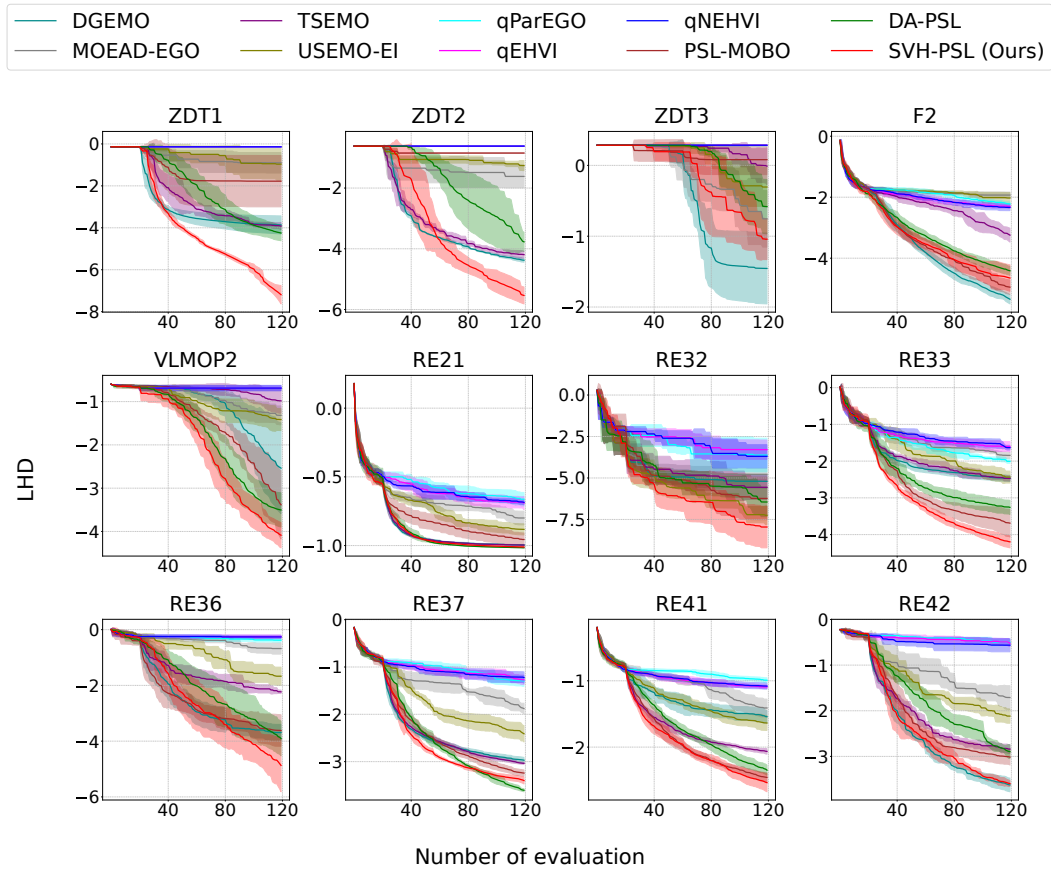


Figure 3: Mean Log Hypervolume Differences between the truth Pareto Front and the learned Pareto Front with respect to the number of expensive evaluations on all MOBO algorithms.

and achieving strong results in real-world scenarios. Notably, the integration of the SVGD algorithm enhances convergence speed compared to the conventional PSL-MOBO method while also providing better diversity on the Pareto set than the DA-PSL method.

Analysis of Local Kernel Benefits in Complex Pareto Fronts. Figure 4 demonstrates the effectiveness of local kernel integration in SVH-PSL, especially in solving real-world problems with complex Pareto fronts (RE37, RE41, and RE42). SVH-PSL achieves a consistently lower LHD compared to SVH-PSL without the local kernel. Moreover, SVH-PSL with the local kernel exhibits a smaller variance, indicating that the local kernel adapts effectively to the complex Pareto front. This results in more stable and consistent solutions across multiple evaluations.

Ablation Studies

The Trade-Off Between Exploration and Exploitation. In Formula 14, the second term acts as the repulsion component, with the parameter α serving as a trade-off between exploration and exploitation. As shown in Figure 5, α significantly impacts performance. When α is small (left), the variance is high, indicating that insufficient repulsion makes it difficult




Figure 4: Illustration of SVH-PSL performance with and without local kernel integration in real-world problems with a complex PF.

for the particles to escape pseudo-optimal points, leading to instability. Conversely, a larger α (right) encourages exploration, driving the particles toward the optimal solution. **High Dimensional.** Figure 6 shows that SVH-PSL performs well in high-dimensional conditions. Although optimization problems become increasingly challenging as dimensionality rises, SVH-PSL maintains its effectiveness. However, it also exhibits higher variance.

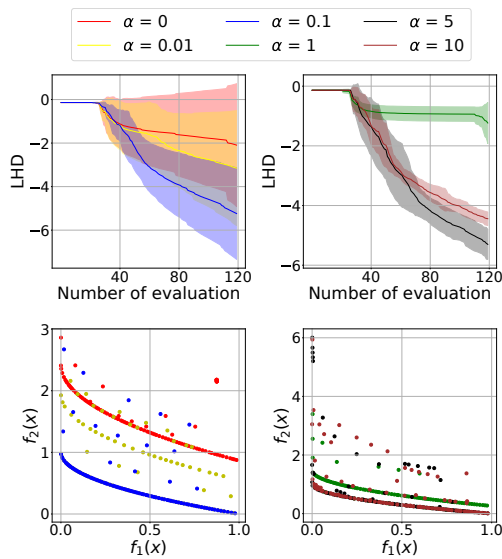


Figure 5: Impact of α on the trade-off between exploration and exploitation in ZDT1 (LHD & Pareto front).

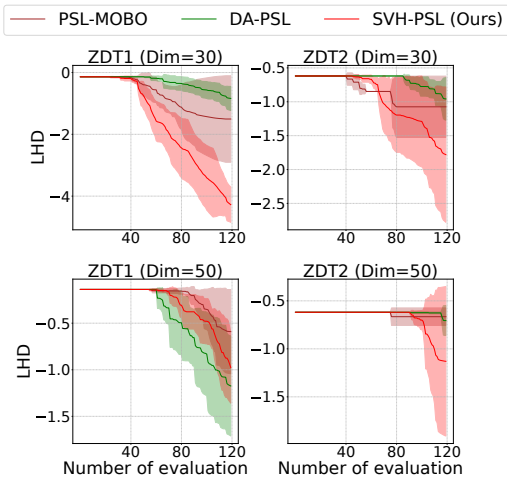


Figure 6: Performance comparison of high-dimensional ZDT1 and ZDT2 problems.

Related Work

Multi-objective Bayesian Optimization. Conventional Multi-objective Bayesian optimization (MOBO) methods have primarily concentrated on locating singular or limited sets of solutions. To achieve a diverse array of solutions catering to varied preferences, scalarization functions have emerged as a prevalent approach. Notably, Paria, Kandasamy, and Póczos (2020) adopts a strategy of scalarizing the Multi-objective problem into a series of single-objective ones, integrating random preference vectors during optimization to yield a collection of diverse solutions. Meanwhile, Abdolshah et al. (2019) delves into distinct regions on the Pareto Front through preference-order constraints, grounded in the Pareto Stationary equation.

Alternatively, evolutionary and genetic algorithms have also played a role in furnishing diversified solution sets, as observed in Zhang et al. (2009) which concurrently tackles a range of surrogate scalarized subproblems within the MOEA/D framework (Zhang and Li 2007). Additionally, Bradford, Schweidtmann, and Lapkin (2018) adeptly combines Thompson Sampling with Hypervolume Improvement to facilitate the selection of successive candidates. Complementing these approaches, Konakovic Lukovic, Tian, and Matusik (2020) employs a well-crafted local search strategy coupled with a specialized mechanism to actively encourage the exploration of diverse solutions. A notable departure from conventional methods, Belakaria et al. (2020) introduces an innovative uncertainty-aware search framework, orchestrating the optimization of input selection through surrogate models. This framework efficiently addresses MOBO problems by discerning and scrutinizing promising candidates based on measures of uncertainty.

Pareto Set Learning or PSL emerges as a proficient means of approximating the complete Pareto front, the set of optimal solutions in multi-objective optimization (MOO) problems, employing Hypernetworks (Chauhan et al. 2023). This involves the utilization of hypernetworks to estimate the relationship between arbitrary preference vectors and their corresponding Pareto optimal solutions. In the realm of known functions, effective solutions for MOO problems have been demonstrated by Navon et al. (2021); Hoang et al. (2023); Tuan et al. (2024) while in the context of combinatorial optimization, (Lin, Yang, and Zhang 2022) have made notable contributions. Notably, Lin et al. (2022) introduced a pioneering approach to employing PSL in the optimization of multiple black-box functions, PSL-MOBO, leveraging surrogate models to learn preference mapping, making it the pioneer method in exploring Pareto Set learning for the black-box multi-objective optimization task. However, PSL-MOBO is hindered by its reliance on Gaussian processes for training the Pareto set model, leading to issues of inadequacy and instability in learning the Pareto front.

Stein Variational Gradient Descent is a particle-based inference method introduced by (Liu and Wang 2016). This approach has gained significant attention due to its ability to approximate complex probability distributions through deterministic sampling methods, making it a powerful alternative to traditional Markov Chain Monte Carlo (MCMC) techniques (Neal 2012; Hoffman, Gelman et al. 2014). Recent research has utilized SVGD in the context of multi-objective optimization (MOO) challenges (Liu, Tong, and Liu 2021; Phan et al. 2022).

Conclusion

This study presents a novel methodology for learning the Pareto set, utilizing Stein Variational Gradient Descent in conjunction with a Hypernetwork. This framework proves to be effective in approximating the complete Pareto Front while simultaneously optimizing several conflicting black-box objectives. The findings from our experiments further indicate that the method effectively acquires a diverse Pareto set, which facilitates robust exploration and exploitation.

Acknowledgments

This research was funded by Vingroup Innovation Foundation (VINIF) under project code VINIF.2024.DA113.

References

- Abdolshah, M.; Shilton, A.; Rana, S.; Gupta, S.; and Venkatesh, S. 2019. Multi-objective Bayesian optimisation with preferences over objectives. *arXiv:1902.04228*.
- Ai, Q.; Liu, S.; He, L.; and Xu, Z. 2023. Stein variational gradient descent with multiple kernels. *Cognitive Computation*, 15(2): 672–682.
- Attia, P. M.; Grover, A.; Jin, N.; Severson, K. A.; Markov, T. M.; Liao, Y.-H.; Chen, M. H.; Cheong, B.; Perkins, N.; Yang, Z.; Herring, P. K.; Aykol, M.; Harris, S. J.; Braatz, R. D.; Ermon, S.; and Chueh, W. C. 2020. Closed-loop optimization of fast-charging protocols for batteries with machine learning. *Nature*, 578: 397–402.
- Belakaria, S.; Deshwal, A.; Jayakodi, N. K.; and Doppa, J. R. 2020. Uncertainty-aware search framework for multi-objective Bayesian optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 10044–10052.
- Bradford, E.; Schweidtmann, A. M.; and Lapkin, A. 2018. Efficient multiobjective optimization employing Gaussian processes, spectral sampling and a genetic algorithm. *Journal of global optimization*, 71(2): 407–438.
- Chauhan, V. K.; Zhou, J.; Lu, P.; Molaei, S.; and Clifton, D. A. 2023. A Brief Review of Hypernetworks in Deep Learning. *arXiv preprint arXiv:2306.06955*.
- Daulton, S.; Balandat, M.; and Bakshy, E. 2020. Differentiable expected hypervolume improvement for parallel multi-objective Bayesian optimization. *Advances in Neural Information Processing Systems*, 33: 9851–9864.
- Daulton, S.; Balandat, M.; and Bakshy, E. 2021. Parallel bayesian optimization of multiple noisy objectives with expected hypervolume improvement. *Advances in Neural Information Processing Systems*, 34: 2187–2200.
- Deb, K.; and Srinivasan, A. 2006. Innovization: Innovating design principles through optimization. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation*, 1629–1636.
- Désidéri, J.-A. 2012. Multiple-gradient descent algorithm (MGDA) for multiobjective optimization. *Comptes Rendus Mathématique*, 350(5-6): 313–318.
- He, C.; Zhang, Y.; Gong, D.; and Ji, X. 2023. A review of surrogate-assisted evolutionary algorithms for expensive optimization problems. *Expert Systems with Applications*, 217: 119495.
- Hoang, L. P.; Le, D. D.; Tuan, T. A.; and Thang, T. N. 2023. Improving pareto front learning via multi-sample hypernetworks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 7875–7883.
- Hoffman, M. D.; Gelman, A.; et al. 2014. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.*, 15(1): 1593–1623.
- Hou, M.; Chen, F.; and Pei, Y. 2024. Optimization of geometric parameters of ejector for fuel cell system based on multi-objective optimization method. *International Journal of Green Energy*, 21(2): 228–243.
- Knowles, J. 2006. ParEGO: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1): 50–66.
- Konakovic Lukovic, M.; Tian, Y.; and Matusik, W. 2020. Diversity-guided multi-objective bayesian optimization with batch evaluations. *Advances in Neural Information Processing Systems*, 33: 17708–17720.
- Laumanns, M.; and Ocenasek, J. 2002. Bayesian optimization algorithms for multi-objective optimization. In *International Conference on Parallel Problem Solving from Nature*, 298–307. Springer.
- Le, D. D.; and Lauw, H. W. 2017. Indexable bayesian personalized ranking for efficient top-k recommendation. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, 1389–1398.
- Lee, S. H.; Li, Y.; Ke, J.; Yoo, I.; Zhang, H.; Yu, J.; Wang, Q.; Deng, F.; Entis, G.; He, J.; et al. 2024. Parrot: Pareto-optimal Multi-Reward Reinforcement Learning Framework for Text-to-Image Generation. *arXiv preprint arXiv:2401.05675*.
- Lin, X.; Yang, Z.; and Zhang, Q. 2022. Pareto Set Learning for Neural Multi-Objective Combinatorial Optimization. In *International Conference on Learning Representations*.
- Lin, X.; Yang, Z.; Zhang, X.; and Zhang, Q. 2022. Pareto set learning for expensive multi-objective optimization. *Advances in Neural Information Processing Systems*, 35: 19231–19247.
- Liu, Q.; and Wang, D. 2016. Stein variational gradient descent: A general purpose bayesian inference algorithm. *Advances in neural information processing systems*, 29.
- Liu, X.; Tong, X.; and Liu, Q. 2021. Profiling pareto front with multi-objective stein variational gradient descent. *Advances in Neural Information Processing Systems*, 34: 14721–14733.
- Lu, Y.; Li, B.; and Zhou, A. 2024. Are You Concerned about Limited Function Evaluations: Data-Augmented Pareto Set Learning for Expensive Multi-Objective Optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 14202–14210.
- Milojkovic, N.; Antognini, D.; Bergamin, G.; Faltings, B.; and Musat, C. 2019. Multi-gradient descent for multi-objective recommender systems. *arXiv preprint arXiv:2001.00846*.
- Navon, A.; Shamsian, A.; Fetaya, E.; and Chechik, G. 2021. Learning the Pareto Front with Hypernetworks. In *International Conference on Learning Representations*.
- Neal, R. M. 2012. MCMC using Hamiltonian dynamics. *arXiv preprint arXiv:1206.1901*.
- Paria, B.; Kandasamy, K.; and Póczos, B. 2020. A flexible framework for multi-objective bayesian optimization using random scalarizations. In *Uncertainty in Artificial Intelligence*, 766–776. PMLR.
- Phan, H.; Tran, N.; Le, T.; Tran, T.; Ho, N.; and Phung, D. 2022. Stochastic multiple target sampling gradient descent. *Advances in neural information processing systems*, 35: 22643–22655.
- Schweikard, A.; Glosner, G. D.; Bodduluri, M.; Murphy, M. J.; and Adler, J. R. 2000. Robotic motion compensation for respiratory movement during radiosurgery. *Computer aided surgery : official journal of the International Society for Computer Aided Surgery*, 5 4: 263–77.
- Swersky, K.; Snoek, J.; and Adams, R. P. 2013. Multi-task bayesian optimization. *Advances in neural information processing systems*, 26.
- Tanabe, R.; and Ishibuchi, H. 2020. An easy-to-use real-world multi-objective optimization problem suite. *Applied Soft Computing*, 89: 106078.
- Tuan, T. A.; Hoang, L. P.; Le, D. D.; and Thang, T. N. 2024. A framework for controllable pareto front learning with completed scalarization functions and its applications. *Neural Networks*, 169: 257–273.

Van Veldhuizen, D. A.; and Lamont, G. B. 1999. Multiobjective evolutionary algorithm test suites. In *Proceedings of the 1999 ACM symposium on Applied computing*, 351–357.

Zhang, Q.; and Li, H. 2007. MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on evolutionary computation*, 11(6): 712–731.

Zhang, Q.; Liu, W.; Tsang, E.; and Virginas, B. 2009. Expensive multiobjective optimization by MOEA/D with Gaussian process model. *IEEE Transactions on Evolutionary Computation*, 14(3): 456–474.

Zitzler, E.; and Thiele, L. 1999. Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach. *IEEE transactions on Evolutionary Computation*, 3(4): 257–271.