ADAPTIVE REFERENCE GEOMETRY FOR PARETO SET LEARNING IN EXPENSIVE MULTI-OBJECTIVE OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Pareto set learning (PSL) learns a continuous mapping from preference vectors to Pareto-optimal solutions, providing an appealing approach to expensive multi-objective optimization under limited evaluation budgets. Most existing scalarization-based PSL methods typically estimate reference points from the evaluated archive. We show that this simple design choice can induce an initialization-dependent reference geometry that restricts Pareto-front exploration: when the archive covers only a local portion of the front, the resulting scalarization can bias the learned solutions toward the observed region and make unexplored trade-offs difficult to reach. We refer to this failure as the reference trap. To address it, we propose Adaptive Reference Geometry for PSL, instantiated through Adaptive Reverse Tchebycheff (AdaRevTche). Instead treating the reference as a fixed archive statistic, AdaRevTche jointly learns the Pareto-set mapping and a bounded worse-side reference, allowing the scalarization geometry to adapt as optimization proceeds. We establish an exact result showing how the reference point restricts the reachable region in two-objective Tchebycheff scalarization, and analyze the corresponding geometric behavior for problems with more objectives. Results on 12 real-world and synthetic problems demonstrate the effectiveness of our AdaRevTche. Notably, compared with the SoTA PSL baseline, CDM-PSL, AdaRevTche reduces the hypervolume difference on RE41 by 51.8%.

1 INTRODUCTION

Many scientific and engineering applications, including materials discovery (Solomou et al., 2018; Gopakumar et al., 2018; Nambiar et al., 2022), battery system design (Gaonkar et al., 2022; Zhang et al., 2023), and electromagnetic device optimization (Easum et al., 2018; Xiao et al., 2018), involve multiple conflicting objectives. However, evaluating these objectives often requires costly physical experiments or high-fidelity simulations (Knowles, 2006; Zhang et al., 2010; Chugh et al., 2019). Therefore, a fundamental challenge is to perform multi-objective optimization with a limited evaluation budget, referred to as expensive multi-objective optimization. Its seeks to efficiently discover Pareto-optimal trade-off solutions despite expensive evaluations.

To address expensive multi-objective optimization, multi-objective Bayesian optimization (MOBO) combines probabilistic surrogate models with acquisition and batch-selection strategies to efficiently allocate limited evaluations (Shahriari et al., 2016; Bradford et al., 2018; Konaković Luković et al., 2020; Daulton et al., 2020b). Beyond sequentially selecting individual solutions, Pareto set learning (PSL) aims to directly learn a continuous mapping from preference vectors to Pareto-optimal solutions (Navon et al., 2021; Lin et al., 2022). This enables the generation of diverse preference-conditioned candidates without repeatedly solving separate scalarized optimization problems.

Recent PSL methods (Lu et al., 2024; Li et al., 2025; Nguyen et al., 2025) improve data efficiency and solution diversity through techniques such as data augmentation, diffusion modeling, and Stein interactions. Nevertheless, most PSL methods rely on scalarization, such as weighted Tchebycheff scalarization or achievement scalarizing functions, to map preference vectors to Pareto-optimal solutions (Wierzbicki, 1982b; Emmerich & Deutz, 2018; Knowles, 2006). These typically require an ideal or aspiration reference point, which in PSL is commonly estimated from the evaluated archive

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

rather than the unknown Pareto front (Lin et al., 2022; Lu et al., 2024; Nguyen et al., 2025). Importantly, the reference point does not only specify a target in the objective space; it also shapes the geometry of the scalarized optimization problem. When the initial archive provides only partial or biased observations of the Pareto front, the estimated reference point may become inaccurate, causing distorted scalarized objectives and restricting the discovery of unexplored Pareto regions.

This limitation is not merely caused by inaccurate surrogate models or ineffective optimization. A initial archive may induce a biased reference geometry that confines the reachable scalarized optima to a limited region, even when the true Pareto front is available and scalarized subproblems are solved exactly. We identify this initialization-sensitive geometric failure as the *reference trap*.

To address this issue, we propose an adaptive reference geometry for Pareto set learning, instantiated by our Adaptive Reverse Tchebycheff PSL (AdaRevTche). Instead of relying on an ideal reference estimated from the current archive, AdaRevTche devises a learnable worse-side reference and reverses the conventional Tchebycheff formulation to encourage exploration toward unexplored regions of the Pareto front. We further devise an adaptive objective-wise reference adjustment strategy that updates the reference together with the Pareto set model, enabling preference-conditioned solutions to progressively expand beyond the observed front. Together, these designs alleviate the dependence on initial archives and achieve broader Pareto front exploration under limited evaluations. Our contributions are summarized as:

1) We identify the ***reference trap*** in Tchebycheff-based PSL, revealing that empirically estimated ideal references can limit Pareto front exploration. We **theoretically analyze** that Pareto solutions reachable by Tchebycheff scalarization is bounded by the estimated reference, even with exact objective evaluations and globally optimized subproblems.

2) We propose AdaRevTche, a PSL framework that adaptively **reshapes the reference geometry** to alleviate the reference trap and expand the region reachable by scalarized optimization.

3) Experiments on 12 real-world and synthetic problems demonstrate the effectiveness of our AdaRevTche. It consistently achieves lower hypervolume differences than the PSL-MOBO and non-PSL baselines across all considered tasks. Against the latest PSL baseline, CDM-PSL, AdaRevTche **reduces the hypervolume difference** on RE41 **by 51.8%**.

2 Related Work

Multi-objective Bayesian optimization. MOBO combines probabilistic surrogates with acquisition and batch selection rules for expensive evaluations (Shahriari et al., 2016; Chugh et al., 2019). TSEMO combines spectral sampling, Thompson sampling, hypervolume, and a genetic algorithm (Bradford et al., 2018); DGEMO promotes diversity over a piecewise continuous Pareto set approximation (Konaković Luković et al., 2020). Pareto set learning methods further model a continuous map from preferences to their solutions, with data augmentation, diffusion modeling, or Stein interactions used to improve training and diversity (Navon et al., 2021; Lu et al., 2024; Li et al., 2025; Nguyen et al., 2025). However, these methods mainly improve the surrogate, acquisition, or diversity mechanism after the candidate region has been specified. They do not address the possibility that an archive-induced reference makes valuable trade-offs unreachable before evaluation. Our method's bounded adaptive reference changes both the PSL-generated candidate scope and the HVI geometry, allowing the optimizer to explore previously inaccessible regions.

Scalarization methods. Scalarization converts a vector objective into preference-indexed single-objective subproblems and is a standard strategy in multi-objective optimization (Emmerich & Deutz, 2018). Linear scalarization can miss nonconvex trade-offs, whereas weighted Tchebycheff and achievement scalarizing functions use deviation from an ideal, utopia, or aspiration reference to target them (Wierzbicki, 1982b; Emmerich & Deutz, 2018). ParEGO and qNParEGO sample augmented Tchebycheff subproblems over Bayesian-optimization rounds (Knowles, 2006; Chugh et al., 2019), while PSL-MOBO learns a continuous preference-to-solution map through scalarized surrogate losses (Lin et al., 2022). These approaches typically treat the ideal or aspiration reference as fixed or as a direct archive estimate; under a limited budget, that estimate can be local and its active bottleneck can restrict the reachable part of the front. Recent work has also considered preference-based expensive multi-objective optimization without relying on an ideal point (Zhao et al., 2025). Our proposed AdaRevTche instead learns a bounded worse-side reference redistribu-

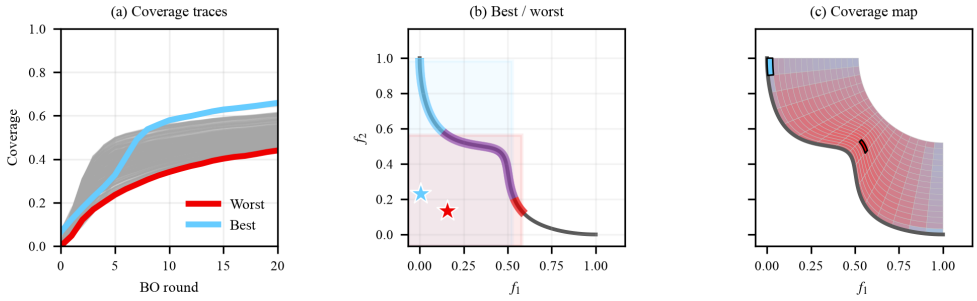

Figure 1: An exact-front reference-trap diagnostic on the convex polar problem, a two-objective synthetic problem with an analytically known Pareto front given by $d = 0$, while every point with $d > 0$ is dominated by its counterpart at $d = 0$.

tion jointly with the preference-to-solution map, therefore it changes the reachable candidate scope rather than merely changing the preferred direction.

3 Proposed Method

3.1 Problem Statement and Motivation

Many real-world applications, such as materials discovery (Solomou et al., 2018; Kang et al., 2025; Gopakumar et al., 2018; Biswas et al., 2021), chemical synthesis (Nambiar et al., 2022; Zhang et al., 2024), battery and energy-system design (Gaonkar et al., 2022; Cho et al., 2025), aerospace engineering (Zhang et al., 2023; Irshad et al., 2023; Zhang et al., 2021; Chen et al., 2023), and antenna or electromagnetic device design (Easum et al., 2018; de Villiers et al., 2017; Xiao et al., 2018; Lim et al., 2015; Wei et al., 2022; Zeng et al., 2024), involve conflicting objectives. Expensive simulations, including full-wave electromagnetic, CFD, and finite-element simulations, make exhaustive optimization impractical; such problems are known as expensive multi-objective optimization problems (Lin et al., 2022). We formulate the problem as

$$\min_{x \in \mathcal{X}} f(x) = (f_1(x), \ldots, f_m(x)), \tag{1}$$

where $\mathcal{X}$ denotes the decision space and m is the number of objectives. Due to the limited evaluation budget, the goal is to efficiently approximate the Pareto front with a small number of objective evaluations.

Among scalarization methods, Tchebycheff scalarization is a widely adopted approach for multi-objective optimization (Emmerich & Deutz, 2018), where it (Knowles, 2006) recovers diverse Pareto optimal solutions by varying the preference vector. Given a preference vector $\lambda \in \Delta_{++}^{m-1} := \{\lambda \in \mathbb{R}_{++}^m : \mathbf{1}^\top \lambda = 1\}$ and a reference point z, the Tchebycheff scalarization is defined as

$$T(y; \lambda, z) = \max_{i=1,\ldots,m} \lambda_i (y_i - z_i). \tag{2}$$

For simplicity, we denote the objective vector $f(x)$ as y when discussing the objective space. Based on Eq. (2), Pareto set Learning (PSL) learns a a set model for mapping preferences to their solutions

$$g_\theta : \Delta^{m-1} \to \mathcal{X}, \tag{3}$$

which directly generates candidate solutions corresponding to different preferences and enables efficient Pareto front approximation with a limited evaluation budget.

Reference Dependence of Tchebycheff Scalarization. We reveal substantial variation in Pareto coverage across different initial archives in Fig. 1. Even with an exact Pareto front oracle and globally solved Tchebycheff subproblems, different initializations result in significantly different coverage trajectories and final Pareto coverage. This suggests that the instability is not caused by surrogate approximation residue or inefficient optimization but by the reference dependent geometry of Tchebycheff scalarization.

In practice, the reference point is commonly estimated from the evaluated archive instead of the unknown true Pareto geometry. A widely used (Lin et al., 2022) choice is the empirical ideal point:

$$\hat{z}_i^* = \min_{(x_j, y_j) \in \mathcal{D}} y_{j,i}, \tag{4}$$

where $\mathcal{D}$ denotes the evaluated archive.

However, an initial archive obtained from a limited number of random evaluations may only cover a portion of the Pareto front. Since the empirical reference is estimated from this archive, it inherits the initialization bias and changes the geometry of the Tchebycheff scalarization. Similar reference point effects have been observed in evolutionary optimization Sato (2014a); Wu et al. (2023b), and inverted or reverse scalarizations have been proposed as remedies Sato (2014a); Jiang & Yang (2016). We provide a formal characterization of this issue, showing how the empirical reference controls not only preference trade-offs but also the scope explored by scalarization.

Empirical Reachable Scope and Reference Trap. To characterize how the empirical reference restricts scalarized optimization, we define the archive induced scope as

$$\mathcal{B}_t^{\text{scope}} = \prod_{i=1}^{m} \left(-\infty, \max_{(x_j, y_j) \in \mathcal{D}_t} y_{j,i} \right]. \tag{5}$$

For any positive preference, the empirical ideal point and archive incumbent induce an exact Tchebycheff sublevel box. Every global Tchebycheff minimizer lies inside this box, while every negative Tchebycheff subgradient outside it points strictly inward, hence its worse-side boundary cannot expand across oracle rounds.

Proposition 3.1 (Reference-induced reachable scope)**.** *For each $\lambda \in \Delta_{++}^{m-1}$, let $\tau_t(\lambda) = \min_{y \in \mathcal{Y}_t} T_t(y; \lambda)$ and $\mathcal{A}_t(\lambda) = \{y : T_t(y; \lambda) \leq \tau_t(\lambda)\}$. Every global minimizer of $T_t(\cdot; \lambda)$ lies in $\mathcal{Y} \cap \mathcal{A}_t(\lambda)$, and every negative subgradient outside $\mathcal{A}_t(\lambda)$ points strictly toward this box. When $m = 2$, $\mathcal{A}_t(\lambda) \subseteq \mathcal{B}_t^{\text{scope}}$ for every λ, so archive-retaining oracle updates cannot expand $\mathcal{B}_t^{\text{scope}}$.*

The above result is an objective-space geometric statement. It does not by itself imply descent or convergence in the decision space. The formal proof and the higher-dimensional extension are provided in the supplementary A. Thus, an archive-induced empirical reference can confine scalarized optima to regions determined by the current archive, preventing subsequent evaluations from expanding toward uncovered trade-offs. Different initial archives can induce different reachable regions and Pareto coverage (Fig. 1(b)); we call this archive-dependent restriction the *reference trap*.

This observation motivates an adaptive mechanism that expands exploration beyond the reference trap while preserving the set-model optimization capability of PSL.

3.2 Adaptive Reverse Tchebycheff PSL

To escape the reference trap, we propose a reference geometry adaptation framework for PSL, termed Adaptive Reverse Tchebycheff PSL. Instead of relying on the empirical ideal reference used by conventional Tchebycheff scalarization, our method proposes a dynamically updated worse-side reference and jointly optimizes the reference geometry with the Pareto set mapping. By adapting the reference point toward unexplored regions, the proposed framework enlarges the reachable scope of scalarized solutions while preserving PSL's ability to map preferences to their solutions.

Reverse Scalarization with a Worse-Side Reference. The conventional Tchebycheff scalarization encourages solutions to approach an ideal-side reference point by minimizing the weighted objective deviations. However, under the reference trap, the empirical ideal reference is estimated from the evaluated archives, and the scalarization geometry is anchored within the explored scope. As a result, the optimization objective mainly pulls solutions toward the existing boundary rather than encouraging expansion into unexplored regions.

To overcome this limitation, we reverse the reference geometry by placing the reference point outside the Pareto front on the worse side of the objective space. Instead of minimizing the distance to an ideal reference, we maximize the normalized margin between a solution and the worse-side reference through the reverse Tchebycheff scalarization:

$$S_{\text{rev}}(\hat{y}; \lambda, \hat{r}) = \min_{i=1,\dots,m} \frac{\hat{r}_i - \hat{y}_i}{\lambda_i} + \rho \sum_{i=1}^{m} (\hat{r}_i - \hat{y}_i), \tag{6}$$

where $\hat{r}$ denotes emperical worse-side reference point satisfying $\hat{r}_i \geq y_i$ for $i = 1, \ldots, m, y_i \in \mathcal{D}_t$, and $\rho > 0$ is a small tie-breaking coefficient. Following (Lin et al., 2022), we use the second term to better distinguish among different solutions with identical bottlenecks by favoring larger overall margins while preserving the max-min margin criterion.

Adaptive Reference Update. Although the reverse Tchebycheff scalarization introduces a worse-side reference geometry, a fixed reference point cannot adapt to evolving Pareto structures. As the optimization proceeds, the reference margin may become insufficient for unexplored regions, limiting further expansion of the reachable scope. A straightforward solution is to directly optimize the worse-side reference point; however, unconstrained reference updates may lead to excessive movements and distort the scalarization geometry. Therefore, we propose a bounded adaptive reference mechanism that learns how to adjust the reference margin.

Specifically, we parameterize the worse-side reference using learnable parameters $\mathbf{a} \in \mathbb{R}^m$ around the nadir $\hat{r}^{(t)}$:

$$\hat{r}_i^{(t)}(\mathbf{a}) = \hat{r}_i^{(t)} + \max\left\{0,\ B\left([\mathrm{softmax}(\mathbf{a})]_i - \frac{1}{m}\right)\right\} \tag{7}$$

where B controls the adaptive range around the worse-side anchor. This parameterization redistributes a bounded reference margin across objectives and clips the negative movement to keep the anchor worse than all evaluated archives, allowing the reference geometry to adapt to the current Pareto structure.

Joint Optimization with PSL.

Given sampled preferences $\{\lambda^{(q)}\}_{q=1}^{Q}$, we jointly optimize the PSL mapping parameters θ and the learnable reference parameters $\mathbf{a}$. For objective i, we use the Gaussian process lower confidence bound (LCB) $\tilde{f}_i^{(t)}(x) = \mu_i^{(t)}(x) - 0.1\sigma_i^{(t)}(x)$. The optimization objective, with regularization weight β, is

$$\min_{\theta,a} -\frac{1}{Q}\sum_{q=1}^{Q} S_{\mathrm{rev}}(\hat{y}^{(q)}; \lambda^{(q)}, \hat{r}^{(t)}(\mathbf{a})) + \beta\mathcal{R}(\mathbf{a}), \tag{8}$$

where

$$\mathcal{R}(\mathbf{a}) = B^2 \left\| \mathrm{softmax}(\mathbf{a}) - \frac{\mathbf{1}}{m} \right\|_2^2 \tag{9}$$

regularizes the reference adjustment by encouraging a diverse allocation of the reference displacement across objectives and avoiding excessive movement along a single objective direction. The gradients provided by S_{rev} jointly train the set model g_θ and the adaptation parameters $\mathbf{a}$ in order to simultaneously learn preference conditioned solutions and adapt the reference geometry.

After optimization, the learned PSL set model generates a candidate set. We use the GP surrogate models to estimate the objective values of these candidates and select the next batch by greedy hypervolume improvement with the same adaptive reference point. Only the selected candidates are evaluated with the true objectives and added to the archive. The complete procedure is summarized in Algorithm 1.

3.3 Analysis of Reachable Scope Expansion

Geometric Analysis of Reachable Scope Expansion. A natural question is whether the reference trap can be resolved by directly learning the ideal reference point. Since the limitation originates from the empirical reference estimated from limited observations, one may expect that adapting the ideal reference could recover the missing Pareto regions. However, the ideal-side geometry of conventional Tchebycheff scalarization does not systematically expand the worse-side boundary. Specifically, for

$$T(y; \lambda, z) = \max_i \lambda_i(y_i - z_i), \tag{10}$$

let $I_T = \arg\max_i \lambda_i(y_i - z_i)$ denote the active coordinates. The subdifferential with respect to the ideal reference is

$$\partial_z T(y; \lambda, z) = \mathrm{conv}\left\{-\lambda_i e_i : i \in I_T\right\}. \tag{11}$$

When the active coordinate is unique, $I_T = \{k\}$, this reduces to $\nabla_z T = -\lambda_k e_k$.

Algorithm 1 Pipeline of our AdaRevTche.

Require: Initial archive $\mathcal{D}_0 = \{(x_j, y_j)\}_{j=1}^{N_0}$ with $y_j = f(x_j)$, number of rounds T, batch size b, number of sampled preferences Q, training iterations I, candidate number C, adaptive range B

for $t = 0, \ldots, T-1$ **do**
  Fit independent GP posterior models $\{(\mu_i^{(t)}, \sigma_i^{(t)})\}_{i=1}^{m}$ on $\mathcal{D}_t$ and define $\tilde{f}_i^{(t)}(x) = \mu_i^{(t)}(x) - 0.1\sigma_i^{(t)}(x)$
  Set the worse-side anchor $\hat{r}_i^{(t)} = \max_{(x_j,y_j)\in\mathcal{D}_t} y_{j,i}$, $i = 1, \ldots, m$
  Initialize the PSL parameters θ and set the reference logits $\mathbf{a} = \mathbf{0}$
  for $s = 0, \ldots, I-1$ **do**
    Sample $\{\lambda^{(j)}\}_{j=1}^{Q} \sim \mathrm{Dirichlet}(\mathbf{1})$
    Generate $x^{(j)} = g_\theta(\lambda^{(j)})$ and $\hat{y}^{(j)} = \tilde{f}^{(t)}(x^{(j)})$
    Compute the adaptive reference $\hat{r}^{(t)}(\mathbf{a})$ according to Eq. (7)
    Update θ and $\mathbf{a}$ by one optimizer step according to Eq. (8)
  end for
  Sample $\{\lambda^{(c)}\}_{c=1}^{C} \sim \mathrm{Dirichlet}(\mathbf{1})$ and construct $\mathcal{C}_t = \{g_\theta(\lambda^{(c)})\}_{c=1}^{C}$
  Select b candidates $\mathcal{X}_t^{\mathrm{eval}} \subset \mathcal{C}_t$ by greedy hypervolume improvement using $\hat{r}^{(t)}(\mathbf{a})$
  Evaluate $\mathcal{X}_t^{\mathrm{eval}}$ and update $\mathcal{D}_{t+1} = \mathcal{D}_t \cup \{(x, f(x)) : x \in \mathcal{X}_t^{\mathrm{eval}}\}$
end for
return Final evaluated archive $\mathcal{D}_T$

Therefore, optimizing the ideal reference only shifts the scalarization boundary toward the active objective, while the induced geometry contracts the attainable range in other objectives. Even with exact Pareto front access under an oracle setting, training the adaptive ideal reference component reduces the empirical scope rather than expanding it.

In contrast, our reverse Tchebycheff scalarization places the reference point beyond the Pareto front on the worse side of the objective space:

$$S_{\mathrm{rev}}(y; \lambda, r) = \min_i \frac{r_i - y_i}{\lambda_i} + \rho \sum_i (r_i - y_i). \tag{12}$$

Let

$$I_{\mathrm{R}} = \arg\min_i \frac{r_i - y_i}{\lambda_i} \tag{13}$$

denote the active bottleneck coordinates. The superdifferential with respect to the reference is

$$\partial_r S_{\mathrm{rev}} = \mathrm{conv}\left\{ \frac{1}{\lambda_i} e_i + \rho \mathbf{1} : i \in I_{\mathrm{R}} \right\}. \tag{14}$$

When the active coordinate is unique, $I_{\mathrm{R}} = \{k_{\mathrm{R}}\}$, this reduces to $\nabla_r S_{\mathrm{rev}} = \lambda_{k_{\mathrm{R}}}^{-1} e_{k_{\mathrm{R}}} + \rho \mathbf{1}$. Thus, maximizing the reverse scalarization moves the active worse-side coordinate outward more strongly than the others, enlarging the solution-reference margin and enabling expansion toward previously uncovered regions. This provides an explanation for why worse-side reference adaptation can alleviate the reference trap; further experiments are provided in the supplementary material C.

4 Experiments

4.1 Experimental Setup

Benchmarks. We evaluate on 11 widely used benchmarks with two to six objectives, including the widely used synthetic problems ZDT1(Zitzler et al., 2000), F2 (Lin et al., 2022), and VLMOP3 (van Veldhuizen & Lamont, 1999), and the real-world engineering design RE suite (Tanabe & Ishibuchi, 2020) (RE21, RE31, RE35, RE36, RE37, RE41, RE42, and RE61). To isolate and validate the reference trap studied in this work, we further develop a synthetic problem *convex polar*. These problems are deliberately designed without additional optimization challenges, allowing us to directly examine the impact of reference selection on Pareto front coverage under limited evaluation budgets. More details are provided in the supplementary B.

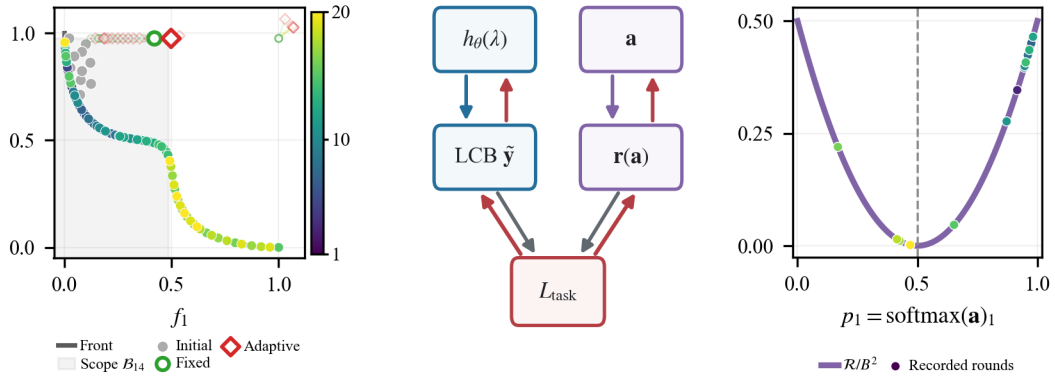

(a) Reference trace (b) Shared-loss updates (c) Regularization

Figure 2: Adaptive-reference mechanism on convex polar. (a) The adapted reference expands toward the previously inaccessible region, which greedy HVI selection subsequently explores. (b) The shared task loss sends gradients to the PSL and reference branches; $\nabla_{\mathbf{a}} L_{\text{task}} = J_{\mathbf{r},\mathbf{a}}^{\mathsf{T}} \nabla_{\mathbf{r}} L_{\text{task}}$. (c) The normalized two-objective penalty is $\mathcal{R}/B^2 = 2(p_1 - 1/2)^2$; dots are recorded allocations.

Baselines. We compare our method against PSL-MOBO (Lin et al., 2022), Stein variational hypernetwork PSL (SVH-PSL), TSEMO (Bradford et al., 2018), qNParEGO(Daulton et al., 2020a), DGEMO (Konaković Luković et al., 2020), data-augmented PSL (DA-PSL) (Lu et al., 2024), qLogEHVI (Daulton et al., 2020b) and CDM-PSL (Li et al., 2025). All baselines are implemented using their official codebases. It is worth noting that only PSL-MOBO and SVH-PSL, like our method, explicitly learn the PSL mapping in Eq. (3), while other methods optimize Pareto front quality. Together, these baselines cover representative and recent state-of-the-art approaches in Pareto set learning, diversity-guided MOBO, decomposition-based MOBO, and hypervolume-based batch BO.

Evaluation Metrics. We evaluate all methods using hypervolume (HV), computed in the objective space (Zitzler & Thiele, 1999). Since raw HV values are difficult to compare and visualize due to differences in scale, we report the logarithm of them, following the common practice in this area (Lin et al., 2022; Lu et al., 2024).: $G_t = \log\big(\max\{H^* - \bar{H}_t, \epsilon\}\big)$, where H^* is the analytic or densely approximated ground-truth HV and $\bar{H}_t$ is the mean HV across runs. This transformation preserves the HV improvement trend, with lower values indicating better performance. Following (Nguyen et al., 2025), ribbons show run-to-run dispersion using $G_t \pm \log(1 + \sigma_{H,t} / \max H^* - \bar{H}_t, \epsilon)$, where $\sigma_{H,t}$ is the HV standard deviation across runs. They are only for visualization and do not represent confidence intervals. More details are contained in supplementary E.

Implementation Details. Following (Lin et al., 2022), we use a two-layer PSL map with 256 ReLU units and train it for 1,000 steps per round with $Q = 10$ preference vectors and $\rho = 0.01$. The regularization parameters are set to $\beta = 0.1$ with an adaptive range $B = 1$. The learning rates are 10^{-3} for the map parameters and 10^{-2} for the free reference logits. Each batch is selected from 1,000 candidates by greedy hypervolume improvement maximization as (Lin et al., 2022). For a fair comparison, all methods use the same run-specific 10-point Latin-hypercube archive and perform 20 Bayesian optimization rounds with a batch size of 5, resulting in 110 total evaluations. We conduct 10 independent runs for each method. We fit independent Matérn-$5/2$ Gaussian processes and assign identical evaluation budgets to all baselines. All hyperparameters are fixed across problems.

4.2 Comparisons with State-of-the-arts

We evaluate AdaRevTche on 12 tasks, including 8 engineering design problems with names beginning with 'RE-', the synthetic zdt1, F2 and VLMOP3 benchmarks, and a deliberately constructed diagnostic problem. We compare AdaRevTche with the original PSL-MOBO, SVH-PSL, DA-PSL, DGEMO and other previously mentioned methods under identical initialization settings and evaluation budgets. The original DGEMO is available only for the two and three objective problems. Figure 3 reports the final round log HV difference over 10 independent runs. Detailed statistics of the results are provided in the supplementary material D.

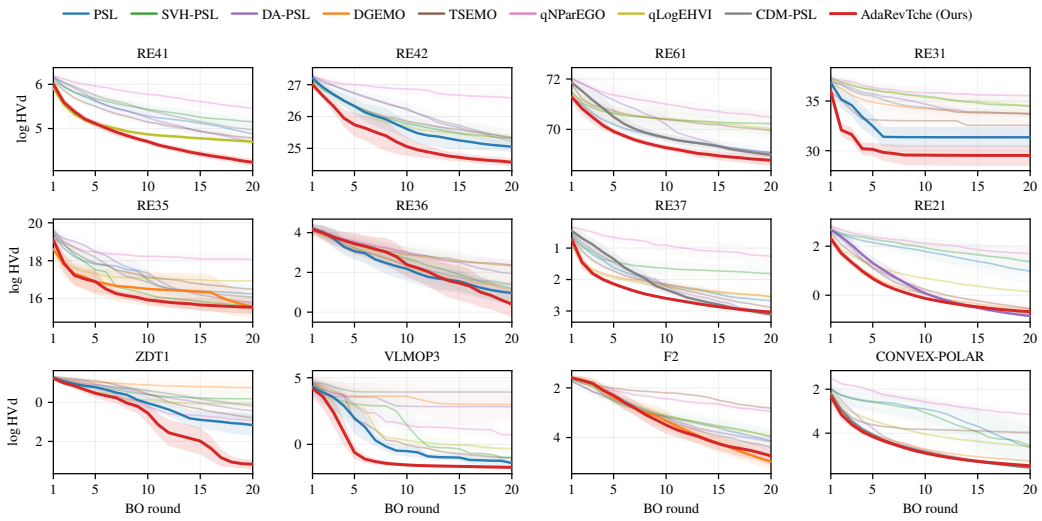

Figure 3: Log HV difference on 12 problems over 10 runs (lower is better). DGEMO is shown on the nine tasks supported by its released code. The proposed method and the strongest available competitor at the final round are opaque in each panel; other methods are de-emphasized but retain fixed colors across panels. Ribbons show the dispersion defined below Eq. (4.1).

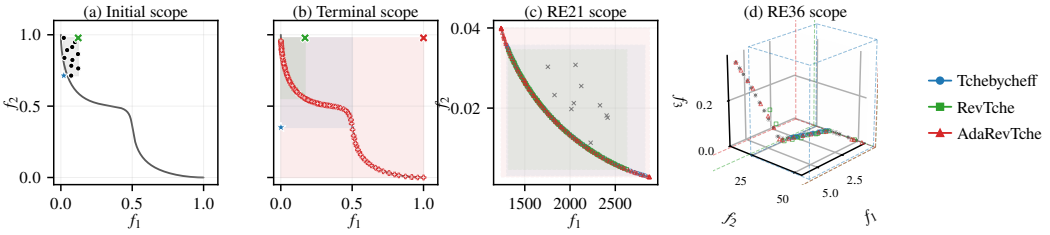

Figure 4: Progressive mechanism ablation. Panels (a)–(b) show the initial and terminal scopes on convex polar, while panels (c)–(d) show the terminal scopes on RE21 and RE36. Adaptive-reference scalarization achieves the broadest terminal coverage across all three problems.

On some relatively simple, low dimensional problems, such as F2 and RE21, AdaRevTche shows mixed performance relative to existing methods. Methods that outperform it in these cases do not explicitly learn the set model in Eq. (3), but instead directly optimize Pareto front quality, which is more closely aligned with the HV-based evaluation. Similar results on the custom diagnostic problem Convex-Polar, designed to validate the reference trap hypothesis, are therefore expected to be higher for non-Tchebycheff methods such as DA-PSL and DGEMO. More importantly, AdaRevTche consistently outperforms the original PSL-MOBO, demonstrating the benefit of adaptive reference adjustment within the preference conditioned PSL paradigm.

The main advantage of AdaRevTche is most evident on the more challenging real-world engineering problems RE41, RE42, and RE61, which involve higher-dimensional objective spaces and increasingly complex Pareto fron structures. By learning an adaptive reference point, AdaRevTche enlarges the reachable scope and improves exploration in insufficiently covered regions. This expanded search scope enables the method to discover diverse trade-off solutions that may remain inaccessible to methods using a fixed reference point. It achieves the lowest final mean HV difference among the available methods on all three tasks, reducing the HV difference relative to CDM-PSL by 51.82% on RE41, 52.35% on RE42, and 19.69% on RE61. These results demonstrate the particular value of adaptive reference adjustment as increasing objective dimensionality makes comprehensive Pareto front coverage more difficult.

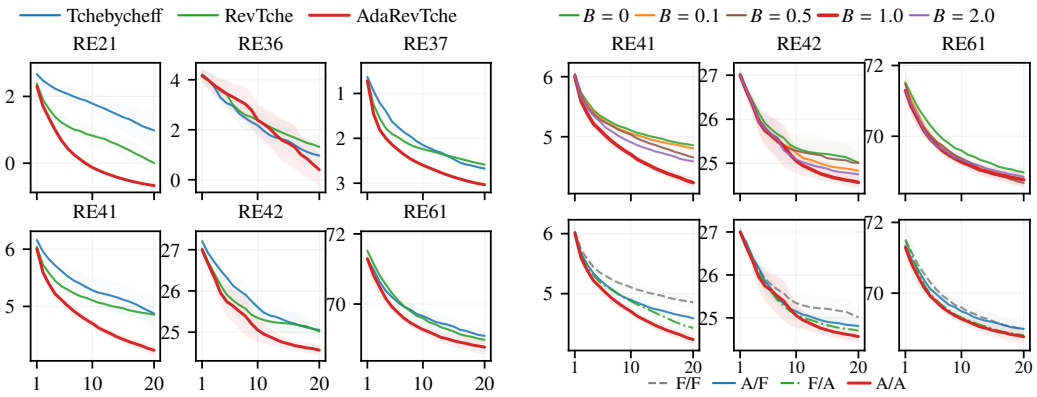

Figure 5: Ablations of adaptive reference geometry. (a, left) Same-budget comparison of Tchebycheff, fixed reverse Tchebycheff (RevTche), and AdaRevTche on six RE tasks. (b, right) Reference ablations on RE41, RE42, and RE61: the top row varies the adaptation range B, while the bottom row compares fixed (F) and adaptive (A) references for PSL training and HVI selection.

4.3 Further Experiments

Demonstration of the Reachable Scope of Different Scalarizations. We compare Tchebycheff (Tche), fixed-reference reverse Tchebycheff (RevTche), and AdaRevTche under localized initialization. Tche and RevTche remain confined to limited front regions, while adaptive reference updates progressively expand the reachable scope and recover uncovered trade-offs. The same behavior is observed on RE21 and RE36 under random initialization: fixed references restrict exploration, whereas AdaRevTche expands the reachable scope and improves front coverage.

Performance Comparisons of Different Scalarizations. Figure 5(a) provides comparisons on 6 engineering design problems. The adaptive method has the lowest final log HV difference on all six tasks. Fixed reverse scalarization is not uniformly sufficient, while adaptive reference updates enlarge the reachable scope during PSL training and subsequent HVI selection. Relative to fixed reverse scalarization, the final raw HV difference is reduced by approximately 49.1%, 59.8%, 36.3%, 46.3%, 36.2%, and 19.1% on RE21, RE36, RE37, RE41, RE42, and RE61, respectively. These indicate dynamic reference adjustment contributes to the improvement achieved by AdaRevTche.

Effect of Adaptive Range B. In the top row of Figure 5(b), we vary the adaptive range $B \in \{0, 0.1, 0.5, 1.0, 2.0\}$. The best value is task-dependent: $B = 1$ performs best on RE41 and RE42, whereas $B = 0.5$ performs best on RE61. When B is too small, the reference moves only slightly at each round and may not expand the exploration scope quickly enough within the fixed budget. A large value can move the reference too aggressively, causing substantial changes and reducing the effectiveness of the learned set model. We use $B = 1.0$ in the main experiments.

PSL-Training and HVI-Selection Reference Ablation. The bottom row of Figure 5(b) compares four configurations. F/F and A/A reuse the recorded fixed-reference and adaptive runs; A/F adapts only the PSL-training reference, whereas F/A adapts only the HVI-selection reference. Either individual change improves final mean hypervolume over F/F, with paired 95% confidence intervals excluding zero. A/A has the highest final mean hypervolume, although its gain on RE61 is within paired uncertainty. The negative interaction on RE41 shows the two gains are not simply additive.

5 Conclusion

We identify a reference trap in Tchebycheff-based Pareto set learning, where the archive-based ideal point can bias scalarized solutions. The proposed AdaRevTche addresses this by jointly learning the Pareto set and a worse-side reference, while adaptively expanding the reference via bounded objectivewise outward adaptation. Across 12 real-world and synthetic tasks, our AdaRevTche consistently improves Pareto front coverage, outperforms PSL-state-of-the-art methods.

AI USE STATEMENT

Generative AI tools are used to assist with language editing and improving the clarity and presentation of parts of the manuscript and visualization. They are not used to generate experimental results, or make scientific conclusions. All AI-assisted content is carefully reviewed and revised by the authors, who take full responsibility for the accuracy and integrity of the manuscript.

REPRODUCIBILITY STATEMENT

We provide detailed information to facilitate reproduction of our results. The supplementary material includes the complete theoretical proofs, benchmark definitions, experimental protocols, randomization and initialization settings, evaluation budgets, model architectures, optimization hyperparameters, surrogate-model configurations, evaluation metrics, and software and hardware environments. All methods are evaluated using matched initializations and identical evaluation budgets where applicable. The benchmark definitions, reference fronts, generated archives, training scripts, baseline implementations, and plotting utilities are maintained with the submission source and will be released publicly upon acceptance.

ETHICS STATEMENT

This work studies a general-purpose algorithm for expensive multi-objective optimization. The experiments are conducted on synthetic problems and publicly available benchmark problems and do not involve human participants, personal data, or private datasets. We are not aware of any immediate ethical concerns specific to the proposed method. As with general-purpose optimization methods, applications in safety-critical or socially sensitive domains may require additional domain-specific constraints and ethical considerations.

REFERENCES

Heinz H Bauschke and Patrick L Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. CMS Books in Mathematics/Ouvrages de Mathematiques de la SMC. Springer, New York, NY, 2011 edition, April 2011.

Arpan Biswas, Anna N Morozovska, Maxim Ziatdinov, Eugene A Eliseev, and Sergei V Kalinin. Multi-objective bayesian optimization of ferroelectric materials with interfacial control for memory and energy storage applications. *J. Appl. Phys.*, 130(20):204102, November 2021.

Eric Bradford, Artur M. Schweidtmann, and Alexei A. Lapkin. Efficient multiobjective optimization employing gaussian processes, spectral sampling and a genetic algorithm. *Journal of Global Optimization*, 71(2):407–438, 2018. doi: 10.1007/s10898-018-0609-2.

Haoying Chen, Yifan Wang, and Haibo Zhang. Multi-objective optimization design of adaptive cycle engine with serpentine 2-D exhaust system based on infrared stealth. *Aerospace*, 10(10): 858, September 2023.

Hanbin Cho, Hugon Lee, and Seunghwa Ryu. Sequential multi-objective bayesian optimization of air-cooled battery thermal management system with spoiler integration. *J. Energy Storage*, 114 (115586):115586, April 2025.

Tinkle Chugh, Karthik Sindhya, Jussi Hakanen, and Kaisa Miettinen. A survey on handling computationally expensive multiobjective optimization problems with evolutionary algorithms. *Soft Comput.*, 23(9):3137–3166, May 2019.

Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Differentiable expected hypervolume improvement for parallel multi-objective bayesian optimization. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS '20, Red Hook, NY, USA, 2020a. Curran Associates Inc. ISBN 9781713829546.

Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Differentiable expected hypervolume improvement for parallel multi-objective bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 33, pp. 9851–9864, 2020b.

Dirk I L de Villiers, Ivo Couckuyt, and Tom Dhaene. Multi-objective optimization of reflector antennas using kriging and probability of improvement. In *2017 IEEE International Symposium on Antennas and Propagation & USNC/URSI National Radio Science Meeting*. IEEE, July 2017.

John A Easum, Jogender Nagar, Pingjuan L Werner, and Douglas H Werner. Efficient multiobjective antenna optimization with tolerance analysis through the use of surrogate models. *IEEE Trans. Antennas Propag.*, 66(12):6706–6715, December 2018.

Michael T M Emmerich and André H Deutz. A tutorial on multiobjective optimization: fundamentals and evolutionary methods. *Nat. Comput.*, 17(3):585–609, May 2018.

Ashwin Gaonkar, Homero Valladares, Andres Tovar, Likun Zhu, and Hazim El-Mounayri. Multi-objective bayesian optimization of lithium-ion battery cells for electric vehicle operational scenarios. *Electronic Materials*, 3(2):201–217, May 2022.

Abhijith M Gopakumar, Prasanna V Balachandran, Dezhen Xue, James E Gubernatis, and Turab Lookman. Multi-objective optimization for materials discovery via adaptive design. *Sci. Rep.*, 8 (1):3738, February 2018.

F Irshad, S Karsch, and A Döpp. Multi-objective and multi-fidelity bayesian optimization of laser-plasma acceleration. *Phys. Rev. Res.*, 5(1), January 2023.

Shouyong Jiang and Shengxiang Yang. An improved multiobjective optimization evolutionary algorithm based on decomposition for complex pareto fronts. *IEEE Trans. Cybern.*, 46(2):421–437, February 2016.

Sukheon Kang, Hyeonbin Moon, Seonho Shin, Mahmoud Mousavi, Hyokyung Sung, and Seunghwa Ryu. Design of auxetic metamaterial for enhanced low cycle fatigue life and negative poisson's ratio through multi-objective bayesian optimization. *Mater. Des.*, 252(113798):113798, April 2025.

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 Konaković Luković, Yunsheng Tian, and Wojciech Matusik. Diversity-guided multi-objective bayesian optimization with batch evaluations. In *Advances in Neural Information Processing Systems*, volume 33, pp. 17708–17720, 2020.

Bingdong Li, Zixiang Di, Yongfan Lu, Hong Qian, Feng Wang, Peng Yang, Ke Tang, and Aimin Zhou. Expensive multi-objective bayesian optimization based on diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 27063–27071, 2025. doi: 10.1609/aaai.v39i25.34913.

Dong-Kuk Lim, Dong-Kyun Woo, Han-Kyeol Yeo, Sang-Yong Jung, Jongsuk Ro, and Hyun-Kyo Jung. A novel surrogate-assisted multi-objective optimization algorithm for an electromagnetic machine design. *IEEE Trans. Magn.*, 51(3):1–4, March 2015.

Xi Lin, Zhiyuan Yang, Xiaoyuan Zhang, and Qingfu Zhang. Pareto set learning for expensive multi-objective optimization. In *Advances in Neural Information Processing Systems*, volume 35, pp. 19231–19247, 2022.

Yongfan Lu, Bingdong Li, and Aimin Zhou. 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, pp. 14202–14210, 2024. doi: 10.1609/aaai.v38i13.29331.

Anirudh M K Nambiar, Christopher P Breen, Travis Hart, Timothy Kulesza, Timothy F Jamison, and Klavs F Jensen. Bayesian optimization of computer-proposed multistep synthetic routes on an automated robotic flow platform. *ACS Cent. Sci.*, 8(6):825–836, June 2022.

Aviv Navon, Aviv Shamsian, Ethan Fetaya, and Gal Chechik. Learning the pareto front with hypernetworks. In *International Conference on Learning Representations*, 2021.

Minh-Duc Nguyen, Phuong Mai Dinh, Quang-Huy Nguyen, Long P. Hoang, and Dung D. Le. Improving pareto set learning for expensive multi-objective optimization via stein variational hypernetworks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 19677–19685, 2025. doi: 10.1609/aaai.v39i18.34167.

R T Rockafellar and Roger J-B Wets. *Variational Analysis*. Grundlehren der mathematischen Wissenschaften. Springer, Berlin, Germany, 1 edition, November 1997.

Hiroyuki Sato. Inverted PBI in MOEA/D and its impact on the search performance on multi and many-objective optimization. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*, New York, NY, USA, July 2014a. ACM.

Hiroyuki Sato. Inverted PBI in MOEA/D and its impact on the search performance on multi and many-objective optimization. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*, New York, NY, USA, July 2014b. ACM.

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. *Proc. IEEE Inst. Electr. Electron. Eng.*, 104(1):148–175, January 2016.

Alexandros Solomou, Guang Zhao, Shahin Boluki, Jobin K Joy, Xiaoning Qian, Ibrahim Karaman, Raymundo Arróyave, and Dimitris C Lagoudas. Multi-objective bayesian materials discovery: Application on the discovery of precipitation strengthened NiTi shape memory alloys through micromechanical modeling. *Mater. Des.*, 160:810–827, December 2018.

Ralph E Steuer and Eng-Ung Choo. An interactive weighted tchebycheff procedure for multiple objective programming. *Math. Program.*, 26(3):326–344, October 1983.

Ryoji Tanabe and Hisao Ishibuchi. An easy-to-use real-world multi-objective optimization problem suite. *Applied Soft Computing*, 89:106078, 2020. ISSN 1568-4946. doi: https://doi.org/10.1016/j.asoc.2020.106078. URL https://www.sciencedirect.com/science/article/pii/S1568494620300181.

David A. van Veldhuizen and Gary B. Lamont. Multiobjective evolutionary algorithm test suites. In *Proceedings of the 1999 ACM Symposium on Applied Computing*, SAC '99, pp. 351–357, New York, NY, USA, 1999. Association for Computing Machinery. ISBN 1581130864. doi: 10.1145/298151.298382. URL https://doi.org/10.1145/298151.298382.

Yongfeng Wei, Guangfei Qi, Yanxing Wang, Ningchaoran Yan, Yongliang Zhang, and Linping Feng. Efficient microwave filter design by a surrogate-model-assisted decomposition-based multi-objective evolutionary algorithm. *Electronics (Basel)*, 11(20):3309, October 2022.

Andrzej P Wierzbicki. A mathematical basis for satisficing decision making. *Math. Model.*, 3(5): 391–405, 1982a.

Andrzej P. Wierzbicki. A mathematical basis for satisficing decision making. *Mathematical Modelling*, 3(5):391–405, 1982b. doi: 10.1016/0270-0255(82)90038-0.

Yin Wu, Ruihao Zheng, and Zhenkun Wang. Decomposition-based multi-objective evolutionary algorithm with model-based ideal point estimation. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 768–776, New York, NY, USA, July 2023a. ACM.

Yin Wu, Ruihao Zheng, and Zhenkun Wang. Decomposition-based multi-objective evolutionary algorithm with model-based ideal point estimation. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 768–776, New York, NY, USA, July 2023b. ACM.

S Xiao, G Q Liu, K L Zhang, Y Z Jing, J H Duan, P Di Barba, and J K Sykulski. Multi-objective pareto optimization of electromagnetic devices exploiting kriging with lipschitzian optimized expected improvement. *IEEE Trans. Magn.*, 54(3):1–4, March 2018.

Yunjia Zeng, Xianming Qing, and Michael Yan-Wah Chia. A wideband circularly polarized antenna with a nonuniform metasurface designed via multiobjective bayesian optimization. *IEEE Antennas Wirel. Propag. Lett.*, 23(6):1739–1743, June 2024.

Jiyizhe Zhang, Naoto Sugisawa, Kobi C Felton, Shinichiro Fuse, and Alexei A Lapkin. Multi-objective bayesian optimisation using q -noisy expected hypervolume improvement (q NEHVI) for the Schotten–Baumann reaction. *React. Chem. Eng.*, 9(3):706–712, 2024.

Qingfu Zhang, Wudong Liu, Edward Tsang, and Botond Virginas. Expensive multiobjective optimization by MOEA/D with gaussian process model. *IEEE Trans. Evol. Comput.*, 14(3):456–474, June 2010.

Saile Zhang, Qingzhen Yang, Rui Wang, and Xufei Wang. Multi-objective bayesian optimization design of elliptical double serpentine nozzle. *Aerospace*, 11(1):48, December 2023.

Yiming Zhang, Sayan Ghosh, Thomas Vandeputte, and Liping Wang. Bayesian optimization for multi-objective high-dimensional turbine aero design. In *Volume 9B: Structures and Dynamics — Fatigue, Fracture, and Life Prediction; Probabilistic Methods; Rotordynamics; Structural Mechanics and Vibration*. American Society of Mechanical Engineers, June 2021.

Peipei Zhao, Liping Wang, and Qicang Qiu. Preference-based expensive multi-objective optimization without using an ideal point. *Complex Intell. Syst.*, 11(7), July 2025.

Eckart Zitzler and Lothar Thiele. Multiobjective evolutionary algorithms: A comparative case study and the strength pareto approach. *IEEE Transactions on Evolutionary Computation*, 3(4):257–271, 1999.

Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary Computation*, 8(2):173–195, 2000. doi: 10.1162/106365600568202.

SUPPLEMENTARY MATERIALS

This supplementary material is organized as follows.

In Section A, we provide a formal statement and proof of Proposition 3.1 in Section 3 of the main paper.

In Section B, we depict more details about benchmark problems in 4.1 of the main paper.

In Section C, we offer more results on adaptive Tchebycheff reference in Section 3.3 of the main paper.

In Section D, we report more results about hypervolume comparisons in Section 4.2 in the main paper.

In Section E, we provide more records about implementations and reproducibility in Section 4.1 in the main paper.

A FORMAL STATEMENT AND PROOF OF PROPOSITION 3.1 OF SECTION 3

This section develops the geometric preliminaries of archive-induced attraction and then establishes the reference-trap proposition stated in Section 3.

A.1 PRELIMINARY RESULTS

Setting and assumptions. Consider

$$\min_{x\in\mathcal{X}} f(x) = \big(f_1(x),\dots,f_m(x)\big), \qquad \mathcal{Y} = f(\mathcal{X}) \subseteq \mathbb{R}^m, \qquad m \geq 2. \tag{15}$$

At round t, let $\mathcal{D}_t = \{(x_j, y_j)\}_{j=1}^{N_t}$, where $y_j = f(x_j)$, and set $\mathcal{Y}_t = \{y_j : (x_j, y_j) \in \mathcal{D}_t\}$. Throughout, $\mathcal{Y}_t$ is assumed finite, nonempty, and contained in $\mathcal{Y}$. Define

$$\hat{z}^*_{t,i} = \min_{y\in\mathcal{Y}_t} y_i, \qquad \overline{y}_{t,i} = \max_{y\in\mathcal{Y}_t} y_i, \qquad \mathcal{B}_t^{\text{scope}} = \prod_{i=1}^{m}(-\infty, \overline{y}_{t,i}]. \tag{16}$$

Let $\Delta_{++}^{m-1} = \{\lambda \in \mathbb{R}^m_{++} : \mathbf{1}^\top\lambda = 1\}$. For $\lambda \in \Delta_{++}^{m-1}$, following (Wierzbicki, 1982a; Steuer & Choo, 1983), we define

$$T_t(y;\lambda) = T(y;\lambda,\hat{z}^*_t) = \max_{1\leq i\leq m} \lambda_i\big(y_i - \hat{z}^*_{t,i}\big). \tag{17}$$

For a nonempty $\Lambda \subseteq \Delta_{++}^{m-1}$, suppose that $T_t(\cdot;\lambda)$ attains its minimum over $\mathcal{Y}$ for every $\lambda \in \Lambda$, and define $\mathcal{R}_t^T(\Lambda) = \bigcup_{\lambda\in\Lambda} \arg\min_{y\in\mathcal{Y}} T_t(y;\lambda)$. An *archive-retaining oracle update* is an update of the form $\mathcal{Y}_{t+1} = \mathcal{Y}_t \cup \mathcal{N}_t$, where $\mathcal{N}_t \subseteq \mathcal{R}_t^T(\Lambda)$.

Proposition A.1 (Bi-objective containment)**.** *Assume $m = 2$, and let $\Lambda \subseteq \Delta_{++}^1$ satisfy the attainment assumption above. Then*

$$\mathcal{R}_t^T(\Lambda) \subseteq \mathcal{Y} \cap \mathcal{B}_t^{\text{scope}}. \tag{18}$$

Consequently, every archive-retaining oracle update satisfies $\mathcal{B}_{t+1}^{\text{scope}} = \mathcal{B}_t^{\text{scope}}$.

Proof. Fix $\lambda \in \Lambda$. Choose $a^{(1)}, a^{(2)} \in \mathcal{Y}_t$ such that $a_i^{(i)} = \hat{z}^*_{t,i}$; these points exist because $\mathcal{Y}_t$ is finite and nonempty. If $y_2 > \overline{y}_{t,2}$, then

$$T_t(a^{(1)};\lambda) \leq \lambda_2(\overline{y}_{t,2} - \hat{z}^*_{t,2}) < \lambda_2(y_2 - \hat{z}^*_{t,2}) \leq T_t(y;\lambda),$$

so y cannot be a global minimizer. Interchanging the two coordinates excludes $y_1 > \overline{y}_{t,1}$. Thus every minimizer lies in $\mathcal{B}_t^{\text{scope}}$, and taking the union over $\lambda \in \Lambda$ proves (18). Archive retention preserves the old maxima, while the new points do not exceed them. □

Remark 1. *The preceding containment is a two-dimensional manifestation of a more general geometry. For each preference, the best archived scalarized value defines an ideal-side box; negative Tchebycheff subgradients outside that box point strictly toward it.*

Definition 1 (Preference-conditioned attraction box)**.** *Fix t, a finite nonempty archive $\mathcal{Y}_t$, and $\lambda \in \Delta_{++}^{m-1}$. Define*

$$\tau_t(\lambda) = \min_{y \in \mathcal{Y}_t} T_t(y;\lambda), \qquad b_{t,i}(\lambda) = \hat{z}^*_{t,i} + \frac{\tau_t(\lambda)}{\lambda_i}, \qquad \mathcal{A}_t(\lambda) = \prod_{i=1}^{m} (-\infty, b_{t,i}(\lambda)]. \tag{19}$$

The set $\mathcal{A}_t(\lambda)$ is the preference-conditioned attraction box. *Since $\lambda_i > 0$ for every i,*

$$y \in \mathcal{A}_t(\lambda) \quad \Longleftrightarrow \quad T_t(y;\lambda) \le \tau_t(\lambda). \tag{20}$$

Following (Rockafellar & Wets, 1997; Bauschke & Combettes, 2011), for a closed convex set C, let $\Pi_C(y)$ denote the Euclidean projection of y onto C. For a vector v, let $[v]_+$ denote its coordinate-wise positive part, and let ∂_y denote the convex subdifferential with respect to y.

Theorem 2 (Archive-induced attraction)**.** *Fix t, a finite nonempty archive $\mathcal{Y}_t \subseteq \mathcal{Y}$, and $\lambda \in \Delta_{++}^{m-1}$.*

1. *For every $y \notin \mathcal{A}_t(\lambda)$ and every $q \in \partial_y T_t(y;\lambda)$,*

$$\big\langle y - \Pi_{\mathcal{A}_t(\lambda)}(y), -q \big\rangle = -\big(T_t(y;\lambda) - \tau_t(\lambda)\big) < 0. \tag{21}$$

   Hence $-q$ is a strict descent direction for $\frac{1}{2}\operatorname{dist}^2(y, \mathcal{A}_t(\lambda))$.

2. *If, in addition, $T_t(\cdot;\lambda)$ attains its minimum over $\mathcal{Y}$, then*

$$\arg\min_{y \in \mathcal{Y}} T_t(y;\lambda) \subseteq \mathcal{Y} \cap \mathcal{A}_t(\lambda). \tag{22}$$

Proof. Let $I_t(y;\lambda) = \arg\max_i \lambda_i(y_i - \hat{z}^*_{t,i})$. Every $q \in \partial_y T_t(y;\lambda)$ has the form

$$q = \sum_{i \in I_t(y;\lambda)} \alpha_i \lambda_i e_i, \qquad \alpha_i \ge 0, \qquad \sum_{i \in I_t(y;\lambda)} \alpha_i = 1.$$

If $y \notin \mathcal{A}_t(\lambda)$, then (20) gives $T_t(y;\lambda) > \tau_t(\lambda)$. For every active i, $\lambda_i[y_i - b_{t,i}(\lambda)]_+ = T_t(y;\lambda) - \tau_t(\lambda)$, while $y - \Pi_{\mathcal{A}_t(\lambda)}(y) = [y - b_t(\lambda)]_+$. Therefore

$$\big\langle y - \Pi_{\mathcal{A}_t(\lambda)}(y), -q \big\rangle = -\sum_{i \in I_t(y;\lambda)} \alpha_i \lambda_i [y_i - b_{t,i}(\lambda)]_+ = -\big(T_t(y;\lambda) - \tau_t(\lambda)\big).$$

This proves (21). The descent statement follows because the squared distance to a closed convex set has gradient $y - \Pi_{\mathcal{A}_t(\lambda)}(y)$.

For the second claim, take $y_t^*(\lambda) \in \arg\min_{y \in \mathcal{Y}} T_t(y;\lambda)$. Since $\mathcal{Y}_t \subseteq \mathcal{Y}$, $T_t(y_t^*(\lambda);\lambda) \le \tau_t(\lambda)$. Equation (20) now yields $y_t^*(\lambda) \in \mathcal{A}_t(\lambda)$. □

Lemma 3 (Attraction-box non-expansion criterion)**.** *Fix t, a finite nonempty archive $\mathcal{Y}_t \subseteq \mathcal{Y}$, and $\lambda \in \Delta_{++}^{m-1}$. Let $r_{t,i} = \overline{y}_{t,i} - \hat{z}^*_{t,i}$, and define*

$$J_t(\lambda) = \{i \in \{1,\dots,m\} : \tau_t(\lambda) \le \lambda_i r_{t,i}\}. \tag{23}$$

Then $J_t(\lambda) \neq \varnothing$, and every $y \in \mathcal{A}_t(\lambda)$ satisfies $y_i \le \overline{y}_{t,i}$ for all $i \in J_t(\lambda)$. Moreover,

$$J_t(\lambda) = \{1,\dots,m\} \quad \Longleftrightarrow \quad \tau_t(\lambda) \le \min_i \lambda_i r_{t,i} \quad \Longleftrightarrow \quad \mathcal{A}_t(\lambda) \subseteq \mathcal{B}_t^{\text{scope}}. \tag{24}$$

Proof. Every $y \in \mathcal{Y}_t$ satisfies $0 \le y_i - \hat{z}^*_{t,i} \le r_{t,i}$, so $\tau_t(\lambda) \le \max_i \lambda_i r_{t,i}$. Hence every maximizer of $i \mapsto \lambda_i r_{t,i}$ belongs to $J_t(\lambda)$, and $J_t(\lambda) \neq \varnothing$.

For each i, positivity of λ_i gives

$$b_{t,i}(\lambda) \le \overline{y}_{t,i} \quad \Longleftrightarrow \quad \tau_t(\lambda) \le \lambda_i r_{t,i} \quad \Longleftrightarrow \quad i \in J_t(\lambda).$$

The coordinate-wise assertion and (24) follow immediately. □

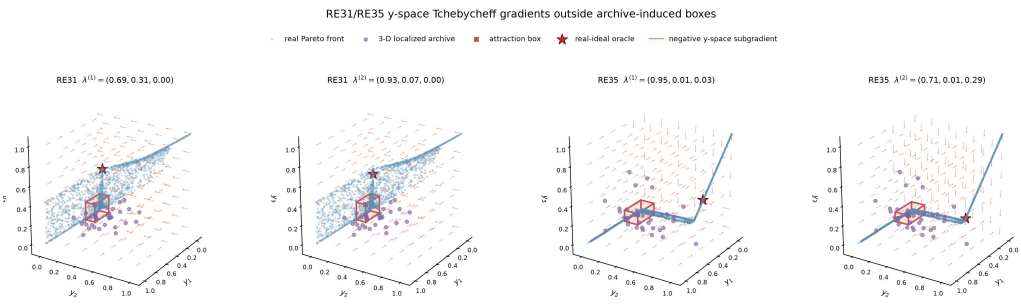

Figure 6: Archive-induced attraction on RE31 and RE35, illustrating the strict inward-pointing property in Theorem 2. The four horizontal panels show two preferences for each problem. Blue points denote the real Pareto front, purple points the dispersed three-dimensional archive, and the red wireframe with translucent faces the preference-conditioned attraction box. Orange arrows are fixed-length negative y-space Tchebycheff subgradients evaluated on a regular exterior lattice; their directions show that exterior points are driven toward the corresponding attraction box. The opaque red star marks the real-ideal oracle outside the box, highlighting the region not covered by the archive-induced attraction geometry.

A.2 Reference Trap

In this subsection, we provide the complete statement and a graphical illustration of Proposition 3.1 in the main paper.

Proposition A.2 (Reference trap)**.** *Fix t, a finite nonempty archive $\mathcal{Y}_t \subseteq \mathcal{Y}$, and a nonempty $\Lambda \subseteq \Delta_{++}^{m-1}$. Assume that $T_t(\cdot;\lambda)$ attains its minimum over $\mathcal{Y}$ for every $\lambda \in \Lambda$. Then*

$$\mathcal{R}_t^T(\Lambda) \subseteq \mathcal{Y} \cap \bigcup_{\lambda\in\Lambda} \mathcal{A}_t(\lambda), \tag{25}$$

and, for each λ, every negative subgradient outside $\mathcal{A}_t(\lambda)$ points strictly toward that box.

Furthermore, every oracle minimizer $y_t^*(\lambda)$ satisfies $y_{t,i}^*(\lambda) \le \overline{y}_{t,i}$ for every $i \in J_t(\lambda)$, where $J_t(\lambda) \neq \varnothing$. If

$$\tau_t(\lambda) \le \min_i \lambda_i\big(\overline{y}_{t,i} - \hat{z}_{t,i}^*\big) \qquad \text{for every } \lambda \in \Lambda, \tag{26}$$

then every archive-retaining oracle update satisfies $\mathcal{B}_{t+1}^{\text{scope}} = \mathcal{B}_t^{\text{scope}}$. When $m = 2$, condition (26) holds automatically; hence Proposition A.1 is exactly the bi-objective special case of this proposition.

Proof. Theorem 2 (ii), applied to each $\lambda \in \Lambda$, gives (25); part (i) gives the inward-pointing statement. Combining (22) with Lemma 3 gives the protected-coordinate conclusion.

If (26) holds, then Lemma 3 yields $\mathcal{A}_t(\lambda) \subseteq \mathcal{B}_t^{\text{scope}}$ for every $\lambda \in \Lambda$. Theorem 2 (ii) therefore gives $\mathcal{R}_t^T(\Lambda) \subseteq \mathcal{B}_t^{\text{scope}}$. An archive-retaining update adds no point beyond the old coordinate-wise maxima, so its worse-side scope is invariant.

It remains to verify the two-dimensional specialization. Fix $i \in \{1,2\}$, let $k \neq i$, and choose $a^{(k)} \in \mathcal{Y}_t$ with $a_k^{(k)} = \hat{z}_{t,k}^*$. Since there is only one nonideal coordinate,

$$\tau_t(\lambda) \le T_t(a^{(k)};\lambda) = \lambda_i(a_i^{(k)} - \hat{z}_{t,i}^*) \le \lambda_i(\overline{y}_{t,i} - \hat{z}_{t,i}^*).$$

This holds for both coordinates, so (26) is automatic when $m = 2$. The resulting containment and scope invariance are precisely the conclusions of Proposition A.1. □

Remark 2 (Objective-space statement)**.** *All sets, subgradients, and updates above are defined in objective space, that is, in y-space. An inward direction $-q$ in y-space need not be realized by a descent step in x, or in the parameters of a Pareto-set model, without additional assumptions on the relevant objective or surrogate Jacobian. The results are geometric statements, not convergence claims in decision or parameter space.*

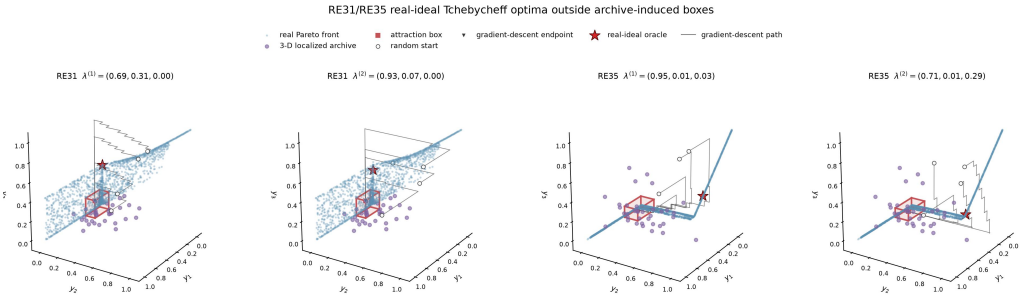

Figure 7: Reference-trap behavior from Proposition A.2 on RE31 and RE35. The four horizontal panels show two preferences for each problem. Blue points denote the real Pareto front, purple points the dispersed archive, and the red wireframe with translucent faces the archive-induced attraction box. Thin gray paths start from random exterior points and follow negative Tchebycheff subgradients to marked endpoints on the box; the resulting trajectories illustrate how the attraction boxes constrain scalarized optimization to archive-induced regions. Open circles and downward triangles denote the starts and endpoints, respectively. The opaque red star marks the real-ideal oracle outside every displayed box, illustrating the uncovered region that the reference trap prevents these trajectories from reaching.

Table 1: Problem inventory and fixed objective-space references for the main-result comparison. All objectives are minimized; n is the number of decision variables, m is the number of objectives, H^* is the ground-truth hypervolume, and the fixed reference is used only for reporting.

Problem	Family	n	m	Real H^*	Reference
F2	synthetic	6	2	0.871	$(1.1, 1.1)$
VLMOP3	synthetic	2	3	672	$(11, 66, 1.1)$
ZDT1	synthetic	30	2	11.8	$(1.1, 11)$
RE21	RE engineering	4	2	60.2	$(3175.01, 0.044)$
RE31	RE engineering	3	3	1.06×10^{17}	$(550.003, 9.071 \times 10^6, 2.130 \times 10^7)$
RE35	RE engineering	7	3	2.75×10^9	$(7808.40, 1865.56, 437.095)$
RE36	RE engineering	4	3	96.4	$(6.5241, 61.6, 0.3913)$
RE37	RE engineering	4	3	1.46	$(1.1022, 1.2073, 1.2032)$
RE41	RE engineering	7	4	1.13×10^3	$(47.0448, 4.8700, 14.4005, 10.3942)$
RE42	RE engineering	6	4	9.13×10^{11}	$(-597.074, 1.739 \times 10^4, 5.715 \times 10^3, 14.329)$
RE61	RE engineering	3	6	4.90×10^{31}	$(8.398 \times 10^4, 1.485 \times 10^3, 3.139 \times 10^6, 1.730 \times 10^7, 3.814 \times 10^5, 1.032 \times 10^5)$
Convex polar	diagnostic	2	2	0.900	$(1.1, 1.1)$

B More Details about Benchmark Problems in Section 4.1

This section summarizes the benchmark problems and fixed objective-space references used in main paper's section 4.1, and then details the Convex Polar problem introduced to isolate the reference-trap phenomenon.

B.1 Widely Used Synthetic and Real-World Problems

The main comparison includes three widely used synthetic benchmarks—F2, VLMOP3, and ZDT1—and eight real-world engineering problems from the RE suite: RE21, RE31, RE35, RE36, RE37, RE41, RE42, and RE61. Their dimensions, ground-truth hypervolumes, and fixed reporting references are summarized in Table 1. All objectives are minimized. For the RE problems, the original physical bounds are retained, and constraint violations are represented as additional minimization objectives.

B.2 Proposed Convex Polar Problem

Convex Polar. Let $x = (u, d) \in [0, 1]^2$, $\theta = \pi u/2$, $A_{\mathrm{v}} = 0.24$, $w_{\mathrm{v}} = 0.20$, and $R_{\mathrm{in}} = 0.48$. We define

$$
\begin{aligned}
r_{\mathrm{v}}(\theta) &= 1 - A_{\mathrm{v}} \exp\left[-\frac{1}{2}\left(\frac{\theta - \pi/4}{w_{\mathrm{v}}}\right)^2\right] \sin^2(2\theta), \\
r_{\mathrm{v}}(\theta, d) &= r_{\mathrm{v}}(\theta) - d\big(r_{\mathrm{v}}(\theta) - R_{\mathrm{in}}\big), \\
f_{\mathrm{v}}(u, d) &= \big(1 - r_{\mathrm{v}}(\theta, d)\cos\theta,\ 1 - r_{\mathrm{v}}(\theta, d)\sin\theta\big).
\end{aligned} \tag{27}
$$

Its exact Pareto front is obtained at $d = 0$:

$$
\mathcal{P}_{\mathrm{v}} = \left\{\big(1 - r_{\mathrm{v}}(\theta)\cos\theta,\ 1 - r_{\mathrm{v}}(\theta)\sin\theta\big) : 0 \le \theta \le \frac{\pi}{2}\right\}. \tag{28}
$$

Every point with $d > 0$ is dominated by the point with the same θ and $d = 0$. The Pareto front is therefore available analytically and can be sampled directly, without surrogate approximation or numerical multi-objective optimization. Convex Polar is deliberately simple: it introduces no additional optimization difficulty, so differences in front coverage can be attributed directly to the reference geometry.

For the exact-front diagnostic, we use a 16×16 grid of localized initialization cells, with ten independent draws per cell. At each of the 20 rounds, candidates are generated directly from a dense representation of $\mathcal{P}_{\mathrm{v}}$, and five are selected by greedy hypervolume improvement. Coverage is measured as

$$
\mathrm{Cov}_\tau(\mathcal{A}, \mathcal{P}) = \frac{1}{|\mathcal{P}|}\left|\left\{p \in \mathcal{P} : \min_{a \in \mathcal{A}} \|\widetilde{p} - \widetilde{a}\|_2 \le \tau\right\}\right|, \qquad \tau = 0.025, \tag{29}
$$

where the tilde denotes coordinate-wise min–max normalization using the dense exact front.

C More Results on Adaptive Tchebycheff Reference

To complement the reachable-scope analysis in Section 3.3 of the main paper, we consider whether the reference trap can be alleviated by directly learning the ideal-side reference of conventional Tchebycheff scalarization (Wu et al., 2023a). We call this variant *AdaTche* and compare it with fixed Tchebycheff and AdaRevTche under the same RE21 initialization, surrogate, preference samples, candidate budget, and batch-selection rule.

For $T(y; \lambda, z) = \max_i \lambda_i(y_i - z_i)$, let $I_T = \arg\max_i \lambda_i(y_i - z_i)$. Its reference subdifferential is $\partial_z T = \mathrm{conv}\{-\lambda_i e_i : i \in I_T\}$. If the active coordinate is unique, $I_T = \{k\}$, a gradient-descent step gives $z^+ = z + \eta\lambda_k e_k$.

Let $c = T(y; \lambda, z)$ and $u_i = z_i + c/\lambda_i$, so that the corresponding sublevel set is $\prod_i(-\infty, u_i]$. While coordinate k remains active, the update satisfies

$$
u_k^+ = u_k, \qquad u_j^+ = u_j - \frac{\eta\lambda_k^2}{\lambda_j} < u_j \quad (j \ne k).
$$

Thus, learning the ideal-side reference preserves the active intercept but contracts the inactive ones. Active-set changes may redistribute this contraction, but do not systematically expand the worse-side boundary. In contrast, for the reverse scalarization, reference ascent moves the active worse-side coordinate outward more strongly because $\partial S_{\mathrm{rev}}/\partial r_{k_{\mathrm{R}}} = 1/\lambda_{k_{\mathrm{R}}} + \rho$, whereas $\partial S_{\mathrm{rev}}/\partial r_j = \rho$ for $j \ne k_{\mathrm{R}}$ (Sato, 2014b). This difference explains why adapting an ideal-side reference alone does not fully remove the reference trap.

Figure 8 reports ten paired runs on RE21 using 20 BO rounds and the coverage metric in Eq. (29). Fixed Tchebycheff and AdaTche achieve mean final HV values of 48.63 and 52.31, with coverages of 0.36 and 0.52, respectively. AdaRevTche achieves a mean final HV of 59.64 and full coverage. These results confirm that adapting the ideal-side reference is beneficial, but the worse-side adaptation of AdaRevTche provides substantially stronger scope expansion.

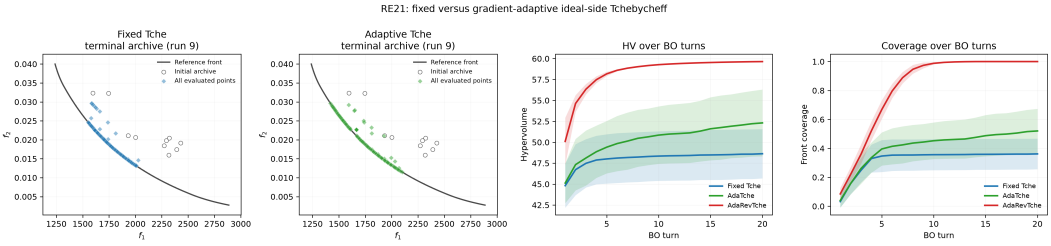

Figure 8: Comparison on RE21 under paired initializations and identical evaluation settings. The first two panels show representative terminal archives, followed by the mean HV and front-coverage trajectories over ten runs. Shaded regions indicate one standard deviation.

Table 2: Final hypervolume (HV), reported as mean with standard deviation shown below, over independent runs. Values are rounded to three significant digits. Boldface indicates the best mean for each problem.

Problem	AdaRevTche	PSL	SVH-PSL	DA-PSL	DGEMO	TSEMO	qNParEGO	qLogEHVI	CDM-PSL
F2	0.863 $\pm$ 0.00404	0.855 $\pm$ 0.00878	0.853 $\pm$ 0.00692	0.856 $\pm$ 0.00295	**0.864** $\pm$ **0.0024**	0.811 $\pm$ 0.0152	0.819 $\pm$ 0.00646	0.852 $\pm$ 0.0031	0.859 $\pm$ 0.0035
VLMOP3	**672** $\pm$ **0.0225**	671 $\pm$ 0.0744	671 $\pm$ 0.286	655 $\pm$ 50.4	652 $\pm$ 49.7	671 $\pm$ 0.0986	670 $\pm$ 1.09	671 $\pm$ 0.204	621 $\pm$ 77
RE21	59.6 $\pm$ 0.0148	57.5 $\pm$ 1.94	56.2 $\pm$ 2.89	**59.7** $\pm$ **0.0273**	59.6 $\pm$ 0.177	59.6 $\pm$ 0.0371	54.7 $\pm$ 1.81	59 $\pm$ 0.132	59.7 $\pm$ 0.0344
RE31	**1.06e17** $\pm$ **1.25e13**	1.06e17 $\pm$ 7.61e13	1.05e17 $\pm$ 1.6e15	1.06e17 $\pm$ 2.2e14	1.06e17 $\pm$ 9.12e14	1.06e17 $\pm$ 3.96e14	1.04e17 $\pm$ 2.72e15	1.05e17 $\pm$ 6.66e14	1.06e17 $\pm$ 2.1e14
RE35	2.74e09 $\pm$ 2.39e06	2.73e09 $\pm$ 7.09e06	2.74e09 $\pm$ 3.74e06	2.74e09 $\pm$ 2.38e06	**2.74e09** $\pm$ **3.91e06**	2.73e09 $\pm$ 2.58e06	2.68e09 $\pm$ 2.56e07	2.72e09 $\pm$ 5.06e06	2.74e09 $\pm$ 4.3e06
RE36	**95** $\pm$ **1.43**	93.8 $\pm$ 1.98	93.2 $\pm$ 4.93	89.5 $\pm$ 6.38	93.4 $\pm$ 3.08	85.8 $\pm$ 3.08	85.9 $\pm$ 4.26	86.4 $\pm$ 3.99	92.5 $\pm$ 4.45
RE37	1.41 $\pm$ 0.00194	1.39 $\pm$ 0.00768	1.29 $\pm$ 0.0453	1.41 $\pm$ 0.00768	1.38 $\pm$ 0.0118	1.4 $\pm$ 0.00576	1.17 $\pm$ 0.101	1.38 $\pm$ 0.00529	**1.41** $\pm$ **0.00403**
RE41	**1.06e03** $\pm$ **6.04**	1e03 $\pm$ 47.7	961 $\pm$ 24.1	1.02e03 $\pm$ 10.3	–	1.01e03 $\pm$ 5.68	900 $\pm$ 46.2	1.02e03 $\pm$ 4.42	990 $\pm$ 15.8
RE42	**8.66e11** $\pm$ **5.78e09**	8.37e11 $\pm$ 1.42e10	8.14e11 $\pm$ 3.93e10	8.23e11 $\pm$ 2.31e10	–	8.11e11 $\pm$ 8.76e09	5.61e11 $\pm$ 9.8e10	8.13e11 $\pm$ 9.22e09	8.15e11 $\pm$ 2.98e10
RE61	**4.83e31** $\pm$ **1.83e29**	4.8e31 $\pm$ 3.01e29	4.59e31 $\pm$ 2.37e30	4.8e31 $\pm$ 1.79e29	–	4.66e31 $\pm$ 5.47e29	4.49e31 $\pm$ 9.95e29	4.64e31 $\pm$ 8.18e29	4.81e31 $\pm$ 1.14e29
ZDT1	**11.7** $\pm$ **0.0118**	11.5 $\pm$ 0.229	10.6 $\pm$ 0.642	11.1 $\pm$ 0.278	9.64 $\pm$ 1.36	11.3 $\pm$ 0.0909	11.4 $\pm$ 0.169	10.8 $\pm$ 0.287	10.9 $\pm$ 0.25
Convex Polar	0.896 $\pm$ 0.000137	0.891 $\pm$ 0.0156	0.89 $\pm$ 0.0181	0.896 $\pm$ 0.000225	0.895 $\pm$ 0.000317	0.881 $\pm$ 0.00228	0.857 $\pm$ 0.00974	0.89 $\pm$ 0.00067	**0.897** $\pm$ **0.000145**

D More Results about Hypervolume Comparisons

In Section 4.2 of the main paper, we report the hypervolume comparison across methods using line plots. This section provides the corresponding numerical results, including the final mean and standard deviation over ten runs for each method–problem pair.

E More Details about Implementations and Reproducibility

This section records the information requested by the reproducibility checklist. The executable training, baseline, benchmark, and plotting scripts are maintained with the submission source; they are not reproduced inline in this appendix and will be released publicly upon acceptance.

E.1 Randomness, Initialization, and Evaluation Budget

The base random seed is 42. Ten independent run seeds are generated with NumPy's `SeedSequence` API using `spawn(10)`; the resulting child sequences control the initial archive, preference sampling, and candidate-generation randomness for the corresponding run. The PSL map uses the shared initialization seed 11 so that paired comparisons do not introduce an additional network-initialization difference. The same run-specific initial archive is supplied to all methods in a paired comparison. Each archive contains ten Latin-hypercube points, followed by $T = 20$ Bayesian-optimization rounds with batch size $b = 5$, for a total of $10 + 20 \times 5 = 110$ true objec-

Table 3: Fixed configuration for the reported AdaRevTche experiments.

Component	Setting
PSL map	two-layer MLP, 256 ReLU units per layer
PSL optimizer	Adam, learning rate 10^{-3}, 1,000 steps per BO round
Preferences	$q = 10$ Dirichlet$(\mathbf{1})$ vectors per update
LCB coefficient	$\kappa = 0.1$
Reverse tie breaker	$\rho = 0.01$
Reference regularization	$\beta = 0.1$
Adaptive range	$B = 1.0$
Reference-logit optimizer	Adam, learning rate 10^{-2}
Candidate pool	$C = 1{,}000$ preference-conditioned candidates
Batch selection	greedy HVI maximization with the adaptive reference
Surrogates	independent Matérn-$5/2$ GPs, one per objective

tive evaluations per run. No early stopping or budget-dependent checkpoint selection is used in the reported main results.

E.2 Model and optimization configuration

Table 3 lists the fixed configuration used for AdaRevTche. Preferences are sampled from Dirichlet$(\mathbf{1})$; ten preferences are resampled at each PSL update. The 1,000 Adam updates are performed on every BO round, with no early stopping. The Gaussian-process lower confidence bound used during PSL training and candidate scoring is $\tilde{f}_i(x) = \mu_i(x) - 0.1\sigma_i(x)$.

The worse-side training anchor at BO round t is the observed archive nadir

$$\hat{r}_i^{(t)} = \max_{(x_j, y_j) \in \mathcal{D}_t} y_{j,i}.$$

The adaptive reference is

$$\hat{r}_i^{(t)}(\mathbf{a}) = \hat{r}_i^{(t)} + \max\left\{0,\ B\left([\mathrm{softmax}(\mathbf{a})]_i - \frac{1}{m}\right)\right\},$$

so it never moves below the observed-nadir anchor. The logits are initialized at $\mathbf{a} = \mathbf{0}$ at every BO round and are not carried over between rounds. The fixed Tchebycheff and fixed-reverse ablations use the same network, optimizer, preference, surrogate, candidate, and evaluation-budget settings, with only the scalarization/reference-adaptation component changed.

E.3 Baselines and implementation scope

We compare against PSL-MOBO, SVH-PSL, DA-PSL, DGEMO, TSEMO, qNParEGO, qLogEHVI, and CDM-PSL. Baseline implementations use the authors' released implementations or the corresponding repository implementations with their documented final settings, while sharing the ten-point initialization, twenty-round budget, batch size, and ten-run protocol. DGEMO is reported only on the two- and three-objective tasks supported by its released Delaunay triangulation implementation.

E.4 Metrics, variation, and statistical reporting

HV is computed in the physical objective space using the fixed, problem-specific evaluation references in Table 1. For an analytic or densely approximated ground-truth HV H^*, the plotted metric is

$$G_t = \log\left(\max\{H^* - \overline{H}_t, \epsilon\}\right),$$

where $\overline{H}_t$ is the mean HV over the ten runs and ϵ is a numerical floor. Lower values are better. Final tables report the mean $\pm$ standard deviation over runs; the ribbons show run-to-run dispersion and are not confidence intervals.

E.5 SOFTWARE, HARDWARE, AND DATA AVAILABILITY

The recorded environment uses Python 3.11.15, PyTorch 2.12.1, BoTorch 0.18.1, GPyTorch 1.15.2, NumPy 2.4.6, SciPy 1.17.1, `pymoo` 0.6.2, pandas 3.0.3, and Matplotlib 3.11.0.

The experiments were run on an Apple M3 Max MacBook Pro with 48 GB memory under macOS 26.3.1. It is worth noting that the MPS backend on this device does not support ‘float64‘; therefore, we adapted the numerical precision of SVH-PSL accordingly.

The benchmark definitions, dense reference fronts, generated archives, and plotting scripts are stored in the submission repository. No external data preprocessing is required: benchmark evaluation and input normalization are performed by the repository's benchmark and transformation utilities. The RE suite and standard synthetic benchmarks are publicly available literature benchmarks; the convex-polar problem is fully specified in this appendix, so no private dataset or preprocessing data are required.