MOGF/D: DECOMPOSITION AND COOPERATION IN GRADIENT FLOW FOR PARETO SET LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Gradient-based multi-objective optimization aims to construct a finite set of solutions representing diverse Pareto trade-offs, but individual gradient descent does not by itself determine how these solutions should organize as a population. Existing methods can generate multiple trade-offs through decomposition, preference adaptation, or particle interaction, but they do not explicitly separate per-subproblem descent from population-level cooperation. To address this, we propose a *normal-tangent cooperation* mechanism that makes population interaction gradient-orthogonal to scalar descent, which preserves first-order decrease and, after a finite formation phase, leads to an exact Lyapunov identity and conditional convergence to Pareto-critical accumulation points. Experiments on synthetic MOO benchmarks and real learning tasks show that the proposed method remains competitive on simple continuous fronts and achieves substantially larger set-quality gains on disconnected and multimodal Pareto problems.

1 INTRODUCTION

Multi-objective optimization (MOO) is central to learning problems with conflicting goals, including multi-task learning, fair classification, and search ranking (Sener & Koltun, 2018; Padh et al., 2021; Carmel et al., 2020). Rather than a single compromise, it seeks a set of Pareto solutions that supports different downstream preferences (Chen et al., 2026). Gradient-based methods make this search practical in high-dimensional models (Chen et al., 2025), but improving individual solutions is only part of the problem. A population must also represent different trade-offs (Zhang et al., 2024; Hotegni & Peitz, 2026). Otherwise, many well-optimized solutions may provide nearly the same choice. This distinction becomes especially important on disconnected fronts or multimodal Pareto sets, where useful trade-offs can lie in different regions or decision-space basins (Li & Zhang, 2009; Ren et al., 2024).

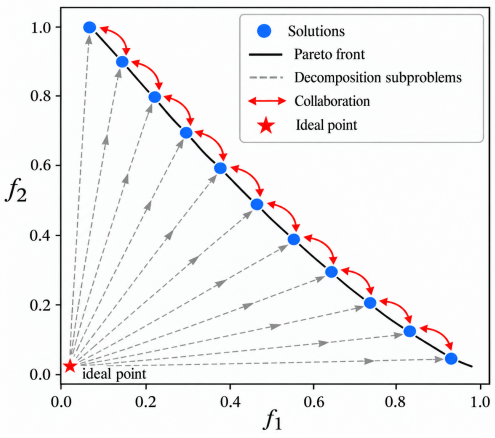

Figure 1: Schematic of decomposition and cooperation: preferences identify trade-offs, while interactions organize their distribution.

Decomposition assigns different preferences to scalar subproblems that can be optimized in parallel (Lin et al., 2019; 2024). Shared parameterizations, adaptive references, and interacting particles can improve their collective coverage (Navon et al., 2021; Chen & Kwok, 2022; Ren et al., 2024). Yet a force that spreads solutions can also oppose a subproblem's descent. Figure 1 illustrates the two roles: preferences specify which trade-offs to seek, whereas cooperation coordinates how the population represents them. The question is whether different subproblems can cooperate without sacrificing their own descent structure.

We propose **M**ulti-**O**bjective **G**radient **F**low based on **D**ecomposition (MOGF/D), which gives decomposition and cooperation distinct roles. Once the scalar subproblems are frozen, each solution

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.

combines projected descent with a feasible cooperative direction orthogonal to its scalar gradient. Cooperation can then redirect motion without changing the scalar subproblem's instantaneous decrease. The key insight is that population-level coordination need not compete with individual descent: decomposition drives progress within each subproblem, while cooperation uses the remaining freedom to organize the solution set. This normal–tangent construction yields an exact Lyapunov identity for the frozen dynamics.

Descent-preserving cooperation cannot by itself cross barriers that require temporary ascent. We therefore first form the population through a finite phase of preference adaptation, interactions, and archive guidance, and then freeze the scalarization before applying normal–tangent dynamics. This two-stage design separates population formation from the subsequent descent analysis; Appendix B explains the transition. It also gives a gradient-flow interpretation of the decomposition–cooperation principle underlying MOEA/D (Zhang & Li, 2007).

The main contributions of this paper are summarized as follows:

1. We introduce a normal-tangent gradient-flow construction that separates decomposition-driven descent from population-level cooperation. The cooperative component is gradient-orthogonal and therefore preserves the first-order decrease of each frozen scalar subproblem.
2. We develop MOGF/D with a finite population-formation phase followed by frozen normal-tangent dynamics. For the frozen phase, we establish an exact continuous-time Lyapunov identity and conditional discrete convergence to Pareto-critical accumulation points.
3. We evaluate MOGF/D on synthetic MOO benchmarks and real learning problems. It remains competitive on simple continuous fronts and achieves substantially larger set-quality gains on disconnected and multimodal Pareto problems, with further results on multi-task learning.

2 Decomposition and Cooperation in Gradient Flow

2.1 Preference-Conditioned Scalar Subproblems

We consider a multi-objective optimization problem with m continuously differentiable objectives,

$$\min_{x\in\mathcal{X}} F(x) = \big(f_1(x), \dots, f_m(x)\big), \tag{1}$$

where $\mathcal{X} \subset \mathbb{R}^d$ is nonempty, compact, and convex. Our goal is not to identify a single compromise, but to construct a finite set of solutions representing different Pareto trade-offs. Following the first-order multi-objective descent framework (Fliege & Svaiter, 2000), a feasible point $x \in \mathcal{X}$ is Pareto critical if there exists $\lambda \in \Delta^{m-1}$ such that

$$0 \in \sum_{j=1}^{m} \lambda_j \nabla f_j(x) + N_{\mathcal{X}}(x), \tag{2}$$

where $N_{\mathcal{X}}(x)$ is the normal cone of $\mathcal{X}$ at x. Geometrically, (2) means that the objective gradients can balance one another, possibly together with a reaction from the feasible-set boundary, so that no feasible first-order direction can simultaneously improve all objectives.

Decomposing Pareto trade-offs. A natural way to obtain multiple trade-offs is to assign different preferences to different scalar subproblems. We associate solution x_i with an interior preference $w_i \in \operatorname{int} \Delta^{m-1}$. Since objectives may have different units and numerical ranges, we first measure their dimensionless deviations from an ideal point $z \in \mathbb{R}^m$:

$$y_j(x) = \frac{f_j(x) - z_j}{s_j}, \qquad j = 1, \dots, m, \tag{3}$$

where $s \in \mathbb{R}^m_{++}$ contains positive scales.

For a given preference w_i, the quantities $w_{ij} y_j(x)$ measure the preference-weighted deviations from the ideal point. A Tchebycheff scalarization minimizes the largest of these deviations (Miettinen, 1998). Its geometry is particularly intuitive: the sublevel set $\max_j w_{ij} y_j(x) \le c$ corresponds in the normalized objective space to a weighted box growing outward from the ideal point. As c increases, this box eventually touches the attainable objective set; changing w_i changes the shape of the box and

therefore the trade-off at which this contact occurs. Unlike a linear weighted sum, this construction does not rely on supporting hyperplanes and therefore does not inherit the convex-front restriction (Miettinen, 1998).

The hard maximum, however, is nonsmooth when several weighted deviations are simultaneously active. We therefore use its log-sum-exp smoothing (Lin et al., 2024) and define

$$\phi_i(x) = \tau \log \left(\sum_{j=1}^{m} \exp \left(\frac{w_{ij} y_j(x)}{\tau} \right) \right) + \nu \sum_{j=1}^{m} w_{ij} y_j(x), \tag{4}$$

where $\tau > 0$ is the smoothing temperature and $\nu \geq 0$ is an augmentation coefficient. The first term can be viewed as a smooth, rounded version of the Tchebycheff box. When τ is small, the objective with the largest weighted deviation dominates the scalar potential. For larger τ, several objectives contribute more smoothly. The augmented term supplies an additional positive contribution from every objective and is useful for retaining all objectives in the scalar descent.

The scalar gradient $\nabla \phi_i$ is a strictly positive combination of the objective gradients, so constrained stationarity of ϕ_i implies Pareto criticality of the original problem. The full gradient expression and proof are given in Lemma C.3 (Appendix C.3).

Hence, different preferences define different scalar force fields and provide a principled mechanism for driving individual solutions toward different Pareto-critical trade-offs. Decomposition alone, however, only specifies how each solution should descend; it does not specify how a finite collection of solutions should organize itself as a population. The next subsection addresses this missing cooperative mechanism.

2.2 Descent-Preserving Cooperative Gradient Flow

The decomposed potentials in (4) specify how individual solutions move toward different Pareto trade-offs, but optimizing them independently provides no direct control over how the finite solution set is organized. A natural remedy is to introduce a population-level velocity that allows different solutions to interact. Directly adding such an interaction to the scalar descent direction, however, may alter the descent property of the corresponding subproblem. We therefore seek a cooperative motion that can change the trajectory of a solution while leaving its scalar decrease unchanged to first order.

Projected scalar descent. Let $g_i = \nabla \phi_i(x_i)$ denote the gradient of the i-th preference-conditioned potential. Without constraints, $-g_i$ is its steepest descent direction. On the feasible set $\mathcal{X}$, only the component pointing along feasible first-order directions can be followed. We therefore define

$$n_i := \mathrm{Proj}_{T_{\mathcal{X}}(x_i)}(-g_i), \tag{5}$$

where $T_{\mathcal{X}}(x_i)$ denotes the tangent cone of $\mathcal{X}$ at x_i. The vector n_i can be viewed as the effective downhill force after removing the component blocked by the feasible-set boundary. By the tangent–normal decomposition (Moreau, 1965), $-g_i = n_i + \zeta_i$ for some $\zeta_i \in N_{\mathcal{X}}(x_i)$ with $\langle n_i, \zeta_i \rangle = 0$, and hence

$$\langle g_i, n_i \rangle = -\|n_i\|_2^2. \tag{6}$$

Thus, $\|n_i\|_2$ is both a projected first-order stationarity measure and the source of scalar dissipation.

Gradient-orthogonal cooperation. Suppose that b_i is a raw population-level signal that suggests how solution i should move relative to the other solutions. Its concrete construction is deferred to Section 3. To use this signal without changing the scalar descent in (6), we retain only a feasible component orthogonal to g_i.

At a constrained point, not every tangent-cone direction permits motion in both signs. We therefore introduce the lineality space

$$L_{\mathcal{X}}(x) := T_{\mathcal{X}}(x) \cap [-T_{\mathcal{X}}(x)], \tag{7}$$

which contains the feasible two-sided directions. For box constraints, $L_{\mathcal{X}}(x)$ is exactly the subspace of free coordinates. Geometrically, these are directions along which a solution can slide locally without pushing against an active boundary.

Let

$$g_L = \mathrm{Proj}_{L_{\mathcal{X}}(x)}\, g, \qquad b_L = \mathrm{Proj}_{L_{\mathcal{X}}(x)}\, b.$$

We define the regularized tangent map

$$T_{\varepsilon,x}(g)b := \frac{\|g_L\|_2^2 b_L - \langle g_L, b_L\rangle g_L}{\|g_L\|_2^2 + \varepsilon^2}, \qquad \varepsilon > 0. \tag{8}$$

The numerator removes from the cooperative signal the component parallel to the scalar gradient, while the denominator regularizes this operation when the projected gradient is small. Consequently,

$$T_{\varepsilon,x}(g)b \in L_{\mathcal{X}}(x), \qquad \langle g, T_{\varepsilon,x}(g)b\rangle = 0,$$

and the map vanishes continuously as $g_L \to 0$. Since the gradient is normal to a scalar level set, the resulting direction can be interpreted as a first-order sliding motion along the level geometry of ϕ_i: it can steer the solution while doing zero first-order work on the scalar potential.

The cooperative velocity should also vanish as the solution approaches stationarity rather than dominating the descent motion. We therefore define

$$\widetilde{t}_i = T_{\varepsilon,x_i}(g_i)b_i, \qquad t_i = \min\left\{1, \frac{\rho\|n_i\|_2}{\|\widetilde{t}_i\|_2}\right\}\widetilde{t}_i, \tag{9}$$

with $t_i = 0$ when $\widetilde{t}_i = 0$. The parameter $\rho \geq 0$ controls the relative strength of cooperation:

$$\|t_i\|_2 \leq \rho\|n_i\|_2.$$

Hence the cooperative steering cannot arbitrarily dominate the dissipative motion, and it disappears automatically when $n_i = 0$. The resulting velocity is

$$v_i = n_i + t_i. \tag{10}$$

The composite velocity is feasible, $v_i \in T_{\mathcal{X}}(x_i)$, and satisfies

$$\langle g_i, v_i\rangle = -\|n_i\|_2^2. \tag{11}$$

Thus, cooperation can steer a solution without changing its first-order scalar decrease. The full descent-angle bound is given in Lemma C.4 (Appendix C.4).

This separation becomes particularly transparent in continuous time. Consider the frozen dynamics

$$\dot{x}_i(t) = n_i(x_i(t)) + t_i(t). \tag{12}$$

Proposition 2.1 (Lyapunov identity)**.** *Along every feasible absolutely continuous solution of* (12)*,*

$$\frac{\mathrm{d}}{\mathrm{d}t}\phi_i(x_i(t)) = -\|n_i(x_i(t))\|_2^2 \leq 0 \tag{13}$$

for almost every t.

The integrated identity and proof are given in Appendix C.5.

This viewpoint also gives a gradient-flow interpretation of the decomposition–cooperation principle underlying MOEA/D (Zhang & Li, 2007): decomposition determines the individual optimization dynamics, while cooperation couples the subproblems at the population level. The construction above is not a continuous-time limit of the evolutionary operators used by classical MOEA/D; rather, it is a differentiable realization of the same design principle.

The frozen flow, however, only explains how an already formed population can cooperate without sacrificing descent. It does not by itself specify how to obtain a favorable finite population or how to escape population configurations that require temporary non-tangential motion. Section 3 addresses this issue through a finite population-formation stage.

3 FROM COOPERATIVE FLOW TO A FINITE-BUDGET ALGORITHM

The cooperative flow in Section 2.2 preserves descent once the scalar subproblems are fixed, but this restriction can be too conservative at the beginning of optimization. A population initialized in an unfavorable region may remain crowded, miss important trade-offs, or be trapped in decision-space basins that cannot be left by descent-neutral motion alone. Rather than enforcing one fixed Lyapunov structure throughout the entire search, we therefore separate the optimization into two stages: a finite *formation* stage that prepares the population, followed by the frozen cooperative flow analyzed in Section 2.2.

During the first K iterations, the scalar subproblems and the population are allowed to adapt. In particular, the normalization and preferences may change, population interaction discourages local crowding, and a nondominated archive provides guidance toward underrepresented regions. These mechanisms address three complementary difficulties—poor subproblem allocation, local population concentration, and uncovered trade-off regions—without requiring the transient dynamics to follow a fixed scalar potential. All formation-only non-tangential moves and all changes to (w_i, z, s) terminate after the prescribed horizon K. The detailed preference updates, interaction forces, archive construction, and finite energy guard are given in Appendix B.

This finite-prefix construction has a simple consequence. If the formation stage produces a feasible state x_i^K, then the subsequent frozen recursion starts from x_i^K with fixed scalarization parameters. Hence the formation stage can change the initial condition, and therefore the basin entered by the frozen dynamics, but it does not alter any asymptotic statement that holds for every feasible frozen initial state. This observation is useful because it separates the two roles of the algorithm: formation is used to improve the population configuration, whereas convergence is supplied by the frozen descent-preserving flow.

We next discretize the frozen dynamics. Tangency is a first-order property, while a finite step is affected by curvature and projection onto $\mathcal{X}$. For $k \geq K$, after computing $v_i^k = n_i^k + t_i^k$, we therefore take a projected trial displacement

$$\Delta x_i^k(\alpha) := \mathrm{Proj}_{\mathcal{X}}\left(x_i^k + \alpha \frac{v_i^k}{\|v_i^k\|_2}\right) - x_i^k, \tag{14}$$

and use per-solution Armijo-type backtracking (Armijo, 1966) to accept a step satisfying

$$\phi_i(x_i^k + \Delta x_i^k) \leq \phi_i(x_i^k) + c\langle g_i^k, \Delta x_i^k\rangle + \epsilon_k, \qquad \sum_{k=K}^{\infty} \epsilon_k < \infty, \tag{15}$$

where $c \in (0, 1)$. The summable error allows small finite-step deviations from the ideal continuous-time level geometry, while preventing these errors from accumulating indefinitely.

The continuous descent identity alone does not guarantee convergence of the projected finite-step scheme: after projection and backtracking, the accepted displacement must still retain a nonvanishing amount of the descent predicted by the continuous flow. Accordingly, we assume that whenever $n_i^k \neq 0$,

$$-\langle g_i^k, \Delta x_i^k\rangle \geq \kappa \|n_i^k\|_2^2 \tag{16}$$

for some constant $\kappa > 0$. This is a finite-step nondegeneracy condition rather than a consequence of tangency alone; sufficient alignment and step-length conditions are given in Appendix D.

Theorem 3.1 (Convergence of the frozen discrete dynamics)**.** *Assume that $\mathcal{X}$ is compact and convex, the objective gradients are Lipschitz, the frozen preferences and scales remain strictly positive, and the accepted displacements satisfy* (15)–(16)*. Let*

$$E_k = \sum_{\ell=k}^{\infty} \epsilon_\ell, \qquad V_i^k = \phi_i(x_i^k) + E_k.$$

Then

$$V_i^{k+1} \leq V_i^k - c\kappa \|n_i^k\|_2^2, \qquad \sum_{k=K}^{\infty} \|n_i^k\|_2^2 < \infty, \qquad \|n_i^k\|_2 \to 0. \tag{17}$$

Every accumulation point of solution i is a constrained stationary point of its frozen scalar subproblem and, by Lemma C.3, a Pareto-critical point of the original multi-objective problem.

Table 1: IGD ↓ on synthetic benchmarks (mean ± sample SD, six seeds). Parentheses give zero-based ranks from unrounded means; bold and underlined values mark the best and second best. Avg. rank averages all seven problems.

Dataset	ParetoMTL	COSMOS	PHN-LS	PHN-EPO	ARGMO-kernel	ARGMO-HV	Particle-WFR	MOGF/D
Simple								
ZDT1	0.08566 (5) ±0.00443	0.04290 (4) ±0.00999	0.01296 (2) ±0.00051	0.01888 (3) ±0.00148	0.31867 (6) ±0.00474	0.32363 (7) ±0.00850	<u>0.00767</u> (1) ±0.00018	**0.00747** (0) ±0.00002
ZDT2	0.30445 (3) ±0.00154	0.32208 (5) ±0.31488	0.21720 (2) ±0.30407	0.31028 (4) ±0.32822	0.38856 (7) ±0.24720	0.37353 (6) ±0.26337	<u>0.00757</u> (1) ±0.00002	**0.00754** (0) ±0.00001
Disconnected PF								
ZDT3	0.31126 (2) ±0.07434	0.54106 (7) ±0.32819	0.37085 (3) ±0.32155	0.41686 (4) ±0.31004	0.53468 (5) ±0.24413	0.54094 (6) ±0.25486	<u>0.01862</u> (1) ±0.00247	**0.00968** (0) ±0.00007
DTLZ7	1.14500 (7) ±0.04412	0.51711 (3) ±0.23176	0.48004 (2) ±0.16724	0.68501 (6) ±0.23400	0.56294 (4) ±0.19786	0.58212 (5) ±0.18807	<u>0.06414</u> (1) ±0.00415	**0.03744** (0) ±0.00050
Complex PS								
LZ09-F2	0.06700 (2) ±0.01204	0.37694 (4) ±0.13903	0.32997 (3) ±0.08754	0.39009 (5) ±0.10113	0.39439 (6) ±0.04974	0.39545 (7) ±0.04127	<u>0.01259</u> (1) ±0.00056	**0.00980** (0) ±0.00233
LZ09-F7	1.68238 (2) ±0.79680	<u>0.75701</u> (1) ±0.28456	4.02054 (7) ±1.59799	3.66362 (6) ±1.85613	1.75754 (3) ±1.85184	1.78026 (4) ±1.42159	2.82260 (5) ±0.15275	**0.03202** (0) ±0.02470
LZ09-F9	0.07539 (2) ±0.01000	0.55767 (5) ±0.11739	0.47104 (3) ±0.12185	0.52688 (4) ±0.16661	0.57992 (6) ±0.08146	0.60662 (7) ±0.13836	<u>0.00833</u> (1) ±0.00025	**0.00811** (0) ±0.00067
Avg. rank	3.29	4.14	3.14	4.57	5.29	6.00	<u>1.57</u>	**0.00**

Theorem 3.1 closes the design loop. The finite formation stage determines where the population enters the frozen dynamics, while the normal–tangent construction determines how the solutions cooperate thereafter without changing their asymptotic first-order target. The theorem concerns the convergence of individual frozen trajectories. It does not imply that different preferences converge to different points or that the limiting particles cover the entire Pareto front. The quality of the returned finite set is therefore evaluated empirically in Section 4.

4 Experiments

We evaluate MOGF/D on analytic benchmarks with known Pareto geometry and high-dimensional multi-objective learning tasks. We use inverted generational distance (IGD ↓) (Ishibuchi et al., 2015) to assess proximity and coverage against known fronts, and hypervolume (HV ↑) (Guerreiro et al., 2021) to evaluate returned-set quality when the true front is unavailable. Appendices F.1 and F.7 give the metric definitions and reproducibility details.

4.1 Pareto-Set Approximation under Increasing Geometric Complexity

We first evaluate MOGF/D on a collection of analytic multi-objective benchmarks with known Pareto geometry, including the ZDT suite (Zitzler et al., 2000), the DTLZ suite (Deb et al., 2002), and the LZ09 problems (Li & Zhang, 2009). ZDT1 and ZDT2 provide simple continuous Pareto fronts, whereas ZDT3 and DTLZ7 introduce disconnected front components; the LZ09 problems further test nonlinear decision-variable coupling and, in the case of LZ09-F7, multiple local Pareto sets. Together, these problems allow us to move progressively from ordinary front fitting to settings in which population organization and basin allocation become increasingly important.

We compare against representative gradient-based and population-based Pareto learning methods: (1) ParetoMTL (Lin et al., 2019), which decomposes the MOO problem into preference-guided sub-problems; (2) COSMOS (Ruchte & Grabocka, 2021), which learns a scalable preference-conditioned approximation of the Pareto front; (3) PHN-LS and PHN-EPO (Navon et al., 2021), which use

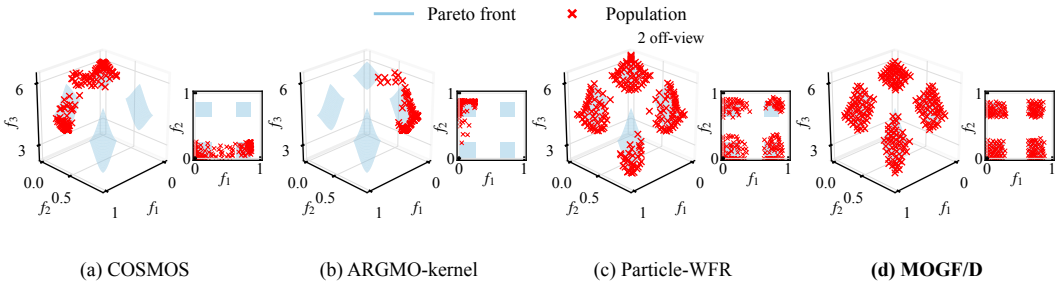

Figure 2: Results on DTLZ7. Light blue shows the analytic Pareto front, and red crosses show the returned populations. Each method is shown in three dimensions alongside its (f_1, f_2) projection. Full eight-method results are in Figure 9.

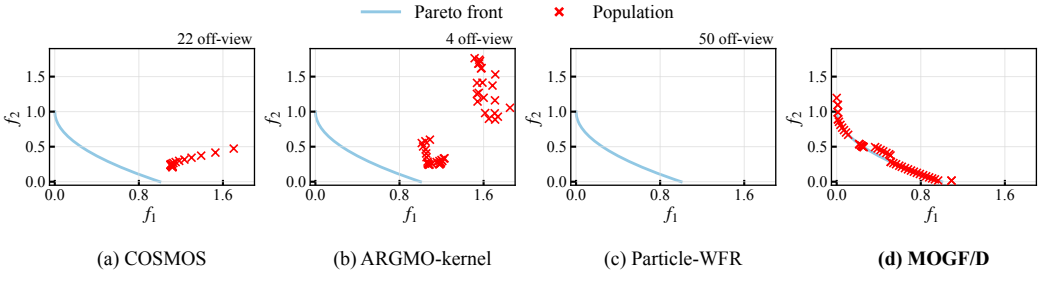

Figure 3: Results on LZ09-F7. Light blue shows the analytic Pareto front, and red crosses show the returned populations. Off-view labels count solutions outside the shared axis limits. Full eight-method results are in Figure 11.

hypernetworks to generate preference-dependent solutions with different scalarization schemes; (4) ARGMO-kernel and ARGMO-HV (Chen & Kwok, 2022), which adapt reference vectors during optimization to improve the distribution of Pareto solutions; and (5) Particle-WFR (Ren et al., 2024), an interacting-particle method based on Wasserstein–Fisher–Rao gradient flow for approximating complex Pareto fronts. Appendix F.2 gives the baseline descriptions and experimental settings, with 30 F7 variables for Particle-WFR and 10 for the other methods.

We present the quantitative results on the synthetic problems in Table 1, with two representative difficult cases visualized in Figures 2 and 3. ParetoMTL poorly approximates the fronts of DTLZ7 and F7 (Table 1); its front plots are therefore deferred to Appendix F.3. MOGF/D achieves the lowest mean IGD on all seven benchmarks and the best overall average rank. More importantly, the magnitude of the improvement changes systematically with the geometry of the problem. On the simple continuous fronts ZDT1 and ZDT2, MOGF/D performs nearly on par with Particle-WFR, indicating that the normal–tangent construction does not sacrifice ordinary Pareto-front fitting. The difference becomes substantially larger on the disconnected ZDT3 and DTLZ7 problems. In particular, Figure 2 shows that several preference-conditioned methods concentrate on only part of the disconnected front or produce many off-front solutions; Particle-WFR reaches all major components but remains visibly dispersed around them, whereas MOGF/D concentrates its population around all four Pareto components. A similar pattern becomes even more pronounced on the LZ09 problems, where the difficulty lies in the geometry of the Pareto set in the decision space. MOGF/D improves substantially on LZ09-F2, remains close to Particle-WFR on LZ09-F9, and shows the largest separation on the multimodal LZ09-F7 problem. As illustrated in Figure 3, most competing methods either miss the Pareto set almost entirely or concentrate on a restricted region, while the population produced by MOGF/D closely follows the full Pareto set. These results suggest that the benefit of cooperation is modest when independent scalar descent already suffices to trace a simple front, but becomes much more pronounced when a finite population must be organized across disconnected Pareto regions or multiple decision-space basins. Appendix F.3 provides the complementary synthetic HV results and full eight-method front plots for all seven benchmarks.

Table 2: Validation-selected test HV ↑ (mean ± sample SD, three seeds). Parentheses give zero-based ranks from unrounded means; bold and underlined values mark the best and second best. MSLR values are in units of 10^{-4}.

Dataset	ParetoMTL	COSMOS	PHN-LS	PHN-EPO	ARGMO-kernel	ARGMO-HV	Particle-WFR	MOGF/D
Adult (2-obj)	0.61999 (7) ±0.00431	0.66966 (6) ±0.00400	0.67120 (4) ±0.00072	0.67095 (5) ±0.00050	0.67165 (3) ±0.00056	0.67239 (2) ±0.00236	<u>0.67470</u> (1) ±0.00054	**0.67488** (0) ±0.00085
Credit (2-obj)	0.55999 (4) ±0.00202	0.55816 (6) ±0.00335	<u>0.56179</u> (1) ±0.00176	0.55735 (7) ±0.00211	0.55900 (5) ±0.00340	0.56050 (2) ±0.00065	0.56035 (3) ±0.00038	**0.56238** (0) ±0.00078
COMPAS (2-obj)	0.77850 (5) ±0.01359	0.80869 (2) ±0.00256	0.79859 (3) ±0.01156	<u>0.80964</u> (1) ±0.00564	0.74192 (7) ±0.11342	0.74381 (6) ±0.11572	0.78250 (4) ±0.00154	**0.81865** (0) ±0.00504
MSLR (6-obj)	7.61323 (2) ±0.29563	4.30206 (4) ±0.60729	4.16405 (5) ±0.62188	4.56538 (3) ±0.17235	3.42511 (7) ±0.82046	3.97654 (6) ±0.64447	<u>8.07166</u> (1) ±0.14919	**8.76222** (0) ±0.20108
Avg. rank	4.50	4.50	3.25	4.00	5.50	4.00	<u>2.25</u>	**0.00**

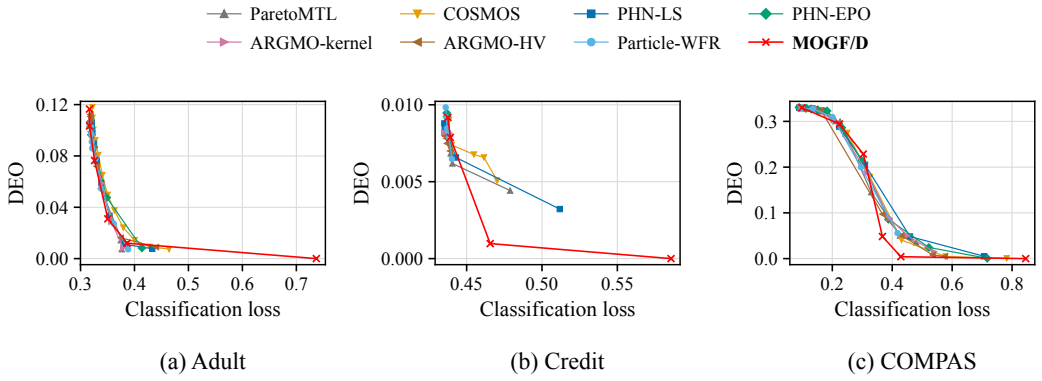

(a) Adult (b) Credit (c) COMPAS

Figure 4: Test trade-offs on (a) Adult, (b) Credit, and (c) COMPAS. Curves connect nondominated solutions from validation-selected checkpoints; each method displays its median-test-HV seed.

Unlike the disconnected-front problems above, F7 contains multiple local Pareto sets, so reaching a well-distributed population may require moving between distinct decision-space basins. This is precisely the situation for which the finite formation stage in Section 3 was introduced. During formation, preferences, population interactions, and archive guidance are allowed to reshape the population before the scalar subproblems are frozen; for F7, we additionally use finite single-coordinate proposals to facilitate basin exploration. The proposals stop before freezing, so the F7 experiments evaluate the complete two-stage method, including finite basin exploration. The frozen normal–tangent flow then preserves the descent structure of the resulting scalar subproblems. Appendix F.4 reports controlled component removals and paired IGD comparisons.

4.2 Results on Pareto Multi-Task Learning

We next evaluate MOGF/D on high-dimensional multi-objective learning problems. For fairness classification, we consider Adult (Becker & Kohavi, 1996), Credit (Yeh & Lien, 2009), and COMPAS (Angwin et al., 2016), where the two objectives are classification loss and a differentiable difference-of-equality-of-opportunity (DEO) loss (Ren et al., 2024). We further consider the six-objective learning-to-rank task on MSLR-WEB10K (Qin & Liu, 2013), following the multi-objective construction used by Particle-WFR (Ren et al., 2024). All methods return a finite set of trade-off solutions, and model selection is based on validation HV. Appendices F.5 and F.6 give the task definitions, data splits, architectures, optimization settings, and selection rules.

The quantitative results are summarized in Table 2, and the fairness trade-offs are visualized in Figure 4. MOGF/D obtains the highest mean test HV on all three fairness datasets, although the

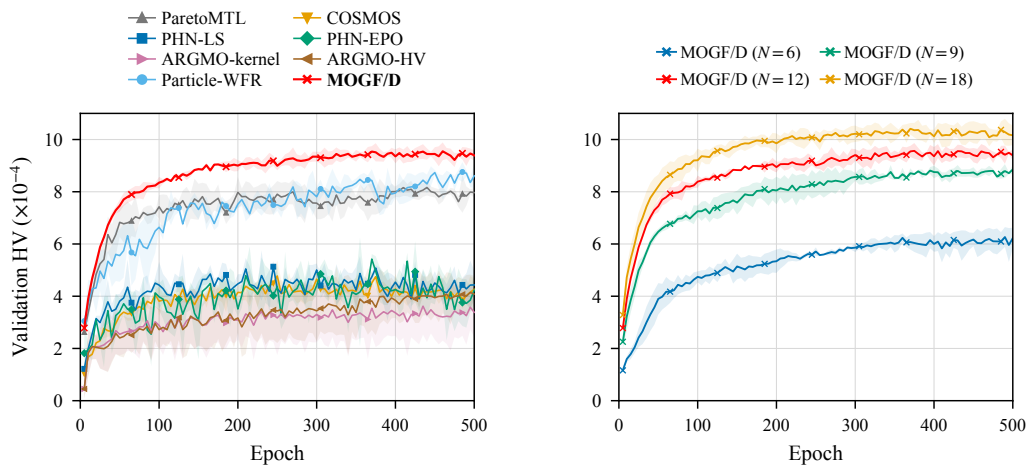

(a) Method comparison ($N = 12$)

(b) MOGF/D population size

Figure 5: Validation HV on MSLR-WEB10K: (a) eight methods with $N = 12$; (b) MOGF/D with $N = 6, 9, 12, 18$. Lines and shading show the mean and sample SD over three seeds.

magnitude of the gain varies considerably. On Adult and Credit, the strongest methods are already close and MOGF/D improves over the best baseline by only about 0.03% and 0.11%, respectively. The corresponding fronts largely overlap, indicating near-parity in these two cases. The difference is clearer on COMPAS, where MOGF/D improves the best baseline by about 1.11% and extends the returned trade-off set toward very low DEO while retaining competitive classification loss. Thus, as in the synthetic study, the gains characterize the quality of the returned trade-off set.

The six-objective MSLR problem provides a more demanding test of population-level optimization. As shown in Table 2, MOGF/D achieves the highest test HV, improving over Particle-WFR by about 8.56%. Figure 5(a) further shows that its validation HV rises rapidly and converges to the highest level among the compared methods for $N = 12$, suggesting that the final test advantage is consistent with the training trajectory. Figure 5(b) shows an additional population-level effect: within the tested sizes, increasing the population from $N = 6$, 9, 12 and 18 progressively improves validation HV. This result is consistent with the central motivation of MOGF/D—a richer population can represent more joint trade-offs, provided that its solutions can be organized cooperatively. Across three seeds, mean test HV at validation-selected checkpoints rises from 5.732×10^{-4} with $N = 6$ to 9.601×10^{-4} with $N = 18$. All sizes use 500 training epochs, with larger populations requiring more model updates per epoch. The population-study protocol and its test results are given in Appendix F.6 and Table 14.

5 Conclusion and Future Work

In this paper, we have proposed a gradient-flow view of Pareto set learning through decomposition and cooperation. MOGF/D separates decomposition-driven descent from population-level cooperation using normal–tangent dynamics, which preserves first-order descent and admits a Lyapunov analysis and conditional convergence guarantee. A finite formation stage further helps organize the population before the dynamics are frozen. Experiments show that MOGF/D remains competitive on simple fronts and achieves larger set-quality gains on disconnected, multimodal, and real learning problems. Overall, our results provide a dynamical interpretation of the decomposition–cooperation principle for gradient-based Pareto set learning. Future work will extend normal–tangent dynamics to stochastic optimization, with particular attention to descent preservation under noisy gradients. We will also investigate adaptive choices of the formation horizon and population size to improve trade-off coverage under a fixed computational budget.

AI USE STATEMENT

Generative AI tools were used to assist with manuscript editing, technical consistency checking, and experimental code development. All AI-assisted content and code were reviewed and verified by the authors. The authors take full responsibility for all claims, results, and materials presented in this work.

ETHICS STATEMENT

This work develops a general methodology for multi-objective optimization and does not involve new data collection or direct interaction with human participants. The learning experiments use established benchmark datasets, including fairness classification tasks on Adult, Credit, and COMPAS. These experiments are intended solely to evaluate multi-objective optimization methods; in particular, the reported fairness trade-offs should not be interpreted as recommendations for deployment or decision making in high-stakes applications.

REPRODUCIBILITY STATEMENT

The problem formulation and MOGF/D algorithm are described in Sections 2 and 3, with construction details and complete theoretical proofs provided in Appendices B–E. Section 4 presents the experiments, while Appendix F specifies the evaluation metrics, benchmark definitions, baseline implementations, hyperparameters, optimization budgets, data preprocessing and splits, random seeds, checkpoint-selection rules, and additional results. Appendix F.7 further documents the provenance and scope of the reported experimental records.

REFERENCES

Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine Bias. ProPublica, 2016. URL https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing. Published May 23, 2016.

Larry Armijo. Minimization of functions having Lipschitz continuous first partial derivatives. *Pacific Journal of Mathematics*, 16(1):1–3, 1966. doi: 10.2140/pjm.1966.16.1. URL https://msp.org/pjm/1966/16-1/pjm-v16-n1-p01-s.pdf.

Barry Becker and Ronny Kohavi. Adult. UCI Machine Learning Repository, 1996. URL https://archive.ics.uci.edu/dataset/2/adult. Dataset.

Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th international conference on Machine learning*, pp. 129–136. ACM, 2007. doi: 10.1145/1273496.1273513. URL https://dl.acm.org/doi/10.1145/1273496.1273513.

David Carmel, Elad Haramaty, Arnon Lazerson, and Liane Lewin-Eytan. Multi-Objective Ranking Optimization for Product Search Using Stochastic Label Aggregation. In *Proceedings of The Web Conference 2020*, pp. 373–383. ACM, 2020. doi: 10.1145/3366423.3380122. URL https://dl.acm.org/doi/10.1145/3366423.3380122.

Edward Chen, Natalie Dullerud, Thomas Niedermayr, Elizabeth Kidd, Ransalu Senanayake, Pang Wei Koh, Sanmi Koyejo, and Carlos Guestrin. Modeling Multi-Objective Tradeoffs with Monotonic Utility Functions. In *Proceedings of The 29th International Conference on Artificial Intelligence and Statistics*, volume 300 of *Proceedings of Machine Learning Research*, pp. 4753–4761. PMLR, 2026. URL https://proceedings.mlr.press/v300/chen26i.html.

Lisha Chen, Quan Xiao, Ellen Hidemi Fukuda, Xinyi Chen, Kun Yuan, and Tianyi Chen. Efficient First-Order Optimization on the Pareto Set for Multi-Objective Learning under Preference Guidance. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 9443–9486. PMLR, 2025. URL https://proceedings.mlr.press/v267/chen25bw.html.

Weiyu Chen and James Kwok. Multi-Objective Deep Learning with Adaptive Reference Vectors. In *Advances in Neural Information Processing Systems*, volume 35, pp. 32723–32735. Curran Associates, Inc., 2022. doi: 10.52202/068431-2371. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/d313b4a8c88eba7f0542c489899cec77-Paper-Conference.pdf.

K. Deb, L. Thiele, M. Laumanns, and E. Zitzler. Scalable multi-objective optimization test problems. In *Proceedings of the 2002 Congress on Evolutionary Computation. CEC'02 (Cat. No.02TH8600)*, volume 1, pp. 825–830. IEEE, 2002. doi: 10.1109/cec.2002.1007032. URL https://doi.org/10.1109/cec.2002.1007032.

Jörg Fliege and Benar Fux Svaiter. Steepest descent methods for multicriteria optimization. *Mathematical Methods of Operations Research*, 51(3):479–494, 2000. doi: 10.1007/s001860000043. URL https://link.springer.com/article/10.1007/s001860000043.

Andreia P. Guerreiro, Carlos M. Fonseca, and Luís Paquete. The Hypervolume Indicator: Computational Problems and Algorithms. *ACM Computing Surveys*, 54(6):1–42, 2021. doi: 10.1145/3453474. URL https://dl.acm.org/doi/10.1145/3453474.

Sture Holm. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics*, 6(2):65–70, 1979. URL https://www.jstor.org/stable/4615733.

Sedjro Salomon Hotegni and Sebastian Peitz. SPREAD: Sampling-based Pareto front Refinement via Efficient Adaptive Diffusion. In *International Conference on Learning Representations*, volume 2026, pp. 115284–115315, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/file/bb74d3d956cc960f693716b3595593d7-Paper-Conference.pdf.

Hisao Ishibuchi, Hiroyuki Masuda, Yuki Tanigaki, and Yusuke Nojima. Modified Distance Calculation in Generational Distance and Inverted Generational Distance. In *Evolutionary Multi-Criterion Optimization*, volume 9019 of *Lecture Notes in Computer Science*, pp. 110–125. Springer, 2015. doi: 10.1007/978-3-319-15892-1_8. URL https://link.springer.com/chapter/10.1007/978-3-319-15892-1_8.

Rolf Jagerman. PyTorchLTR. Software and documentation, 2020. URL https://pytorchltr.readthedocs.io/en/stable/.

Kalervo Järvelin and Jaana Kekäläinen. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems*, 20(4):422–446, 2002. doi: 10.1145/582415.582418. URL https://dl.acm.org/doi/10.1145/582415.582418.

Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, 2015. URL https://arxiv.org/abs/1412.6980.

Hui Li and Qingfu Zhang. Multiobjective Optimization Problems With Complicated Pareto Sets, MOEA/D and NSGA-II. *IEEE Transactions on Evolutionary Computation*, 13(2):284–302, 2009. doi: 10.1109/tevc.2008.925798. URL https://ieeexplore.ieee.org/document/4633340/.

Xi Lin, Hui-Ling Zhen, Zhenhua Li, Qing-Fu Zhang, and Sam Kwong. Pareto Multi-Task Learning. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/685bfde03eb646c27ed565881917c71c-Paper.pdf.

Xi Lin, Xiaoyuan Zhang, Zhiyuan Yang, Fei Liu, Zhenkun Wang, and Qingfu Zhang. Smooth Tchebycheff Scalarization for Multi-Objective Optimization. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 30479–30509. PMLR, 2024. URL https://proceedings.mlr.press/v235/lin24y.html.

Xingchao Liu, Xin Tong, and Qiang Liu. Profiling Pareto Front With Multi-Objective Stein Variational Gradient Descent. In *Advances in Neural Information Processing Systems*, volume 34, pp. 14721–14733. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/7bb16972da003e87724f048d76b7e0e1-Paper.pdf.

Debabrata Mahapatra and Vaibhav Rajan. Multi-Task Learning with User Preferences: Gradient Descent with Controlled Ascent in Pareto Optimization. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 6597–6607. PMLR, 2020. URL https://proceedings.mlr.press/v119/mahapatra20a.html.

Kaisa Miettinen. *Nonlinear Multiobjective Optimization*, volume 12 of *International Series in Operations Research & Management Science*. Springer, 1998. doi: 10.1007/978-1-4615-5563-6. URL https://link.springer.com/book/10.1007/978-1-4615-5563-6.

J.J. Moreau. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société Mathématique de France*, 93:273–299, 1965. doi: 10.24033/bsmf.1625. URL https://numdam.org/articles/10.24033/bsmf.1625/.

Aviv Navon, Aviv Shamsian, Ethan Fetaya, and Gal Chechik. Learning the Pareto Front with Hypernetworks. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=NjF772F4ZZR.

Kirtan Padh, Diego Antognini, Emma Lejal-Glaude, Boi Faltings, and Claudiu Musat. Addressing fairness in classification with a model-agnostic multi-objective algorithm. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, volume 161 of *Proceedings of Machine Learning Research*, pp. 600–609. PMLR, 2021. URL https://proceedings.mlr.press/v161/padh21a.html.

Tao Qin and Tie-Yan Liu. Introducing LETOR 4.0 Datasets. *CoRR*, abs/1306.2597, 2013. URL https://arxiv.org/abs/1306.2597.

Yinuo Ren, Tesi Xiao, Tanmay Gangwani, Anshuka Rangi, Holakou Rahmanian, Lexing Ying, and Subhajit Sanyal. Multi-objective Optimization via Wasserstein-Fisher-Rao Gradient Flow. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 3862–3870. PMLR, 2024. URL https://proceedings.mlr.press/v238/ren24b.html.

Michael Ruchte and Josif Grabocka. Scalable Pareto Front Approximation for Deep Multi-Objective Learning. In *2021 IEEE International Conference on Data Mining (ICDM)*, pp. 1306–1311. IEEE, 2021. doi: 10.1109/icdm51629.2021.00162. URL https://doi.org/10.1109/ICDM51629.2021.00162.

Ozan Sener and Vladlen Koltun. Multi-Task Learning as Multi-Objective Optimization. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/432aca3a1e345e339f35a30c8f65edce-Paper.pdf.

Mao Ye and Qiang Liu. Pareto navigation gradient descent: a first-order algorithm for optimization in pareto set. In *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, volume 180 of *Proceedings of Machine Learning Research*, pp. 2246–2255. PMLR, 2022. URL https://proceedings.mlr.press/v180/ye22a.html.

I-Cheng Yeh and Che-hui Lien. The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert Systems with Applications*, 36(2): 2473–2480, 2009. doi: 10.1016/j.eswa.2007.12.020. URL https://doi.org/10.1016/j.eswa.2007.12.020.

Qingfu Zhang and Hui Li. MOEA/D: A Multiobjective Evolutionary Algorithm Based on Decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6):712–731, 2007. doi: 10.1109/tevc.2007.892759. URL https://ieeexplore.ieee.org/document/4358754.

Xiaoyuan Zhang, Genghui Li, Xi Lin, Yichi Zhang, Yifan Chen, and Qingfu Zhang. Gliding over the Pareto Front with Uniform Designs. In *Advances in Neural Information Processing Systems*, volume 37, pp. 2215–2245. Curran Associates, Inc., 2024. doi: 10.52202/079017-0072. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/040d3b6af368bf71f952c18da5713b48-Paper-Conference.pdf.

Yawen Zhou, Ke Xue, and Chao Qian. MGDA/D: A Multiple Gradient Descent Algorithm Based on Decomposition. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 455–458. ACM, 2025. doi: 10.1145/3712255.3726613. URL https://doi.org/10.1145/3712255.3726613.

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. URL https://doi.org/10.1162/106365600568202.

A Related Work

A.1 Decomposition and preference-conditioned learning

Decomposition constructs multiple trade-offs through preference-defined subproblems, as in MOEA/D and ParetoMTL (Zhang & Li, 2007; Lin et al., 2019), while smooth Tchebycheff scalarization provides a differentiable formulation (Lin et al., 2024). Shared models amortize this optimization: Pareto hypernetworks generate parameters from preferences, and COSMOS conditions the predictive network directly on them (Navon et al., 2021; Ruchte & Grabocka, 2021). To improve front coverage, adaptive-reference methods optimize reference vectors through bilevel learning (Chen & Kwok, 2022), and UMOD uses a learned preference-to-objective model to optimize a max-packing design (Zhang et al., 2024). MGDA/D further combines gradient-based subproblems with neighborhood information to escape local optima (Zhou et al., 2025). MOGF/D shares the decomposition–cooperation principle, but separates finite population formation from a frozen phase in which cooperation preserves each scalar subproblem's first-order descent.

A.2 Population dynamics and Pareto-set optimization

Particle methods organize solutions through coupled dynamics. MOO-SVGD combines multiple-gradient descent with kernel repulsion (Liu et al., 2021), while Particle-WFR uses Langevin and birth–death dynamics to implement a distributional gradient flow with dominance and diversity potentials (Ren et al., 2024). SPREAD extends gradient-guided refinement and repulsion to diffusion sampling (Hotegni & Peitz, 2026). Complementary approaches optimize an additional preference or diversity criterion over Pareto solutions: PNG constrains objective descent (Ye & Liu, 2022), and recent first-order methods use merit functions and penalty reformulations under weak Pareto optimality (Chen et al., 2025). MOGF/D imposes a different structural condition: cooperative motion is orthogonal to the gradient of each frozen scalar potential, so its instantaneous decrease is determined entirely by the descent component. This identity applies after formation, when the scalarization is fixed.

B Finite Population Formation and Construction Details

Section 3 uses a finite formation stage to prepare the population before the descent-preserving dynamics are frozen. This appendix gives the construction details omitted from the main text. The role of these mechanisms is deliberately limited: they improve the finite population configuration, whereas the asymptotic descent guarantee is supplied by the frozen normal–tangent dynamics.

Changing and frozen scalar landscapes. To illustrate the role of this phase, let $\xi_i = (w_i, z, s)$ denote the parameters defining the scalar potential of subproblem i. Schematically, in continuous time,

$$\frac{\mathrm{d}}{\mathrm{d}t}\phi_i(x_i(t); \xi_i(t)) = \langle \nabla_x \phi_i, \dot{x}_i \rangle + \left\langle \nabla_\xi \phi_i, \dot{\xi}_i \right\rangle .$$

During formation, both the solution and its scalar landscape may change, so we do not impose a fixed Lyapunov structure on this transient. After a prescribed horizon K, the scalarization is frozen, $\dot{\xi}_i = 0$, and cooperation is restricted by $\langle \nabla \phi_i, t_i \rangle = 0$, recovering

$$\frac{\mathrm{d}}{\mathrm{d}t}\phi_i(x_i(t)) = -\|n_i\|_2^2.$$

Thus, the finite prefix is used to form the population, while the frozen tail admits a clean dissipative analysis.

Positive and anchored preferences. Each solution starts from a positive reference preference $\bar{w}_i \in \operatorname{int} \Delta^{m-1}$. During formation, the deterministic implementation adapts the preference with a centered simplex-tangent increment. Let $q_i^k \in [0, 1]$ be a crowding gate,

$$\pi_i^k = \operatorname{softmax}\big(-y(x_i^k)/\tau\big), \qquad H = I - \frac{\mathbf{1}\mathbf{1}^\top}{m}.$$

The update is

$$w_i^{k+1} = R_{\epsilon_w}\left(w_i^k + \gamma q_i^k H\big[\pi_i^k - w_i^k + \lambda_w(\bar{w}_i - w_i^k)\big]\right), \qquad [R_{\epsilon_w}(u)]_j = \frac{\max\{u_j, \epsilon_w\}}{\sum_r \max\{u_r, \epsilon_w\}}. \tag{18}$$

The centered increment changes the relative preference while preserving the simplex tangent direction before renormalization, and the anchor $\bar{w}_i$ discourages uncontrolled drift. The clipping–renormalization map keeps every finite preference update positive. The deterministic implementation uses $\epsilon_w = 10^{-4}$ as a numerical positivity floor before renormalization; learning implementations may instead project onto a simplex with a configured positive floor.

Proposition B.1 (Positive frozen preferences)**.** *Starting from positive preferences, every finite sequence of updates* (18) *remains in* $\operatorname{int}\Delta^{m-1}$*. After the finite formation stage, the frozen collection therefore admits a common constant* $\omega > 0$ *such that* $w_{ij} \geq \omega$ *for all solutions and objectives.*

Proof. For every finite input vector, clipping at $\epsilon_w > 0$ produces strictly positive coordinates with a finite positive sum. Renormalization therefore produces an interior simplex point. Induction gives positivity throughout the finite adaptation phase. At the freezing time there are only finitely many positive coordinates, so their minimum is positive; call it ω. Since the preferences are then fixed, the same lower bound holds for all subsequent iterations. □

The anchor in (18) penalizes drift but does not mathematically prevent two preferences from becoming equal, nor does it guarantee distinct limiting solutions. No disjoint preference-cell condition is imposed.

Visited archive and local gap targets. Preference adaptation changes which trade-offs individual scalar subproblems emphasize, but it does not directly identify uncovered regions of the currently observed front. We therefore maintain a visited nondominated archive $\mathcal{A}^k$ in objective space. For $m = 2$, candidate gaps are formed by adjacent archive points after sorting by the first objective. For $m \geq 3$, we form a nearest-neighbor graph in archive-normalized objective space. For an eligible edge (a, b), write

$$d_{ab} = \|y_a - y_b\|_2$$

and let $\widetilde{d}$ denote the median positive nearest-neighbor distance. We only consider edges satisfying

$$\widetilde{d} < d_{ab} \leq \chi\widetilde{d}, \qquad \chi > 1. \tag{19}$$

The lower inequality identifies an unusually large local gap, while the upper inequality filters very long edges that are more likely to connect unrelated regions. This is a local geometric filter, not a certificate that the edge lies on a single connected component of the true Pareto front.

For an eligible edge, the candidate decision point is the feasible midpoint

$$\widehat{x} = \operatorname{Proj}_{\mathcal{X}}\left(\frac{x_a + x_b}{2}\right).$$

It is activated only when $F(\widehat{x})$ is not dominated by the visited archive and the local two-edge surrogate improves:

$$G_{ab}(\widehat{x}) := \frac{\big[d_{ab}^2 - \|y(\widehat{x}) - y_a\|_2^2 - \|y_b - y(\widehat{x})\|_2^2\big]_+}{\max\{d_{ab}^2, \epsilon_G\}} > 0. \tag{20}$$

The numerator compares the original squared edge length with the two squared edges obtained by inserting the candidate. Thus G_{ab} measures a directly observable local reduction in the gap surrogate; it does not require access to the unknown true Pareto front.

Proposition B.2 (Local archive-gap certificate)**.** *Every activated target* $\widehat{x}$ *is nondominated with respect to the visited archive and strictly reduces the local two-edge squared-gap surrogate. More precisely,*

$$\frac{d_{ab}^2 - \|y(\widehat{x}) - y_a\|_2^2 - \|y_b - y(\widehat{x})\|_2^2}{d_{ab}^2} \geq G_{ab}(\widehat{x}) > 0.$$

For $m \geq 3$*, no edge with* $d_{ab} > \chi\widetilde{d}$ *can be activated.*

Proof. Activation requires that $F(\widehat{x})$ is not dominated by any objective vector in $\mathcal{A}^k$, proving the first statement. Let

$$S = \|y(\widehat{x}) - y_a\|_2^2 + \|y_b - y(\widehat{x})\|_2^2, \qquad D = \max\{d_{ab}^2, \epsilon_G\}.$$

Since $G_{ab}(\widehat{x}) > 0$, the positive-part operator in (20) is active and

$$G_{ab}(\widehat{x})D = d_{ab}^2 - S > 0.$$

Hence $S < d_{ab}^2$. Moreover,

$$\frac{d_{ab}^2 - S}{d_{ab}^2} = G_{ab}(\widehat{x})\frac{D}{d_{ab}^2} \geq G_{ab}(\widehat{x}),$$

because $D \geq d_{ab}^2$. The final statement follows directly from the eligibility condition (19). □

The certificate is intentionally local. It does not imply monotonic IGD improvement, discovery of an unseen Pareto component, or preservation of true-front connectivity; nor does candidate attraction guarantee that an actual solution reaches the midpoint.

Interaction and coverage signals. The raw population signal b_i used in Section 2.2 combines a crowding signal and, when available, an archive-guided coverage signal. Let $J_y(x) \in \mathbb{R}^{m\times d}$ denote the Jacobian of the normalized objective map. We use the Gaussian crowding potential

$$U_i(y_1, \ldots, y_N) := \frac{1}{N-1}\sum_{q\neq i} \exp\left(-\frac{\|y_i - y_q\|_2^2}{2h^2}\right). \tag{21}$$

This potential is largest when solutions coincide. Descending it therefore generates a repulsive decision-space signal

$$b_i^{\text{int}} := -J_y(x_i)^\top \nabla_{y_i} U_i = J_y(x_i)^\top \frac{1}{h^2(N-1)}\sum_{q\neq i} K_{iq}(y_i - y_q), \qquad K_{iq} = \exp\left(-\frac{\|y_i - y_q\|_2^2}{2h^2}\right). \tag{22}$$

An activated archive target $\widehat{x}$ generates

$$b_i^{\text{cov}} := -J_y(x_i)^\top \big(F(x_i) - F(\widehat{x})\big). \tag{23}$$

Treating the target and scales as fixed, this is the negative gradient of

$$\frac{1}{2}\sum_{j=1}^{m} \frac{(f_j(x_i) - f_j(\widehat{x}))^2}{s_j}.$$

Thus the first signal discourages local crowding, whereas the second attracts a solution toward a locally underrepresented objective-space region. The raw cooperative signal is

$$b_i = \lambda_{\text{int}} b_i^{\text{int}} + \lambda_{\text{cov}} q_i G_{ab}(\widehat{x}) b_i^{\text{cov}}. \tag{24}$$

During the frozen phase, Section 2.2 maps this raw signal into the gradient-orthogonal component through (8)–(9). During formation, a projected candidate may instead use the raw signal directly.

Transient energy guard and freezing. Formation deliberately permits non-tangential motion because its purpose is to reshape the population, not to preserve one fixed Lyapunov function while (w_i, z, s) are changing. A finite energy guard accepts a formation candidate only when the increase of the current scalar potential stays within the prescribed interaction–coverage budget; otherwise the algorithm falls back to the decomposition descent direction. All formation-only non-tangential moves and every update to (w_i, z, s) terminate at the prescribed horizon K.

Proposition B.3 (Finite-prefix neutrality)**.** *Suppose all formation-only non-tangential moves and all changes to (w_i, z, s) terminate after a finite K, and formation produces $x_i^K \in \mathcal{X}$. Then every asymptotic conclusion proved for the frozen recursion from an arbitrary feasible initial state applies from x_i^K. The formation phase can therefore change the initial state, and hence the basin entered by the frozen dynamics, but not the asymptotic implication of the frozen analysis.*

Proof. Fix a solution and denote the complete frozen update map, including its accepted backtracking step, by R_k. For $k \geq K$,

$$x^{k+1} = R_k(x^k)$$

with fixed scalarization parameters. Define $\bar{x}^r = x^{K+r}$ and $\bar{R}_r = R_{K+r}$. Then

$$\bar{x}^{r+1} = \bar{R}_r(\bar{x}^r), \qquad \bar{x}^0 = x^K \in \mathcal{X}.$$

All assumptions used by the frozen analysis are properties of this tail recursion and are unchanged by the re-indexing. Adding or removing finitely many prefix iterates does not change convergence of the tail, summability of its residuals, or its accumulation points. Hence formation can select x^K and thereby affect the basin from which the frozen recursion starts, but it cannot alter any asymptotic implication proved for every feasible frozen initial state. □

For LZ09-F7 only, the reported formation stage additionally proposes Gaussian perturbations of one randomly selected decision coordinate per solution. A clipped proposal is accepted only when the current decomposed scalar value strictly decreases. This experiment-specific basin-hopping mechanism is disabled before freezing and is not part of the continuous normal–tangent theory.

Finite returned set. At a finite stopping time, the synthetic experiments return a fixed-size selection from the nondominated union of the visited archive and final solutions. Archived earlier iterates are not covered by a theorem about limiting solution positions, so the archive and readout rule are part of the empirical finite-budget method rather than the asymptotic convergence claim.

C Proofs for the Decomposition and Cooperative Flow

The main arguments rely on two reusable analytical facts. The first concerns the smooth augmented Tchebycheff scalarization; the second is the orthogonal decomposition of a vector into tangent and normal components. We use the normal-cone convention

$$N_{\mathcal{X}}(x) = \{\zeta \in \mathbb{R}^d : \langle \zeta, y - x \rangle \leq 0 \text{ for every } y \in \mathcal{X}\}.$$

C.1 Properties of the Smooth Augmented Scalarization

Define

$$\ell_\tau(u) = \tau \log \sum_{j=1}^{m} \exp(u_j/\tau).$$

Lemma C.1 (Properties of log-sum-exp and the frozen scalar potential)**.** *For $\tau > 0$, the map $\ell_\tau : \mathbb{R}^m \to \mathbb{R}$ satisfies:*

1. *$\nabla \ell_\tau(u) = a(u)$ with*
$$a_j(u) = \frac{\exp(u_j/\tau)}{\sum_r \exp(u_r/\tau)} > 0.$$
2. *It is convex, strictly increasing in every coordinate, translation equivariant, and one-Lipschitz from ℓ_∞ to $\mathbb{R}$:*
$$\ell_\tau(u + c\mathbf{1}) = \ell_\tau(u) + c, \qquad |\ell_\tau(u) - \ell_\tau(v)| \leq \|u - v\|_\infty.$$
3. *Under the frozen assumptions in Theorem 3.1, every ϕ_i is continuously differentiable on a neighborhood of $\mathcal{X}$, bounded below on $\mathcal{X}$, and has a Lipschitz gradient on $\mathcal{X}$. If every f_j is convex, then ϕ_i is convex. Regardless of convexity, ϕ_i is strictly increasing in every objective value whenever $w_{ij} > 0$.*

Proof. Direct differentiation gives

$$\frac{\partial \ell_\tau(u)}{\partial u_j} = a_j(u) > 0, \qquad \nabla^2 \ell_\tau(u) = \frac{1}{\tau}\{\mathrm{Diag}(a) - aa^\top\}.$$

For every $q \in \mathbb{R}^m$,

$$q^\top \{\mathrm{Diag}(a) - aa^\top\} q = \sum_j a_j q_j^2 - \left(\sum_j a_j q_j \right)^2 \geq 0,$$

because the right-hand side is the variance of the values q_j under probability vector a. Thus ℓ_τ is convex, while positivity of every partial derivative gives strict coordinatewise monotonicity.

Translation equivariance follows by factoring $\exp(c/\tau)$ from the defining sum. Let $D = \|u - v\|_\infty$. The componentwise inequalities $u \leq v + D\mathbf{1}$ and $v \leq u + D\mathbf{1}$, together with monotonicity and translation equivariance, yield

$$\ell_\tau(u) \leq \ell_\tau(v) + D, \qquad \ell_\tau(v) \leq \ell_\tau(u) + D,$$

which proves the Lipschitz bound.

For the regularity claim, compactness of $\mathcal{X}$ and continuity of the objective gradients imply that every ∇f_j is bounded on $\mathcal{X}$. Hence every f_j, normalized objective y_j, and weighted input $w_{ij} y_j$ is Lipschitz. The softmax coefficients are Lipschitz on the compact image of $\mathcal{X}$. Therefore the coefficients

$$\beta_{ij}(x) = (a_{ij}(x) + \nu) \frac{w_{ij}}{s_j}$$

are Lipschitz, and

$$\nabla \phi_i(x) = \sum_j \beta_{ij}(x) \nabla f_j(x)$$

has a Lipschitz gradient on $\mathcal{X}$. Continuity on a compact set gives a finite lower bound.

If each f_j is convex, then each $u_{ij}(x) = w_{ij}(f_j(x) - z_j)/s_j$ is convex. Convexity and coordinatewise monotonicity of ℓ_τ preserve convexity under this componentwise composition, and the augmented term is a positive weighted sum of convex functions. Finally,

$$\frac{\partial \phi_i}{\partial f_j} = (a_{ij}(x) + \nu) \frac{w_{ij}}{s_j} > 0,$$

which proves strict objectivewise monotonicity. □

C.2 Tangent–Normal Geometry

Lemma C.2 (Tangent–normal decomposition)**.** *Let $\mathcal{X} \subset \mathbb{R}^d$ be nonempty, closed, and convex. For $x \in \mathcal{X}$, $g \in \mathbb{R}^d$, and*

$$n = \mathrm{Proj}_{T_{\mathcal{X}}(x)}(-g),$$

there is a unique $\zeta \in N_{\mathcal{X}}(x)$ such that

$$-g = n + \zeta, \qquad \langle n, \zeta \rangle = 0. \tag{25}$$

Consequently,

$$\langle g, n \rangle = -\|n\|_2^2. \tag{26}$$

If $t \in L_{\mathcal{X}}(x)$, then $\langle \zeta, t \rangle = 0$. Moreover, the graph of the normal-cone map is closed: $x_r \to x$, $\zeta_r \to \zeta$, and $\zeta_r \in N_{\mathcal{X}}(x_r)$ imply $\zeta \in N_{\mathcal{X}}(x)$.

Proof. The tangent cone is a closed convex cone and its polar is the normal cone. Moreau's orthogonal cone decomposition (Moreau, 1965) applied to $-g$ gives the unique decomposition (25). Taking its inner product with n proves (26).

If $t \in L_{\mathcal{X}}(x)$, then both t and $-t$ belong to the tangent cone. Polarity gives $\langle \zeta, t \rangle \leq 0$ and $\langle \zeta, -t \rangle \leq 0$, hence $\langle \zeta, t \rangle = 0$.

For closedness, fix any $y \in \mathcal{X}$. Since $\langle \zeta_r, y - x_r \rangle \leq 0$, passing to the limit gives $\langle \zeta, y - x \rangle \leq 0$. This holds for every $y \in \mathcal{X}$, so $\zeta \in N_{\mathcal{X}}(x)$. □

C.3 Proof of the Positive-Gradient Decomposition

This interpretation also appears directly in the gradient of ϕ_i. The log-sum-exp term assigns a softmax coefficient to each weighted objective: objectives with larger preference-scaled deviations exert a stronger influence on the descent direction.

Lemma C.3 (Positive-gradient decomposition)**.** *For any* $w_i \in \operatorname{int} \Delta^{m-1}$,

$$g_i(x) := \nabla \phi_i(x) = \sum_{j=1}^{m} \left(a_{ij}(x) + \nu\right) \frac{w_{ij}}{s_j} \nabla f_j(x), \tag{27}$$

where

$$a_{ij}(x) = \frac{\exp(w_{ij} y_j(x)/\tau)}{\sum_{r=1}^{m} \exp(w_{ir} y_r(x)/\tau)}. \tag{28}$$

Every coefficient in (27) *is strictly positive. Consequently, constrained stationarity of* ϕ_i *implies Pareto criticality of the original multi-objective problem.*

Lemma C.3 gives the key link between decomposition and the original MOO problem. The scalar gradient is not an arbitrary search direction: it is a strictly positive mixture of the objective gradients. One may therefore view $-\nabla \phi_i$ as the resultant force produced by all objectives, with the softmax coefficients dynamically emphasizing the objectives that are currently most violated under preference w_i. At a constrained stationary point, this resultant force is either zero in the interior or balanced by the boundary normal force. After normalization of its positive coefficients, this is exactly the Pareto-criticality condition in (2).

Proof of Lemma C.3. Fix i and define

$$u_{ij}(x) = w_{ij} \frac{f_j(x) - z_j}{s_j}.$$

By the chain rule and Lemma C.1,

$$\nabla \ell_\tau(u_i(x)) = \sum_{j=1}^{m} a_{ij}(x) \frac{w_{ij}}{s_j} \nabla f_j(x),$$

whereas

$$\nabla \left(\nu \sum_{j=1}^{m} w_{ij} y_j(x) \right) = \sum_{j=1}^{m} \nu \frac{w_{ij}}{s_j} \nabla f_j(x).$$

Adding the two identities gives the gradient expression in Lemma C.3. Since $a_{ij}(x) > 0$, $w_{ij} > 0$, and $s_j > 0$, every coefficient

$$\beta_{ij}(x) = \left(a_{ij}(x) + \nu\right) \frac{w_{ij}}{s_j}$$

is strictly positive.

Suppose now that

$$0 \in \nabla \phi_i(x) + N_{\mathcal{X}}(x).$$

Then some $\zeta \in N_{\mathcal{X}}(x)$ satisfies

$$\sum_j \beta_{ij}(x) \nabla f_j(x) + \zeta = 0.$$

Let

$$B_i = \sum_j \beta_{ij}(x) > 0, \qquad \lambda_{ij} = \beta_{ij}(x)/B_i.$$

Then $\lambda_i \in \operatorname{int} \Delta^{m-1}$. Since the normal cone is a cone, $\zeta/B_i \in N_{\mathcal{X}}(x)$, and division by B_i gives

$$0 \in \sum_{j=1}^{m} \lambda_{ij} \nabla f_j(x) + N_{\mathcal{X}}(x),$$

which is exactly the Pareto-criticality condition (2). □

C.4 Proof of Descent-Preserving Cooperation

Lemma C.4 (Descent-preserving tangential cooperation)**.** *For every nonstationary solution,* $v_i \in T_{\mathcal{X}}(x_i)$ *and*

$$\langle g_i, v_i \rangle = -\|n_i\|_2^2, \qquad -\left\langle g_i, \frac{v_i}{\|v_i\|_2} \right\rangle \geq \frac{\|n_i\|_2}{\sqrt{1+\rho^2}} > 0. \tag{29}$$

Lemma C.4 separates the two roles in the dynamics. The normal component n_i determines the first-order decrease of the scalar subproblem, whereas the tangent component t_i can change the direction of motion without contributing to this decrease. The second inequality further ensures that, after normalization, the composite velocity retains a quantitative downhill component even in the presence of cooperation.

Proof of Lemma C.4. We suppress the solution index. Because g_L and b_L lie in the linear space $L_{\mathcal{X}}(x)$, the numerator of (8), and hence $\widetilde{t}$, also lies in this space. Orthogonal projection onto the lineality space gives

$$\langle g, b_L \rangle = \langle g_L, b_L \rangle, \qquad \langle g, g_L \rangle = \|g_L\|_2^2.$$

Therefore

$$\langle g, \widetilde{t} \rangle = \frac{\|g_L\|_2^2 \langle g_L, b_L \rangle - \langle g_L, b_L \rangle \|g_L\|_2^2}{\|g_L\|_2^2 + \varepsilon^2} = 0.$$

The scalar cap in (9) preserves both membership in the lineality space and orthogonality to g, while ensuring

$$\|t\|_2 \leq \rho \|n\|_2.$$

In particular, $n = 0$ implies $t = 0$.

By Lemma C.2,

$$-g = n + \zeta, \qquad \zeta \in N_{\mathcal{X}}(x), \qquad \langle n, \zeta \rangle = 0.$$

Since $t \in L_{\mathcal{X}}(x)$, $\langle \zeta, t \rangle = 0$. Taking the inner product of $-g = n + \zeta$ with t and using $\langle g, t \rangle = 0$ gives $\langle n, t \rangle = 0$. The tangent cone is a convex cone, so $n \in T_{\mathcal{X}}(x)$ and $t \in L_{\mathcal{X}}(x) \subseteq T_{\mathcal{X}}(x)$ imply $v = n + t \in T_{\mathcal{X}}(x)$.

Finally,

$$\langle g, v \rangle = \langle g, n \rangle + \langle g, t \rangle = -\|n\|_2^2.$$

Moreover, orthogonality of n and t and the relative cap imply

$$\|v\|_2^2 = \|n\|_2^2 + \|t\|_2^2 \leq (1+\rho^2)\|n\|_2^2.$$

For $n \neq 0$,

$$-\left\langle g, \frac{v}{\|v\|_2} \right\rangle = \frac{\|n\|_2^2}{\|v\|_2} \geq \frac{\|n\|_2}{\sqrt{1+\rho^2}},$$

which proves the result. □

C.5 Proof of the Lyapunov Identity

Under the conditions of Proposition 2.1, the integrated identity is

$$\phi_i(x_i(T)) + \int_0^T \|n_i(x_i(t))\|_2^2 \, \mathrm{d}t = \phi_i(x_i(0)). \tag{30}$$

Equation (30) gives an energy interpretation of the cooperative flow: the decrease of the frozen scalar potential is exactly the accumulated normal dissipation, while the cooperative motion contributes no additional first-order energy change. In this sense, decomposition and cooperation play complementary dynamical roles: decomposition supplies the dissipative motion of each subproblem, whereas cooperation transports the population through descent-neutral directions.

Proof of Proposition 2.1. By Lemma C.1, the frozen potential ϕ_i is continuously differentiable. For every feasible absolutely continuous trajectory, the chain rule gives, for almost every t,

$$\frac{\mathrm{d}}{\mathrm{d}t}\phi_i(x_i(t)) = \langle \nabla\phi_i(x_i(t)), \dot{x}_i(t)\rangle\,.$$

Substituting the flow (12) and applying Lemma C.4,

$$\frac{\mathrm{d}}{\mathrm{d}t}\phi_i(x_i(t)) = \langle g_i, n_i + t_i\rangle = -\|n_i\|_2^2 \le 0.$$

Integrating from 0 to T gives

$$\phi_i(x_i(T)) + \int_0^T \|n_i(x_i(t))\|_2^2\,\mathrm{d}t = \phi_i(x_i(0)),$$

which is (30). Since ϕ_i is bounded below on the compact feasible set, the total normal dissipation is finite along every infinite feasible trajectory. □

D Discrete Descent and Convergence

The continuous normal–tangent construction controls the instantaneous descent angle. A finite projected step, however, may be shortened or rotated by projection and backtracking. The main-text condition (16) therefore isolates the exact finite-step property needed by the discrete convergence proof.

D.1 A Sufficient Finite-Step Condition

Lemma D.1 (Sufficient finite-step condition)**.** *Fix a frozen iteration with $n \neq 0$ and let $d = \Delta x(\alpha)$ be an accepted projected displacement. Suppose that for constants $\sigma > 0$ and $a > 0$,*

$$-\langle g, d\rangle \ge \sigma\|n\|_2\|d\|_2, \qquad \|d\|_2 \ge a\|n\|_2. \tag{31}$$

Then (16) *holds with $\kappa = \sigma a$. If, in addition, $\nabla\phi$ is L_ϕ-Lipschitz and the first inequality in* (31) *holds, every displacement satisfying*

$$\|d\|_2 \le \frac{2(1-c)\sigma}{L_\phi}\|n\|_2 \tag{32}$$

satisfies the strict Armijo condition $\phi(x+d) \le \phi(x) + c\langle g, d\rangle$.

Proof. Multiplying the two inequalities in (31) gives

$$-\langle g, d\rangle \ge \sigma a\|n\|_2^2,$$

which is (16) with $\kappa = \sigma a$.

Lipschitz continuity of $\nabla\phi$ gives the descent lemma

$$\phi(x+d) \le \phi(x) + \langle g, d\rangle + \frac{L_\phi}{2}\|d\|_2^2.$$

By the alignment condition and (32),

$$\frac{L_\phi}{2}\|d\|_2^2 \le (1-c)\sigma\|n\|_2\|d\|_2 \le -(1-c)\langle g, d\rangle.$$

Substitution yields

$$\phi(x+d) \le \phi(x) + c\langle g, d\rangle.$$

For the unprojected normalized composite direction, Lemma C.4 gives $\sigma = 1/\sqrt{1+\rho^2}$. The main-text nondegeneracy condition therefore asks the accepted projected displacement to retain a uniform product of alignment and effective length after projection and backtracking. □

D.2 Proof of the Discrete Convergence Theorem

Proof of Theorem 3.1. Fix a solution i and suppress the index. The proof has four steps: quasi-Lyapunov descent, residual summability, constrained stationarity of every accumulation point, and conversion from scalar stationarity to Pareto criticality.

Step 1: quasi-Lyapunov descent. The accepted condition (15) and (16) imply

$$\phi(x^{k+1}) \le \phi(x^k) - c\kappa\|n^k\|_2^2 + \epsilon_k. \tag{33}$$

Since $\sum_{k=K}^{\infty} \epsilon_k < \infty$, define

$$E_k = \sum_{\ell=k}^{\infty} \epsilon_\ell, \qquad V^k = \phi(x^k) + E_k.$$

Because $E_{k+1} = E_k - \epsilon_k$, adding E_{k+1} to (33) gives

$$V^{k+1} \le V^k - c\kappa\|n^k\|_2^2. \tag{34}$$

Hence $\{V^k\}$ is nonincreasing. By Lemma C.1, ϕ is bounded below on the compact feasible set, and $E_k \ge 0$. Therefore V^k converges to a finite limit. Since $E_k \to 0$, the scalar values $\phi(x^k) = V^k - E_k$ converge to the same limit.

Step 2: residual summability. Summing (34) from $k = K$ to M gives

$$c\kappa \sum_{k=K}^{M} \|n^k\|_2^2 \le V^K - V^{M+1}.$$

Letting $M \to \infty$ yields

$$\sum_{k=K}^{\infty} \|n^k\|_2^2 < \infty,$$

and therefore $\|n^k\|_2 \to 0$.

Step 3: stationarity of every accumulation point. Compactness of $\mathcal{X}$ gives at least one accumulation point. Let $x^{k_r} \to x^\star$ be an arbitrary convergent subsequence. By Lemma C.2,

$$\zeta^{k_r} = -g(x^{k_r}) - n^{k_r} \in N_{\mathcal{X}}(x^{k_r}).$$

Continuity of $g = \nabla\phi$ and $n^{k_r} \to 0$ imply

$$\zeta^{k_r} \to -g(x^\star).$$

The closed-graph property of the normal-cone map then gives

$$-g(x^\star) \in N_{\mathcal{X}}(x^\star),$$

or equivalently

$$0 \in \nabla\phi(x^\star) + N_{\mathcal{X}}(x^\star).$$

Because the convergent subsequence was arbitrary, every accumulation point is a constrained stationary point of the frozen scalar subproblem.

Step 4: Pareto criticality. Lemma C.3 expresses $\nabla\phi(x^\star)$ as a strictly positive combination of the objective gradients. Normalizing those positive coefficients converts the preceding stationarity inclusion into (2). Hence every accumulation point is Pareto critical for the original multi-objective problem. □

E Conditional Extensions

The main text deliberately separates the basic Pareto-critical convergence statement from stronger claims that require additional assumptions. We collect two such extensions here.

E.1 Convex Optimality and a Local PL Rate

Corollary E.1 (Convex optimality and local PL convergence)**.** *If every f_j is convex, every accumulation point supplied by Theorem 3.1 is Pareto optimal. If, in a frozen basin, the projected Polyak–Łojasiewicz condition*

$$\frac{1}{2}\|n_i(x)\|_2^2 \geq \mu\big(\phi_i(x) - \phi_i^\star\big) \tag{35}$$

holds for some $\mu > 0$, then strict Armijo yields a local linear decrease of the scalar gap. With summable errors, the gap obeys a linear contraction plus a vanishing convolution of the error sequence.

Proof. If every f_j is convex, Lemma C.1 implies that the frozen potential ϕ_i is convex. Let $x^\star$ be an accumulation point from Theorem 3.1. There exists $\zeta^\star \in N_{\mathcal{X}}(x^\star)$ such that

$$\nabla\phi_i(x^\star) + \zeta^\star = 0.$$

For every $x \in \mathcal{X}$, convexity and the normal-cone inequality give

$$\phi_i(x) \geq \phi_i(x^\star) + \langle \nabla\phi_i(x^\star), x - x^\star\rangle = \phi_i(x^\star) - \langle \zeta^\star, x - x^\star\rangle \geq \phi_i(x^\star).$$

Thus $x^\star$ is a global minimizer of the frozen scalar subproblem. If some feasible x dominated $x^\star$, strict objectivewise monotonicity of ϕ_i would imply $\phi_i(x) < \phi_i(x^\star)$, a contradiction. Hence $x^\star$ is Pareto optimal.

For the rate statement, write

$$\delta_k = \phi_i(x_i^k) - \phi_i^\star.$$

Combining the quasi-descent inequality with (35) gives

$$\delta_{k+1} \leq (1 - 2c\kappa\mu)\delta_k + \epsilon_k \leq (1-\eta)\delta_k + \epsilon_k, \qquad \eta = \min\{2c\kappa\mu, 1/2\} \in (0,1).$$

For every starting index K_0 inside the basin,

$$\delta_k \leq (1-\eta)^{k-K_0}\delta_{K_0} + \sum_{r=K_0}^{k-1}(1-\eta)^{k-1-r}\epsilon_r. \tag{36}$$

When $\epsilon_r = 0$, this is a linear contraction. For summable nonnegative errors, $\epsilon_r \to 0$ and the convolution in (36) vanishes, so $\delta_k \to 0$. □

E.2 From Empirical to Population Trade-offs

The optimization results above concern the objective functions supplied to the optimizer. For a learning problem, let $\widehat{R}_j$ and R_j denote empirical and population risks. The following result separates statistical estimation from the optimization analysis while retaining the same frozen scalarization.

Proposition E.2 (Uniform transfer of the augmented decomposition)**.** *Suppose that, with probability at least $1-\delta$,*

$$\sup_{x\in\mathcal{X}} |R_j(x) - \widehat{R}_j(x)| \leq \varepsilon_j, \qquad j = 1, \dots, m. \tag{37}$$

Use the same frozen (w_i, z, s, τ, ν) to define the empirical and population versions of (4)*, and let*

$$D_i := \max_j \frac{w_{ij}\varepsilon_j}{s_j} + \nu\sum_{j=1}^m \frac{w_{ij}\varepsilon_j}{s_j}. \tag{38}$$

On the event (37)*,*

$$\sup_{x\in\mathcal{X}} |\phi_i^{\mathrm{pop}}(x) - \widehat{\phi}_i(x)| \leq D_i.$$

If

$$\widehat{\phi}_i(\widehat{x}_i) \leq \inf_{x\in\mathcal{X}} \widehat{\phi}_i(x) + \eta_i,$$

then

$$\phi_i^{\mathrm{pop}}(\widehat{x}_i) \leq \inf_{x\in\mathcal{X}} \phi_i^{\mathrm{pop}}(x) + \eta_i + 2D_i. \tag{39}$$

In particular, there is no $x \in \mathcal{X}$ and $\gamma > 0$ satisfying

$$R_j(x) \leq R_j(\widehat{x}_i) - \gamma\frac{s_j}{w_{ij}} \quad \textit{for every } j, \qquad (1+\nu m)\gamma > \eta_i + 2D_i. \tag{40}$$

Proof. Define the weighted normalized population and empirical inputs

$$u_j^{\mathrm{pop}}(x) = \frac{w_{ij}(R_j(x) - z_j)}{s_j}, \qquad \widehat{u}_j(x) = \frac{w_{ij}(\widehat{R}_j(x) - z_j)}{s_j}.$$

The event (37) implies, uniformly in x,

$$\|u^{\mathrm{pop}}(x) - \widehat{u}(x)\|_\infty \le \max_j \frac{w_{ij}\varepsilon_j}{s_j}.$$

By the ℓ_∞-Lipschitz property of log-sum-exp in Lemma C.1, the difference between the two smooth Tchebycheff terms is bounded by the right-hand side. For the augmented terms,

$$\left| \nu \sum_j \frac{w_{ij}}{s_j} \big(R_j(x) - \widehat{R}_j(x)\big) \right| \le \nu \sum_j \frac{w_{ij}\varepsilon_j}{s_j}.$$

Adding the two bounds proves

$$\sup_x |\phi_i^{\mathrm{pop}}(x) - \widehat{\phi}_i(x)| \le D_i.$$

For every $x \in \mathcal{X}$,

$$\phi_i^{\mathrm{pop}}(\widehat{x}_i) \le \widehat{\phi}_i(\widehat{x}_i) + D_i \le \widehat{\phi}_i(x) + \eta_i + D_i \le \phi_i^{\mathrm{pop}}(x) + \eta_i + 2D_i.$$

Taking the infimum over x proves (39).

Finally, suppose an x and $\gamma > 0$ satisfy (40). Multiplying each component inequality by w_{ij}/s_j gives

$$u_j^{\mathrm{pop}}(x) \le u_j^{\mathrm{pop}}(\widehat{x}_i) - \gamma \qquad \text{for every } j.$$

Translation equivariance and monotonicity of log-sum-exp imply that the smooth Tchebycheff term decreases by at least γ. The augmented term decreases by at least $\nu m\gamma$. Hence

$$\phi_i^{\mathrm{pop}}(x) \le \phi_i^{\mathrm{pop}}(\widehat{x}_i) - (1 + \nu m)\gamma,$$

which contradicts (39) whenever $(1 + \nu m)\gamma > \eta_i + 2D_i$. □

This transfer statement is conditional. Neither the uniform risk event (37) nor near-global empirical scalar optimality follows from nonconvex stationarity. It should therefore not be interpreted as a generalization theorem for the neural-network experiments.

F Experimental Details and Additional Results

This section gives the evaluation protocol, implementation settings, and additional results for the experiments in Section 4. We first specify the metrics and synthetic problems, then report the remaining front plots and controlled component removals. Finally, we describe the fairness and ranking tasks and clarify the scope of the comparisons. All numerical results below are obtained from the saved experimental outputs.

F.1 Metrics and Reporting Protocol

Inverted generational distance. IGD measures how well the returned set represents a known front by averaging the distance from reference Pareto points to their nearest returned solutions. It therefore reflects both proximity to the front and uncovered regions.

Let S be the returned objective vectors and $P^\star$ a finite reference set on the analytic Pareto front. The reported IGD (Ishibuchi et al., 2015) is

$$\mathrm{IGD}(P^\star, S) = \frac{1}{|P^\star|} \sum_{y \in P^\star} \min_{s \in \mathrm{ND}(S)} \|y - s\|_2, \tag{41}$$

where $\mathrm{ND}(S)$ is the nondominated subset used by the evaluation code, with coordinate tolerance 10^{-10}. Distances are computed in the original objective coordinates, without normalization. For

ZDT1/2 and LZ09-F2/F7/F9, the reference set contains 2,000 points sampled uniformly in $f_1 \in [0, 1]$. For ZDT3, we sample 400 points in each of the five intervals listed below, giving 2,000 points in total. For DTLZ7, the retained evaluator constructs a 44×44 grid in $(f_1, f_2) \in [0, 1]^2$, evaluates the analytic $g = 1$ surface, and removes dominated grid points, leaving 484 reference points. DTLZ7 IGD uses this fixed 484-point reference. The reference sets are shared by methods and seeds within each problem.

Hypervolume. HV measures the objective-space volume dominated by the returned set relative to a fixed reference point. It summarizes convergence and trade-off coverage without requiring the true Pareto front, making it suitable for the real learning tasks.

For minimization objectives and a fixed reference vector r, we compute (Guerreiro et al., 2021)

$$\mathrm{HV}(S; r) = \mathrm{vol}_m \left(\bigcup_{s \in S:\, s \le r} \prod_{j=1}^{m} [s_j, r_j] \right). \tag{42}$$

The reference is $(1.1, 1.1)$ for all two-objective synthetic problems, $(1.1, 1.1, 6.6)$ for DTLZ7, and $(1, 1)$ for fairness classification. HV is evaluated in the original objective coordinates; points outside the reference box contribute zero volume. For MSLR, the stored objectives are negative NDCG@10 values and the minimization reference is 0^6. Equivalently, HV measures the volume covered between the origin and the returned positive NDCG vectors. Method comparisons use a fixed reference across seeds within each task.

Fill distance. To complement average reference-to-set distance, the component analysis also reports the worst uncovered distance,

$$h(\mathcal{P}, S) = \sup_{y \in \mathcal{P}} \min_{s \in S} \|y - s\|_2, \tag{43}$$

following the coverage measure used by Zhang et al. (2024). Here $\mathcal{P}$ is the continuous analytic front and S contains all 50 returned points, including dominated points. For ZDT3, the numerical evaluation treats its five connected components separately and takes the largest distance. The saved evaluations bound numerical uncertainty below 10^{-7} in the original objective coordinates.

Aggregation and selection. Synthetic results use six seeds, numbered 0–5; learning results use three seeds, numbered 0–2. We report the arithmetic mean and sample standard deviation, with denominator $n - 1$. Performance ranks are zero-based and computed from unrounded means: lower IGD or higher HV receives a better rank. Exact ties share a rank; boldface and underlining identify rank 0 and rank 1. The average rank in the synthetic tables averages all seven problems for the eight displayed set-valued methods. The ranked comparison comprises eight set-valued methods. Synthetic metrics use the terminal returned sets. For the learning tables, the whole-population checkpoint is selected by maximum validation HV, and the corresponding test output is reported. Appendix F.7 describes the historical test traces.

F.2 Synthetic Benchmarks and Implementations

Problems. Table 3 lists the dimensions, population sizes, and MOGF/D budgets. ZDT1 and ZDT2 have fronts $f_2 = 1 - \sqrt{f_1}$ and $f_2 = 1 - f_1^2$, respectively. ZDT3 uses $f_2 = 1 - \sqrt{f_1} - f_1 \sin(10\pi f_1)$ on the five intervals

$$[0, 0.0830015349], \quad [0.1822287280, 0.2577623634], \quad [0.4093136748, 0.4538821041],$$
$$[0.6183967944, 0.6525117038], \quad [0.8233317983, 0.8518328654].$$

The DTLZ7 reference surface is $f_3 = 6 - f_1(1 + \sin(3\pi f_1)) - f_2(1 + \sin(3\pi f_2))$; only its nondominated patches form the true front. Figures display analytic curves and surfaces; IGD uses the discrete references in Appendix F.1.

For the LZ09 problems, we use the box parameterization of the retained implementation of Li & Zhang (2009). Write $t = x_1$ and $z_j = 2x_j - 1$ for $j = 2, \ldots, d$. Define $J_1 = \{j \ge 2 : j \text{ is odd}\}$ and $J_2 = \{j \ge 2 : j \text{ is even}\}$. The implemented objectives are

$$f_1(x) = t + \frac{2}{|J_1|} \sum_{j \in J_1} q_j(x), \qquad f_2(x) = b(t) + \frac{2}{|J_2|} \sum_{j \in J_2} q_j(x). \tag{44}$$

Table 3: Synthetic settings: dimension d, objective count m, population size N, and first fully frozen iteration K (zero-based). All domains are $[0, 1]^d$; each problem uses seeds 0–5.

Problem	d	m	N	Updates	K	Returned
ZDT1	30	2	50	5001	3251	50
ZDT2	30	2	50	5001	4901	50
ZDT3	30	2	50	5001	4901	50
DTLZ7	30	3	200	1001	921	200
LZ09-F2	30	2	50	5001	3251	50
LZ09-F7	10	2	50	5001	3501	50
LZ09-F9	30	2	50	5001	3501	50

For F2 and F9, $q_j(x) = [z_j - \sin(6\pi t + j\pi/d)]^2$; $b(t) = 1 - \sqrt{t}$ for F2 and $b(t) = 1 - t^2$ for F9. For F7, let $a_j = (d + 3j - 8)/[2(d - 2)]$ and $e_j = z_j - t^{a_j}$. Then $q_j(x) = 4e_j^2 - \cos(8\pi e_j) + 1$ and $b(t) = 1 - \sqrt{t}$. The global Pareto set is obtained by setting every residual to zero. The oscillatory F7 penalty additionally produces local basins, although its global objective-space front is continuous. This distinguishes the complex decision-space geometry of F7 from the disconnected objective-space fronts of ZDT3 and DTLZ7.

Baseline overview. We compare against representative gradient-based and population-based Pareto learning methods: (1) ParetoMTL (Lin et al., 2019), which decomposes the MOO problem into preference-guided subproblems; (2) COSMOS (Ruchte & Grabocka, 2021), which learns a scalable preference-conditioned approximation of the Pareto front; (3) PHN-LS and PHN-EPO (Navon et al., 2021), which use hypernetworks to generate preference-dependent solutions with different scalarization schemes; (4) ARGMO-kernel and ARGMO-HV (Chen & Kwok, 2022), which adapt reference vectors during optimization to improve the distribution of Pareto solutions; and (5) Particle-WFR (Ren et al., 2024), an interacting-particle method based on Wasserstein–Fisher–Rao gradient flow for approximating complex Pareto fronts.

Baseline implementations. We use the original constrained minimum-norm update of ParetoMTL (Lin et al., 2019), with analytic-objective adapters for ZDT/DTLZ and LZ09. The two-objective LZ09 adapter uses an equivalent two-dimensional minimum-norm solver. COSMOS, PHN-LS, PHN-EPO, ARGMO-kernel, and ARGMO-HV are run through the vendored Particle-WFR framework (Ren et al., 2024); the two PHN variants retain their linear and EPO (Mahapatra & Rajan, 2020) solvers, and the two ARGMO variants retain kernel and HV reference adaptation. Particle-WFR's synthetic runs follow its direct visualization dynamics, rather than the generic learning-task runner. Table 4 reports the actual budgets. Update units follow each implementation: a preference-conditioned network, a whole population, or one independently optimized preference. Their computational costs therefore vary by method.

Table 4: Synthetic update budgets. The two-objective column covers ZDT1–3 and LZ09-F2/F7/F9. Methods return 50 solutions in two objectives and 200 on DTLZ7.

Method	Update / optimizer	2-objective	DTLZ7
ParetoMTL	Constrained min-norm step	5,000	1,000
COSMOS	Adam, 10^{-4}	5,000	1,000
PHN-LS / PHN-EPO	Adam, 10^{-4}	5,000	1,000
ARGMO-kernel / ARGMO-HV	Adam, 10^{-4}	5,000	1,000
Particle-WFR	Adam, source settings	5,000	1,000
MOGF/D	Projected step + backtracking	5,000	1,000

ParetoMTL uses step sizes 0.10, 0.05, and 0.03 on ZDT1/2/3; 0.04 on DTLZ7; and 0.10, 0.01, and 0.05 on LZ09-F2/F7/F9. Its initial feasibility phase is bounded by $\lfloor 0.2T \rfloor$ iterations, where T is the outer budget. DTLZ7 ParetoMTL uses six completed runs with an inner cap of 250 iterations in both minimum-norm routines and stopping tolerance 10^{-6}. The rerun corrects the inner-loop counter

while retaining 200 preferences, 1,001 outer updates, dimension 30, and the original initialization. The ZDT results use the original implementation.

The direct Particle-WFR implementation uses Adam (Kingma & Ba, 2015) with learning rate 5×10^{-4} for the two-objective problems and 8×10^{-3} for DTLZ7. For two objectives, its source coefficients for Coulomb and dominance forces are 5×10^{-4} and 0.6, with noise scale 10^{-3} and birth–death scale 20. For DTLZ7, these coefficients are 4×10^{-4} and 1, with the same noise scale and birth–death scale 0.1. The retained LZ09 driver initializes 30 decision variables also for F7, whereas the other F7 methods use 10. COSMOS and both PHN variants use Dirichlet concentration 0.1 for the two-objective tasks and $(0.5, 0.1, 0.5)$ for DTLZ7; COSMOS uses cosine penalty coefficient 0.5. ARGMO uses reference learning rate 5×10^{-4} and warm-up 500 in 2D, and 10^{-3} and warm-up 200 for DTLZ7, with one reference adaptation and constraint constant 0.1.

MOGF/D dynamics. Tables 5 and 6 specify the stored configurations. The synthetic implementation differentiates the analytic objectives and uses float32 population states. It projects updates onto the box and uses per-solution backtracking after the finite exploration phase. The initial step is multiplied by a cosine schedule that decreases to the listed minimum-step ratio. Normalization and preference adaptation stop at their listed iteration boundaries; K is their maximum together with the exploration boundary. Before normalization freezes, the reference and scale are obtained from the current population and archive; evaluation metrics still use the original objective coordinates. The implemented scalar potential is $\tau\log\sum_j\exp(w_{ij}\widehat{f}_j/\tau)$, corresponding to $\nu=0$ in the augmented family of Section 2. The tangent regularizer is 10^{-4}, the stationarity tolerance is 10^{-6}, and the boundary tolerance is 2×10^{-5}. Backtracking uses Armijo coefficient 10^{-4}, shrink factor 0.5, at most 16 trials, zero additive acceptance tolerance, and unit direction scaling. Nonzero summable errors have exponent 1.25.

Two-objective preferences follow the source simplex rule, while DTLZ7 uses a simplex lattice. All tasks use inverse preference-to-weight conversion.

The nondominated archive retains objective values and their associated decisions, with the capacity and update interval in Table 6. Gap proposals use adjacent archive points in 2D and nearest-neighbor edges in higher dimensions, rejecting edges above the configured gap threshold. Each run returns a size-N selection from its own population and archive. Maximin, ordered quantile, and hybrid arc/coordinate quantile readouts are identified separately in the table. For F7 only, each active proposal perturbs one randomly chosen decision coordinate by Gaussian noise with scale 0.21926258, clips it to the box, and accepts it when it decreases the current scalar value. Proposals occur every two iterations during the finite exploration prefix and are disabled afterwards. The F7 configuration uses an unguarded finite exploration prefix, with proposals disabled before the frozen phase.

Table 5: MOGF/D dynamics settings. LZ09-F2 shares the ZDT1 settings; decimal values are rounded for display.

Parameter	ZDT1	ZDT2	ZDT3	DTLZ7	LZ09-F7	LZ09-F9
Initial step	0.076858	0.096125	0.22306	0.15179	0.21001	0.075084
Minimum step ratio	0.019196	0.50042	0.017799	0.024197	0.021716	0.84057
τ	0.28544	0.097309	0.08538	0.40861	0.38466	0.084663
λ_{cov}	0.1532	1.8715	0.033942	2.2142	0.3025	1.0213
λ_{int}	2.0503	1.2776	0.92765	1.4705	0.11585	1.7708
Gaussian bandwidth	0.047416	0.035659	0.16071	0.044213	0.032882	0.096745
Interaction energy budget	0.023346	0.0049217	0.021296	0.001539	0.0045572	0.0051451
Interaction warm-up	536	443	150	225	700	200
Preference rate	0.015947	0.023602	0	0.0026859	0.035596	0
Preference anchor	4.6287	7.4037	0	1.947	4.9553	0
Preference stop	500	4851	0	901	3501	0
Normalization stop	3251	4901	4901	921	3501	3501
Exploration stop	0	4851	4851	901	3501	3501
Tangent cap ρ	1	1	1	1	0.5	4
Error amplitude	0	0.001	0.001	0	0	0.001

Table 6: Archive, preference, and readout settings. LZ09-F2 shares the ZDT1 settings; intervals and warm-ups are in iterations.

Parameter	ZDT1	ZDT2	ZDT3	DTLZ7	LZ09-F7	LZ09-F9
Archive capacity	200	200	200	200	200	300
Archive update interval	1	1	1	1	1	2
Gap proposal interval	10	10	10	1	10	1
Archive warm-up	36	114	64	220	5	100
Gap factor χ	7.9818	6.1868	3.8231	1.3292	9.5293	785.14
Stretch rate	2.0678	0.57756	0.57654	1.8528	2.0678	0.57756
Inverse-weight floor	0.001	0.001	0.001	0.001	0.075694	0.010674
Preference logit scale	1	1	1	1	2.0281	0.79163
Preference logit bias	0	0	0	0	0.077929	0.0075741
Arc mixing weight	0.75	0.75	0.75	0.75	0.8991	0.75
Readout	Hybrid	Hybrid	Maximin	Maximin	Hybrid	Ordered
Raw finite coverage	No	Yes	Yes	Yes	No	Yes
Finite guard	Yes	Yes	Yes	Yes	No	Yes

F.3 Additional Synthetic Fronts

Figures 6–12 show all eight methods on the seven synthetic problems, including the full comparisons corresponding to Figures 2 and 3. For each method and problem, we sort the six runs by increasing unrounded IGD and display rank 3 of 6, with ascending seed number breaking ties. This lower-median rule selects one complete run per method and problem. The plots show returned populations, including dominated points, against the analytic front. Each problem uses common axis limits across methods, and off-view counts identify points omitted by those limits. All metrics use the complete saved outputs. The three LZ09 problems share continuous objective-space fronts with ZDT1 or ZDT2, despite their different decision-space geometry. These figures visualize approximation in objective space.

Table 7 gives complementary HV results for the same terminal sets as Table 1. MOGF/D has the largest mean HV on the seven problems under the stated references. HV measures dominated volume under the fixed reference, while IGD averages reference-to-set distances after nondominated filtering. Zero F7 HV entries correspond to zero dominated volume under reference $(1.1, 1.1)$.

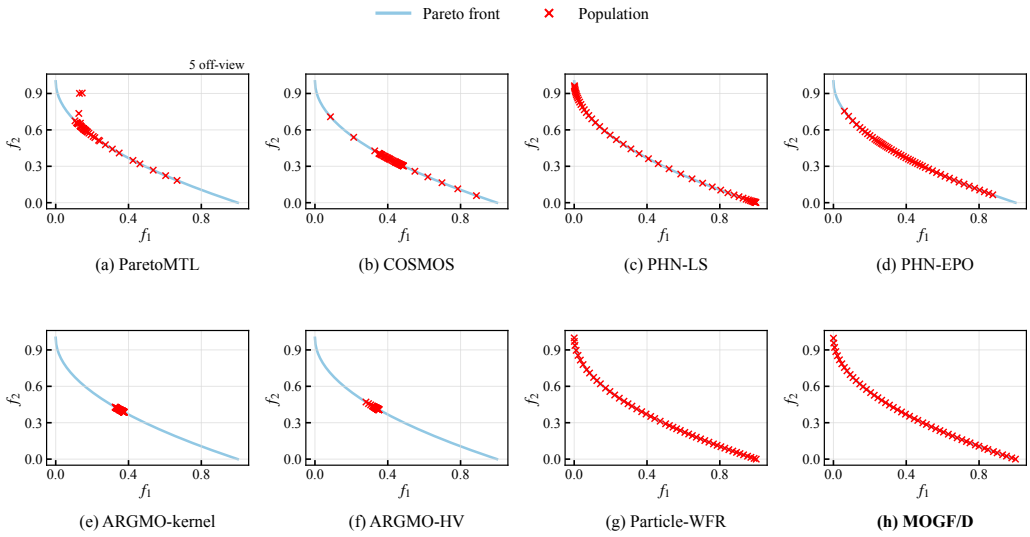

Figure 6: Results on ZDT1. Light blue shows the analytic Pareto front, and red crosses show the returned solutions. Each method displays its lower-median IGD run among six seeds, with shared axis limits.

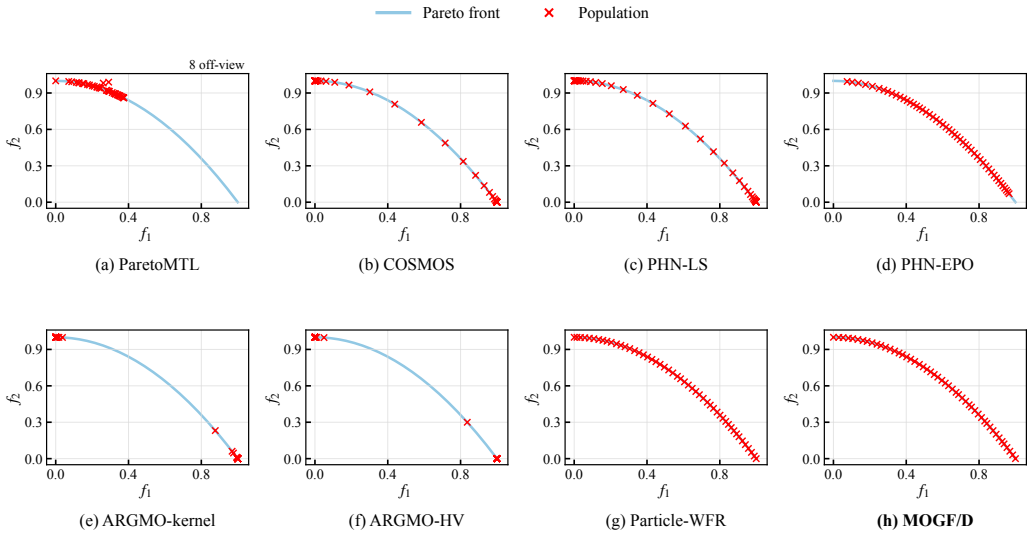

Figure 7: Results on ZDT2, whose Pareto front is continuous and concave. Colors and representative-run selection follow Figure 6.

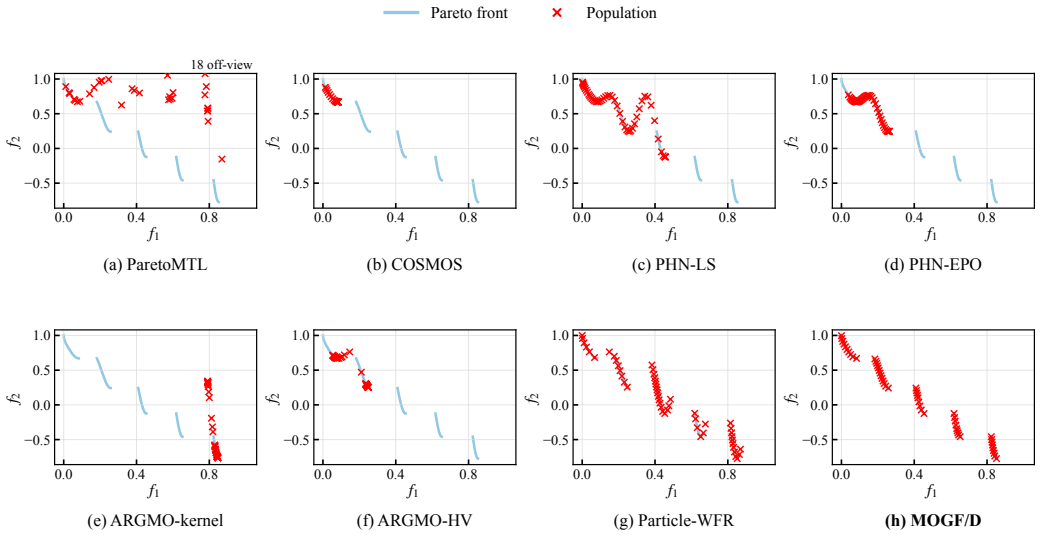

Figure 8: Results on ZDT3, showing its five disconnected Pareto-front components. Colors and representative-run selection follow Figure 6.

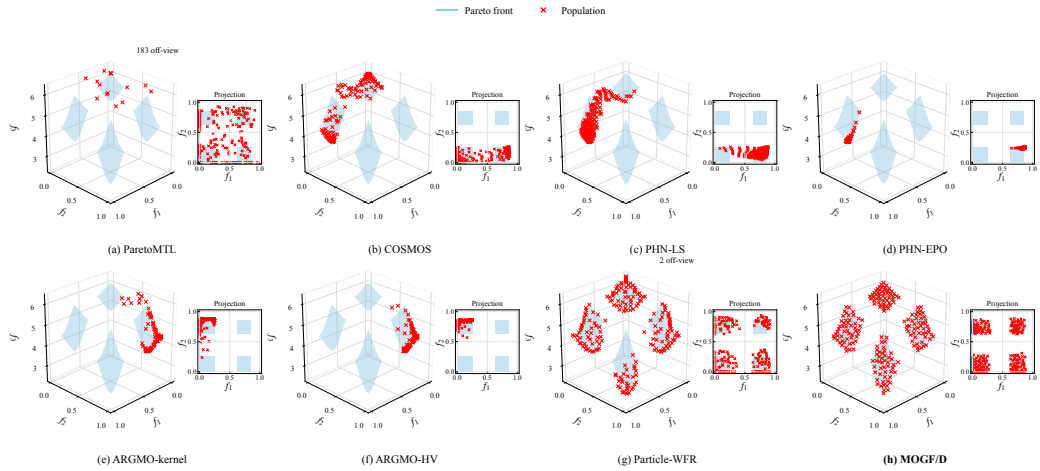

Figure 9: Full eight-method results on DTLZ7, with three-dimensional views and (f_1, f_2) projections. Colors and representative-run selection follow Figure 6.

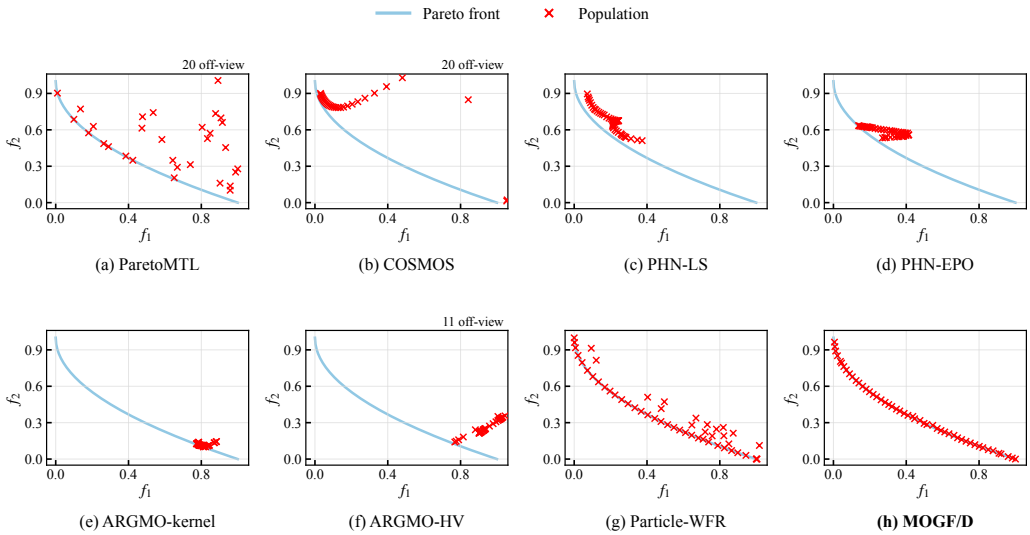

Figure 10: Results on LZ09-F2, with a continuous Pareto front and nonlinear decision-variable coupling. Colors and representative-run selection follow Figure 6.

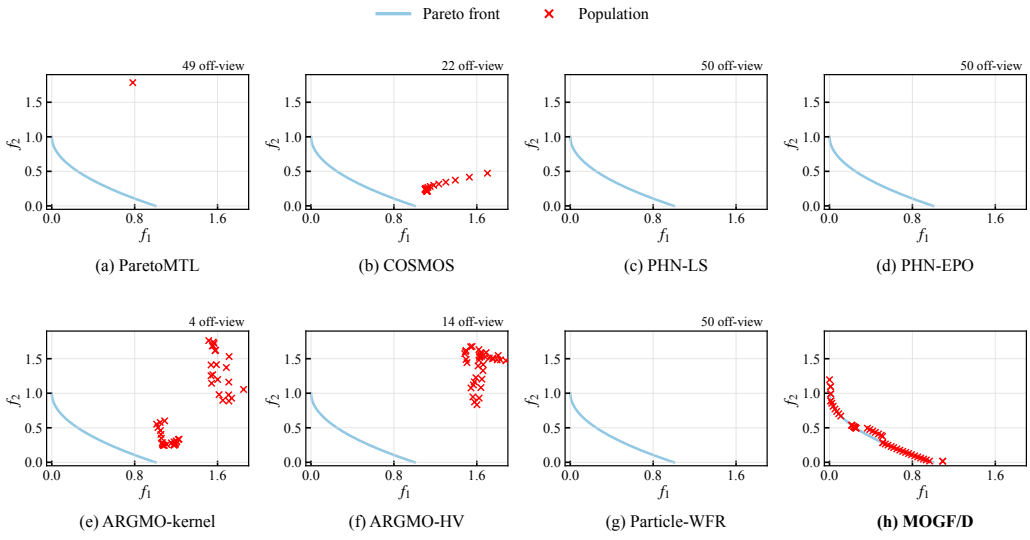

Figure 11: Full eight-method results on LZ09-F7. Colors and representative-run selection follow Figure 6; off-view counts use the shared limits of Figure 3.

Table 7: HV ↑ on synthetic benchmarks (mean ± sample SD, six seeds). Parentheses give zero-based ranks from unrounded means; bold and underlined values mark the best and second best. Avg. rank averages all seven problems.

Dataset	ParetoMTL	COSMOS	PHN-LS	PHN-EPO	ARGMO-kernel	ARGMO-HV	Particle-WFR	MOGF/D
Simple								
ZDT1	0.78168 (5) ±0.00304	0.80931 (4) ±0.01482	0.85909 (2) ±0.00080	0.84434 (3) ±0.00174	0.55286 (6) ±0.00534	0.55261 (7) ±0.01244	<u>0.86658</u> (1) ±0.00031	**0.86716** (0) ±0.00002
ZDT2	0.22824 (7) ±0.00077	0.30385 (4) ±0.21236	0.37880 (2) ±0.20871	0.31331 (3) ±0.22393	0.24107 (6) ±0.15278	0.25116 (5) ±0.16006	<u>0.53351</u> (1) ±0.00011	**0.53371** (0) ±0.00002
Disconnected PF								
ZDT3	0.84614 (4) ±0.08475	0.68064 (5) ±0.35183	0.90025 (2) ±0.37575	0.84757 (3) ±0.36304	0.59269 (7) ±0.17828	0.66583 (6) ±0.25099	<u>1.32018</u> (1) ±0.00244	**1.32627** (0) ±0.00011
DTLZ7	1.15055 (7) ±0.03528	1.92085 (2) ±0.23438	1.74146 (3) ±0.28138	1.27735 (6) ±0.38489	1.69172 (4) ±0.31791	1.60559 (5) ±0.35772	<u>2.57159</u> (1) ±0.01065	**2.63293** (0) ±0.00205
Complex PS								
LZ09-F2	0.77040 (2) ±0.02327	0.42140 (7) ±0.15845	0.45085 (3) ±0.15867	0.44347 (4) ±0.15727	0.43430 (5) ±0.10516	0.42633 (6) ±0.09030	<u>0.85982</u> (1) ±0.00083	**0.86216** (0) ±0.00369
LZ09-F7	0.01510 (3) ±0.02937	0.04110 (2) ±0.06445	0.00000 (5) ±0.00000	0.00000 (5) ±0.00000	<u>0.04722</u> (1) ±0.04000	0.00323 (4) ±0.00791	0.00000 (5) ±0.00000	**0.82680** (0) ±0.03895
LZ09-F9	0.41896 (2) ±0.01649	0.14142 (5) ±0.04732	0.15550 (4) ±0.04596	0.16099 (3) ±0.06179	0.13455 (6) ±0.03504	0.11535 (7) ±0.06139	<u>0.53176</u> (1) ±0.00116	**0.53183** (0) ±0.00073
Avg. rank	4.29	4.14	3.00	3.86	5.00	5.71	<u>1.57</u>	**0.00**

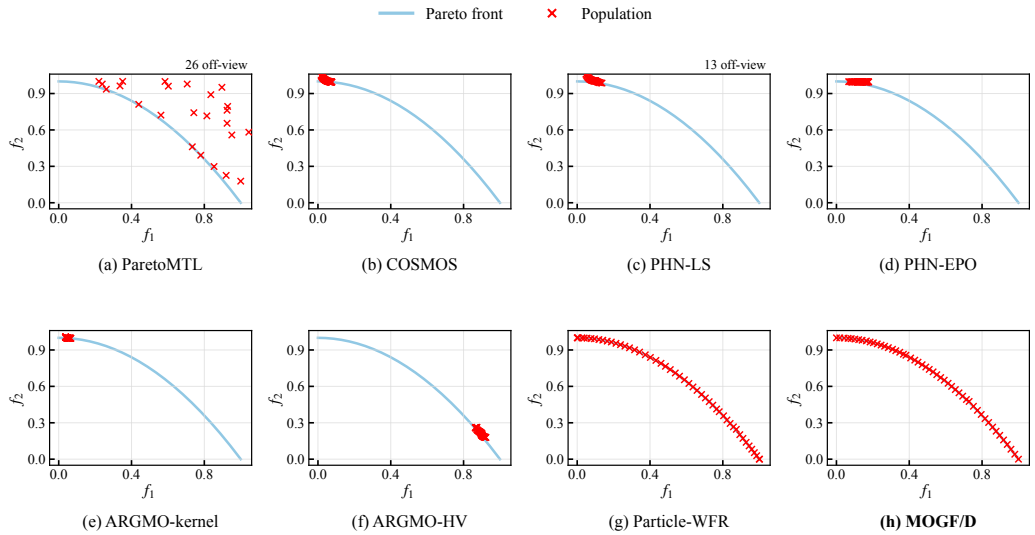

Figure 12: Results on LZ09-F9, with a continuous concave front and a nonlinear Pareto set. Colors and representative-run selection follow Figure 6.

F.4 Controlled Component Analysis

We examine the contributions of the implemented population mechanisms by removing one component at a time from the selected full configuration. The study contains 48 valid float32 runs: three variants on ZDT3 and five on F7, each with seeds 0–5, 50 returned points, and 5,001 updates. The full runs reproduce the MOGF/D outputs used in the main synthetic comparison. Each removal retains

the remaining parameters of the full configuration. Removing interaction sets its force strength to zero; removing gap guidance sets only the coverage-force strength to zero and retains the archive and final readout. Fixed preferences set the preference rate to zero, and removing coordinate proposals disables their interval and perturbation scale. For ZDT3, fixing preferences and removing coordinate proposals leave the full configuration unchanged; dashes identify these redundant controls. HV and fill distance are evaluated on the same saved outputs.

Table 8: Component removals on ZDT3 and LZ09-F7 (mean ± sample SD, six matched seeds). Each run returns 50 solutions after 5,001 updates with the remaining parameters fixed. Dashes mark redundant controls.

Variant	IGD ↓	HV ↑	Fill distance ↓
ZDT3			
Full MOGF/D	0.00968 ± 0.00007	1.32627 ± 0.00011	0.02941 ± 0.00033
Without interaction	0.00975 ± 0.00007	1.32628 ± 0.00013	0.03404 ± 0.00832
Without gap guidance	0.01017 ± 0.00066	1.32598 ± 0.00048	0.04818 ± 0.03082
Fixed preferences	—	—	—
Without coordinate proposals	—	—	—
LZ09-F7			
Full MOGF/D	0.03202 ± 0.02470	0.82680 ± 0.03895	0.10458 ± 0.01577
Without interaction	0.03828 ± 0.01024	0.82108 ± 0.01545	0.13850 ± 0.02834
Without gap guidance	0.03364 ± 0.01493	0.82319 ± 0.02451	0.14206 ± 0.09390
Fixed preferences	0.04591 ± 0.01859	0.80524 ± 0.03042	0.10353 ± 0.01703
Without coordinate proposals	0.17524 ± 0.00870	0.58721 ± 0.01343	0.25156 ± 0.01344

Table 9: Paired IGD comparisons over six seeds. W/T/L counts lower/equal/higher IGD for full MOGF/D. Exact one-sided Wilcoxon tests assess whether removal increases IGD; Holm adjustment covers all six comparisons.

Problem	Removal	W/T/L	p	p_{Holm}
ZDT3	No gap guidance	5/0/1	0.03125	0.15625
ZDT3	No interaction	4/0/2	0.10938	0.43750
LZ09-F7	Fixed preferences	4/0/2	0.21875	0.65625
LZ09-F7	No coordinate proposals	6/0/0	0.01562	0.09375
LZ09-F7	No gap guidance	4/0/2	0.28125	0.65625
LZ09-F7	No interaction	4/0/2	0.28125	0.65625

Table 8 shows that removing F7 coordinate proposals increases mean IGD from 0.03202 to 0.17524, a factor of 5.47. The full method has lower IGD in all six matched seeds for this comparison. Fixed preferences increase mean IGD to 0.04591, with lower IGD for the full method in four of six pairs. Interaction and gap guidance have smaller, seed-dependent effects; on ZDT3, removing interaction changes mean IGD by only about 0.69%. The metrics capture different effects: ZDT3 HV is almost unchanged without interaction, while fixed preferences on F7 yield slightly smaller mean fill distance and worse mean IGD. The clearest F7 improvement is associated with coordinate-based basin exploration.

Table 9 reports post-hoc paired comparisons from unrounded per-seed IGD. Holm-adjusted p values (Holm, 1979) range from 0.09375 to 0.65625; the six-of-six F7 coordinate-proposal comparison has $p_{\mathrm{Holm}} = 0.09375$. The study assesses the listed component removals at the selected full configuration using terminal outputs.

F.5 Fairness Classification

Data and models. Table 10 gives the fixed split sizes and prediction networks. All methods use the same split indices across seeds; the seed controls optimization randomness rather than a new data

partition. Adult preprocessing replaces missing categorical values with the source's fixed values and groups education and marital-status categories. Credit encodes the designated education, marital-status, and payment-history fields as categorical variables. COMPAS applies the source screening filters, converts jail duration to hours, and predicts two-year recidivism. Categorical features are one-hot encoded, and the target and sensitive sex field are removed from the predictor inputs. The retained loaders fit the feature standardizer on the entire preprocessed dataset before applying the fixed splits. All methods share this preprocessing convention.

Table 10: Fairness datasets and ReLU network architectures. Each task uses two objectives, sex as the sensitive attribute, and 12 trade-off solutions.

Dataset	Features	Train	Validation	Test	Architecture
Adult	88	34188	4885	9769	88–60–25–1
Credit	90	21000	3000	6000	90–60–25–1
COMPAS	20	4319	618	1235	20–60–25–1

Training objectives and evaluation. For a minibatch B of size b, let a_i be a model logit, $p_i = \sigma(a_i)$ its predicted probability, $y_i \in \{0, 1\}$ the label, and $s_i \in \{0, 1\}$ the sensitive-group code. Male and female are coded as 0 and 1, respectively. The two training objectives are

$$f_1(B) = -\frac{1}{b} \sum_{i \in B} \big[y_i \log p_i + (1 - y_i) \log(1 - p_i) \big], \tag{45}$$

$$f_2^{\text{soft}}(B) = \frac{1}{b} \left| \sum_{i \in B : y_i = 1, s_i = 0} \tanh(p_i) - \sum_{i \in B : y_i = 1, s_i = 1} \tanh(p_i) \right|. \tag{46}$$

This reproduces the source hyperbolic-tangent surrogate (Ren et al., 2024) with coefficient 1; the ReLU applied to probabilities in the code has no effect because $p_i \geq 0$. Validation and test replace $\tanh(p_i)$ by $\mathbf{1}[a_i > 0]$, yielding the stored hard-DEO score. Both versions use the total batch size as the denominator. The hard-DEO score is a batch-scaled difference in true-positive counts between groups. The evaluator averages the per-batch BCE and hard-DEO values with equal batch weights, including the final smaller batch. HV and the plotted fairness fronts use these evaluation scores; the soft surrogate is used for training.

Optimization and selection. Each method returns 12 models or preference-conditioned solutions after a budget of 50 epochs, with batch size 512 and validation evaluation every five epochs. The selected checkpoint maximizes whole-set validation HV under reference $(1, 1)$. Tables 11–13 give the saved optimization settings. The ParetoMTL fairness implementation adapts its original feasibility and constrained minimum-norm updates to the same loaders, networks, and two losses. It trains one model per preference with Adam at learning rate 10^{-3}; the first two epochs are reserved for feasibility and are included in the 50-epoch budget. The retained update-equivalence checks compare the adapted two-task directions with the original implementation. Its nine completed runs use the same three seeds and population-level validation selection as the comparison.

The target-weight temperature equals τ. MOGF/D uses direct minibatch Adam steps with tangent projection and relative norm caps. COSMOS uses Dirichlet concentration 0.5 on Adult/COMPAS, $(0.1, 0.5)$ on Credit, and 1 on MSLR; its cosine penalty is 0.01 for fairness and 1 for MSLR. ARGMO reference learning rates are 0.005 (kernel) and 0.0005 (HV) on fairness, and 0.0005 for both on MSLR. Both use a one-epoch warm-up and one reference adaptation, with constraint constants 0.01 for fairness and 10^{-4} for MSLR.

For Figure 4, each method displays the median-test-HV seed among three runs after validation-based checkpoint selection. Dominated points are removed separately within each displayed method, and adjacent nondominated points are connected only as a visual guide. The learning table uses all three seeds, regardless of this display choice. The selected MOGF/D fairness checkpoints occur during formation or at its boundary.

Table 11: MOGF/D parameters for the learning tasks. All population-size runs share the MSLR configuration.

Parameter	Adult	Credit	COMPAS	MSLR
Adam learning rate	0.0015075	0.0024061	0.0039788	0.0028207
τ	0.14748	0.067552	0.025168	0.11673
λ_{cov}	0.0060267	0.009282	0.0036481	0.031058
λ_{int}	0.023417	0.017878	0.06988	0.024805
Gaussian bandwidth	0.082369	0.24169	0.11596	0.67935
Preference rate	0.00032469	0.0020186	0.0042526	0.0056514
Preference anchor	0.1	0.1	0.1	0.1
Formation epochs	46	25	45	50
Archive warm-up	10	9	4	5
Preference floor	0.046808	0.077766	0.021817	0.024292
Preference power	1	1	0.71844	1
Tangent cap	1	1	1	1

Table 12: Adam learning rates for the learning baselines. PHN-LS and PHN-EPO use linear and EPO subsolvers, respectively; learning rates remain constant.

Method	Fairness	MSLR
ParetoMTL	10^{-3}	10^{-3}
COSMOS	10^{-3}	10^{-3}
PHN-LS / PHN-EPO	10^{-4}	10^{-4}
ARGMO-kernel / ARGMO-HV	10^{-3}	10^{-3}
Particle-WFR	10^{-3}	10^{-3}

Table 13: Particle-WFR learning parameters, using a Gaussian interaction kernel.

Dataset	α_2	β	γ	M	width
Adult	0.079644	0.00015949	6.3659e-06	4.7603	3.0426
Credit	0.055476	0.00031537	4.3653e-05	5.8213	1.7546
COMPAS	0.10906	0.00024267	2.4327e-05	1.2375	3.7991
MSLR	100	1e-05	0	30	0.0001

F.6 MSLR-WEB10K

Data and objectives. We use the official Fold 1 train/validation/test partition of MSLR-WEB10K (Qin & Liu, 2013) through the retained PyTorchLTR loader (Jagerman, 2020), with its normalization enabled. Following the Particle-WFR task construction (Ren et al., 2024), the first 131 of the 136 features form the input. The remaining five signals—QS, QS2, QUCC, UCC, and UDT—together with the relevance label define six ranking objectives. A prediction network has architecture 131–32–1 with ReLU in its hidden layer. For query q with documents D_q, let a_{qi} be a predicted score and $\ell_{qi}^{(j)}$ the label for objective j. Define $P_{qi} = \exp(a_{qi})/\sum_{k\in D_q}\exp(a_{qk})$ and $Q_{qi}^{(j)} = \exp(\ell_{qi}^{(j)})/\sum_{k\in D_q}\exp(\ell_{qk}^{(j)})$. The implemented ListNet loss (Cao et al., 2007) for objective j is

$$L_j = -\frac{1}{|B|}\sum_{q\in B}\sum_{i\in D_q} Q_{qi}^{(j)}\log(P_{qi}+10^{-10}), \tag{47}$$

where padded documents are masked out. Evaluation uses NDCG@10 (Järvelin & Kekäläinen, 2002) on each objective, averaged within each loader batch and then across batches with equal batch weights. The six-dimensional HV uses negative NDCG values with zero reference, as defined in Appendix F.1. Tables express this six-dimensional HV in units of 10^{-4}.

Main comparison. All eight methods use a budget of 500 epochs, batch size 512 queries, population size $N = 12$, and evaluation every five epochs. The checkpoint maximizing validation HV is used for the reported test HV. MOGF/D and ParetoMTL use explicit populations of ranking networks; the other baselines retain their source-specific population or preference-conditioned parameterization. The ParetoMTL adapter generalizes the original constrained direction to six objectives, using axis directions supplemented by deterministic maximin Dirichlet directions, with two feasibility epochs included in the budget. The baseline learning rates and MOGF/D parameters are given in Tables 11 and 12. ARGMO-HV seed 2 uses the complete rerun, selected at epoch 485 by validation HV. Figure 5(a) displays current-checkpoint validation HV, averaged over three seeds with sample-SD bands.

Population-size study. We additionally use the completed MOGF/D runs at $N \in \{6, 9, 12, 18\}$, with three seeds per size. Only population size changes in the training configuration: the optimizer, previously selected parameters, 500-epoch budget, evaluation interval, and formation horizon of 50 epochs are retained. Table 14 summarizes the final and best validation HV and the selected-checkpoint test HV. Figure 5(a) and the eight-method table use the original $N = 12$ cohort; panel (b) uses independent $N = 12$ reruns. For $N = 9$ seed 2, we retain the first complete, hash-verified run.

Table 14: MOGF/D population-size study on MSLR-WEB10K (mean $\pm$ sample SD, three seeds; HV in units of 10^{-4}). Final validation HV is measured at epoch 500; test HV uses the validation-best checkpoint.

N	Final validation HV	Best validation HV	Selected test HV
6	$6.225_{\pm 0.398}$	$6.336_{\pm 0.322}$	$5.732_{\pm 0.308}$
9	$8.853_{\pm 0.174}$	$8.927_{\pm 0.115}$	$8.149_{\pm 0.240}$
12	$9.393_{\pm 0.215}$	$9.691_{\pm 0.203}$	$8.809_{\pm 0.194}$
18	$10.160_{\pm 0.375}$	$10.480_{\pm 0.247}$	$9.601_{\pm 0.260}$

Across these tested sizes, mean final validation HV increases from 6.225×10^{-4} at $N = 6$ to 10.160×10^{-4} at $N = 18$. Mean selected-checkpoint test HV also increases, from 5.732×10^{-4} to 9.601×10^{-4}. Across the tested sizes, larger populations improve returned-set quality while increasing the number of models and computation per epoch. Checkpoint selection and the convergence figure use validation HV. Because the NDCG implementation randomly breaks score ties, independent test reevaluations of the same checkpoint can differ slightly; the table uses the original selected-test record, with checkpoint identity and fixed-random-state replay checked separately.

F.7 Reproducibility and Scope

The retained records identify each problem, method, seed, output set, configuration, and selected checkpoint, together with file hashes. Means, sample SDs, ranks, and the additional summaries are computed from per-seed outputs rather than from rounded table entries. Reported results use the completed cap-corrected DTLZ7 ParetoMTL cohort, completed fairness runs, and 48 float32 component-analysis runs.

Synthetic tuning initially used seeds 0–2, which are included in the six-run summaries; later tuning used separate seeds. ParetoMTL initializes each preference with the nominal seed plus preference index, producing overlapping random streams across nominal seeds. The reported SDs summarize the six saved runs. Method-specific budgets and the F7 dimensions are listed in Appendix F.2.

For learning tasks, reported checkpoints are selected by validation HV. MOGF/D's selected-test output is evaluated after restoring that checkpoint. Some baseline implementations also retain test histories, and the population-size extension evaluates saved checkpoints on test after training; model selection throughout follows validation HV. Appendix F.5 specifies the shared preprocessing and evaluation conventions.

The neural experiments use minibatch gradients and Adam with numerical tangent projection and relative norm caps. The deterministic convergence result in Section 3 applies to the projected Armijo recursion. Learning results measure empirical returned-set performance under the stated stochastic training protocol.