SURFACE: PREFERENCE DISTRIBUTION MATCHING ON PARETO FRONTS VIA SCALARIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Constructing a finite Pareto representation requires deciding where a limited number of solutions should be placed along a continuous Pareto front (PF). A desired preference distribution provides a direct way to specify this allocation by assigning more solutions to some PF regions and fewer to others. Scalarization generates Pareto solutions by varying objective weights and solving single-objective problems. However, a distribution of scalarization weights generally does not translate into the same distribution of Pareto solutions along the PF, because the weight-to-solution map is nonlinear and unknown in advance. We introduce SURFACE (***S**calarization **U**pdates for **R**edistributing **F**inite **A**pproximations with **C**hart **E**stimation*), which adaptively updates scalarization weights to match a desired PF distribution. SURFACE estimates how each region of the weight grid corresponds to an area of the PF from the current scalarized solutions, then redistributes the weights accordingly. The method supports parallel, warm-started, inexact optimization with standard single-objective solvers. For smooth positive preference distributions, we prove that the error decreases geometrically to an $\mathcal{O}(N^{-2/(M-1)})$ finite-sample level, with the same rate for the mismatch between the desired and achieved PF allocation, where N is the number of adaptive weights and M is the number of objectives. Experiments across diverse benchmarks and comparisons with nine baselines consistently demonstrate strong performance.

1 INTRODUCTION

Multi-objective optimization (MOO) is increasingly important in modern artificial intelligence (AI), including multimodal learning (Wei & Hu, 2024, Saha et al., 2025), large language model (LLM) alignment (Shi et al., 2024), and agentic AI systems (Zhang et al., 2026). Consider variable $\mathbf{x}$ in feasible domain $\mathcal{X} \subseteq \mathbb{R}^n$, MOO targets a vector of objectives $f_m : \mathcal{X} \to \mathbb{R}$:

$$\min_{\mathbf{x}\in\mathcal{X}} \mathbf{f}(\mathbf{x}) := [f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_M(\mathbf{x})]^\top . \tag{1}$$

Unlike single-objective optimization, generally no single solution is best in every objective. Therefore, (1) seeks some *Pareto-optimal* solution $\mathbf{x}^\star \in \mathcal{X}$ such that no $\mathbf{x} \in \mathcal{X}$ satisfies $f_m(\mathbf{x}) \le f_m(\mathbf{x}^\star)$ for all m, with strict inequality for some m. The collection of such solutions is called the *Pareto set* $\mathcal{X}_{\mathrm{PF}}$. Its image $\mathcal{F}_{\mathrm{PF}} := \mathbf{f}(\mathcal{X}_{\mathrm{PF}})$ forms the *Pareto front* (PF) (Miettinen et al., 1999, Ehrgott, 2005).

Preference information is essential in MOO because Pareto optimality alone does not determine which tradeoffs matter most to a decision maker (Pajasmaa et al., 2025). Under standard regularity conditions, the PF forms an $(M-1)$-dimensional manifold of tradeoff solutions (Hillermeier, 2001). This motivates two classical paradigms. *A priori* methods incorporate preferences before optimization to target preferred Pareto solutions (Chen et al., 2024, 2025), but require preferences to be specified in advance. *A posteriori* methods instead construct a representative PF approximation for later selection (Zhang et al., 2024, Boehmer & Wittmann, 2026), but must represent a continuous PF with finitely many solutions. Hence, the central challenge is to represent a continuous PF with finitely many Pareto solutions while preserving the desired preference, leading to our first question:

(Q1) How should finitely many Pareto solutions be distributed across a continuous PF to reflect a desired preference distribution?

Existing MOO methods provide limited control over how the resulting Pareto solutions are distributed along the PF. MGDA-type methods seek Pareto-stationary solutions (Désidéri, 2012, Sener

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.

& Koltun, 2018), preference-guided methods target specified tradeoffs or regions (Lin et al., 2019, Chen et al., 2024, 2025), and geometric, continuation, and finite-design methods promote well-spread Pareto representations (Das & Dennis, 1998, Zhang & Gao, 2006, Pereyra, 2009, Pereyra et al., 2013, Zhang et al., 2024, Boehmer & Wittmann, 2026). Scalarization offers a particularly natural route to distributional control by generating Pareto solutions through tunable weights (Ehrgott & Ruzika, 2008, Marler & Arora, 2010). Among them, Tchebycheff Scalarization (TS) (Steuer & Choo, 1983) is especially useful for PF construction because it can recover Pareto solutions on nonconvex fronts that some other scalarizations may miss (Ehrgott, 2005, Lin et al., 2024):

$$\mathbf{x}^{\star}_{\mathrm{TS}}(\mathbf{w}) \in \operatorname*{argmin}_{\mathbf{x}\in\mathcal{X}} \max_{m=1,\dots,M} \frac{f_m(\mathbf{x}) - z_m}{w_m}, \quad z_m < \min_{\mathbf{x}\in\mathcal{X}} f_m(\mathbf{x}), \quad \mathbf{w} \in \mathrm{relint}(\Delta_M) \tag{2}$$

Here $\mathbf{z}$ is a componentwise lower reference point, which may be chosen as the ideal objective vector. Yet geometric distortion in the weight-to-PF map generally makes the induced PF distribution differ from the prescribed weight distribution (Jiang et al., 2026), motivating our second question:

(Q2) How can scalarization weights be adapted to compensate for weight-to-PF distortion and match a desired PF preference distribution?

We answer these questions with SURFACE (***S**calarization **U**pdates for **R**edistributing **F**inite **A**pproximations with **C**hart **E**stimation*), a geometry-aware framework for matching finite Pareto representations to a desired preference distribution. As illustrated in Figure 1, SURFACE estimates the PF geometry from a finite grid of weights and updates the weights to match the desired preference distribution over the PF. Specifically, it estimates the PF distribution induced by the weights by measuring the area of the PF patch corresponding to each cell of the weight grid. It then updates the weights so that the resulting Pareto solutions move toward the desired preference distribution. Our contributions are:

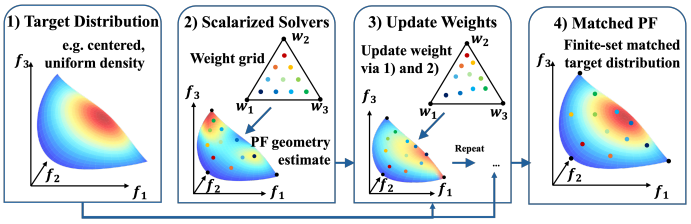

Figure 1: An overview of SURFACE. Cooler regions indicate lower target density, while warmer regions indicate higher density and a denser allocation of Pareto solutions. SURFACE estimates the weight-to-PF geometry and updates the scalarization weights to match this target distribution.

C1) Section 2 shows how a desired preference distribution over the PF maps to a distribution of scalarization weights via the weight-to-PF geometry, and gives sufficient conditions for Tchebycheff scalarization to yield a smooth, nondegenerate PF parameterization.

C2) In Section 3, we introduce SURFACE, which adaptively updates scalarization weights towards a desired preference coverage under finite budget. It is compatible with off-the-shelf optimizers, supports warm-started inexact solves, and solves all scalarized subproblems in parallel.

C3) We provide a rigorous convergence analysis in Section 4, showing that the weight updates contract the preference mismatch linearly down to a tight finite-resolution floor of $\mathcal{O}(N^{-2/(M-1)})$.

C4) Experiments in Section 5 empirically corroborate the effectiveness of SURFACE and verify our theoretical claims: across 9 baselines and benchmarks including ZDT/DTLZ and MO-Gym, SURFACE consistently achieves strong, steady performance.

Related work. Existing approaches to finite PF representation mainly target well-spread solutions through different design criteria. Adaptive scalarization adjusts parameters toward more even objective-space spacing (Eichfelder, 2009, Zhang & Gao, 2006, Saini et al., 2025), but may require local sensitivity information (Eichfelder, 2009) or sequential parameter updates (Zhang & Gao, 2006, Saini et al., 2025). Geometric methods construct additional objective-space directions or constraints (Das & Dennis, 1998, Messac et al., 2003, Erfani & Utyuzhnikov, 2011), while continuation methods trace the PF sequentially using local derivative or sensitivity information (Hillermeier, 2001, Pereyra et al., 2013, Bolten et al., 2021). Evolutionary methods adapt weights or reference points according to sparsity and diversity (Zhang & Li, 2007, Qi et al., 2014, Jain & Deb, 2013, Cheng et al., 2016), and finite-design methods use covering, packing, subset-selection, or repulsion objectives (Blank et al., 2020, Borghi, 2023, Zhang et al., 2024, Shang et al., 2023). Preference-conditioned models instead learn mappings from preferences to Pareto solutions (Navon et al., 2021,

Lin et al., 2022, Ye et al., 2024). SURFACE differs by taking the desired PF distribution as an explicit input, reconstructing the local PF geometry from parallel scalarized solves, and adapting the weights to match that distribution without differentiating through the optimizer. SURF (Jiang et al., 2026) considers the special case of uniform arc-length allocation for two objectives and uses a cumulative density function (CDF) update. SURFACE addresses general prescribed distributions and higher-dimensional PFs through reconstructed PF geometry and multidimensional mesh redistribution. In addition, the theory of SURFACE covers a broader class of scalarization methods beyond linear scalarization. Appendix A.1 provides detailed comparisons.

Notation. Let $\Delta_M := \{\mathbf{w} \in \mathbb{R}^M : w_m \geq 0, \sum_{m=1}^M w_m = 1\}$ denote the weight simplex. We write D for the Jacobian with respect to the displayed parameter coordinates, $\|\cdot\|$ for the Euclidean norm, and $\|\cdot\|_\infty$ for the uniform norm. For $k \geq 0$, $C^{k,1}$ denotes functions with k continuous derivatives and Lipschitz kth derivative, with $\|\cdot\|_{C^k}$ the supremum norm through order k. We use Vol_{M-1} for intrinsic weight-space volume. Recurring symbols are collected in Appendix A.2.

2 Pareto-Front Geometry from Scalarization

2.1 Scalarization-Induced PF Geometry and Preference Matching

Varying scalarization weights generates different Pareto outcomes and, under suitable regularity, can parameterize a region of the PF (Miettinen et al., 1999, Ehrgott, 2005, Hillermeier, 2001):

$$\mathbf{w} \mapsto \mathbf{x}^\star(\mathbf{w}) \mapsto \mathbf{f}_{\mathrm{PF}}(\mathbf{w}) := \mathbf{f}(\mathbf{x}^\star(\mathbf{w})) \in \mathcal{F}_{\mathrm{PF}}, \quad \mathbf{w} \in \Delta_M. \tag{3}$$

Section 2.2 formalizes this for $\mathbf{x}^\star(\mathbf{w}) = \mathbf{x}^\star_{\mathrm{TS}}(\mathbf{w})$ in (2) and Appendix B.4 provides an alternative analysis for weighted-sum scalarization. However, a distribution over $\mathbf{w} \in \Delta_M$ is generally distorted under the weight-to-PF map $\mathbf{f}_{\mathrm{PF}}$ (Jiang et al., 2026).As illustrated in Figure 2, a simplex cell in weight space maps to a PF patch whose surface area depends on the local geometry of $\mathbf{f}_{\mathrm{PF}}$. This distortion is characterized by the local *PF-area density* and the induced area of a simplex cell $\tau \subseteq \Delta_M$:

$$J(\mathbf{w}) := \frac{1}{\sqrt{M}} \sqrt{\det\big(D\mathbf{f}_{\mathrm{PF}}(\mathbf{w})^\top D\mathbf{f}_{\mathrm{PF}}(\mathbf{w})\big)}. \tag{4}$$

This geometry provides a direct way to encode preferences over the PF. Let $\vartheta : \mathcal{F}_{\mathrm{PF}} \to \mathbb{R}_+$ be an (unnormalized) *preference density* with respect to PF surface area (e.g., $\vartheta = 1$ corresponds to uniform surface-area preference). If J were available, a PF distribution with density proportional to ϑ could be realized by sampling weights from

$$\rho_\vartheta(\mathbf{w}) := \frac{\vartheta(\mathbf{f}_{\mathrm{PF}}(\mathbf{w}))J(\mathbf{w})}{\int_{\Delta_M} \vartheta(\mathbf{f}_{\mathrm{PF}}(\mathbf{v}))J(\mathbf{v})\mathrm{d}\,\mathrm{Vol}_{M-1}(\mathbf{v})}. \tag{5}$$

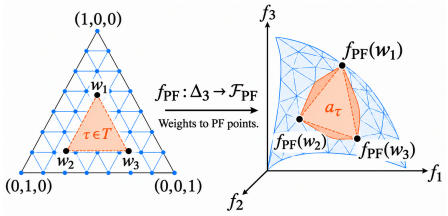

Figure 2: Illustration of weight-to-PF geometry. A simplex cell τ in weight space maps to a PF patch $\mathbf{f}_{\mathrm{PF}}(\tau)$. Its surface area depends on $J(\mathbf{w})$, so equal-sized weight cells can represent different amounts of PF area.

However, evaluating J requires the sensitivity $D\mathbf{f}_{\mathrm{PF}}(\mathbf{w})$, which may be unavailable or costly to obtain by differentiating through the scalarized optimization. Under a finite budget, we instead use only sampled pairs $(\mathbf{w}, \mathbf{f}_{\mathrm{PF}}(\mathbf{w}))$ to estimate the represented PF areas and match the preference.

Lemma 1 (Preference Matching)**.** *Let $\mathbf{f}_{\mathrm{PF}}$ be injective and full-rank, and let $W = \{\mathbf{w}_n\}_{n=1}^N \subseteq \Delta_M$ define a mesh $\mathcal{T}(W)$ of the weight simplex. For each $\tau \in \mathcal{T}(W)$, let $a_\tau := \mathbf{f}_{\mathrm{PF}}(\tau)$ denote the corresponding PF patch. Define $\mathcal{R}_\vartheta(W) := \max_{\tau\in\mathcal{T}(W)} \big||\mathcal{T}(W)| \int_\tau \rho_\vartheta(\mathbf{w})\mathrm{d}\mathbf{w} - 1\big|$. Then $\mathcal{R}_\vartheta(W) = 0$ if and only if all PF patches $\{a_\tau\}_{\tau\in\mathcal{T}(W)}$ carry equal preference-weighted area.*

The proof of Lemma 1 is provided in Appendix B.1. For $\vartheta = 1$, $\mathcal{R}_\vartheta(W)$ measures equal PF surface-area allocation. When $M = 2$, this reduces to equal arc-length allocation Jiang et al. (2026).

2.2 Regular PF Charts from TS

In this section, we specialize the weight-to-PF map to TS and establish its PF coverage and regular chart properties. For $\mathbf{w} \in \mathrm{relint}(\Delta_M)$, TS in (2) admits the equivalent epigraph formulation (6), under which we impose standard and checkable regularity conditions.

Assumption 1 (Regular full-ray TS regime)**.** *Fix $0 < \delta < 1/M$. Suppose $\mathcal{X}$ is compact and $\mathbf{z}$ is utopian, i.e., $z_m < \min_{\mathbf{x}\in\mathcal{X}} f_m(\mathbf{x})$ for every m. For every $\mathbf{w} \in \Delta_M^\delta$, the problem*

$$\min_{\mathbf{x}\in\mathcal{X}, s\in\mathbb{R}} \quad s \quad \textit{s.t.} \quad f_m(\mathbf{x}) - z_m \le s w_m, \quad m = 1, \dots, M, \tag{6}$$

has a unique solution $(\mathbf{x}^\star_{\mathrm{TS}}(\mathbf{w}), s^\star(\mathbf{w}))$. At this solution, all M ray inequalities are active, the linear independence constraint qualification (LICQ) holds, every active inequality has a positive multiplier, and the second-order sufficient condition (SOSC) holds on the critical cone.

Compactness and a utopian $\mathbf{z}$ are standard in TS (Ehrgott, 2005). Full activity, LICQ, positive multipliers, and SOSC are directly checkable from the ray residuals, active-constraint Jacobian, multiplier signs, and reduced Hessian (Fiacco, 1976, Bonnans & Shapiro, 2013). These conditions ensure a smooth, nonfolding TS parameterization over Δ_M^δ. When the standard Tchebycheff formulation is nonsmooth, smooth Tchebycheff scalarizations provide a related alternative (Lin et al., 2024). Appendix B.3 gives concrete sufficient regimes for uniqueness and these regularity conditions. Classical TS completeness separately gives PF coverage of the represented positive-gap rays without requiring PF convexity (Miettinen et al., 1999, Ehrgott, 2005).

Proposition 1 (PF coverage and injectivity of TS)**.** *Under Assumption 1, TS maps Δ_M^δ injectively onto the corresponding PF patch $\mathbf{f}_{\mathrm{PF}}(\Delta_M^\delta)$. Moreover, every Pareto point whose normalized positive-gap ray lies in Δ_M^δ is represented.*

The injectivity statement follows from full ray activity, which associates each represented PF point with a unique positive ray from $\mathbf{z}$. Additionally, we impose the following standard smooth-data conditions for parametric sensitivity (Fiacco, 1976, Bonnans & Shapiro, 2013).

Assumption 2 (Shared feasible-set and smooth-data conditions)**.** *The feasible set $\mathcal{X} = \{\mathbf{x} : \mathbf{c}_1(\mathbf{x}) \le \mathbf{0},\ \mathbf{c}_2(\mathbf{x}) = \mathbf{0}\}$ is nonempty and closed. Every f_m and every component of $\mathbf{c}_1$ is $C^{3,1}$ on a common open convex neighborhood of $\mathcal{X}$, each component of $\mathbf{c}_1$ is convex there, and $\mathbf{c}_2$ is affine.*

These conditions are directly checkable in many applications, including DTLZ2 and a strongly convex regime considered in Appendix B.3. They yield the following regularity result.

Proposition 2 (Regular PF chart from TS)**.** *Suppose Assumptions 1 and 2 hold, with $\Delta_M^\delta := \{\mathbf{w} \in \Delta_M : w_m \ge \delta,\ m = 1, \dots, M\}$. Then the TS-induced PF chart satisfies (i) **smoothness**: $\mathbf{x}^\star_{\mathrm{TS}}$ is uniformly Lipschitz and face-wise $C^{2,1}$ on Δ_M^δ, with compatible one-sided extensions, and $\mathbf{f}_{\mathrm{PF}} = \mathbf{f} \circ \mathbf{x}^\star_{\mathrm{TS}}$ is face-wise $C^{2,1}$; (ii) **nonfolding**: $\mathbf{f}_{\mathrm{PF}}$ is injective, has full tangent rank, and there exist $0 < c_P \le C_P < \infty$ such that*

$$c_P\|\mathbf{w} - \mathbf{v}\| \le \|\mathbf{f}_{\mathrm{PF}}(\mathbf{w}) - \mathbf{f}_{\mathrm{PF}}(\mathbf{v})\| \le C_P\|\mathbf{w} - \mathbf{v}\|, \quad \mathbf{w}, \mathbf{v} \in \Delta_M^\delta; \tag{7}$$

*and (iii) **nondegenerate area**: J is face-wise $C^{1,1}$ with compatible one-sided traces and satisfies $\underline{J} \le J(\mathbf{w}) \le \overline{J}$ for some $0 < \underline{J} \le \overline{J} < \infty$ and all $\mathbf{w} \in \Delta_M^\delta$.*

The proof of Proposition 2 is given in Appendix B.2. Thus TS induces a smooth, nonfolding PF chart with uniformly nondegenerate area distortion. Combined with Lemma 1, these properties yield $\rho_\vartheta(\mathbf{w}) \propto \vartheta(\mathbf{f}_{\mathrm{PF}}(\mathbf{w}))J(\mathbf{w})$, with $\rho_1(\mathbf{w}) \propto J(\mathbf{w})$ recovering uniform PF surface-area coverage.

3 The SURFACE Algorithm

Building on Section 2, which identifies the target weight density $\rho_\vartheta \propto (\vartheta \circ \mathbf{f}_{\mathrm{PF}})J$ for the prescribed PF preference, we propose SURFACE (***S**calarization **U**pdates to **R**edistribute **F**inite **A**pproximations through **C**hart **E**stimation*), an general framework operates on a chosen scalarization with adaptive weights: it estimates the target density from finitely many scalarized solutions and redistributes the weights through mass equidistribution (MEQ). We use TS as the running specialization below, while Appendix B.4 gives the analogous construction for weighted-sum scalarization. For the TS convergence analysis, the adaptive weights and MEQ are restricted to any fixed Δ_M^δ from Proposition 2; the extremal references remain auxiliary boundary anchors.

Following classical PF-generation methods that use individual-objective optima to characterize objective-space geometry (Das & Dennis, 1998, Messac et al., 2003), SURFACE uses M single-objective extremal references $\tilde{\mathbf{f}}_m \approx \mathbf{f}(\mathbf{x}^\star_m)$, where $\mathbf{x}^\star_m \in \arg\min_{\mathbf{x}\in\mathcal{X}} f_m(\mathbf{x})$ for $m \in [M]$. These

Algorithm 1 SURFACE

Require: Number of adaptive weights N, single-objective extremal references $\{\tilde{\mathbf{f}}_m\}_{m=1}^M$, preference profile ϑ, initial solver states $\{\tilde{\mathbf{x}}_n^{(-1)}\}_{n=1}^N$, damping $\alpha \in (0,1]$, and outer rounds T

1: **for** $t = 0, \dots, T-1$ **do**
2: $\quad \{\mathbf{w}_n^{(t)}\}_{n=1}^N \leftarrow \mathrm{MEQ}_N(\rho_t)$ ▷ Initialize ρ_0 as uniform on the chosen weight domain
3: $\quad \tilde{\mathbf{x}}_n^{(t)} \leftarrow \textsc{InnerSolver}\left(\mathbf{w}_n^{(t)}, \mathbf{x}_{\mathrm{init}} = \tilde{\mathbf{x}}_n^{(t-1)}\right)$ in parallel for $n \in [N]$
4: $\quad$ **Optional:** refine $\tilde{\mathbf{f}}_m \leftarrow \mathbf{f}(\mathrm{argmin}_{\mathbf{x}\in\mathcal{X}} f_m(\mathbf{x}))$ in parallel for $m \in [M]$
5: $\quad$ Interpolate P_t from $\{(\mathbf{w}_n^{(t)}, \mathbf{f}(\tilde{\mathbf{x}}_n^{(t)}))\}_{n=1}^N \cup \{(\mathbf{e}_m, \tilde{\mathbf{f}}_m)\}_{m=1}^M$
6: $\quad$ Update ρ_{t+1} by (9)
7: **end for**

Ensure: $\{(\mathbf{w}_n^{(T-1)}, \mathbf{f}(\tilde{\mathbf{x}}_n^{(T-1)}))\}_{n=1}^N$ together with the final extremal references

references $\{\tilde{\mathbf{f}}_m\}_{m=1}^M$ characterize the PF boundary, while N adaptive weights $\{\mathbf{w}_n^{(t)}\}_{n=1}^N$ sample its interior geometry. The smooth and nonfolding chart characterized in Proposition 2 allows these samples to track the PF without ambiguity.

At outer round t, SURFACE applies InnerSolver in parallel on the scalarized objectives to obtain

$$\tilde{\mathbf{x}}_n^{(t)} = \textsc{InnerSolver}\left(\mathbf{w}_n^{(t)}, \mathbf{x}_{\mathrm{init}} = \tilde{\mathbf{x}}_n^{(t-1)}\right) \approx \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t)}), \quad n = 1, \dots, N. \tag{8}$$

Here, one may use gradient-based optimization such as Adam for differentiable scalarized objectives (Kingma & Ba, 2015), or policy optimization such as PPO for reinforcement learning (Schulman et al., 2017). Here, warm-starting from the preceding round allows the scalarized problems to be solved inexactly while keeping refining, as detailed in the next section. The references $\{\tilde{\mathbf{f}}_m\}_{m=1}^M$ may be held fixed when accurate estimates are available or refined similarly in parallel.

Using the N returned weight-objective pairs $\{(\mathbf{w}_n^{(t)}, \mathbf{f}(\tilde{\mathbf{x}}_n^{(t)}))\}_{n=1}^N$ together with the M single-objective extremal pairs $\{(\mathbf{e}_m, \tilde{\mathbf{f}}_m)\}_{m=1}^M$, SURFACE interpolates the sampled PF geometry to obtain a smooth chart approximation P_t. We use local quadratically augmented radial basis function (RBF) interpolation combined through a partition-of-unity construction (Wendland, 2005). The extremal references are auxiliary boundary anchors and are not part of the regular TS chart or transported mesh analyzed on Δ_M^δ; because the reconstruction is local with compact support, they do not affect P_t on Δ_M^δ for sufficiently fine meshes. The resulting P_t can be differentiated throughout the weight domain, so replacing $\mathbf{f}_{\mathrm{PF}}$ by P_t in (4) gives the estimated area factor J_{P_t} and hence the local weight-to-PF distortion without differentiating through the scalarized optimizer. SURFACE then combines this reconstructed geometry with the prescribed PF preference and updates the target weight density with damping factor $\alpha \in (0,1]$:

$$\rho_{t+1}(\mathbf{w}) := (1-\alpha)\rho_t(\mathbf{w}) + \alpha \frac{\vartheta(P_t(\mathbf{w}))J_{P_t}(\mathbf{w})}{\int \vartheta(P_t(\mathbf{v}))J_{P_t}(\mathbf{v})\mathrm{d}\mathbf{v}}. \tag{9}$$

For $\vartheta = 1$, this reduces to the reconstructed PF-area density and targets uniform PF surface-area coverage. Appendix C.1 gives the interpolation details.

Finally, SURFACE converts ρ_{t+1} into the N adaptive weights for the next round through mass equidistribution (MEQ) (Budd & Williams, 2009). Ideally, the free weights are placed so that every induced weight-space cell carries the same ρ_{t+1}-mass. If $\{\tau_k(\mathbf{w}_{1:N})\}_{k=1}^K$ denotes the simplicial cells induced by the N free weights on the chosen weight domain, this target can be written as

$$\{\mathbf{w}_n^{(t+1)}\}_{n=1}^N \in \mathrm{MEQ}(\rho_{t+1}) := \underset{\{\mathbf{w}_n\}_{n=1}^N}{\mathrm{argmin}} \max_{k=1,\dots,K} \left| K \int_{\tau_k(\mathbf{w}_{1:N})} \rho_{t+1}(\mathbf{w})\mathrm{d}\mathbf{w} - 1 \right|. \tag{10}$$

Thus regions with larger ρ_{t+1} receive smaller cells and consequently denser scalarization weights. Rather than solving (10) directly, SURFACE realizes this equal-mass target by transporting an initially equal-volume tessellation. For $M = 2$, the transport is given exactly by inverse cumulative-distribution transport Jiang et al. (2026). For $M \geq 3$, we use Moser transport (Bruveris et al., 2018), which preserves the simplex boundary while redistributing its cells toward equal ρ_{t+1}-mass. Appendix C.2.1 gives the construction, with numerical implementation details in Appendix C.2.2.

Algorithm 1 summarizes the iteration. This algorithm requires $(T)N$ scalarized solves, with extremal references either supplied or updated in parallel. The dominant cost is typically INNERSOLVER, since reconstruction and MEQ operate only in the $(M-1)$-dimensional weight space and require no additional scalarized solves.

4 CONVERGENCE ANALYSIS

This section establishes the convergence of SURFACE by tracking the three approximations in Section 3: finite PF reconstruction, numerical MEQ, and inexact scalarized solves. We use TS as the scalarization and work on any fixed Δ_M^δ from Proposition 2; throughout this section, the density in (5) is normalized over Δ_M^δ. We impose the following condition on the preference density.

Assumption 3 (Regular preference)**.** *The preference density $\vartheta : \mathcal{F}_{\mathrm{PF}} \to \mathbb{R}_+$ extends to a $C^{2,1}$ function on an open neighborhood containing $\mathcal{F}_{\mathrm{PF}}$ and $P_t(\Delta_M^\delta)$ for all t, with $0 < \underline{\vartheta} \le \vartheta(\mathbf{y}) \le \overline{\vartheta} < \infty$ and $\|\vartheta\|_{C^{2,1}} \le B_\vartheta$ for some constants $\underline{\vartheta}, \overline{\vartheta}, B_\vartheta$.*

For finite PF reconstruction, Quadratic RBF interpolation (Wendland, 2005) gives second-order C^1 accuracy, while inexact PF samples incur an additional error amplified by one inverse mesh scale. To characterize the target distortion, define the preference-weighted condition number $\kappa_\vartheta := \overline{\vartheta}\,\overline{J}/(\underline{\vartheta}\,\underline{J})$.

Lemma 2 (Reconstruction and density accuracy)**.** *Suppose Assumptions 1–3 hold. Then*

$$\left\| \frac{\vartheta(P_t)J_{P_t}}{\int \vartheta(P_t)J_{P_t}} - \rho_\vartheta \right\|_\infty \le C_t N^{-\frac{2}{M-1}} + C N^{\frac{1}{M-1}} \max_{n\in[N]} \left\| \tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}^\star_{\mathrm{TS}}(\mathbf{w}_n^{(t)}) \right\|. \tag{11}$$

Here C and the round-dependent coefficient C_t are both bounded.

With exact scalarized solves, Lemma 2 gives density error $C_t N^{-\frac{2}{M-1}}$. Under the initial uniform weighting, the bound on C_t can be loose when κ_ϑ is large. As MEQ adapts the weights toward the target distribution, this geometric imbalance is reduced and C_t becomes $\mathcal{O}(1)$. Thus κ_ϑ affects the transient reconstruction, while the converged finite-resolution floor is $\mathcal{O}(N^{-\frac{2}{M-1}})$. For inexact solves, the second term captures the inverse-mesh amplification of inner optimization error. Appendix C.1 gives the reconstruction analysis. We next control the numerical MEQ error. Given ρ_{t+1}, the Moser transport redistributes the weights so that the represented cells have approximately equal ρ_{t+1}-mass.

Proposition 3 (Finite MEQ stability)**.** *Suppose Assumptions 1–3 hold. The finite Moser transport used by SURFACE achieves relative cell-mass error $\mathcal{O}(N^{-2/(M-1)})$. Moreover, the redistributed weights remain sufficiently well separated for stable reconstruction at the next round.*

Thus numerical MEQ has the same $\mathcal{O}(N^{-2/(M-1)})$ resolution as PF reconstruction while preserving the mesh regularity needed by Lemma 2; Appendix C.2 gives the construction and proof. We finally quantify the effect of solving each scalarized problem for only K inner steps. Following standard local convergence and sensitivity theory (Bonnans & Shapiro, 2013) and the finite-step analysis of SURF (Jiang et al., 2026), let η_K denote a K-step contraction factor such that $\|\mathbf{x}_K - \mathbf{x}^\star_{\mathrm{TS}}(\mathbf{w})\| \le \eta_K \|\mathbf{x}_0 - \mathbf{x}^\star_{\mathrm{TS}}(\mathbf{w})\|$. For linearly convergent methods, η_K decreases geometrically in K, while under the corresponding local error-bound or quadratic-growth conditions, sublinear guarantees give $\eta_K = \mathcal{O}(K^{-p})$. Since damping limits the motion of successive weights, α trades outer progress against the required inner accuracy. We therefore define the coupled density–solver error

$$\mathcal{E}_t := \|\rho_t - \rho_\vartheta\|_\infty + C_0 \left(1 + \alpha N^{1/(M-1)}\right) \max_{n\in[N]} \left\| \tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}^\star_{\mathrm{TS}}(\mathbf{w}_n^{(t)}) \right\|. \tag{12}$$

Theorem 1 (SURFACE convergence with inexact solves)**.** *Suppose Assumptions 1–3 hold, let $N^{1/(M-1)} = \Omega(\kappa_\vartheta^{3/2})$, and let $\bar{\alpha} = \Omega(N^{-\frac{3}{2(M-1)}})$ with $0 < \alpha \le \bar{\alpha} < 1$. Suppose the initial solves satisfy* (48) *and the warm-started K-step* INNERSOLVER *has contraction factor η_K satisfying $\bar{\alpha}^2 \eta_K = \mathcal{O}(N^{-\frac{3}{M-1}})$ with a sufficiently small hidden constant. Then, for $t = 0, \dots, T-1$,*

$$\mathcal{E}_t \le (1 - \alpha/4)^t \mathcal{E}_0 + \mathcal{O}\left(N^{-\frac{2}{M-1}}\right), \quad \widehat{\mathcal{R}}_\vartheta(W^{(t)}) \le C\mathcal{E}_t + \mathcal{O}\left(N^{-\frac{2}{M-1}}\right). \tag{13}$$

Here $W^{(t)} := \{\mathbf{w}_n^{(t)}\}_{n=1}^N$ is the transported weight mesh on Δ_M^δ at round t, and $\widehat{\mathcal{R}}_\vartheta(W^{(t)})$ denotes the target-mass residual on the complete cells transported by MEQ, defined in (24)*.*

Damping $\bar{\alpha}$	Contraction η_K	Inner K	Outer T	Use When
$\Theta(1)$	$\mathcal{O}(N^{-\frac{3}{M-1}})$	$\mathcal{O}(\log N)$	$\mathcal{O}(\log N)$	Linear/superlinear; $p > 2$
$\Theta(N^{-\frac{1}{M-1}})$	$\mathcal{O}(N^{-\frac{1}{M-1}})$	$\mathcal{O}(\log N)$	$\mathcal{O}(N^{\frac{1}{M-1}} \log N)$	$p = 2$ / moderate K
$\Theta(N^{-\frac{3}{2(M-1)}})$	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(N^{\frac{3}{2(M-1)}} \log N)$	$p < 2$ / fixed K

Table 1: Representative outer–inner regimes implied by $\bar{\alpha}^2 \eta_K = \mathcal{O}(N^{-\frac{3}{M-1}})$. The K and T columns use $\alpha = \Theta(\bar{\alpha})$ and linear inner convergence $\eta_K \le q^K$. In the last column, p refers to sublinear convergence $\eta_K = \mathcal{O}(K^{-p})$.

Appendix D gives the convergence proof. The requirement $N^{1/(M-1)} = \Omega(\kappa_\vartheta^{3/2})$ controls the condition-dependent transient, while MEQ redistribution brings C_t to $\mathcal{O}(1)$. Table 1 summarizes the resulting outer–inner tradeoff: larger $\bar{\alpha}$ is preferable for linear, superlinear, or sufficiently fast sublinear inner convergence, whereas stronger damping reduces the required inner accuracy and can make a fixed inner budget sufficient. At target accuracy ϵ, for fixed κ_ϑ, choosing $N = \Theta(\epsilon^{-(M-1)/2})$ sufficiently large to satisfy $N^{1/(M-1)} = \Omega(\kappa_\vartheta^{3/2})$ gives the $\mathcal{O}(\epsilon)$ resolution floor. If $\eta_K = \mathcal{O}(K^{-p})$, the three regimes require total repeated inner work $\widetilde{\mathcal{O}}(\epsilon^{-3/(2p)})$, $\widetilde{\mathcal{O}}(\epsilon^{-1/2-1/(2p)})$, and $\widetilde{\mathcal{O}}(\epsilon^{-3/4})$, respectively. Thus the first regime is preferable for $p > 2$, the three have the same leading ϵ-dependence at $p = 2$, and the fixed-budget regime is preferable for $p < 2$.

5 Experiments

We begin with controlled preference matching and geometric distortion, then evaluate uniform PF coverage on standard benchmarks and multi-objective reinforcement learning policy problems. Additional protocols, experiments, and analysis are in Appendix E.

5.1 Controlled Geometry and Preference Matching

This section considers an analytic three-objective problem:

$$\min_{\mathbf{x}\in\Delta_3} \mathbf{f}(\mathbf{x}) = \exp\left(\kappa \left\|\mathbf{x} - \tfrac{1}{3}\mathbf{1}\right\|_2^2\right)\mathbf{x}, \quad \kappa \ge 0. \tag{14}$$

Using TS, it gives $\mathbf{f}_{\mathrm{PF}}(\mathbf{w}) = \exp(\kappa\|\mathbf{w} - \frac{1}{3}\mathbf{1}\|_2^2)\mathbf{w}$ and larger κ produced more ill-conditioned PF with larger $\overline{J}/\underline{J} = \frac{\max_{\mathbf{w}\in\Delta_3}\|\mathbf{f}_{\mathrm{PF}}(\mathbf{w})\|_2}{\min_{\mathbf{w}\in\Delta_3}\|\mathbf{f}_{\mathrm{PF}}(\mathbf{w})\|_2} = \Theta(\exp(\kappa)\kappa)$. Setting $\kappa = 1$, we first illustrate how SURFACE controls *where* finite PF resolution is allocated through ϑ. We consider the uniform preference $\vartheta = 1$ and a focused preference $\vartheta = \vartheta_{\mathrm{focus}}(\mathbf{y}) = \exp\left(-\frac{\|\mathbf{y}-\mathbf{f}_{\mathrm{PF}}([0.55,0.30,0.15]^\top)\|_2^2}{2\sigma^2}\right)$, where σ is 15% of the PF diameter. As shown in Figure 3, using $N = 63$, $\alpha = 1$, and $T = 5$ on (14) with $\kappa = 1$, SURFACE under the focused preference concentrates Pareto solutions around the target region, while SURFACE with $\vartheta = 1$ redistributes samples toward uniform PF surface-area coverage. In contrast, fixed Uniform weighting inherits the uncontrolled weight-to-PF distortion and concentrates toward the PF center.

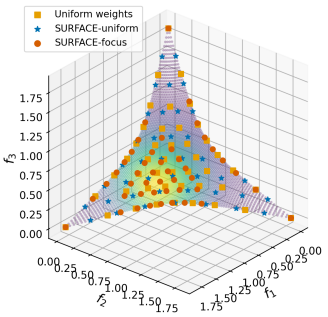

Figure 3: Uniform weighting versus SURFACE with $\vartheta = 1, \vartheta_{\mathrm{focus}}$. Warmer colors indicate higher $\vartheta_{\mathrm{focus}}$.

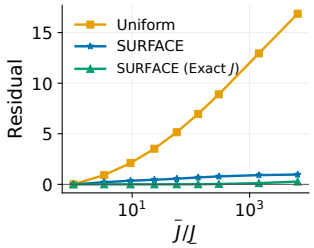

Figure 4: Residual $\mathcal{R}_\vartheta(W_T)$ versus PF distortion. SURFACE (Exact J) is a reference using analytic density.

We fix $\vartheta = 1$ and vary $\kappa \in \{0, 0.5, 1, 1.5, 2, 2.5, 3, 4, 5\}$ to examine increasingly distorted PFs. As shown in Figure 4, with $N = 63$ adaptive weights, $\alpha = 0.2$, and $T = 5$ rounds, the distortion ratio $\overline{J}/\underline{J}$ increases from 1 to 6.58×10^3. Uniform weighting deteriorates to $\mathcal{R}_\vartheta = 16.866$, whereas SURFACE remains below 0.969 throughout, corroborating Theorem 1 that the target-mass error need not scale with the scalarization-induced PF distortion. For reference, SURFACE (Exact J) replaces the reconstructed

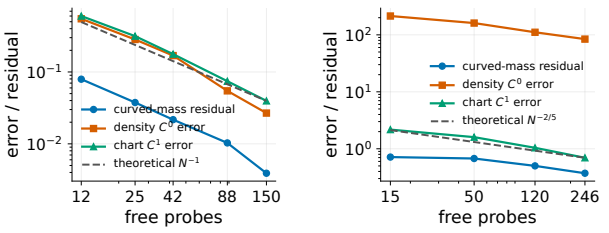

Figure 5: Finite-resolution errors versus the number of free probes N for $M = 3$ (left) and $M = 6$ (right). The decay is consistent with the predicted $\mathcal{O}(N^{-1})$ and $\mathcal{O}(N^{-2/5})$ rates.

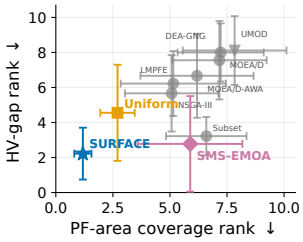

Figure 6: Mean ranks on PF coverage and HV over nine ZDT/DTLZ tasks. Error bars show standard deviations. Complete results are in Appendix E.3.

Metric	SMS-EMOA	NSGA-III	LMPFE	Subset	UMOD	Uniform	SURFACE
ZDT2, $M = 2$, $N = 9$							
$\hat{\mathcal{R}}_\vartheta \downarrow$	.480 ± .018	.589 ± .000	.881 ± .374	.491 ± .000	.812 ± .047	.514	**.397**
IGD ↓	.039 ± .000	.038 ± .000	.048 ± .013	.039 ± .000	.041 ± .000	.037	**.036**
Fill↓	.126 ± .003	.089 ± .000	.156 ± .068	.122 ± .000	.101 ± .003	.087	**.073**
ΔHV ↓	**.069 ± .000**	.076 ± .000	.097 ± .025	.075 ± .000	.084 ± .001	.078	.074
ZDT4, $M = 2$, $N = 9$							
$\hat{\mathcal{R}}_\vartheta \downarrow$	.418 ± .006	.611 ± .003	.572 ± .033	.760 ± .005	2.124 ± .972	.505	**.316**
IGD ↓	.037 ± .000	.043 ± .000	.038 ± .000	.040 ± .000	.108 ± .060	.037	**.035**
Fill↓	.090 ± .001	.150 ± .000	**.088 ± .004**	.173 ± .001	.324 ± .155	.106	.096
ΔHV ↓	.049 ± .000	.056 ± .001	.053 ± .002	.050 ± .000	.196 ± .122	.048	**.045**
DTLZ3, $M = 2$, $N = 9$							
$\hat{\mathcal{R}}_\vartheta \downarrow$	.509 ± .017	.616 ± .006	.492 ± .016	.624 ± .001	1.471 ± .399	.543	**.402**
IGD ↓	.046 ± .000	.041 ± .000	.041 ± .001	.048 ± .001	.085 ± .022	.040	**.038**
Fill↓	.143 ± .001	.099 ± .000	.096 ± .006	.148 ± .001	.233 ± .052	.097	**.078**
ΔHV ↓	**.070 ± .006**	.088 ± .003	.092 ± .007	.079 ± .013	.289 ± .113	.085	.074
DTLZ3, $M = 3$, $N = 18$							
$\hat{\mathcal{R}}_\vartheta \downarrow$	1.490 ± .402	1.179 ± .001	2.546 ± 1.447	1.588 ± .001	3.120 ± .202	1.078	**.693**
IGD ↓	.166 ± .005	.137 ± .000	.210 ± .062	.162 ± .000	.204 ± .022	.132	**.129**
Fill↓	.342 ± .020	.272 ± .000	.479 ± .137	.333 ± .000	.476 ± .050	.265	**.227**
ΔHV ↓	**.134 ± .001**	.162 ± .000	.258 ± .071	.139 ± .002	.266 ± .021	.175	.162
DTLZ4, $M = 3$, $N = 18$							
$\hat{\mathcal{R}}_\vartheta \downarrow$	20.000 ± .000	20.000 ± .000	4.698 ± 8.556	8.954 ± 10.083	20.000 ± .000	1.078	**.693**
IGD ↓	.943 ± .000	.943 ± .000	.296 ± .361	.475 ± .427	.943 ± .000	.132	**.129**
Fill↓	1.414 ± .000	1.414 ± .000	.523 ± .498	.765 ± .592	1.414 ± .000	.265	**.227**
ΔHV ↓	.848 ± .000	.848 ± .000	.307 ± .302	.422 ± .389	.848 ± .000	.175	**.162**

Table 2: Benchmark results for representative baselines. Entries for stochastic baselines are mean ± standard deviation over five seeds. SURFACE and Uniform are deterministic. Complete results for all ten methods are reported in Appendix E.3.

area density with the analytic J while keeping the same transport procedure, reducing the residual to 0.276 and indicating the remaining error after removing finite PF reconstruction. Figure 5 further shows that the finite-resolution errors decrease with N at rates consistent with the theoretical $\mathcal{O}(N^{-1})$ for $M = 3$ and $\mathcal{O}(N^{-2/5})$ for $M = 6$ in Theorem 1.

5.2 Uniform Pareto-Front Coverage on ZDT and DTLZ Benchmarks

We further evaluate SURFACE on ZDT1/2/4 and DTLZ2/3/4 (Zitzler et al., 2000, Deb et al., 2002) towards the uniform PF coverage preference $\vartheta = 1$. We compare SURFACE with *Uniform weighting* (Zhang & Li, 2007), *UMOD* (Zhang et al., 2024), *LMPFE* (Tian et al., 2022), *DEA-GNG* (Liu et al., 2019), *Subset selection* (Shang et al., 2023), *NSGA-III* (Jain & Deb, 2013), *SMS-EMOA* (Beume et al., 2007), *MOEA/D* (Zhang & Li, 2007), and *MOEA/D-AWA* (Qi et al., 2014), with matched returned cardinality. SURFACE and Uniform use the same scalarized solver, while the remaining baselines are evaluated end-to-end. We use target-mass residuals as the primary distribution-matching metrics, together with area-weighted IGD (Ishibuchi et al., 2015), fill distance (Zhang et al., 2024), and hypervolume (Beume et al., 2007, Shang et al., 2020). SURFACE and Uniform use the same scalarized optimization procedure and differ only in the choice

Metric	SMS-EMOA	NSGA-III	MOEA/D	MOEA/D-AWA	UMOD	Uniform	SURFACE
Deep Sea Treasure							
$\mathcal{R}_\vartheta \downarrow$	$.943 \pm .379$	$.838 \pm .279$	$1.296 \pm .201$	$1.156 \pm .299$	$1.138 \pm .371$	1.069	**.249**
IGD $\downarrow$	$.346 \pm .038$	$.335 \pm .030$	$.416 \pm .056$	$.387 \pm .067$	$.383 \pm .071$	.077	**.026**
Fill$\downarrow$	$.822 \pm .086$	$.797 \pm .072$	$.945 \pm .079$	$.895 \pm .104$	$.883 \pm .120$	.255	**.105**
HV$\uparrow$	$1.031 \pm .178$	$\mathbf{1.093 \pm .142}$	$.812 \pm .131$	$.909 \pm .187$	$.911 \pm .208$	1.041	1.043
Resource Gathering							
$\mathcal{R}_\vartheta \downarrow$	$1.806 \pm .055$	$1.848 \pm .052$	$1.830 \pm .042$	$1.783 \pm .117$	$1.848 \pm .053$	.454	**.392**
IGD $\downarrow$	$2.808 \pm .054$	$1.961 \pm .732$	4.722 ± 1.119	$1.216 \pm .245$	4.447 ± 1.846	.095	**.083**
Fill$\downarrow$	$3.288 \pm .037$	$2.377 \pm .800$	5.289 ± 1.236	$1.623 \pm .249$	4.956 ± 1.891	.200	**.154**
HV$\uparrow$	$1.174 \pm .243$	$.951 \pm .331$	$1.057 \pm .154$	$\mathbf{1.279 \pm .084}$	$1.150 \pm .250$	.114	.114
Fruit Tree							
$\mathcal{R}_\vartheta \downarrow$	$1.485 \pm .141$	$1.258 \pm .047$	$.956 \pm .028$	$1.138 \pm .061$	$1.254 \pm .087$	1.023	**.943**
IGD $\downarrow$	$.377 \pm .032$	$.262 \pm .014$	$.263 \pm .010$	$.284 \pm .017$	$.307 \pm .015$	.255	**.249**
Fill$\downarrow$	$.642 \pm .045$	$.662 \pm .008$	$.787 \pm .092$	$.556 \pm .141$	$.637 \pm .064$	.392	**.381**
HV$\uparrow$	$.035 \pm .006$	$.069 \pm .002$	$\mathbf{.078 \pm .002}$	$.078 \pm .002$	$.059 \pm .005$	.068	.069

Table 3: MO-Gym policy results. $\mathcal{R}_\vartheta$ is the independently evaluated PF-area target-mass imbalance, and Fill is the PF-area fill distance. Stochastic entries are mean $\pm$ sample standard deviation over five seeds; Uniform and SURFACE are deterministic. Bold marks the best mean in each row.

of scalarization weights, allowing us to directly evaluate the effect of SURFACE's adaptive weight redistribution.

Table 2 reports a preview of the results, while complete results for all ten methods are provided in Appendix E.3. consistently give strong performance. Across the 11 settings, SURFACE achieves the best mean rank, as visualized in Figure 6, where the coverage metric averages the first three PF-representation metrics. Compared with Uniform, SURFACE improves all four metrics in every primary comparison, with median reductions of 35.7% in PF-area imbalance, 3.64% in IGD, 14.2% in fill distance, and 7.21% in normalized hypervolume gap. These results show that adapting the scalarization weights to the PF geometry produces a more uniform finite representation while maintaining competitive objective-space quality.

5.3 Multi-Objective Policy Optimization

We additionally evaluate SURFACE's compatibility beyond TS by implementing it with weighted-sum scalarization on Deep Sea Treasure, Resource Gathering, and Fruit Tree from MO-Gymnasium (Felten et al., 2023). Under hidden strong convexity, weighted-sum scalarization is sufficient to recover the full PF, including the regularized reinforcement-learning formulation considered in SURF (Jiang et al., 2026); Appendix B.4 establishes the corresponding regular PF chart for this regime. Each scalarized problem is solved by a KL-regularized soft-value optimizer. With the uniform PF target $\vartheta \equiv 1$, SURFACE and Uniform use the same deterministic inner solver, while the remaining baselines use their standard evolutionary policy optimizers. As shown in Table 3, SURFACE consistently improves both PF-area residuals and area-weighted IGD over Uniform on all three environments while maintaining comparable hypervolume. Experimental settings, additional resolution studies, and policy diagnostics are provided in Appendix E.4.

6 Conclusion

This paper introduced SURFACE, a geometry-aware framework for matching a prescribed preference distribution on the PF by estimating the weight-to-PF chart from finitely many scalarized solutions and redistributing scalarization weights via mass equidistribution. Theoretically, we showed that Tchebycheff scalarization induces a smooth, nonfolding, and nondegenerate PF chart under verifiable regularity conditions, and proved that SURFACE contracts the coupled density–solver error geometrically to an $\mathcal{O}(N^{-2/(M-1)})$ finite-resolution floor, with the same rate for target-mass imbalance. Experiments across ZDT/DTLZ benchmarks and MO-Gymnasium policy tasks against nine baselines confirm consistent gains in preference matching and Pareto representation quality while remaining compatible with off-the-shelf solvers and inexact warm-started optimization. A natural future direction is to mitigate the rapid growth of the required number of adaptive weights N with the number of objectives M through adaptive sparse weight placement or hierarchical mesh refinement.

AI USE STATEMENT

We used generative AI tools to help polish the writing and to double-check parts of the mathematical proofs. All final arguments and proofs were reviewed and verified by the authors.

REPRODUCIBILITY STATEMENT

Implementation details for SURFACE and the experimental settings are provided in Appendix E. The code is provided in the supplementary material.

REFERENCES

Axel Abels, Diederik Roijers, Tom Lenaerts, Ann Nowé, and Denis Steckelmacher. Dynamic weights in multi-objective deep reinforcement learning. In *International conference on machine learning*, pp. 11–20. PMLR, 2019.

Lucas Nunes Alegre, Ana Bazzan, and Bruno C Da Silva. Optimistic linear support and successor features as a basis for optimal policy transfer. In *International conference on machine learning*, pp. 394–413. PMLR, 2022.

Fernanda Beltrán, Oliver Cuate, and Oliver Schütze. The pareto tracer for general inequality constrained multi-objective optimization problems. *Mathematical and Computational Applications*, 25(4):80, 2020.

Claude Berge. *Topological spaces: Including a treatment of multi-valued functions, vector spaces and convexity*. Oliver & Boyd, 1877.

Nicola Beume, Boris Naujoks, and Michael Emmerich. Sms-emoa: Multiobjective selection based on dominated hypervolume. *European journal of operational research*, 181(3):1653–1669, 2007.

Julian Blank, Kalyanmoy Deb, Yashesh Dhebar, Sunith Bandaru, and Haitham Seada. Generating well-spaced points on a unit simplex for evolutionary many-objective optimization. *IEEE Transactions on Evolutionary Computation*, 25(1):48–60, 2020.

Niclas Boehmer and Maximilian T Wittmann. Picking a representative set of solutions in multiobjective optimization: axioms, algorithms, and experiments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 16717–16725, 2026.

Matthias Bolten, Onur Tanil Doganay, Hanno Gottschalk, and Kathrin Klamroth. Tracing locally pareto-optimal points by numerical integration. *SIAM Journal on Control and Optimization*, 59 (5):3302–3328, 2021.

J Frédéric Bonnans and Alexander Shapiro. *Perturbation analysis of optimization problems*. Springer Science & Business Media, 2013.

Giacomo Borghi. Repulsion dynamics for uniform pareto front approximation in multi-objective optimization problems. *PAMM*, 23(1):e202200285, 2023.

Giacomo Borghi, Michael Herty, and Lorenzo Pareschi. An adaptive consensus based method for multi-objective optimization with uniform pareto front approximation. *Applied Mathematics & Optimization*, 88(2):58, 2023.

Martins Bruveris, Peter Michor, Adam Parusiński, and Armin Rainer. Moser's theorem on manifolds with corners. *Proceedings of the American Mathematical Society*, 146(11):4889–4897, 2018.

Chris J Budd and JF Williams. Moving mesh generation using the parabolic monge–ampère equation. *SIAM Journal on Scientific Computing*, 31(5):3438–3465, 2009.

Lisha Chen, AFM Saif, Yanning Shen, and Tianyi Chen. Ferero: A flexible framework for preference-guided multi-objective learning. *Advances in Neural Information Processing Systems*, 37:18758–18805, 2024.

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 *Forty-second International Conference on Machine Learning*, 2025.

Ran Cheng, Yaochu Jin, Markus Olhofer, and Bernhard Sendhoff. A reference vector guided evolutionary algorithm for many-objective optimization. *IEEE transactions on evolutionary computation*, 20(5):773–791, 2016.

Bernard Dacorogna and Jürgen Moser. On a partial differential equation involving the jacobian determinant. *Annales de l'Institut Henri Poincaré C, Analyse non linéaire*, 7(1):1–26, 1990.

Indraneel Das and John E Dennis. Normal-boundary intersection: A new method for generating the pareto surface in nonlinear multicriteria optimization problems. *SIAM journal on optimization*, 8 (3):631–657, 1998.

Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation*, 6(2): 182–197, 2002.

Jean-Antoine Désidéri. Multiple-gradient descent algorithm (mgda) for multiobjective optimization. *Comptes Rendus. Mathématique*, 350(5-6):313–318, 2012.

Qiang Du, Max D Gunzburger, and Lili Ju. Constrained centroidal voronoi tessellations for surfaces. *SIAM Journal on Scientific Computing*, 24(5):1488–1506, 2003.

Matthias Ehrgott. *Multicriteria optimization*. Springer, 2005.

Matthias Ehrgott and Stefan Ruzika. Improved ε-constraint method for multiobjective programming. *Journal of optimization Theory and Applications*, 138(3):375–396, 2008.

Gabriele Eichfelder. An adaptive scalarization method in multiobjective optimization. *SIAM Journal on Optimization*, 19(4):1694–1718, 2009.

Tohid Erfani and Sergei V Utyuzhnikov. Directed search domain: a method for even generation of the pareto frontier in multiobjective optimization. *Engineering Optimization*, 43(5):467–484, 2011.

Florian Felten, Lucas N Alegre, Ann Nowe, Ana Bazzan, El Ghazali Talbi, Grégoire Danoy, and Bruno C da Silva. A toolkit for reliable benchmarking and research in multi-objective reinforcement learning. *Advances in Neural Information Processing Systems*, 36:23671–23700, 2023.

Anthony V Fiacco. Sensitivity analysis for nonlinear programming using penalty methods. *Mathematical programming*, 10(1):287–311, 1976.

Claus Hillermeier. Generalized homotopy approach to multiobjective optimization. *Journal of Optimization Theory and Applications*, 110(3):557–583, 2001.

Weizhang Huang and Robert D Russell. *Adaptive moving mesh methods*. Springer Science & Business Media, 2010.

Hisao Ishibuchi, Hiroyuki Masuda, and Yusuke Nojima. A study on performance evaluation ability of a modified inverted generational distance indicator. In *Proceedings of the 2015 annual conference on genetic and evolutionary computation*, pp. 695–702, 2015.

Himanshu Jain and Kalyanmoy Deb. An evolutionary many-objective optimization algorithm using reference-point based nondominated sorting approach, part II: Handling constraints and extending to an adaptive approach. *IEEE Transactions on evolutionary computation*, 18(4):602–622, 2013.

Liuyuan Jiang, Chentong Huang, and Lisha Chen. Surf: Steering the scalarization weight to uniformly traverse the pareto front. *arXiv preprint arXiv:2605.20619*, 2026.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

Xi Lin, Hui-Ling Zhen, Zhenhua Li, Qing-Fu Zhang, and Sam Kwong. Pareto multi-task learning. *Advances in neural information processing systems*, 32, 2019.

Xi Lin, Zhiyuan Yang, Xiaoyuan Zhang, and Qingfu Zhang. Pareto set learning for expensive multi-objective optimization. *Advances in neural information processing systems*, 35:19231–19247, 2022.

Xi Lin, Xiaoyuan Zhang, Zhiyuan Yang, Fei Liu, Zhenkun Wang, and Qingfu Zhang. Smooth tchebycheff scalarization for multi-objective optimization. In *Forty-first International Conference on Machine Learning*, 2024.

Yiping Liu, Hisao Ishibuchi, Naoki Masuyama, and Yusuke Nojima. Adapting reference vectors and scalarizing functions by growing neural gas to handle irregular pareto fronts. *IEEE Transactions on Evolutionary Computation*, 24(3):439–453, 2019.

R Timothy Marler and Jasbir S Arora. The weighted sum method for multi-objective optimization: new insights. *Structural and multidisciplinary optimization*, 41(6):853–862, 2010.

Benjamin Martin, Alexandre Goldsztejn, Laurent Granvilliers, and Christophe Jermann. On continuation methods for non-linear bi-objective optimization: towards a certified interval-based approach. *Journal of Global Optimization*, 64(1):3–16, 2016.

Achille Messac, Amir Ismail-Yahaya, and Christopher A Mattson. The normalized normal constraint method for generating the pareto frontier. *Structural and multidisciplinary optimization*, 25(2): 86–98, 2003.

Kaisa Miettinen et al. *Nonlinear multiobjective optimization*, volume 12. Kluwer Academic Publishers Boston, 1999.

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

Mahboubeh Nezhadmoghaddam, Julio Juárez, Jesús Guillermo Falcón-Cardona, Michael TM Emmerich, and André Deutz. Approximating optimal μ-point distributions over a generalized sphere for riesz s-energy using a gradient descent algorithm. *Information Sciences*, pp. 123470, 2026.

Juuso Pajasmaa, Kaisa Miettinen, and Johanna Silvennoinen. Group decision making in multiobjective optimization: A systematic literature review. *Group Decision and Negotiation*, 34(2): 329–371, 2025.

Victor Pereyra. Fast computation of equispaced pareto manifolds and pareto fronts for multiobjective optimization problems. *Mathematics and Computers in Simulation*, 79(6):1935–1947, 2009.

Victor Pereyra, Michael Saunders, and Jose Castillo. Equispaced pareto front construction for constrained bi-objective optimization. *Mathematical and Computer Modelling*, 57(9-10):2122–2131, 2013.

Yutao Qi, Xiaoliang Ma, Fang Liu, Licheng Jiao, Jianyong Sun, and Jianshe Wu. MOEA/D with adaptive weight adjustment. *Evolutionary computation*, 22(2):231–264, 2014.

Pramit Saha, Felix Wagner, Divyanshu Mishra, Can Peng, Anshul Thakur, David A Clifton, Konstantinos Kamnitsas, and J Alison Noble. F^{3}ocus-federated finetuning of vision-language foundation models with optimal client layer updating strategy via multi-objective meta-heuristics. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20006–20017. IEEE, 2025.

Bhupinder Singh Saini, Hemant Kumar Singh, Babooshka Shavazipour, and Kaisa Miettinen. An efficient iterative approach for uniformly representing pareto fronts. In *International Conference on Evolutionary Multi-Criterion Optimization*, pp. 241–256. Springer, 2025.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. *Advances in neural information processing systems*, 31, 2018.

Ke Shang, Hisao Ishibuchi, Linjun He, and Lie Meng Pang. A survey on the hypervolume indicator in evolutionary multiobjective optimization. *IEEE Transactions on Evolutionary Computation*, 25(1):1–20, 2020.

Ke Shang, Tianye Shu, Hisao Ishibuchi, Yang Nan, and Lie Meng Pang. Benchmarking large-scale subset selection in evolutionary multi-objective optimization. *Information Sciences*, 622: 755–770, 2023.

Ruizhe Shi, Yifang Chen, Yushi Hu, Alisa Liu, Hannaneh Hajishirzi, Noah A Smith, and Simon S Du. Decoding-time language model alignment with multiple objectives. *Advances in Neural Information Processing Systems*, 37:48875–48920, 2024.

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

Ryoji Tanabe and Hisao Ishibuchi. An easy-to-use real-world multi-objective optimization problem suite. *Applied Soft Computing*, 89:106078, 2020.

Ye Tian, Langchun Si, Xingyi Zhang, Kay Chen Tan, and Yaochu Jin. Local model-based pareto front estimation for multiobjective optimization. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 53(1):623–634, 2022.

Yake Wei and Di Hu. Mmpareto: boosting multimodal learning with innocent unimodal assistance. In *International Conference on Machine Learning*, 2024.

Holger Wendland. *Scattered data approximation*, volume 17. Cambridge university press Cambridge, 2005.

Xiangmin Xu, Weizhang Huang, Robert D. Russell, and J. F. Williams. Convergence of de boor's algorithm for the generation of equidistributing meshes. *IMA Journal of Numerical Analysis*, 31 (2):580–596, 2011.

Rongguang Ye, Longcan Chen, Wei-Bin Kou, Jinyuan Zhang, and Hisao Ishibuchi. Pareto front shape-agnostic pareto set learning in multi-objective optimization. In *2024 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 4730–4736. IEEE, 2024.

Boxi Yu and Sergey Utyuzhnikov. Extension of the directed search domain algorithm for multi-objective optimization to higher dimensions. *Annals of Operations Research*, 340(2):1125–1160, 2024.

Guibin Zhang, Haiyang Yu, Kaiming Yang, Bingli Wu, Fei Huang, Yongbin Li, and Shuicheng Yan. Evoroute: Experience-driven self-routing LLM agent systems. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics*, pp. 38213–38225, 2026.

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

Sen Zhang, Jianwei Guo, Hui Zhang, Xiaohong Jia, Dong-Ming Yan, Junhai Yong, and Peter Wonka. Capacity constrained blue-noise sampling on surfaces. *Computers & Graphics*, 55:44–54, 2016.

Wei Hong Zhang and Tong Gao. A min–max method with adaptive weightings for uniformly spaced pareto optimum points. *Computers & structures*, 84(28):1760–1769, 2006.

Xiaoyuan Zhang, Genghui Li, Xi Lin, Yichi Zhang, Yifan Chen, and Qingfu Zhang. Gliding over the pareto front with uniform designs. *Advances in Neural Information Processing Systems*, 37: 2215–2245, 2024.

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

APPENDIX CONTENTS

A Related Work and Notation **14**
A.1 Detailed Related Work 14
A.2 Notation 16
B Proofs for Pareto-Front Geometry **16**
B.1 Proof of Lemma 1 16
B.2 Proof of Proposition 2 17
B.3 Sufficient Regular Regimes 18
B.4 Weighted-Sum Specialization 19
C Details for the SURFACE Algorithm **20**
C.1 Reconstructing the PF Between Sampled Weights 21
C.2 Finite MEQ: Mass Accuracy and Mesh Stability 22
C.2.1 Face-Preserving Moser Transport 22
C.2.2 Numerical MEQ Accuracy 23
C.2.3 Mesh Stability Across SURFACE Rounds 24
D Proofs for the Convergence Analysis **25**
D.1 Exact Scalarized Solves 25
D.2 Inexact Scalarized Solves 26
E Experiment Details **29**
E.1 Experimental Protocol and Benchmark Scope 29
E.2 Controlled Validation 30
E.3 Details of ZDT and DTLZ Benchmark Results 30
E.4 Multi-Objective Policy Optimization 31

A RELATED WORK AND NOTATION

This appendix expands the literature positioning in Section 1 and collects the notation used throughout the analysis.

A.1 DETAILED RELATED WORK

Preference specification and Pareto-set learning. Preference information in MOO is commonly incorporated either *a priori*, before optimization, or *a posteriori*, after constructing a Pareto representation (Pajasmaa et al., 2025). A priori methods use specified preferences to target preferred Pareto solutions or regions (Chen et al., 2024, 2025). Preference-conditioned learning extends this paradigm by learning a map from preferences to Pareto solutions. Pareto HyperNetworks (Navon et al., 2021) generate models conditioned on preference vectors, while Pareto Set Learning (Lin et al., 2022) learns a continuous preference-to-solution map for expensive MOO. Evolutionary Preference Sampling (Ye et al., 2024) adapts the preference distribution used for training, and shape-agnostic Pareto Set Learning (Ye et al., 2024) formulates Pareto-set learning as distribution transformation with a hypervolume objective. A posteriori methods instead construct a finite PF approximation for subsequent selection (Zhang et al., 2024, Boehmer & Wittmann, 2026). SURFACE connects these two paradigms: it specifies the desired allocation *a priori* through a probability distribution over the PF and returns an *a posteriori* finite representation whose resolution follows that distribution.

Scalarization and geometry-based PF coverage. Scalarization generates Pareto solutions by varying preference parameters. Weighted-sum scalarization (Marler & Arora, 2010) covers the full

PF under suitable convexity conditions but can miss Pareto solutions on nonconvex fronts (Miettinen et al., 1999, Ehrgott, 2005). Tchebycheff scalarization (Miettinen et al., 1999, Ehrgott, 2005), achievement scalarization (Ehrgott, 2005), and Pascoletti–Serafini scalarization (Eichfelder, 2009) provide alternative parameterizations under their respective conditions. Standard scalarization determines which Pareto solution corresponds to a parameter but does not control the distribution of a finite parameter set over the PF. In particular, uniformly spaced parameters can produce nonuniformly spaced Pareto points (Das & Dennis, 1998, Zhang & Gao, 2006, Jiang et al., 2026). We therefore use uniform scalarization weights as a baseline for isolating the effect of SURFACE's redistribution.

Several methods explicitly adapt scalarization parameters or objective-space constructions to improve PF coverage. Eichfelder (2009) uses local sensitivity information to adapt Pascoletti–Serafini parameters toward nearly equidistant objective-space points. Zhang & Gao (2006) adjusts min–max weights to generate more uniformly spaced Pareto points, while Saini et al. (2025) iteratively removes directions in densely represented regions and adds directions in sparse regions. Normal-boundary intersection (Das & Dennis, 1998) generates solutions along normal directions from reference points on the convex hull of the individual minima. The normalized normal-constraint method (Messac et al., 2003) normalizes the objectives and generates solutions through moving objective-space constraints. Directed Search Domain methods (Erfani & Utyuzhnikov, 2011) sequentially restrict objective-space regions to place new solutions in sparsely represented areas, with a higher-dimensional extension given by Yu & Utyuzhnikov (2024). These methods primarily target uniform or well-spread representations. Sensitivity-based methods require local solution information (Eichfelder, 2009), several adaptive methods construct new parameters sequentially (Zhang & Gao, 2006, Saini et al., 2025, Erfani & Utyuzhnikov, 2011), and geometric methods modify each optimization subproblem through additional directions or constraints (Das & Dennis, 1998, Messac et al., 2003). These requirements limit parallelization when each Pareto solution requires an expensive stochastic training run. SURFACE instead supports a general prescribed PF distribution, keeps the scalarized optimization problem unchanged, solves the current scalarized problems in parallel, and estimates the redistribution directly from their weight–objective pairs.

Finite Pareto representations: tracing, diversity, and design. Continuation methods construct finite PF representations by following local Pareto geometry. Hillermeier (2001) models Pareto candidates as a differentiable manifold and uses generalized homotopy to obtain a homogeneous discretization. Pereyra (2009), Pereyra et al. (2013) incorporate spacing conditions into continuation procedures, including objective-space arc-length spacing for constrained bi-objective problems. Martin et al. (2016) follows bi-objective branches with interval validation, Bolten et al. (2021) traces locally Pareto-optimal solutions through numerical integration of sensitivity information, and Beltrán et al. (2020) extends Pareto tracing to inequality-constrained problems. SURFACE follows the same principle of using local PF geometry to control solution placement. These continuation methods advance sequentially from previously computed solutions and use local derivative, tangent, or sensitivity information (Hillermeier, 2001, Pereyra et al., 2013, Bolten et al., 2021); SURFACE instead reconstructs the relevant PF geometry from parallel scalarized solves without tracing the PF sequentially.

Evolutionary decomposition and reference-vector methods construct finite representations by controlling search directions or reference points. MOEA/D (Zhang & Li, 2007) decomposes MOO into scalar subproblems associated with a fixed set of weights. MOEA/D-AWA (Qi et al., 2014) removes weights associated with crowded or duplicate solutions and adds weights in sparse PF regions. DEA-GNG (Liu et al., 2019) adapts preference vectors using a growing neural gas representation of the evolving PF. NSGA-III (Jain & Deb, 2013) maintains diversity through reference-point association, while RVEA (Cheng et al., 2016) combines angle-penalized distance with adaptive reference vectors. These methods determine the allocation indirectly through sparsity, occupancy, learned PF structure, or angular diversity. SURFACE instead specifies the desired objective-space allocation directly as a probability distribution and adapts the scalarization weights to match it.

Other methods define representativeness through an explicit finite-set objective. LMPFE (Tian et al., 2022) approximates the PF using multiple local models, while fixed-cardinality subset selection (Shang et al., 2023) selects a representative subset from a larger candidate set. UMOD (Zhang et al., 2024) measures representativeness by fill distance and optimizes a max-packing surrogate with a constant-factor fill-distance guarantee. Borghi (2023) uses binary repulsive potentials for bi-

and tri-objective PFs, while Borghi et al. (2023) combines gradient-free search with energy-based parameter dynamics. Riesz energy has been used to generate arbitrary-cardinality reference sets on the unit simplex (Blank et al., 2020) and to optimize point configurations directly on nonlinear PFs (Nezhadmoghaddam et al., 2026). These methods define the desired representation through approximation, covering, packing, subset selection, or repulsion. SURFACE instead takes the desired PF distribution as an input and equalizes its probability mass across the finite representation, with uniform PF surface area as a special case.

Set-quality indicators provide another, implicit specification of the desired finite representation. Classical indicators combine convergence and spread in different ways (Zitzler et al., 2000). Crowding distance (Deb et al., 2002) favors locally isolated solutions. IGD (Ishibuchi et al., 2015) measures distances to a chosen reference set, while IGD+ (Ishibuchi et al., 2015) modifies this distance to satisfy weak Pareto compliance. Hypervolume (Beume et al., 2007) measures dominated objective-space volume relative to a reference point, and its optimal fixed-cardinality distribution depends on the PF geometry and reference point (Shang et al., 2020). SMS-EMOA (Beume et al., 2007) selects solutions by hypervolume contribution, while indicator-based subset selection (Shang et al., 2023) retains the finite set preferred by the selected criterion. These indicators therefore induce an allocation through a reference set, reference point, or diversity criterion. SURFACE instead specifies the allocation directly through a target PF distribution and measures its realization through target-mass imbalance.

Mesh redistribution for distribution matching. SURFACE draws its redistribution mechanism from numerical mesh generation and geometric transport rather than from a competing MOO method. Adaptive moving meshes (Huang & Russell, 2010, Budd & Williams, 2009, Xu et al., 2011) equidistribute a prescribed monitor density while controlling mesh quality. Prescribed-Jacobian transport (Dacorogna & Moser, 1990) and Moser-type transport on domains with corners (Bruveris et al., 2018) construct smooth maps between positive densities. Centroidal Voronoi methods (Du et al., 2003) and capacity-constrained sampling (Zhang et al., 2016) provide related mechanisms for distributing finite sites. SURFACE follows this mass-equidistribution principle after the prescribed PF preference and the weight-to-PF geometry determine the density to redistribute. For two objectives and a constant preference profile, the construction reduces to SURF's chord-length CDF and inverse-CDF update (Jiang et al., 2026). SURF addresses uniform allocation along a one-dimensional PF, while SURFACE supports nonconstant PF distributions and higher-dimensional PFs through face-preserving transport. For two objectives, a nonconstant profile gives the corresponding preference-weighted arc-length construction.

A.2 Notation

Table 4 collects symbols that recur across sections. The last column points to each symbol's main definition or explanation. Each proof introduces its local notation.

B Proofs for Pareto-Front Geometry

This appendix proves the geometric results in Section 2. We first relate weight-space volume to preference-weighted PF area, then establish Tchebycheff coverage and regularity. We finally give concrete regimes in which the regularity assumptions hold, together with the weighted-sum specialization referenced in Section 2.1.

B.1 Proof of Lemma 1

Proof. For any measurable $A \subseteq \Delta_M$, injectivity and full rank of $\mathbf{f}_{\mathrm{PF}}$ give the area formula

$$\int_A \vartheta(\mathbf{f}_{\mathrm{PF}}(\mathbf{w}))J(\mathbf{w})\mathrm{d}\mathbf{w} = \int_{\mathbf{f}_{\mathrm{PF}}(A)} \vartheta(\mathbf{y})\mathrm{d}\mathcal{H}^{M-1}(\mathbf{y}). \tag{15}$$

Using (5),

$$\int_A \rho_\vartheta(\mathbf{w})\mathrm{d}\mathbf{w} = \frac{\int_{\mathbf{f}_{\mathrm{PF}}(A)} \vartheta(\mathbf{y})\mathrm{d}\mathcal{H}^{M-1}(\mathbf{y})}{\int_{\mathbf{f}_{\mathrm{PF}}(\Delta_M)} \vartheta(\mathbf{y})\mathrm{d}\mathcal{H}^{M-1}(\mathbf{y})}. \tag{16}$$

Table 4: Recurring notation guide.

Symbol	Meaning	Definition
Problem and PF chart		
M	Number of objectives.	p. 1
$\mathbf{x}$, $\mathcal{X}$	Decision variable and feasible set.	p. 1
$\mathbf{f} = (f_1, \ldots, f_M)$	Vector-valued objective.	p. 1
$\mathbf{z}$	Utopian Tchebycheff reference, with $z_m < \min_{\mathbf{x}\in\mathcal{X}} f_m(\mathbf{x})$.	p. 2
$\mathbf{w}$	Scalarization weight.	p. 2
$\mathbf{x}^\star(\mathbf{w})$, $\mathbf{x}^\star_{\mathrm{TS}}(\mathbf{w})$	Generic scalarized optimizer and its Tchebycheff specialization.	p. 3
$\mathbf{f}_{\mathrm{PF}}(\mathbf{w})$, $\mathcal{F}_{\mathrm{PF}}$	Scalarization-induced PF map and Pareto front.	p. 3
Δ_M, Δ_M^δ, δ	Weight simplex, truncated TS domain, and truncation margin $0 < \delta < 1/M$.	p. 4
PF target distribution and finite representation		
$J(\mathbf{w})$	PF-area Jacobian factor of the chart $\mathbf{f}_{\mathrm{PF}}$.	p. 3
ϑ, ρ_ϑ	PF preference density and its normalized pullback to weight space.	p. 3
W, $W^{(t)}$, $\mathcal{T}(W)$	Weight mesh, its round-t instance, and induced cell partition.	p. 3
a_τ	PF patch $a_\tau := \mathbf{f}_{\mathrm{PF}}(\tau)$ represented by cell τ.	p. 3
$\mathcal{R}_\vartheta(W)$	Worst-cell target-mass imbalance over $\mathcal{T}(W)$.	p. 3
N	Number of adaptive weights in SURFACE.	p. 5
SURFACE and inexact solves		
P_t	Reconstructed PF chart at outer round t.	p. 5
ρ_t, α	Outer-round weight density and damping parameter.	p. 5
$\mathcal{E}_t$	Coupled density–solver error at outer round t.	p. 6

For a cell $\tau \in \mathcal{T}(W)$, let $a_\tau := \mathbf{f}_{\mathrm{PF}}(\tau)$. Equation (16) gives

$$\mathcal{R}_\vartheta(W) = 0 \quad \Longleftrightarrow \quad \int_\tau \rho_\vartheta(\mathbf{w})\mathrm{d}\mathbf{w} = \frac{1}{|\mathcal{T}(W)|} \quad \text{for every } \tau \in \mathcal{T}(W), \tag{17}$$

which is equivalent to all PF patches a_τ carrying the same preference-weighted area.

The same identity gives the continuous distributional statement. For any measurable $B \subseteq \mathbf{f}_{\mathrm{PF}}(\Delta_M)$,

$$\Pr\left(\mathbf{f}_{\mathrm{PF}}(\mathbf{w}) \in B\right) = \int_{\mathbf{f}_{\mathrm{PF}}{}^{-1}(B)} \rho_\vartheta(\mathbf{w})\mathrm{d}\mathbf{w} = \frac{\int_B \vartheta(\mathbf{y})\mathrm{d}\mathcal{H}^{M-1}(\mathbf{y})}{\int_{\mathbf{f}_{\mathrm{PF}}(\Delta_M)} \vartheta(\mathbf{y})\mathrm{d}\mathcal{H}^{M-1}(\mathbf{y})}. \tag{18}$$

Thus the pushforward of ρ_ϑ under $\mathbf{f}_{\mathrm{PF}}$ has density proportional to ϑ with respect to PF surface area. □

For later reconstruction arguments, we use the same area formula for a generic chart. For a C^1 map P, let J_P denote the area factor obtained by replacing $\mathbf{f}_{\mathrm{PF}}$ with P in (4). If P is injective on a measurable region A, then $\mathcal{H}^{M-1}(P(A)) = \int_A J_P(\mathbf{w})\mathrm{d}\mathbf{w}$.. Thus $J_P(\mathbf{w})\mathrm{d}\mathbf{w}$ is the PF area represented by an infinitesimal weight-space region.

B.2 Proof of Proposition 2

Proof. Fix $0 < \delta < 1/M$. We first establish smoothness. On Δ_M^δ, the equivalent Tchebycheff objective $\max_{m=1,\ldots,M}(f_m(\mathbf{x}) - z_m)/w_m$ is jointly continuous in $(\mathbf{x}, \mathbf{w})$. Compactness of $\mathcal{X}$ and uniqueness in Assumption 1 therefore imply, by Berge's maximum theorem Berge (1877), that $\mathbf{x}^\star_{\mathrm{TS}}(\mathbf{w})$ and $s^\star(\mathbf{w})$ are continuous on Δ_M^δ. Since the reference point is utopian and every ray constraint is active, $s^\star(\mathbf{w}) > 0$, and compactness gives a uniform positive lower bound. Fix $\mathbf{w} \in \Delta_M^\delta$. Under LICQ, strict complementarity, and SOSC, the active Karush–Kuhn–Tucker (KKT) system of (6) satisfies the standard strong-regularity conditions for parametric nonlinear programming Bonnans & Shapiro (2013). Since the problem data are $C^{3,1}$, the KKT residual is $C^{2,1}$, and the implicit-function theorem yields a locally unique $C^{2,1}$ primal-dual solution mapping Bonnans &

Shapiro (2013). Strict complementarity preserves the active set locally. Continuity and uniqueness of the global solution identify this local mapping with the selected TS solution. Because Δ_M^δ is compactly contained in $\mathrm{relint}(\Delta_M)$, finitely many such neighborhoods cover Δ_M^δ, giving uniform $C^{2,1}$ bounds on a neighborhood of Δ_M^δ in the affine hyperplane $\mathbf{1}^\top \mathbf{w} = 1$. Hence $\mathbf{x}_{\mathrm{TS}}^\star$ and $s^\star$ are uniformly Lipschitz, and $\mathbf{f}_{\mathrm{PF}} = \mathbf{f} \circ \mathbf{x}_{\mathrm{TS}}^\star$ is $C^{2,1}$ on this neighborhood.

For the nonfolding geometry, the equivalent form gives $\mathbf{f}_{\mathrm{PF}}(\mathbf{w}) = \mathbf{z} + s^\star(\mathbf{w})\mathbf{w}$. For any tangent direction $\mathbf{v}$ satisfying $\mathbf{1}^\top \mathbf{v} = 0$,

$$D\mathbf{f}_{\mathrm{PF}}(\mathbf{w})[\mathbf{v}] = s^\star(\mathbf{w})\mathbf{v} + Ds^\star(\mathbf{w})[\mathbf{v}]\mathbf{w}. \tag{19}$$

If $D\mathbf{f}_{\mathrm{PF}}(\mathbf{w})[\mathbf{v}] = 0$, summing the coordinates in (19) gives $Ds^\star(\mathbf{w})[\mathbf{v}] = 0$. Since $s^\star(\mathbf{w}) > 0$, it follows that $\mathbf{v} = 0$, so $\mathbf{f}_{\mathrm{PF}}$ has full tangent rank. Injectivity follows from Proposition 1. The uniform C^1 bound gives $|\mathbf{f}_{\mathrm{PF}}(\mathbf{w}) - \mathbf{f}_{\mathrm{PF}}(\mathbf{v})| \le C_P|\mathbf{w} - \mathbf{v}|$. Full tangent rank, continuity of $D\mathbf{f}_{\mathrm{PF}}$, and compactness give a uniform local lower-Lipschitz bound. For pairs separated from the diagonal, compactness and injectivity imply that $|\mathbf{f}_{\mathrm{PF}}(\mathbf{w}) - \mathbf{f}_{\mathrm{PF}}(\mathbf{v})|/|\mathbf{w} - \mathbf{v}|$ attains a positive minimum. Combining the local and separated-pair bounds yields (7).

Finally, since $\mathbf{f}_{\mathrm{PF}}$ is $C^{2,1}$ on a neighborhood of Δ_M^δ, its derivative is $C^{1,1}$. Full tangent rank and compactness imply that the singular values of the tangent derivative are uniformly bounded above and bounded away from zero. The Gram-determinant map in (4) is continuous on full-rank matrices, so J is continuous on Δ_M^δ. Compactness therefore gives constants $0 < \underline{J} \le \overline{J} < \infty$ such that $\underline{J} \le J(\mathbf{w}) \le \overline{J}$ for all $\mathbf{w} \in \Delta_M^\delta$. This proves the claim. □

B.3 Sufficient Regular Regimes

Proposition 4 (Sufficient conditions for regular full-ray TS)**.** *Suppose Assumption 2 holds, $\mathcal{X}$ is compact, and $\mathbf{z}$ is utopian. Fix $0 < \delta < 1/M$ and assume that, for every $\mathbf{w} \in \Delta_M^\delta$, the following conditions hold.*

(i) *There exists a unique Pareto-optimal $\mathbf{x} \in \mathcal{X}$ satisfying $\frac{f_1(\mathbf{x})-z_1}{w_1} = \cdots = \frac{f_M(\mathbf{x})-z_M}{w_M}$.*

(ii) *At this point, the gradients $\nabla\left(\frac{f_m(\mathbf{x})-z_m}{w_m} - \frac{f_M(\mathbf{x})-z_M}{w_M}\right)$ for $m \in [M]$, together with the gradients of the active feasibility constraints, are linearly independent.*

(iii) *The constrained first-order stationarity condition for* (2) *admits strictly positive coefficients on all M normalized objectives.*

(iv) *For the same first-order multipliers, the corresponding Lagrangian has positive second variation along every nonzero feasible first-order direction for which all M normalized objectives have the same directional derivative.*

Then the restriction of Assumption 1 to Δ_M^δ holds, and Proposition 2 applies on this domain.

Proof. Fix $\mathbf{w} \in \Delta_M^\delta$ and let $\mathbf{x}$ be the unique point in (i). Let s be the common value of the normalized objective gaps. Then

$$f_m(\mathbf{x}) - z_m = sw_m, \quad m = 1, \dots, M, \tag{20}$$

so $(\mathbf{x}, s)$ is feasible for (6) and all ray inequalities are active.

We first show global optimality and uniqueness. If another feasible pair had an epigraph value smaller than s, then

$$f_m(\mathbf{x}') - z_m < f_m(\mathbf{x}) - z_m, \quad m = 1, \dots, M, \tag{21}$$

which would strictly improve every objective and contradict Pareto optimality of $\mathbf{x}$. Hence $(\mathbf{x}, s)$ is globally optimal. If another global solution had the same epigraph value, its objective values would be no larger than those of $\mathbf{x}$. Pareto optimality forces equality in every objective. The normalized objective gaps are then equal at that solution as well, and uniqueness in (i) gives the same decision point. Thus (6) has a unique solution.

We next verify LICQ. Consider any linear dependence among the gradients of the M active ray constraints and the active feasibility constraints in the variables $(\mathbf{x}, s)$. The s component implies

that the coefficients of the ray gradients have weighted sum zero. Eliminating the coefficient of the Mth ray constraint reduces the $\mathbf{x}$ component to a linear dependence among the $M-1$ normalized-gap gradients in part (ii) and the active feasibility-constraint gradients. Part (ii) forces all coefficients to vanish. Hence LICQ holds for (6).

Part (iii) gives strictly positive coefficients on the gradients $\nabla f_m(\mathbf{x})/w_m$ in the max-form stationarity condition. Dividing the coefficient associated with objective m by $w_m > 0$ gives a strictly positive multiplier for the corresponding ray inequality. Since the max-form coefficients sum to one, these ray multipliers satisfy the stationarity condition with respect to s. The stationarity condition with respect to $\mathbf{x}$ is unchanged. Hence every active ray inequality has a positive multiplier.

Finally, full activity and positive multipliers imply that every critical direction of (6) satisfies the linearized equalities of all M ray constraints. After eliminating the epigraph component, its decision component is a feasible first-order direction along which all normalized objectives have the same directional derivative. Conversely, if the decision component vanishes, the linearized ray equalities force the epigraph component to vanish as well. Part (iv) therefore gives strictly positive second variation along every nonzero critical direction. Thus SOSC holds, and all conditions in Assumption 1 follow. □

Remark 1 (DTLZ2)**.** *Standard DTLZ2 (Deb et al., 2002) satisfies Proposition 4 on every fixed Δ_M^δ for a sufficiently small strictly utopian reference $\mathbf{z} = -\epsilon\mathbf{1}$. Its Pareto front is the strictly positive portion of the unit sphere on this restricted weight domain. For sufficiently small $\epsilon > 0$, each ray from $\mathbf{z}$ in a direction $\mathbf{w} \in \Delta_M^\delta$ intersects this Pareto front at exactly one point, and the standard DTLZ2 parameterization gives a unique decision point realizing that intersection. This verifies part* (i)*. On the strictly positive Pareto front, the DTLZ2 objective map has full tangent rank. Since $\mathbf{f}_{\mathrm{PF}}(\mathbf{w})^\top \mathbf{w} > 0$, the radial direction is transverse to the tangent space of the sphere, which gives the linear independence in part* (ii)*. The max-form stationarity condition admits strictly positive coefficients proportional to $w_m f_m(\mathbf{x}^\star_{\mathrm{TS}}(\mathbf{w}))$, which verifies part* (iii)*. Finally, equality of the first-order normalized objective changes eliminates all nonzero tangent motion along the sphere, while the remaining DTLZ2 decision directions enter through strictly positive quadratic growth away from the Pareto front. This gives part* (iv)*. Hence the regular full-ray conditions hold uniformly on every fixed Δ_M^δ for sufficiently small $\epsilon > 0$.*

B.4 Weighted-Sum Specialization

Weighted-sum scalarization provides an alternative regular PF construction under hidden strong convexity. This specialization is also used in the objective-gap estimate of Appendix D.2. Let

$$\mathbf{x}^\star_{\mathrm{LS}}(\mathbf{w}) \in \operatorname*{argmin}_{\mathbf{x}\in\mathcal{X}} \sum_{m=1}^{M} w_m f_m(\mathbf{x}), \quad \mathbf{w} \in \Delta_M. \tag{22}$$

Assumption 4 (Hidden strongly convex weighted-sum regime)**.** *There exist a compact convex set $\mathcal{Y}$ and a $C^{3,1}$ diffeomorphism Ψ from a neighborhood of $\mathcal{Y}$ onto a neighborhood of $\mathcal{X}$ such that $\Psi(\mathcal{Y}) = \mathcal{X}$. Under the change of variables $\mathbf{x} = \Psi(\mathbf{y})$, the set $\mathcal{Y}$ is described by $C^{3,1}$ convex inequalities and affine equalities, and every $f_m \circ \Psi$ is μ-strongly convex on $\mathcal{Y}$ for a common $\mu > 0$. At the minimizer corresponding to every $\mathbf{w} \in \Delta_M$, the linear independence constraint qualification (LICQ) and strict complementarity hold. Moreover, for every nonzero $\mathbf{v}$ satisfying $\mathbf{1}^\top \mathbf{v} = 0$, there exists a feasible tangent direction d such that $d^\top \sum_{m=1}^{M} v_m \nabla(f_m \circ \Psi) \neq 0$, where the gradients are evaluated at the transformed minimizer.*

Thus the objectives need not be convex in the original model parameter $\mathbf{x}$. They become strongly convex only after the smooth change of variables. The last condition ensures that a nonzero change in scalarization weights changes the attained Pareto point to first order.

Proposition 5 (Regular PF chart from weighted sums)**.** *Suppose Assumption 4 holds. Then* (22) *has a unique solution for every $\mathbf{w} \in \Delta_M$, weighted-sum scalarization covers the PF, and $\mathbf{f}_{\mathrm{PF}}(\mathbf{w}) := \mathbf{f}(\mathbf{x}^\star_{\mathrm{LS}}(\mathbf{w}))$ is injective with full tangent rank. Moreover, $\mathbf{f}_{\mathrm{PF}}$ admits a $C^{2,1}$ extension to a neighborhood of Δ_M in the affine hyperplane $\mathbf{1}^\top \mathbf{w} = 1$. There exist constants $0 < c_P \leq C_P < \infty$ such that*

$$c_P \|\mathbf{w} - \mathbf{v}\| \leq \|\mathbf{f}_{\mathrm{PF}}(\mathbf{w}) - \mathbf{f}_{\mathrm{PF}}(\mathbf{v})\| \leq C_P \|\mathbf{w} - \mathbf{v}\|, \quad \mathbf{w}, \mathbf{v} \in \Delta_M. \tag{23}$$

The area factor J is $C^{1,1}$ on Δ_M and satisfies $0 < \underline{J} \leq J(\mathbf{w}) \leq \overline{J} < \infty$.

Proof. Under the change of variables $\mathbf{x} = \Psi(\mathbf{y})$, problem (22) becomes the minimization of $\sum_{m=1}^{M} w_m(f_m \circ \Psi)(\mathbf{y})$ over $\mathcal{Y}$. Since the weights are nonnegative and sum to one, this objective is μ-strongly convex. Hence the minimizer is unique for every $\mathbf{w} \in \Delta_M$, which makes $\mathbf{x}^\star_{\mathrm{LS}}(\mathbf{w})$ well defined.

The transformed objectives are convex on the convex set $\mathcal{Y}$. Standard completeness of weighted-sum scalarization for convex multi-objective problems therefore implies that every Pareto point minimizes a nonnegative weighted sum (Ehrgott, 2005). Conversely, every solution of (22) is Pareto optimal. If a different feasible point dominated it, then either the weighted objective would decrease or the same weighted value would be attained at two different points, contradicting optimality or uniqueness. Since Ψ is one-to-one, Pareto optimality is unchanged by the transformation. Thus weighted-sum scalarization covers the PF.

We next establish smooth dependence on the weights. Strong convexity of the transformed weighted objective and convexity of the active inequality constraints give positive second variation along every nonzero feasible tangent direction, so SOSC holds. Together with LICQ and strict complementarity, the Karush–Kuhn–Tucker (KKT) system satisfies the standard strong-regularity conditions for parametric nonlinear programming (Bonnans & Shapiro, 2013). Since the transformed problem data are $C^{3,1}$, the KKT residual is $C^{2,1}$, and the implicit-function theorem gives a locally unique $C^{2,1}$ primal-dual solution mapping (Bonnans & Shapiro, 2013). Strict complementarity preserves the active set locally. Uniqueness of the global minimizer identifies this local mapping with the weighted-sum solution. Compactness of Δ_M then gives uniform $C^{2,1}$ bounds and a $C^{2,1}$ extension to a neighborhood of Δ_M in the affine hyperplane $\mathbf{1}^\top \mathbf{w} = 1$. Composing with Ψ gives the same regularity for $\mathbf{x}^\star_{\mathrm{LS}}$ and $\mathbf{f}_{\mathrm{PF}}$.

We now prove full tangent rank. Fix $\mathbf{w} \in \Delta_M$ and a nonzero $\mathbf{v}$ satisfying $\mathbf{1}^\top \mathbf{v} = 0$. Let $\dot{\mathbf{y}}$ be the derivative, in direction $\mathbf{v}$, of the transformed minimizer $\Psi^{-1}(\mathbf{x}^\star_{\mathrm{LS}}(\mathbf{w}))$. Suppose $D\mathbf{f}_{\mathrm{PF}}(\mathbf{w})[\mathbf{v}] = 0$. Then $\nabla(f_m \circ \Psi)^\top \dot{\mathbf{y}} = 0$ for every m. Differentiating the KKT conditions and taking the inner product of the stationarity equation with $\dot{\mathbf{y}}$ gives $\dot{\mathbf{y}}^\top H \dot{\mathbf{y}} = 0$, where H is the Lagrangian Hessian. The second-order condition therefore gives $\dot{\mathbf{y}} = 0$. The differentiated stationarity condition then implies that $\sum_{m=1}^{M} v_m \nabla(f_m \circ \Psi)$ is a linear combination of the active constraint gradients. Its inner product with every feasible tangent direction is therefore zero, contradicting Assumption 4. Hence $\mathbf{v} = 0$, and $\mathbf{f}_{\mathrm{PF}}$ has full tangent rank.

To prove injectivity, suppose $\mathbf{f}_{\mathrm{PF}}(\mathbf{w}) = \mathbf{f}_{\mathrm{PF}}(\mathbf{v})$. The corresponding transformed minimizers have the same objective vector, so each attains the optimal weighted value of the other problem. Uniqueness therefore gives the same transformed decision point. Subtracting the two KKT stationarity conditions at this point shows that $\sum_{m=1}^{M}(w_m - v_m)\nabla(f_m \circ \Psi)$ is a linear combination of the active constraint gradients. Since $\mathbf{1}^\top(\mathbf{w} - \mathbf{v}) = 0$, the last condition in Assumption 4 implies $\mathbf{w} = \mathbf{v}$. Thus $\mathbf{f}_{\mathrm{PF}}$ is injective.

The uniform C^1 bound on the compact simplex gives the upper bound in (23). Full tangent rank and continuity of $D\mathbf{f}_{\mathrm{PF}}$ give a uniform local lower bound. For pairs separated from the diagonal, compactness and injectivity imply that $\|\mathbf{f}_{\mathrm{PF}}(\mathbf{w}) - \mathbf{f}_{\mathrm{PF}}(\mathbf{v})\|/\|\mathbf{w} - \mathbf{v}\|$ attains a positive minimum. Combining these two cases gives the lower bound in (23).

Finally, the $C^{2,1}$ regularity of $\mathbf{f}_{\mathrm{PF}}$ implies that $D\mathbf{f}_{\mathrm{PF}}$ is $C^{1,1}$. Full tangent rank and compactness imply that the singular values of the tangent derivative are uniformly bounded above and bounded away from zero. The Gram-determinant map in (4) is smooth on this set of derivatives, so J is $C^{1,1}$ on Δ_M. Compactness then gives $0 < \underline{J} \le J(\mathbf{w}) \le \overline{J} < \infty$ for every $\mathbf{w} \in \Delta_M$. □

C Details for the SURFACE Algorithm

This section gives the reconstruction and redistribution results used by SURFACE. We keep the notation of the main paper. In particular, $\mathbf{f}_{\mathrm{PF}}$ and J denote the exact PF chart and its area factor, P_t and J_{P_t} denote the reconstructed chart and its area factor at round t, ρ_ϑ is the target weight density, ρ_t is the density used to generate the current weights, and $W^{(t)}$ is the corresponding weight mesh. Throughout, we work on any fixed restricted TS weight domain Δ_M^δ from Proposition 2. Its compact inclusion in $\mathrm{relint}(\Delta_M)$ gives the uniform regularity bounds used below; these constants may depend on the fixed δ and need not remain uniform as $\delta \downarrow 0$. Constants denoted by C may

depend on fixed smoothness, preference, interpolation, and chart-regularity bounds, including the fixed bounds entering κ_ϑ, but not on N, T, t, or α. We keep the round-dependent reconstruction coefficient C_t explicit. This separates the approximation order from the current PF resolution.

For the mass analysis, we use one additional cell notation. For a mesh W produced by finite MEQ, let $\mathcal{C}(W)$ denote the complete cells obtained by transporting the fixed equal-volume reference cells. The existing $\mathcal{T}(W)$ remains the straight simplicial mesh through the transported vertices. Define

$$\widehat{\mathcal{R}}_\vartheta(W) := \max_{\tau\in\mathcal{C}(W)} \left| |\mathcal{C}(W)| \int_\tau \rho_\vartheta(\mathbf{w})\,\mathrm{d}\mathbf{w} - 1 \right|. \tag{24}$$

Thus $\mathcal{R}_\vartheta$ in Lemma 1 is evaluated on the straight cells $\mathcal{T}(W)$, whereas $\widehat{\mathcal{R}}_\vartheta$ is evaluated on the complete transported cells $\mathcal{C}(W)$. The two coincide when $M = 2$.

C.1 Reconstructing the PF Between Sampled Weights

SURFACE reconstructs the PF from the weight–objective pairs returned at the current round. We use the quadratically augmented local radial basis function (RBF) interpolation described in Section 3. The local fits use a compactly supported Wendland kernel and are combined by a partition of unity (Wendland, 2005). Proposition 2 provides the $C^{2,1}$ regularity required below, including one-sided extensions near the boundary of Δ_M^δ. To retain the effect of the current PF resolution, define

$$C_t := 1 + N^{2/(M-1)} \max_{\tau\in\mathcal{T}(W^{(t)})} \operatorname{diam}\left(\mathbf{f}_{\mathrm{PF}}(\tau)\right)^2. \tag{25}$$

Thus C_t measures the largest PF patch represented by one current straight cell relative to the nominal resolution $N^{-1/(M-1)}$. It remains bounded at every admissible round, but its initial bound can be loose when uniform scalarization weights resolve a target with large κ_ϑ unevenly.

Lemma 3 (PF reconstruction accuracy)**.** *Suppose Assumptions 1 and 2 hold. Suppose the current weight mesh is quasi-uniform and shape regular on Δ_M^δ and on each relative face, with mesh spacing comparable to $N^{-1/(M-1)}$. If the adaptive scalarized solutions used to construct P_t are exact, then, for all sufficiently large N, the restriction of P_t to Δ_M^δ satisfies*

$$\begin{aligned} \|P_t - \mathbf{f}_{\mathrm{PF}}\|_{C^1} &\le C C_t N^{-2/(M-1)}, \\ \|P_t - \mathbf{f}_{\mathrm{PF}}\|_{C^2} &\le C C_t^{1/2} N^{-1/(M-1)}, \\ \|P_t\|_{C^{2,1}} &\le C. \end{aligned} \tag{26}$$

If the scalarized solutions are inexact, then

$$\|P_t - \mathbf{f}_{\mathrm{PF}}\|_{C^1} \le C C_t N^{-2/(M-1)} + C N^{1/(M-1)} \max_{n\in[N]} \left\| \tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t)}) \right\|, \tag{27}$$

$$\|P_t - \mathbf{f}_{\mathrm{PF}}\|_{C^2} \le C C_t^{1/2} N^{-1/(M-1)} + C N^{2/(M-1)} \max_{n\in[N]} \left\| \tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t)}) \right\|. \tag{28}$$

Moreover, if the maximum decision error is at most $cN^{-3/(M-1)}$ for a fixed sufficiently small c, then P_t has a uniform $C^{2,1}$ bound on Δ_M^δ. The same bounds hold when the reconstruction additionally includes the fixed extremal anchors in Algorithm 1, since these anchors lie a fixed positive distance from Δ_M^δ and therefore do not enter the local compactly supported interpolation patches for sufficiently large N.

Proof. For a cell of the current mesh, the co-Lipschitz bound in Proposition 2 bounds its weight-space diameter by a constant times the diameter of its PF image. Equation (25) therefore bounds the largest local interpolation scale by $CC_t^{1/2}N^{-1/(M-1)}$. On each local patch, take the quadratic Taylor polynomial of $\mathbf{f}_{\mathrm{PF}}$ at any point in the patch. Since $\mathbf{f}_{\mathrm{PF}}$ is $C^{2,1}$, quadratic reproduction and uniform stability of the scaled RBF systems give the stated C^1 and C^2 errors. Bounded overlap of the partition of unity gives the global bounds. The same scaled argument through third derivatives gives the uniform $C^{2,1}$ bound for exact data. The one-sided extensions supplied by Proposition 2 give the same estimates at the boundary.

For inexact values, the interpolation operator is linear in its nodal data. Rescaling a patch to unit diameter shows that one derivative amplifies a nodal perturbation by at most $CN^{1/(M-1)}$, two derivatives by at most $CN^{2/(M-1)}$, and the Lipschitz second-derivative norm by at most $CN^{3/(M-1)}$. Assumption 2 gives a uniform Lipschitz bound for $\mathbf{f}$ near the selected TS solutions. Hence decision error controls the objective-vector perturbation. Adding the nodal perturbation to the exact-data estimates proves (27)–(28). If the decision error is $O(N^{-3/(M-1)})$, the third-derivative amplification is $O(1)$, which preserves the uniform $C^{2,1}$ bound. □

The next result transfers chart error to the density used in (9).

Lemma 4 (Preference-density stability)**.** *Suppose Assumptions 1–3 hold. If P_t is sufficiently close to $\mathbf{f}_{\mathrm{PF}}$ in C^1, then P_t has full tangent rank and*

$$\left\| \frac{\vartheta(P_t) J_{P_t}}{\int_{\Delta_M^\delta} \vartheta(P_t) J_{P_t}} - \rho_\vartheta \right\|_\infty \le C \|P_t - \mathbf{f}_{\mathrm{PF}}\|_{C^1}, \tag{29}$$

$$\left\| \frac{\vartheta(P_t) J_{P_t}}{\int_{\Delta_M^\delta} \vartheta(P_t) J_{P_t}} - \rho_\vartheta \right\|_{C^1} \le C \|P_t - \mathbf{f}_{\mathrm{PF}}\|_{C^2}. \tag{30}$$

If P_t has a uniform $C^{2,1}$ bound, then the normalized density on the left is strictly positive and has a uniform $C^{1,1}$ bound.

Proof. Proposition 2 gives a uniform positive lower bound on the tangent singular values of $D\mathbf{f}_{\mathrm{PF}}$. A sufficiently small C^1 perturbation preserves this bound for DP_t. The Gram-determinant expression in (4) is smooth on the resulting compact full-rank set. Hence

$$\|J_{P_t} - J\|_\infty \le C \|P_t - \mathbf{f}_{\mathrm{PF}}\|_{C^1}, \quad \text{and} \quad \|J_{P_t} - J\|_{C^1} \le C \|P_t - \mathbf{f}_{\mathrm{PF}}\|_{C^2}. \tag{31}$$

Assumption 3 and the chain rule give the corresponding estimates for $\vartheta(P_t) - \vartheta(\mathbf{f}_{\mathrm{PF}})$. The product rule gives the same orders for $\vartheta(P_t)J_{P_t} - \vartheta(\mathbf{f}_{\mathrm{PF}})J$. Both integrands have positive lower and finite upper bounds, so their normalizing integrals are bounded away from zero. Normalization preserves the estimates. The final claim follows from the same chain and product rules using the uniform $C^{2,1}$ bound. □

Combining Lemmas 3 and 4 gives

$$\left\| \frac{\vartheta(P_t) J_{P_t}}{\int_{\Delta_M^\delta} \vartheta(P_t) J_{P_t}} - \rho_\vartheta \right\|_\infty \le C C_t N^{-2/(M-1)} + C N^{1/(M-1)} \max_{n \in [N]} \left\| \tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t)}) \right\|. \tag{32}$$

This is the reconstruction bound stated in Lemma 2. Decision errors of order $N^{-3/(M-1)}$ contribute at the same $N^{-2/(M-1)}$ scale after differentiation.

C.2 Finite MEQ: Mass Accuracy and Mesh Stability

This subsection establishes the two properties of finite MEQ used in the convergence analysis. The first is accurate target-mass redistribution on the complete represented cells $\mathcal{C}(W)$. The second is stability of the straight vertex mesh $\mathcal{T}(W)$ used for PF reconstruction.

C.2.1 Face-Preserving Moser Transport

SURFACE converts the updated density into scalarization weights by transporting a fixed equal-volume reference mesh. The transport preserves every relative face of Δ_M^δ, so the boundary structure of the weight domain is unchanged.

Proposition 6 (Face-preserving MEQ transport)**.** *Let ρ be a positive normalized $C^{1,1}$ density on Δ_M^δ whose lower bound, upper bound, and $C^{1,1}$ norm are fixed. Then a Moser transport gives a face-preserving $C^{1,1}$ diffeomorphism of Δ_M^δ that maps every equal-volume reference cell to a represented cell with equal ρ-mass (Bruveris et al., 2018). The map and its inverse are uniformly Lipschitz, with constants determined only by the stated density bounds.*

If the divergence solve and flow integration are computed to accuracy $\mathcal{O}(N^{-2/(M-1)})$, the finite transport returns $W = \mathrm{MEQ}_N(\rho)$ whose represented cells satisfy

$$\max_{\tau \in \mathcal{C}(W)} \left| |\mathcal{C}(W)| \int_\tau \rho(\mathbf{w})\, \mathrm{d}\mathbf{w} - 1 \right| \le C N^{-2/(M-1)}. \tag{33}$$

The transported vertices depend Lipschitz-continuously on ρ in the uniform norm.

Proof. The density ρ and the uniform density have the same total mass and are bounded away from zero. Moser's construction on manifolds with corners gives a flow whose velocity is tangent to every relative face and whose time-one map, denoted by Φ, satisfies

$$\rho(\Phi(\mathbf{w}))\, \det D\Phi(\mathbf{w}) = \frac{1}{\mathrm{Vol}_{M-1}(\Delta_M^\delta)}. \tag{34}$$

The tangency condition keeps every relative face invariant. Applying the change-of-variables formula to an equal-volume reference cell gives exact mass equidistribution. The common lower, upper, and C^1 density bounds give uniform bounds on the Moser velocity field. Standard stability of ordinary differential equation flows gives uniform Lipschitz bounds for Φ and its inverse and Lipschitz dependence on the input density (Bruveris et al., 2018).

For the finite implementation, the divergence residual and the C^1 flow error perturb (34) by at most the same order. The uniform density and map bounds make this estimate uniform over all reference cells. Running both numerical components to $\mathcal{O}(N^{-2/(M-1)})$ accuracy therefore gives (33). The same flow-stability estimate gives the stated dependence of the transported vertices on the input density. □

For $M = 2$, the face-preserving transport is the monotone inverse cumulative-distribution transform (Jiang et al., 2026). For $M \ge 3$, Proposition 6 gives the corresponding multidimensional construction.

C.2.2 Numerical MEQ Accuracy

The exact Moser transport maps the equal-volume reference cells to represented cells of exactly equal target mass. The finite implementation only needs to approximate this transport at the resolution of the PF reconstruction, namely $\mathcal{O}(N^{-2/(M-1)})$. We therefore compute each numerical component so that its contribution to the transported density is of this order.

The normalizing integral in (9) is evaluated by adaptive quadrature until the resulting normalized density differs by at most $\mathcal{O}(N^{-2/(M-1)})$. For $M = 2$, SURFACE numerically evaluates and inverts the cumulative density, recovering the one-dimensional transport used by SURF (Jiang et al., 2026). For $M \ge 3$, SURFACE solves the Moser divergence equation with face-tangent boundary conditions and integrates the resulting flow (Bruveris et al., 2018). The internal discretization is independent of the returned weight resolution and is refined until the divergence residual and the C^1 flow error are both $\mathcal{O}(N^{-2/(M-1)})$.

This accuracy gives the finite cell-mass rate directly. The numerical errors perturb (34) by $\mathcal{O}(N^{-2/(M-1)})$ uniformly over the weight domain. Each reference cell has volume $\Theta(N^{-1})$, so its absolute transported-mass error is $\mathcal{O}(N^{-1-2/(M-1)})$. The target mass of each represented cell is $\Theta(N^{-1})$. Therefore its relative mass error is $\mathcal{O}(N^{-2/(M-1)})$, which is exactly (33).

After transport, SURFACE recomputes the represented-cell masses by independent quadrature. It is sufficient to evaluate each cell integral to absolute accuracy

$$\mathcal{O}\left(N^{-1-2/(M-1)}\right) = \mathcal{O}\left(N^{-(M+1)/(M-1)}\right). \tag{35}$$

The resulting normalized masses provide a direct certificate for (33). Boundary drift, cell orientation, and cell conditioning are recorded as geometric diagnostics. They are not additional terms in the convergence recursion.

C.2.3 MESH STABILITY ACROSS SURFACE ROUNDS

Mass accuracy concerns the complete represented cells. PF reconstruction uses the straight simplices through the transported vertices. We next show that these simplices remain uniformly suitable for interpolation.

Lemma 5 (Redistribution preserves mesh quality)**.** *Suppose the density passed to* MEQ_N *has common positive lower and upper bounds and a common* $C^{1,1}$ *bound. For all sufficiently large* N*, the transported vertices form a quasi-uniform, shape-regular simplicial mesh of* Δ_M^δ*. The same conclusion holds on each relative face. The constants are independent of the outer round and of* N*.*

Proof. Let Φ denote the exact Moser map for the current density. Proposition 6 gives uniform Lipschitz and inverse-Lipschitz bounds for Φ and a uniform $C^{1,1}$ bound. On the fixed shape-regular reference mesh, standard piecewise-affine interpolation approximates Φ with value error $\mathcal{O}(N^{-2/(M-1)})$ and derivative error $\mathcal{O}(N^{-1/(M-1)})$. For sufficiently large N, the derivative error is smaller than half the uniform lower singular-value bound of $D\Phi$. The affine map on every reference simplex therefore has singular values bounded above and away from zero. Hence the straight transported simplices are uniformly shape regular.

The same bi-Lipschitz bounds map reference vertices separated by order $N^{-1/(M-1)}$ to transported vertices with comparable separation and fill distance. Face preservation gives the same conclusion on every relative face. The numerical transport differs from the exact map by $\mathcal{O}(N^{-2/(M-1)})$, which is smaller than the vertex-separation scale. It therefore does not change the separation, fill-distance, or shape-regularity conclusions for sufficiently large N. □

The preceding results are local to one admissible density. Exact SURFACE updates preserve the required density bounds from one round to the next.

Lemma 6 (Exact round-to-round regularity)**.** *Suppose Assumptions 1–3 hold and the scalarized solutions used to construct* P_t *are exact. Then, for all sufficiently large* N*, every reconstructed chart has full tangent rank and a uniform* $C^{2,1}$ *bound. The normalized reconstructed density in* (9) *has uniform positive lower and upper bounds and a uniform* $C^{1,1}$ *bound. The damped density* ρ_{t+1} *has the same properties,* $W^{(t+1)}$ *satisfies Lemma 5, and* C_t *remains bounded uniformly in* t *and* N*.*

Proof. The uniform initial density generates the fixed reference mesh. Suppose the current mesh satisfies Lemma 5. Its straight cell diameters are $\mathcal{O}(N^{-1/(M-1)})$. The Lipschitz bound for $\mathbf{f}_{\mathrm{PF}}$ in Proposition 2 therefore gives $C_t \le C$. Lemma 3 gives $\|P_t - \mathbf{f}_{\mathrm{PF}}\|_{C^1} = \mathcal{O}(N^{-2/(M-1)})$ and a uniform $C^{2,1}$ bound. For sufficiently large N, Lemma 4 gives full tangent rank and common positive, finite, and $C^{1,1}$ bounds for the normalized reconstructed density. The update (9) is a convex combination of that density and ρ_t, so it preserves the same type of bounds. Proposition 6 and Lemma 5 then give the next regular mesh. Induction proves the result. □

We can now prove the finite-MEQ statement in the main paper without identifying the represented cells with their straight simplicial display.

Proof of Proposition 3. Let ρ be the density passed to finite MEQ. The hypotheses of Proposition 3 give fixed positive lower and upper bounds and a fixed $C^{1,1}$ bound for ρ. Proposition 6 therefore gives

$$\max_{\tau\in\mathcal{C}(\mathrm{MEQ}_N(\rho))} \left| |\mathcal{C}(\mathrm{MEQ}_N(\rho))| \int_\tau \rho(\mathbf{w})\,\mathrm{d}\mathbf{w} - 1 \right| \le CN^{-2/(M-1)}. \tag{36}$$

This is the stated relative cell-mass accuracy on the complete represented cells.

The same density bounds give uniform Lipschitz, inverse-Lipschitz, and $C^{1,1}$ bounds for the transport. Lemma 5 therefore shows that the transported vertices define a quasi-uniform, shape-regular straight simplicial mesh at scale $N^{-1/(M-1)}$. Hence the returned weights remain sufficiently separated for the next PF reconstruction. This proves Proposition 3. □

Proposition 3 controls mass with respect to the density passed to MEQ. The next result transfers this guarantee to the prescribed target density.

Proposition 7 (Target-mass error after MEQ)**.** *Let ρ satisfy the density conditions of Proposition 3 and let $W = \mathrm{MEQ}_N(\rho)$. Then*

$$\widehat{\mathcal{R}}_\vartheta(W) \le C\|\rho - \rho_\vartheta\|_\infty + CN^{-2/(M-1)}. \tag{37}$$

Consequently, if $\|\rho - \rho_\vartheta\|_\infty = \mathcal{O}(N^{-2/(M-1)})$, then

$$\widehat{\mathcal{R}}_\vartheta(W) = \mathcal{O}(N^{-2/(M-1)}) \tag{38}$$

for every $M \ge 2$.

Proof. By (33), every represented cell in $\mathcal{C}(W)$ has ρ-mass $1/|\mathcal{C}(W)|$ up to relative error $\mathcal{O}(N^{-2/(M-1)})$. The bi-Lipschitz transport bounds imply that every represented cell has volume at most $C/|\mathcal{C}(W)|$. Hence, for every represented cell τ,

$$|\mathcal{C}(W)| \left| \int_\tau \big(\rho_\vartheta(\mathbf{w}) - \rho(\mathbf{w})\big)\, \mathrm{d}\mathbf{w} \right| \le C\|\rho - \rho_\vartheta\|_\infty. \tag{39}$$

Adding the finite-MEQ error and maximizing over $\mathcal{C}(W)$ gives (37). □

The same target-mass control bounds the round-dependent reconstruction coefficient. Each represented cell and its corresponding straight simplex share the same transported vertices. Lemma 5 gives comparable diameters for the two cells. The bi-Lipschitz PF chart then gives

$$C_t \le C\left[1 + \left(1 + \widehat{\mathcal{R}}_\vartheta(W^{(t)})\right)^{2/(M-1)}\right]. \tag{40}$$

Indeed, the definition of $\widehat{\mathcal{R}}_\vartheta$ bounds the preference-weighted area of every represented PF patch by $(1+\widehat{\mathcal{R}}_\vartheta(W^{(t)}))/|\mathcal{C}(W^{(t)})|$ times the total preference mass. The positive lower bound on ϑ converts this to a PF-area bound. Shape regularity and the bi-Lipschitz PF chart convert PF area to patch diameter. Since $|\mathcal{C}(W^{(t)})| \asymp N$, (40) follows. The constant depends on the fixed bi-Lipschitz chart and preference bounds. The bound introduces no explicit multiplicative factor $\overline{J}/\underline{J}$. Thus, once target-mass imbalance is controlled, redistribution gives the sharper $C_t = \mathcal{O}(1)$ bound used in the convergence recursion, even when the initial uniform-weight bound is loose because κ_ϑ is large.

The main-paper residual $\mathcal{R}_\vartheta$ uses the straight cells in $\mathcal{T}(W)$. Comparing each straight cell with its corresponding complete transported cell gives

$$\left|\mathcal{R}_\vartheta(W) - \widehat{\mathcal{R}}_\vartheta(W)\right| \le \begin{cases} 0, & M = 2, \\ CN^{-1/(M-1)}, & M \ge 3. \end{cases} \tag{41}$$

This comparison is not used in the convergence result below; it only records the additional straight-cell representation error.

D Proofs for the Convergence Analysis

We now prove the convergence results in Section 4. The preceding section separates the two cell representations used by SURFACE. PF reconstruction uses the straight mesh $\mathcal{T}(W^{(t)})$, while target-mass matching uses the complete represented cells $\mathcal{C}(W^{(t)})$. We first consider exact scalarized solves. We then prove the inexact result by a joint induction that keeps solver accuracy, density regularity, and mesh regularity valid at every round. Throughout this section, $N^{1/(M-1)} = \Omega(\kappa_\vartheta^{3/2})$ is understood with a sufficiently large hidden constant to absorb the fixed resolution thresholds in the reconstruction and transport results above.

D.1 Exact Scalarized Solves

Theorem 2 (Exact-solve convergence of SURFACE)**.** *Suppose Assumptions 1–3 hold, let $N^{1/(M-1)} = \Omega(\kappa_\vartheta^{3/2})$, and suppose all scalarized solutions used to construct P_t are exact. Suppose*

the numerical MEQ implementation satisfies (33)*. Then, for all represented rounds,*

$$\|\rho_t - \rho_\vartheta\|_\infty \le (1-\alpha)^t \|\rho_0 - \rho_\vartheta\|_\infty + C\alpha N^{-2/(M-1)} \sum_{j=0}^{t-1} (1-\alpha)^{t-1-j} C_j$$
$$\le (1-\alpha)^t \|\rho_0 - \rho_\vartheta\|_\infty + C N^{-2/(M-1)}. \tag{42}$$

Moreover,

$$\widehat{\mathcal{R}}_\vartheta(W^{(t)}) \le C(1-\alpha)^t \|\rho_0 - \rho_\vartheta\|_\infty + C N^{-2/(M-1)}, \tag{43}$$
$$\|P_t - \mathbf{f}_{\mathrm{PF}}\|_{C^1} \le C C_t N^{-2/(M-1)}. \tag{44}$$

Proof. Lemma 6 shows that all exact rounds remain in the regular reconstruction and transport regime. Lemmas 3 and 4 give

$$\left\| \frac{\vartheta(P_t) J_{P_t}}{\int_{\Delta_M^\delta} \vartheta(P_t) J_{P_t}} - \rho_\vartheta \right\|_\infty \le C C_t N^{-2/(M-1)}. \tag{45}$$

Using (9),

$$\|\rho_{t+1} - \rho_\vartheta\|_\infty \le (1-\alpha)\|\rho_t - \rho_\vartheta\|_\infty + C\alpha C_t N^{-2/(M-1)}. \tag{46}$$

Iterating this recursion gives the first inequality in (42). Lemma 6 gives $C_j \le C$, so the geometrically weighted sum is uniformly bounded and yields the second inequality.

Keeping the sum explicit records the first-reconstruction effect. A large early C_j, whose bound may depend on κ_ϑ, is geometrically discounted. Equation (40) shows that later coefficients become $\mathcal{O}(1)$ once represented-cell target-mass imbalance is controlled, without an explicit factor $\overline{J}/\underline{J}$. Equation (44) follows directly from Lemma 3. Proposition 7 combined with (42) proves (43). □

Equation (41) separately relates the represented-cell residual in (43) to the straight-cell residual $\mathcal{R}_\vartheta$.

D.2 Inexact Scalarized Solves

Following the finite-step inner-solver analysis of SURF (Jiang et al., 2026), we characterize the work of InnerSolver by a K-step contraction factor η_K. When the solve at round $t+1$ is warm-started from the returned state at round t, we assume

$$\left\| \tilde{\mathbf{x}}_n^{(t+1)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t+1)}) \right\| \le \eta_K \left\| \tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t+1)}) \right\|, \quad n \in [N]. \tag{47}$$

Here η_K decreases as the inner budget K increases. This abstraction includes linearly convergent methods with geometric decay in K and sublinearly convergent methods after converting their optimization guarantee to decision error. In the hidden strongly convex weighted-sum specialization, projected gradient gives a standard realization (Jiang et al., 2026). On the regular TS solution mapping, Assumption 1 gives the local strong regularity and quadratic growth needed for a locally contractive inner method uniformly over Δ_M^δ (Bonnans & Shapiro, 2013).

The damping factor α controls how far the scalarized problems move between outer rounds and therefore trades outer progress against inner work. Compared with the one-dimensional SURF analysis, SURFACE additionally differentiates the reconstructed chart: the proof of Lemma 3 shows that decision error is amplified by $N^{1/(M-1)}$ in C^1, by $N^{2/(M-1)}$ in C^2, and by $N^{3/(M-1)}$ in the Lipschitz second derivative. Consequently, full tangent rank requires the C^1 perturbation to remain small, while the damped density update remains uniformly $C^{1,1}$ when $\alpha N^{3/(M-1)} \max_{n\in[N]} \|\tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t)})\| = \mathcal{O}(1)$. We therefore require the initial warmup accuracy

$$\max_{n\in[N]} \left\| \tilde{\mathbf{x}}_n^{(0)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(0)}) \right\| \le c_x \min\left\{ N^{-1/(M-1)},\, \alpha^{-1} N^{-3/(M-1)} \right\}, \tag{48}$$

and choose the inner budget so that

$$\eta_K = \mathcal{O}\left(\min\left\{ \left(1 + \alpha N^{1/(M-1)}\right)^{-1},\, \alpha^{-2} N^{-3/(M-1)} \right\} \right), \tag{49}$$

where $c_x > 0$ and the hidden constant in (49) are sufficiently small and independent of N, T, and α. The extremal references are outside the analyzed interior chart and do not enter these asymptotic accuracy conditions. Exact solves correspond to $\eta_K = 0$. For the simplified condition in Theorem 1, if $\bar{\alpha} = \Omega(N^{-\frac{3}{2(M-1)}})$ and $0 < \alpha \le \bar{\alpha}$, then $\bar{\alpha}^2\eta_K = \mathcal{O}(N^{-\frac{3}{M-1}})$ with a sufficiently small hidden constant implies (49).

The next lemma shows that (48)–(49) are preserved by the coupled reconstruction–redistribution–solver dynamics.

Lemma 7 (Invariant inexact trajectory)**.** *Suppose Assumptions 1–3 hold,* $0 < \alpha \le \bar{\alpha} < 1$*, and* (48)–(49) *hold. The constant* c_x *and the hidden constant in* (49) *can be chosen independently of* N*,* T*, and* α *such that, for all sufficiently large* N *and every represented round,*

$$\max_{n\in[N]} \left\|\tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t)})\right\| \le c_x \min\left\{N^{-1/(M-1)},\, \alpha^{-1}N^{-3/(M-1)}\right\}, \quad C_t \le C. \tag{50}$$

Moreover, P_t *has full tangent rank,* ρ_t *has common positive lower and upper bounds and a common* $C^{1,1}$ *bound, and* $W^{(t)}$ *satisfies Lemma 5.*

Proof. At $t = 0$, the uniform density generates the fixed reference mesh, so the round coefficient in (25) is bounded, while (48) gives the required warmup accuracy.

Suppose the conclusions hold at round t. Lemma 3 gives

$$\|P_t - \mathbf{f}_{\mathrm{PF}}\|_{C^1} \le CN^{-2/(M-1)} + CN^{1/(M-1)} \max_{n\in[N]} \left\|\tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t)})\right\|. \tag{51}$$

Choosing c_x sufficiently small and then N sufficiently large makes the right-hand side smaller than the uniform tangent-rank margin in Proposition 2, so P_t has full tangent rank. The scaled interpolation argument in the proof of Lemma 3 also gives

$$\|P_t\|_{C^{2,1}} \le C\left(1 + N^{3/(M-1)} \max_{n\in[N]} \left\|\tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t)})\right\|\right). \tag{52}$$

Applying the chain and product estimates in Lemma 4 shows that the normalized reconstructed density has common positive lower and upper bounds and $C^{1,1}$ norm controlled by the same right-hand side. Since (9) is damped,

$$\|\rho_{t+1}\|_{C^{1,1}} \le C + C\alpha N^{3/(M-1)} \max_{n\in[N]} \left\|\tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t)})\right\| \le C. \tag{53}$$

The same convex-combination argument preserves common positive lower and upper bounds. Proposition 3 and Lemma 5 then give a regular next mesh and $C_{t+1} \le C$.

It remains to propagate the solver error. Proposition 6 gives Lipschitz dependence of the transported vertices on the input density, hence

$$\max_{n\in[N]} \|\mathbf{w}_n^{(t+1)} - \mathbf{w}_n^{(t)}\| \le C\|\rho_{t+1} - \rho_t\|_\infty. \tag{54}$$

Using (9) and (32),

$$\|\rho_{t+1} - \rho_t\|_\infty \le C\alpha\left[\|\rho_t - \rho_\vartheta\|_\infty + C_t N^{-2/(M-1)} + N^{1/(M-1)} \max_{n\in[N]} \left\|\tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t)})\right\|\right]. \tag{55}$$

The density bounds already established imply $\|\rho_t - \rho_\vartheta\|_\infty \le C$. Proposition 2 makes $\mathbf{x}_{\mathrm{TS}}^\star$ uniformly Lipschitz, so (47) yields

$$\begin{aligned}&\max_{n\in[N]} \left\|\tilde{\mathbf{x}}_n^{(t+1)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t+1)})\right\| \\ \le&\eta_K\left[\left(1 + C\alpha N^{1/(M-1)}\right) \max_{n\in[N]} \left\|\tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t)})\right\| + C\alpha + C\alpha C_t N^{-2/(M-1)}\right].\end{aligned} \tag{56}$$

The first term is absorbed by the first requirement in (49). If $\alpha \le N^{-2/(M-1)}$, then $\eta_K\alpha \le C\alpha \le CN^{-1/(M-1)}$, while if $\alpha > N^{-2/(M-1)}$, the second requirement in (49) gives $\eta_K\alpha \le C\alpha^{-1}N^{-3/(M-1)}$. Taking the hidden constant in (49) sufficiently small therefore gives

$$\max_{n\in[N]} \left\|\tilde{\mathbf{x}}_n^{(t+1)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t+1)})\right\| \le c_x \min\left\{N^{-1/(M-1)},\, \alpha^{-1}N^{-3/(M-1)}\right\}, \tag{57}$$

which closes the joint induction. □

Lemma 7 therefore justifies the reconstruction and transport estimates below without assuming independently high-accuracy solves at every round. Smaller α permits larger tracking error during the transient because it reduces both the motion of the scalarized problems and the contribution of reconstruction roughness to the next density.

Lemma 8 (One-step coupled contraction)**.** *Under the conditions of Lemma 7, there exist constants $C_0, C > 0$, independent of N, T, and α, such that, for all sufficiently large N,*

$$\mathcal{E}_{t+1} \le (1-\alpha/4)\mathcal{E}_t + C\alpha C_t N^{-2/(M-1)}. \tag{58}$$

Proof. Equation (32) and the density update give

$$\|\rho_{t+1} - \rho_\vartheta\|_\infty \le (1-\alpha)\|\rho_t - \rho_\vartheta\|_\infty + C\alpha C_t N^{-2/(M-1)} \tag{59}$$

$$+ C\alpha N^{1/(M-1)} \max_{n\in[N]} \big\|\tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t)})\big\|. \tag{60}$$

Equations (54)–(55), the Lipschitz continuity of $\mathbf{x}_{\mathrm{TS}}^\star$, and (47) give

$$\max_{n\in[N]} \Big\|\tilde{\mathbf{x}}_n^{(t+1)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t+1)})\Big\| \le \eta_K\Big(\big(1 + C\alpha N^{1/(M-1)}\big) \max_{n\in[N]} \big\|\tilde{\mathbf{x}}_n^{(t)} - \mathbf{x}_{\mathrm{TS}}^\star(\mathbf{w}_n^{(t)})\big\| + C\alpha\|\rho_t - \rho_\vartheta\|_\infty + C\alpha C_t N^{-2/(M-1)}\Big). \tag{61}$$

Multiply (61) by $C_0(1 + \alpha N^{1/(M-1)})$ and add it to (60). Choose C_0 sufficiently large to absorb the direct solver-error term, then take the hidden constant in (49) sufficiently small so that $\eta_K(1 + \alpha N^{1/(M-1)})$ is below a fixed contraction factor $q_0 < 1 - \bar{\alpha}/4$. The cross terms are then absorbed into the two components of $\mathcal{E}_t$, the density component has coefficient at most $1 - \alpha/2$, and the solver component has coefficient at most $q_0 \le 1 - \alpha/4$. The remaining terms are bounded by $C\alpha C_t N^{-2/(M-1)}$, which proves (58). □

proof of Theorem 1. Lemma 7 gives a bounded C_t uniformly over the trajectory. Its initial bound may depend on κ_ϑ, while (40) gives $C_t = \mathcal{O}(1)$ once target-mass imbalance is controlled. Iterating (58) gives

$$\mathcal{E}_t \le (1-\alpha/4)^t\mathcal{E}_0 + C\alpha N^{-2/(M-1)} \sum_{j=0}^{t-1} (1-\alpha/4)^{t-1-j} \le (1-\alpha/4)^t\mathcal{E}_0 + CN^{-2/(M-1)},$$

which proves $\mathcal{E}_t \le (1-\alpha/4)^t\mathcal{E}_0 + CN^{-2/(M-1)}$. Proposition 7 then gives

$$\widehat{\mathcal{R}}_\vartheta(W^{(t)}) \le C\|\rho_t - \rho_\vartheta\|_\infty + CN^{-2/(M-1)} \le C\mathcal{E}_t + CN^{-2/(M-1)},$$

which proves $\widehat{\mathcal{R}}_\vartheta(W^{(t)}) \le C\mathcal{E}_t + CN^{-2/(M-1)}$. □

Theorem 1 makes the outer–inner tradeoff explicit. For fixed κ_ϑ, choose $N = \Theta(\epsilon^{-(M-1)/2})$ large enough to satisfy $N^{1/(M-1)} = \Omega(\kappa_\vartheta^{3/2})$, so that the finite-resolution floor satisfies $N^{-2/(M-1)} = \Theta(\epsilon)$. The three damping choices in Table 1 then become $\alpha = \Theta(1)$, $\Theta(\epsilon^{1/2})$, and $\Theta(\epsilon^{3/4})$, respectively. Since the outer contraction factor is $1 - \Theta(\alpha)$, reducing a constant-size transient to the $\mathcal{O}(\epsilon)$ floor requires $T = \mathcal{O}(\alpha^{-1}\log\epsilon^{-1})$ outer rounds.

For a linearly convergent InnerSolver with $\eta_K \le q^K$ for some $q \in (0,1)$, the first regime requires $\eta_K = \mathcal{O}(\epsilon^{3/2})$ and hence $K = \mathcal{O}(\log\epsilon^{-1})$, giving total repeated inner work $TK = \mathcal{O}(\log^2\epsilon^{-1})$. The intermediate regime requires $\eta_K = \mathcal{O}(\epsilon^{1/2})$, so $K = \mathcal{O}(\log\epsilon^{-1})$ and $TK = \mathcal{O}(\epsilon^{-1/2}\log^2\epsilon^{-1})$. The most strongly damped regime requires only a fixed contraction $\eta_K = \mathcal{O}(1)$, hence fixed K, but needs $T = \mathcal{O}(\epsilon^{-3/4}\log\epsilon^{-1})$ outer rounds. Thus, for linear or faster inner convergence, $\alpha = \Theta(1)$ is asymptotically preferable when additional inner iterations are affordable, while smaller α is useful when the inner budget per outer round is constrained.

For a sublinearly convergent inner method with $\eta_K = \mathcal{O}(K^{-p})$, the three regimes require $K = \Omega(\epsilon^{-3/(2p)})$, $\Omega(\epsilon^{-1/(2p)})$, and $\mathcal{O}(1)$ inner steps per round, respectively. Combining these requirements with the corresponding number of outer rounds gives total repeated inner work $\widetilde{\mathcal{O}}(\epsilon^{-3/(2p)})$,

$\widetilde{\mathcal{O}}(\epsilon^{-1/2-1/(2p)})$, and $\widetilde{\mathcal{O}}(\epsilon^{-3/4})$, where $\widetilde{\mathcal{O}}$ suppresses logarithmic factors. Hence $\alpha = \Theta(1)$ has the best leading ϵ-dependence when $p > 2$, all three regimes have the same leading $\epsilon^{-3/4}$ power when $p = 2$, and $\alpha = \Theta(\epsilon^{3/4}) = \Theta(N^{-3/(2(M-1))})$ is preferable when $p < 2$. The one-time warmup requirements in Table 1 do not change these leading regime comparisons.

This tradeoff follows the same warm-start principle as SURF (Jiang et al., 2026): damping reduces the motion of successive scalarized problems and therefore weakens the contraction required from each inner solve. SURFACE requires stronger damping in the fixed-budget regime because errors in the sampled Pareto points are amplified when differentiating the multidimensional PF reconstruction before forming the MEQ density.

Once the density transient reaches $\mathcal{O}(N^{-2/(M-1)}) = \mathcal{O}(\epsilon)$, (61) further gives a tracking floor of $\mathcal{O}(\eta_K \alpha N^{-2/(M-1)})$. Hence the scalarized problems need not be solved independently to $\mathcal{O}(N^{-3/(M-1)})$ at every outer round: α controls how far the next scalarized problem moves, while K controls how accurately the inner solver tracks its moving solution.

E Experiment Details

This appendix provides the complete experimental protocol, controlled diagnostics, benchmark results, and implementation details supporting Section 5.

E.1 Experimental Protocol and Benchmark Scope

Benchmarks. We evaluate SURFACE across $M \in \{2, 3, 4, 6\}$ objectives on synthetic, real-world, and multi-objective reinforcement-learning problems. The synthetic suite includes ZDT1/2/4 (Zitzler et al., 2000) and DTLZ2/3/4 (Deb et al., 2002). ZDT1/2/4 provide smooth connected bi-objective PFs on the represented interior branch, with ZDT2 changing the PF curvature and ZDT4 additionally introducing multimodality in the decision space. DTLZ2/3/4 share the same smooth full-dimensional spherical PF, while DTLZ3 introduces multimodality and DTLZ4 a biased decision-to-PF parameterization. Proposition 4 in Appendix B.3 verifies the regular full-ray TS conditions for DTLZ2. We further consider the real-world RE21 benchmark (Tanabe & Ishibuchi, 2020) using its empirical reference front, and multi-objective policy problems including Deep Sea Treasure, Resource Gathering, Fruit Tree, and Minecart from MO-Gymnasium (Felten et al., 2023). Together, these experiments span analytic PF geometries, multimodal decision landscapes, empirical real-world fronts, and learned policy spaces.

Inner solvers. The inner solver is chosen according to the underlying scalarized problem and is held fixed between SURFACE and Uniform within each comparison. For the ZDT/DTLZ benchmarks, both methods use the same scalarized solver and returned cardinality; the controlled DTLZ2 resolution studies use exact scalarized solutions. The quadratic inexact-solve experiment instead uses warm-started gradient descent, allowing the effect of controlled inner-solver error to be isolated. For Deep Sea Treasure, Resource Gathering, and Fruit Tree, each scalarized policy problem is solved by the same deterministic KL-regularized soft-value iteration with tolerance 10^{-11}. Minecart uses a common frozen GPI-LS model, with $2{,}016$ evaluation episodes per method. The end-to-end stochastic baselines are evaluated with five seeds at matched returned cardinality.

Metrics. We evaluate finite Pareto representations using complementary measures of distribution matching, representativeness, uniformity, and convergence. Our primary metric is the **finite target-mass residual** $\widehat{\mathcal{R}}_\vartheta(W)$ ($\downarrow$), which measures the maximum relative imbalance of independently evaluated target mass among the represented cells. For $\vartheta = 1$, it measures imbalance in PF surface-area allocation. The **IGD** ($\downarrow$) measures the average distance from $\mathcal{F}_{\mathrm{PF}}$ to the nearest returned objective, while the **fill distance** ($\downarrow$) measures the corresponding worst-case distance, i.e., the covering radius of the returned set (Ishibuchi et al., 2015, Zhang et al., 2024). We further report **hypervolume** HV ($\uparrow$) (Beume et al., 2007, Shang et al., 2020), computed with a fixed task-specific reference point, and its normalized **hypervolume gap** $\Delta\mathrm{HV}$ ($\downarrow$) for cross-task comparison. Together, $\widehat{\mathcal{R}}_\vartheta$ measures target matching, IGD and fill distance measure PF representativeness, and HV measures dominated objective-space coverage.

Baselines. Uniform weighting uses the same scalarized solver and returned cardinality as SURFACE, providing a direct comparison between fixed and geometry-adaptive scalarization weights. We further compare with several classes of finite-PF methods. *MOEA/D* (Zhang & Li, 2007) decomposes the MOO problem using a fixed set of weight vectors, while *MOEA/D-AWA* (Qi et al., 2014) adaptively removes weights in crowded regions and adds weights toward sparse PF regions. *DEA-GNG* (Liu et al., 2019) also adapts preference vectors, using a growing neural gas model to capture the evolving PF geometry. *NSGA-III* (Jain & Deb, 2013) maintains population diversity through reference-point association, whereas *SMS-EMOA* (Beume et al., 2007) uses hypervolume contribution for environmental selection. *LMPFE* (Tian et al., 2022) approximates the PF through multiple local models, and *Subset selection* (Shang et al., 2023) selects a fixed-cardinality representation from a larger candidate set. Finally, *UMOD* (Zhang et al., 2024) directly targets representativeness through a max-packing design, which serves as a surrogate for minimizing PF fill distance. All methods are evaluated at matched returned cardinality, and stochastic baselines use five seeds in the primary synthetic experiments.

E.2 Controlled Validation

Controlled PF distortion. We first use the three-objective radial problem in (14) to isolate the effect of scalarization-induced PF distortion. We vary $\kappa \in \{0, 0.5, 1, 1.5, 2, 2.5, 3, 4, 5\}$ and normalize the analytic density using four independently scrambled Sobol rules with $8{,}192$ points each. The mesh contains 66 returned vertices, including three pinned corners, and 100 cells. SURFACE uses $T = 5$ outer rounds and $\alpha = 0.2$. Across the sweep, the measured distortion ratio $\overline{J}/\underline{J}$ increases from 1 to 6.58×10^3. Uniform weighting reaches $\widehat{\mathcal{R}}_\vartheta = 16.866$, whereas SURFACE remains below 0.969. Supplying the analytic PF-area density to the same transport solver gives $\widehat{\mathcal{R}}_\vartheta = 0.276$, isolating the remaining numerical transport error.

Verification of the $\mathcal{O}(N^{-2/(M-1)})$ floor. Figures 5 corroborate the finite-resolution floor in Lemma 2 and Theorem 1 at $M = 3$ and $M = 6$, respectively. For DTLZ2 with $M = 3$, exact scalarized solves and one undamped update reduce the chart error $\|P_T - \mathbf{f}_{\mathrm{PF}}\|_{C^1}$, density error $\mathcal{E}_T$, and residual $\widehat{\mathcal{R}}_\vartheta(W_T)$ from $(0.2368, 0.3300, 0.04909)$ at $N = 12$ to $(0.02912, 0.01934, 0.00652)$ at $N = 150$, consistent with the predicted $\mathcal{O}(N^{-1})$ scale. The $M = 6$ results show the same decay with increasing N, with the chart-error slope close to the predicted $\mathcal{O}(N^{-2/5})$ scale. Separately, on the analytic quadratic control, exact solves give $\widehat{\mathcal{R}}_\vartheta = 0.00253$ versus 0.320 for Uniform, while 16 warm-started gradient steps reduce the decision error to 1.24×10^{-9} with $\widehat{\mathcal{R}}_\vartheta < 0.00254$, corroborating the inexact-solver term in the convergence analysis.

Finite MEQ validation. We validate the numerical MEQ step by independently evaluating transported cell masses. For $M = 3$, the selected configuration reduces the worst cell-mass defect from 0.0575 to 0.0395. For $M = 4$, across $N \in \{6, 16, 31, 52, 80\}$ free probes, all tested transports remain nondegenerate and the finest-resolution cell-mass defect is 0.0354, with minimum orientation ratio 0.535 and maximum condition number 7.07. The corresponding connected-cell PF-area residuals are 0.160, 0.155, 0.185, 0.207, and 0.207. These results support the finite-resolution MEQ accuracy and mesh regularity required by Proposition 3; implementation details are given in Appendix C.2.2.

Reconstruction accuracy. We test the reconstruction-to-density relation in Lemma 2 by replaying the saved $N \in \{16, 31, 52\}$ results without additional scalarized solves, where errors are measured by root-mean-square (RMS) values over evaluation points. At $N = 31$, radial reconstruction reduces the PF reconstruction error from 0.0646 to 0.0267 and the density error from 0.671 to 0.127; at $N = 52$, it reduces the density error from 0.377 to 0.0942 and the independently evaluated cell-mass defect from 0.0479 to 0.0196, corroborating that improved PF reconstruction yields a more accurate target-density estimate. A fixed piecewise-linear density approximation introduces additional interpolation error at finer resolutions, so we evaluate the reconstructed density directly.

E.3 Details of ZDT and DTLZ Benchmarks Results

We evaluate SURFACE on ZDT1/2/4 and DTLZ2/3/4 across the 11 $M = 2, 3, 4$ tasks, using matched returned cardinality $N + M$ for all methods. SURFACE and Uniform use the same scalarized solver, so their comparison isolates the effect of geometry-aware weight redistribution, while the remaining baselines are evaluated end-to-end. Stochastic baselines are run with five seeds. Table 6 reports the complete task-level results, Table 5 summarizes the mean ranks, and Figure 6 visualizes the tradeoff between uniform PF-area coverage and hypervolume performance.

Method	$\widehat{\mathcal{R}}_\vartheta$	IGD	Fill	ΔHV	Coverage
SURFACE	**1.00**	**1.00**	**1.56**	**2.22**	**1.19**
Uniform	3.22	2.11	2.78	4.56	2.70
NSGA-III	5.44	5.11	4.67	5.67	5.07
LMPFE	5.22	5.22	5.00	6.22	5.15
SMS-EMOA	4.22	6.67	6.78	2.78	5.89
MOEA/D-AWA	6.67	6.11	5.78	6.67	6.19
Subset	6.22	6.44	7.11	3.22	6.59
MOEA/D	7.72	7.22	6.56	7.56	7.17
DEA-GNG	6.89	7.11	7.67	8.00	7.22
UMOD	8.39	8.00	7.11	8.11	7.83

Table 5: Mean ranks over nine ZDT/DTLZ task–dimension groups. Coverage averages the first three PF-representation metrics and excludes hypervolume. Lower is better.

Across the experiments, SURFACE consistently provides the strongest approximation to the uniform PF-area target. Relative to Uniform, it improves the reported PF-area imbalance, IGD, fill distance, and hypervolume gap throughout the primary comparisons, indicating that adapting the scalarization weights to the reconstructed PF geometry improves the finite representation even when the underlying optimizer is unchanged. SURFACE attains the best mean rank in PF-area imbalance and IGD, remains strongest overall in fill distance, and consequently achieves the best aggregate PF-coverage rank in Table 5. The advantage is therefore not tied to a single notion of uniformity: it appears both in how evenly PF area is allocated and in the average and worst-case distances from the represented set to the PF.

The improved coverage also does not require sacrificing objective-space quality. Some evolutionary methods attain better hypervolume on individual problems, reflecting the fact that hypervolume and uniform PF representation emphasize different finite designs. Nevertheless, SURFACE remains competitive on hypervolume across the suite and achieves the best mean ΔHV rank over the nine primary groups. This behavior is summarized in Figure 6. SURFACE lies in the favorable lower-left region, combining the strongest PF-area coverage rank with a strong hypervolume rank, whereas methods that perform well on hypervolume for selected tasks generally exhibit substantially less uniform PF representations.

E.4 Multi-Objective Policy Optimization

We evaluate SURFACE on Deep Sea Treasure, Resource Gathering, and Fruit Tree from MO-Gymnasium (Felten et al., 2023). For each preference $\mathbf{w}$, we optimize

$$F_{\mathbf{w}}(\pi) = -\sum_{m=1}^{M} w_m J_m(\pi) + \beta R_{\mathrm{KL}}(\pi). \tag{62}$$

We use discount factor 0.999, with $\beta = 1.5$ for Deep Sea Treasure and $\beta = 0.5$ for Resource Gathering and Fruit Tree. Soft value iteration terminates at tolerance 10^{-11}. SURFACE uses $\alpha = 0.3$, slot-matched warm starts, cached pinned corners, and a fresh terminal solve. Deep Sea Treasure and Resource Gathering use 30 outer rounds, while Fruit Tree uses ten; the returned cardinalities are 16, 21, and 56, respectively. SURFACE and Uniform use the same deterministic soft-value solver, while the remaining baselines use a five-seed evolutionary-logit policy optimizer. Table 3 reports the resulting comparisons.

We further study the effect of returned-set resolution on the $M = 6$ Fruit Tree task (Felten et al., 2023). Using differentiated policy sensitivities and a common independently sampled PF-area reference, balanced equal-capacity transport achieves decreasing target-mass residuals $\widehat{\mathcal{R}}_\vartheta = 0.2523, 0.1955, 0.1694$ as the returned cardinality increases from $N = 126$ to 252 and 462. These results show that SURFACE continues to improve with finer returned-set resolution in a higher-dimensional objective space.

We further demonstrate the applicability of SURFACE with a different policy-optimization backend on Minecart, the continuous-state, discrete-action benchmark from MO-Gymnasium (Felten et al., 2023), originally introduced by Abels et al. (2019). We use a common frozen GPI-LS policy set (Alegre et al., 2022) for both methods, with $2{,}016$ evaluation episodes per method at temperature

Metric	SMS-EMOA	NSGA-III	LMPFE	Subset	MOEA/D-AWA	MOEA/D	DEA-GNG	UMOD	Uniform	SURFACE
ZDT1, $M = 2$, $N = 9$										
$\widehat{\mathcal{R}}_\vartheta \downarrow$	.431 ± .015	.608 ± .000	.563 ± .060	.769 ± .001	.434 ± .076	.897 ± .239	.516 ± .122	.483 ± .065	.505	**.316**
IGD ↓	.037 ± .000	.042 ± .000	.039 ± .003	.040 ± .000	.039 ± .000	.049 ± .005	.041 ± .002	.037 ± .000	.037	**.035**
Fill↓	.092 ± .002	.150 ± .000	.095 ± .016	.174 ± .000	.105 ± .003	.189 ± .032	.119 ± .009	**.085 ± .008**	.106	.096
ΔHV ↓	.048 ± .000	.055 ± .000	.053 ± .004	.049 ± .000	.052 ± .000	.060 ± .004	.055 ± .003	.049 ± .001	.048	**.045**
Coverage↓	2.67 ± .38	8.33 ± .38	5.33 ± 1.69	8.33 ± .53	4.33 ± .62	10.00 ± .73	7.00 ± 1.22	2.67 ± .96	4.33	**2.00**
ZDT2, $M = 2$, $N = 9$										
$\widehat{\mathcal{R}}_\vartheta \downarrow$	.480 ± .018	.589 ± .000	.881 ± .374	.491 ± .000	.580 ± .044	.917 ± .120	.599 ± .070	.812 ± .047	.514	**.397**
IGD ↓	.039 ± .000	.038 ± .000	.048 ± .013	.039 ± .000	.038 ± .000	.040 ± .000	.039 ± .001	.041 ± .000	.037	**.036**
Fill↓	.126 ± .003	.089 ± .000	.156 ± .068	.122 ± .000	.088 ± .000	.102 ± .012	.106 ± .011	.101 ± .003	.087	**.073**
ΔHV ↓	**.069 ± .000**	.076 ± .000	.097 ± .025	.075 ± .000	.076 ± .001	.095 ± .003	.076 ± .002	.084 ± .001	.078	.074
Coverage↓	5.67 ± .55	4.33 ± .43	9.67 ± 2.83	5.33 ± .61	4.00 ± .43	8.00 ± 1.09	7.00 ± 1.51	7.33 ± .33	2.67	**1.00**
ZDT4, $M = 2$, $N = 9$										
$\widehat{\mathcal{R}}_\vartheta \downarrow$	.418 ± .006	.611 ± .003	.572 ± .033	.760 ± .005	.612 ± .167	1.618 ± .044	.862 ± .739	2.124 ± .972	.505	**.316**
IGD ↓	.037 ± .000	.043 ± .000	.038 ± .000	.040 ± .000	.040 ± .001	.056 ± .000	.049 ± .012	.108 ± .060	.037	**.035**
Fill↓	.090 ± .001	.150 ± .000	**.088 ± .004**	.173 ± .001	.111 ± .006	.281 ± .007	.178 ± .103	.324 ± .155	.106	.096
ΔHV ↓	.049 ± .000	.056 ± .001	.053 ± .002	.050 ± .000	.055 ± .003	.075 ± .007	.065 ± .012	.196 ± .122	.048	**.045**
Coverage↓	2.33 ± .15	6.00 ± .49	3.00 ± .51	6.67 ± .60	5.33 ± .90	9.00 ± .69	8.00 ± 1.59	10.00 ± .33	3.00	**1.67**
DTLZ2, $M = 2$, $N = 9$										
$\widehat{\mathcal{R}}_\vartheta \downarrow$	.503 ± .008	.613 ± .000	.587 ± .065	.623 ± .001	.673 ± .078	.613 ± .000	.691 ± .068	.810 ± .092	.543	**.402**
IGD ↓	.045 ± .000	.041 ± .000	.041 ± .001	.048 ± .000	.041 ± .000	.041 ± .000	.043 ± .003	.043 ± .002	.040	**.038**
Fill↓	.141 ± .001	.099 ± .000	.106 ± .012	.148 ± .000	.099 ± .000	.099 ± .000	.121 ± .021	.102 ± .017	.097	**.078**
ΔHV ↓	**.061 ± .000**	.082 ± .000	.079 ± .004	.063 ± .000	.081 ± .001	.082 ± .000	.085 ± .011	.081 ± .003	.085	.074
Coverage↓	6.67 ± .30	4.00 ± .76	4.67 ± 1.96	9.00 ± .28	6.33 ± .68	5.00 ± .71	8.33 ± 1.12	7.67 ± 1.24	2.33	**1.00**
DTLZ2, $M = 3$, $N = 18$										
$\widehat{\mathcal{R}}_\vartheta \downarrow$	1.312 ± .038	1.178 ± .001	.913 ± .086	1.588 ± .001	4.469 ± 1.499	1.180 ± .000	.896 ± .067	1.457 ± .174	1.078	**.693**
IGD ↓	.164 ± .001	.137 ± .000	.134 ± .003	.162 ± .000	.334 ± .117	.137 ± .000	.132 ± .003	.142 ± .004	.132	**.129**
Fill↓	.334 ± .008	.272 ± .000	.296 ± .027	.333 ± .000	.796 ± .306	.273 ± .000	.319 ± .015	.309 ± .032	.265	**.227**
ΔHV ↓	**.133 ± .001**	.162 ± .000	.169 ± .006	.138 ± .000	.475 ± .093	.162 ± .000	.200 ± .015	.176 ± .003	.175	.162
Coverage↓	8.33 ± .53	4.33 ± .28	4.00 ± 1.01	8.33 ± .37	10.00 ± .00	5.33 ± .24	3.67 ± .71	7.00 ± .87	3.00	**1.00**
DTLZ2 , $M = 4$, $N = 31$										
$\widehat{\mathcal{R}}_\vartheta \downarrow$	5.547 ± 2.009	3.436 ± .001	**1.767 ± .327**	3.771 ± .002	9.975 ± 3.034	3.434 ± .002	1.827 ± .478	4.523 ± .509	3.042	2.498
IGD ↓	.348 ± .043	.228 ± .000	.216 ± .003	.246 ± .000	.510 ± .091	.228 ± .000	**.208 ± .014**	.245 ± .006	.218	.216
Fill↓	.723 ± .091	.391 ± .000	.420 ± .021	.467 ± .000	.975 ± .225	.391 ± .000	.447 ± .023	.467 ± .035	**.380**	.381
ΔHV ↓	.493 ± .050	.183 ± .000	.196 ± .004	**.166 ± .000**	.581 ± .069	.183 ± .000	.266 ± .008	.209 ± .005	.208	.196
Coverage↓	9.00 ± .94	5.00 ± .25	**2.67 ± .61**	7.67 ± .41	10.00 ± .15	4.67 ± .48	3.00 ± 1.26	7.33 ± .41	3.00	**2.67**
DTLZ3, $M = 2$, $N = 9$										
$\widehat{\mathcal{R}}_\vartheta \downarrow$	.509 ± .017	.616 ± .006	.492 ± .016	.624 ± .001	.652 ± .042	.784 ± .200	3.273 ± 3.130	1.471 ± .399	.543	**.402**
IGD ↓	.046 ± .000	.041 ± .000	.041 ± .001	.048 ± .001	.043 ± .002	.047 ± .005	.189 ± .185	.085 ± .022	.040	**.038**
Fill↓	.143 ± .001	.099 ± .000	.096 ± .006	.148 ± .001	.100 ± .002	.107 ± .010	.525 ± .444	.233 ± .052	.097	**.078**
ΔHV ↓	**.070 ± .006**	.088 ± .003	.092 ± .007	.079 ± .013	.106 ± .025	.151 ± .040	.307 ± .231	.289 ± .113	.085	.074
Coverage↓	5.33 ± .43	4.33 ± .40	2.33 ± .78	7.33 ± .77	5.67 ± .88	7.00 ± 1.09	10.00 ± 2.37	9.00 ± .55	3.00	**1.00**
DTLZ3, $M = 3$, $N = 18$										
$\widehat{\mathcal{R}}_\vartheta \downarrow$	1.490 ± .402	1.179 ± .001	2.546 ± 1.447	1.588 ± .001	20.000 ± .000	1.180 ± .002	4.833 ± 1.872	3.120 ± .202	1.078	**.693**
IGD ↓	.166 ± .005	.137 ± .000	.210 ± .062	.162 ± .000	143.225 ± 159.716	.138 ± .000	.359 ± .123	.204 ± .022	.132	**.129**
Fill↓	.342 ± .020	.272 ± .000	.479 ± .137	.333 ± .000	143.717 ± 159.735	.273 ± .001	.828 ± .284	.476 ± .050	.265	**.227**
ΔHV ↓	**.134 ± .001**	.162 ± .000	.258 ± .071	.139 ± .002	.939 ± .083	.172 ± .010	.432 ± .129	.266 ± .021	.175	.162
Coverage↓	5.67 ± .37	3.00 ± .15	7.67 ± 1.10	5.33 ± .37	10.00 ± .00	4.00 ± .15	9.00 ± .28	7.33 ± .55	2.00	**1.00**
DTLZ3, $M = 4$, $N = 31$										
$\widehat{\mathcal{R}}_\vartheta \downarrow$	6.121 ± 2.177	3.436 ± .001	3.977 ± 2.356	3.781 ± .025	34.00 ± .000	3.436 ± .004	7.726 ± 2.906	4.766 ± 1.414	3.042	**2.498**
IGD ↓	.439 ± .079	.228 ± .000	.294 ± .059	.247 ± .001	73.139 ± 161.229	.228 ± .000	.456 ± .089	.309 ± .032	.218	**.216**
Fill↓	.967 ± .186	.391 ± .000	.616 ± .059	.467 ± .000	73.537 ± 161.271	.391 ± .001	.972 ± .244	.653 ± .085	**.380**	.381
ΔHV ↓	.529 ± .044	.184 ± .000	.267 ± .061	**.166 ± .000**	.908 ± .051	.184 ± .000	.498 ± .127	.267 ± .030	.208	.196
Coverage↓	8.00 ± .62	3.17 ± .47	6.00 ± 1.01	5.00 ± .33	10.00 ± .00	3.83 ± .30	9.00 ± 1.64	7.00 ± .64	1.67	**1.33**
DTLZ4, $M = 2$, $N = 9$										
$\widehat{\mathcal{R}}_\vartheta \downarrow$	8.101 ± 4.247	4.368 ± 5.142	8.112 ± 4.222	6.249 ± 5.136	8.140 ± 4.159	10.000 ± .000	8.086 ± 4.279	10.000 ± .000	.543	**.402**
IGD ↓	.606 ± .313	.323 ± .386	.605 ± .316	.467 ± .382	.605 ± .315	.746 ± .000	.605 ± .315	.746 ± .000	.040	**.038**
Fill↓	1.160 ± .569	.625 ± .721	1.149 ± .592	.908 ± .694	1.151 ± .588	1.414 ± .000	1.152 ± .586	1.414 ± .000	.097	**.078**
ΔHV ↓	.602 ± .302	.344 ± .359	.604 ± .296	.468 ± .368	.606 ± .293	.737 ± .000	.604 ± .297	.737 ± .000	.085	**.074**
Coverage↓	7.33 ± 1.84	3.00 ± 1.45	5.67 ± 2.65	4.00 ± 1.47	7.00 ± 1.75	9.17 ± 1.22	6.00 ± 2.52	9.83 ± .52	2.00	**1.00**
DTLZ4, $M = 3$, $N = 18$										
$\widehat{\mathcal{R}}_\vartheta \downarrow$	20.000 ± .000	20.000 ± .000	4.698 ± 8.556	8.954 ± 10.083	1.253 ± .549	16.236 ± 8.417	12.282 ± 10.568	20.000 ± .000	1.078	**.693**
IGD ↓	.943 ± .000	.943 ± .000	.296 ± .361	.475 ± .427	.150 ± .019	.782 ± .360	.615 ± .449	.943 ± .000	.132	**.129**
Fill↓	1.414 ± .000	1.414 ± .000	.523 ± .498	.765 ± .592	.345 ± .048	1.186 ± .511	.959 ± .623	1.414 ± .000	.265	**.227**
ΔHV ↓	.848 ± .000	.848 ± .000	.307 ± .302	.422 ± .389	.187 ± .026	.711 ± .307	.586 ± .358	.848 ± .000	.175	**.162**
Coverage↓	9.00 ± 2.11	8.33 ± 1.03	4.00 ± 1.80	5.00 ± 1.09	3.00 ± 1.32	7.00 ± 1.98	6.00 ± 3.19	9.67 ± .63	2.00	**1.00**

Table 6: Complete results for SURFACE and 9 baselines. Per-task entries are mean ± sample standard deviation over five seeds. Uniform and SURFACE are deterministic and therefore omit standard-deviation terms. Coverage row reports the mean rank over $\widehat{\mathcal{R}}_\vartheta$, IGD, and fill distance.

0.03. Relative to Uniform, SURFACE improves return-space coverage from 3.55 to 2.50, IGD from 0.429 to 0.421, and hypervolume from 507.7 to 516.0.