algorithms

MDPI

Article

Certified Adaptive Triangulation Sampling for Deterministic Pareto-Surface Reconstruction

Massimiliano Caramia

Dipartimento di Ingegneria dell'Impresa, University of Rome Tor Vergata, Via del Politecnico, 1, 00133 Rome, Italy; caramia@dii.uniroma2.it

Abstract

Many deterministic multi-objective optimization methods generate Pareto outcomes by repeatedly solving scalarized subproblems for different preference or reference vectors. When the number of objectives is $m \geq 3$, the resulting samples lie on an $(m-1)$-dimensional Pareto surface in objective space. For tasks such as visualization, trade-off exploration, interactive decision making, and sensitivity analysis, a finite cloud of non-dominated points may be insufficient; one often needs a continuous surrogate of the Pareto surface together with a quantitative control of its reconstruction error. This paper studies the corresponding outer-loop reconstruction problem: how should new reference vectors be selected so as to reconstruct the Pareto surface to a prescribed uniform accuracy while using as few scalarized solves as possible? We propose Certified Adaptive Triangulation Sampling (CATS), a curvature-aware adaptive triangulation method for reconstructing a Pareto surface from an oracle $u \mapsto z(u)$, $u \in \Delta_d$, where $d = m-1$. CATS builds a simplicial mesh over the reference simplex and refines the cell with the largest local interpolation quantity $\eta(\tau) = \frac{1}{2}\big(\max_k M_{\tau,k}\big)\operatorname{diam}(\tau)^2$, where $M_{\tau,k}$ is an upper bound on the Hessian norm of the kth component of the oracle-induced map over τ. This quantity matches the natural error scale of affine interpolation for C^2 maps. The rigorous certified interpretation of CATS applies when the preference-to-Pareto map is single-valued, C^2, and equipped with reliable local Hessian-norm upper bounds. If such bounds are replaced by numerical curvature estimates, the same rule can still be used as an adaptive refinement indicator, but the resulting stopping test is not a formal certificate unless those estimates are themselves validated. Under the certified assumptions, we prove that the stopping condition $\max_\tau \eta(\tau) \leq \varepsilon$ guarantees $\sup_{u\in\Delta_d} \|z(u) - \hat{z}(u)\|_\infty \leq \varepsilon$, and that the oracle complexity of certified simplicial piecewise-affine reconstruction is $\Theta(\varepsilon^{-d/2})$. On the rigorously certified core tests, CATS uses 2.7×–3.8× fewer oracle calls than uniform reference-direction sampling and 1.2×–1.6× fewer than an AWS-inspired patch-area refinement rule. Additional benchmark studies, evaluated with the same interpolation quantity as a practical stopping indicator, show the same qualitative advantage, especially on anisotropic and localized surface geometries.

Keywords: multi-objective optimization; pareto surface; scalarization; adaptive sampling; surface reconstruction; interpolation certificate

Check for updates

Academic Editor: Binhai Zhu

Received: 30 May 2026
Revised: 4 June 2026
Accepted: 9 June 2026
Published: 11 June 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Multi-objective optimization arises naturally in engineering design, industrial systems, robotics, machine learning, and decision-support problems whenever several performance criteria must be optimized simultaneously. In many such applications, the objectives are

genuinely conflicting: improving one performance index may deteriorate another. Examples include reducing energy consumption while preserving productivity in robotic systems, increasing mechanical efficiency while limiting wear and friction in transmission systems, or improving predictive performance while controlling communication and computational complexity in data-driven models. In these settings, the decision-maker is rarely only interested in a single solution; rather, one often needs to understand the structure of the trade-off set in order to identify robust compromises, sensitive regions, and alternative designs.

This need is particularly visible in mechanical and tribological design. In worm gear units and related transmission systems, efficiency, power losses, vibration, friction, wear, and service life are coupled quantities that may lead to competing design objectives. Recent investigations on carbon-nanotube additives in lubricants for worm gear systems show that lubricant formulation can significantly influence efficiency and tribological behavior [1]. Similar motivations occur in robotic and data-driven engineering applications: in trajectory planning for six-axis manipulators, path quality, energy consumption, execution time, and smoothness may be optimized jointly [2]; in machine-learning applications, multi-objective neural architecture search has been used to balance accuracy, communication cost, efficiency, and model complexity, as in centralized federated deep fuzzy neural-network models for epistatic detection [3]. These examples motivate methods that do not merely return isolated non-dominated points, but also help reconstruct the geometry of the trade-off surface for visualization, sensitivity analysis, and decision support.

Formally, a deterministic multi-objective optimization problem can be written as

$$\min_{x \in \mathcal{X}} \; F(x) := (f_1(x), \ldots, f_m(x)) \in \mathbb{R}^m, \tag{1}$$

where $\mathcal{X} \subseteq \mathbb{R}^n$ denotes the feasible set, $F : \mathcal{X} \to \mathbb{R}^m$ is the vector-valued objective map, and $f_i : \mathcal{X} \to \mathbb{R}$ is the ith objective function, for $i = 1, \ldots, m$; see, e.g., recent surveys and reviews [4–7]. Since the objectives are typically conflicting, one does not seek a single global minimizer of all components simultaneously. Instead, one seeks Pareto-optimal solutions, i.e., points $x^\star \in \mathcal{X}$ such that there is no feasible $x \in \mathcal{X}$ that is no worse in all objectives and strictly better in at least one. Standard references on Pareto optimality and scalarization theory include [8,9]. The image of Pareto-optimal points in objective space is called the Pareto front for $m = 2$ and the Pareto surface for $m \geq 3$.

A large part of the deterministic multi-objective optimization literature focuses on generating finite sets of non-dominated points. This is useful and often sufficient when the goal is to provide representative trade-off alternatives. However, for downstream tasks such as interactive decision-making, visualization, sensitivity analysis, and post-optimal trade-off exploration, a finite point cloud may be insufficient. In these cases, one would like to reconstruct a continuous surrogate of the Pareto front or surface, ideally together with a quantitative accuracy guarantee. This leads to the outer-loop question studied in this paper: how should new preference vectors be selected so as to reconstruct the entire Pareto surface to a prescribed uniform accuracy using as few scalarized solves as possible?

Many deterministic multi-objective optimization methods can be viewed as a two-level procedure. The outer loop selects a preference vector, weight vector, or reference direction; the inner loop solves a scalarized optimization problem and returns a Pareto outcome. Thus, for the purpose of studying the sampling and reconstruction problem, we abstract the inner solve as an oracle

$$\mathcal{O} : \; u \in \Delta_d \mapsto z(u) \in \mathbb{R}^m, \tag{2}$$

where $d = m - 1$, Δ_d is the reference simplex, and $z(u)$ is the Pareto outcome returned for the preference vector u. This abstraction is deliberately independent of the particular

scalarization used inside the oracle: the inner solver may be based on any deterministic scalarized formulation. The focus of this paper is not the design of a new scalarization, nor the improvement of an individual Pareto point returned by the inner solver. Rather, the focus is the outer-loop placement of preference vectors for reconstructing the full Pareto surface.

We assume, in the certified part of our analysis, that the oracle induces a single-valued smooth map

$$z : \Delta_d \to \mathbb{R}^m, \qquad u \mapsto z(u), \tag{3}$$

with sufficient regularity, namely $z \in C^2(\Delta_d; \mathbb{R}^m)$. This smooth single-valued setting is the natural regime in which a parametric Pareto surface and a classical interpolation certificate are well defined. Situations involving discontinuities, branch switching, bifurcations, or genuinely set-valued Pareto correspondences require a different mathematical framework and are outside the certified scope of the present paper.

Given finitely many oracle evaluations, we build a conforming triangulation $\mathcal{T}$ of the reference simplex Δ_d and construct the piecewise-affine interpolant $\hat{z}$ that matches z at all vertices of $\mathcal{T}$. The reconstruction error is measured by the uniform norm

$$\|z - \hat{z}\|_\infty := \sup_{u \in \Delta_d} \|z(u) - \hat{z}(u)\|_\infty. \tag{4}$$

The outer-loop task is therefore to choose new vertices of $\mathcal{T}$, i.e., new preference vectors, so that

$$\|z - \hat{z}\|_\infty \le \varepsilon \tag{5}$$

with as few oracle calls as possible.

The proposed method, Certified Adaptive Triangulation Sampling (CATS), is based on one guiding principle: if z is C^2, then the local error of affine interpolation on a simplex τ scales as (local curvature bound on τ) $\times$ $\mathrm{diam}(\tau)^2$. Accordingly, CATS assigns to each simplex $\tau \in \mathcal{T}$ the local quantity

$$\eta(\tau) = \frac{1}{2}\left(\max_k M_{\tau,k}\right)\mathrm{diam}(\tau)^2, \tag{6}$$

where $M_{\tau,k}$ is an upper bound on the spectral norm of the Hessian of the kth component of z over τ. The algorithm repeatedly refines the simplex with the largest value of $\eta(\tau)$ and stops when $\max_{\tau \in \mathcal{T}} \eta(\tau) \le \varepsilon$. Under the stated smoothness and curvature-bound assumptions, this stopping rule implies the global guarantee in (5).

The word "certified" is used in this paper in a conditional sense. The certificate is rigorous only if the quantities $M_{\tau,k}$ are valid upper bounds on the Hessian norms of the oracle-induced map over the corresponding simplex. In a general MOO pipeline, the map $u \mapsto z(u)$ is defined implicitly by a scalarized optimization problem, and the solver may not provide second-order sensitivity information with respect to the preference vector. Therefore, CATS should not be interpreted as a universally plug-in certified outer loop for arbitrary deterministic MOO solvers. Its certified version applies when reliable curvature bounds are available, for example from analytic expressions, validated sensitivity analysis, interval bounds, or conservative model-based estimates. When such bounds are replaced by numerical estimates, the same rule remains useful as an adaptive refinement indicator, but the stopping test is no longer a formal certificate unless those estimates are themselves validated.

This perspective also clarifies the difference between CATS and common sampling practices. Uniform simplex-lattice sampling (see, e.g., [10,11]) is simple and widely used, but it can oversample nearly flat regions while undersampling localized high-curvature

regions. This effect becomes more severe when $m \geq 3$, since the Pareto surface has the dimensions of $d = m - 1$ and the number of uniform samples grows rapidly with the desired resolution. Adaptive methods inspired by patch size in objective space, such as AWS-type methods, improve over purely uniform placement, but patch area is only an indirect proxy for interpolation error: a large patch can be nearly affine, whereas a small patch can be highly curved. Since uniform reconstruction error is controlled by the worst-resolved region, CATS uses a direct interpolation certificate rather than patch size alone.

The guarantee considered in this paper should also be distinguished from other notions of approximation accuracy used in multi-objective optimization. A finite non-dominated point set may provide useful trade-off alternatives without certifying a continuous reconstructed surface. Dominance-based ε-Pareto sets guarantee coverage of the feasible image in an approximate-dominance sense, but do not directly provide a parametric interpolation bound. Convex sandwiching and Benson-type methods certify gaps between inner and outer approximations of the upper image, but they address a different set-approximation objective. CATS instead targets the uniform reconstruction of a smooth parametric Pareto surface over a reference simplex.

The specific contributions of this paper are as follows:

1. A certified outer-loop algorithm for Pareto-surface reconstruction. We propose CATS, an adaptive triangulation method that selects new preference vectors by refining the simplex with the largest Hessian-based interpolation certificate. The method acts only at the outer-loop level: it decides where to query the scalarized solver next, not how the inner scalarized optimization problem is solved.
2. A global uniform error guarantee. We prove that, when valid local curvature bounds are available, the stopping rule $\max_{\tau \in \mathcal{T}} \eta(\tau) \leq \varepsilon$ implies the global reconstruction bound $\|z - \hat{z}\|_\infty \leq \varepsilon$.
3. Rate-optimal oracle complexity. Let $d = m - 1$ be the dimension of the reconstructed Pareto surface. Under standard smoothness and mesh-regularity assumptions, we prove an upper bound $O(\varepsilon^{-d/2})$ on the number of oracle calls required to reach tolerance ε. We also prove a matching minimax lower bound $\Omega(\varepsilon^{-d/2})$ for certified simplicial piecewise-affine reconstructions based on point queries, showing that the exponent is information-theoretically tight.
4. Interpretability and empirical gains. On rigorously certified core tests, CATS uses substantially fewer oracle calls than uniform reference-direction sampling and fewer calls than an AWS-inspired patch-area refinement rule. Additional smooth parametric surface-reconstruction benchmarks and compact DTLZ-style tri-objective tests, evaluated with the same interpolation quantity as a practical stopping indicator, show the same qualitative trend, especially on anisotropic and localized geometries.

Overall, the significance of the work is that it isolates a reconstruction problem that is often implicit in deterministic multi-objective optimization workflows. Existing approaches typically emphasize scalarization design, point generation, or set/image approximation. CATS instead addresses the complementary question of how to place preference vectors when the desired output is a continuous triangulated Pareto surface with a prescribed uniform reconstruction accuracy. This makes the method especially relevant when each oracle call corresponds to an expensive deterministic optimization, simulation, or engineering design evaluation.

The rest of the paper is organized as follows. Section 2 positions CATS among point-generation, triangulated-surrogate, and sandwiching approaches, and clarifies which accuracy notions they certify. Section 3 formalizes the oracle model and the reconstruction objective. Section 4 derives the curvature-based certificate and explains why it is the appro-

priate adaptivity signal for uniform interpolation error. Section 5 presents the algorithm and the baselines. Section 6 proves upper and lower oracle-complexity bounds. Finally, Section 7 provides visual and quantitative reconstructions that connect oracle savings to curvature localization.

2. Related Work

This section positions CATS relative to three established strands of deterministic work: (i) scalarization-based point generation in MOO, (ii) geometric surrogates built from sampled Pareto points, and (iii) certified set/image approximations (dominance-based or convex-analytic). The key message is that CATS is an outer-loop method for certified parametric surface reconstruction—it is not a new scalarization, and it does not target dominance-only or polyhedral-gap notions of accuracy.

General background information on deterministic MOO, Pareto optimality, and scalarization theory can be found in standard references such as Ehrgott [8] and Miettinen [9]. Scalarizations convert the vector problem (1) into a parameterized family of scalar problems whose solutions (under suitable conditions) are Pareto optimal. Classical examples include weighted sums, Pascoletti–Serafini scalarization, ϵ-constraint schemes, and many variants. Our contribution is orthogonal to scalarization choice: we treat the inner scalarized solver as a black-box oracle that maps a preference vector to a Pareto outcome.

2.1. Point-Generation Outer Loops

A common deterministic workflow is: (i) select reference directions u (often on a simplex), (ii) solve one scalarized problem per u, and (iii) return the resulting non-dominated points. Two influential methods in this paradigm are:

- Normal Boundary Intersection (NBI) [10], which generates points by intersecting normal directions with the attainable objective set;
- Normalized Normal Constraint (NNC) [11], which enforces normalized constraint geometry and aims at evenly distributed points.

In practice, both are frequently paired with a uniform simplex lattice of directions. This motivates our UNIFORM baseline (used in the experiments): it captures the widely used, non-adaptive choice of evenly spaced reference vectors.

To adaptively improve point distributions, Kim and de Weck's Adaptive Weighted Sum (AWS) method [12] refines by splitting "large patches" in objective space. AWS directly targets the bottleneck we study—where to place new reference directions—but its refinement signal (patch size) is only an indirect proxy for uniform reconstruction error. In dimensions $m \geq 3$, a patch can be large yet nearly planar (easy to approximate), while a small patch can be highly curved (hard to approximate). This motivates our AREA-MAX baseline in the experiments (AWS-inspired) and clarifies why CATS uses a curvature-driven certificate.

Recent scalarization-based front-construction approaches continue this line by proposing new objective-constraint formulations tailored to Pareto-front generation [13].

2.2. From Point Sets to Continuous Surrogates: Triangulated Approximations

Beyond point sets, several works build continuous geometric surrogates of the Pareto front/surface to support visualization and decision making. A representative example is the triangulation-based approximation approach of Hartikainen–Miettinen–Wiecek [14], which constructs a simplicial complex (often derived from a Delaunay triangulation) over known Pareto outcomes and uses indicators motivated by coverage and dominance consistency.

These approaches share an important theme with CATS: a continuous curve/surface is more useful than a point cloud. The difference is the domain and the refinement signal. Triangulation-in-objective-space methods triangulate the set of sampled outcomes and drive

refinement using indicators derived from those samples. CATS instead triangulates the reference simplex (the parameter domain) and uses an a posteriori interpolation certificate tied to the smoothness of the parametric map $u \mapsto z(u)$. This enables a direct certificate for a uniform parametric reconstruction error of the form

$$\sup_{u \in \Delta_d} \|z(u) - \hat{z}(u)\|_\infty \leq \varepsilon. \tag{7}$$

2.3. Guaranteed-Accuracy Approximations: Dominance-Based ε-Pareto Sets

In deterministic global multi-objective optimization, a standard goal is to compute an ε-Pareto set: a finite set of feasible solutions such that every feasible outcome is approximately dominated (within ε) by some returned point. This notion is fundamental when the goal is to cover the feasible image rather than reconstruct a smooth manifold. Representative deterministic approaches include non-uniform covering methods [15] and branch-and-bound variants [16]. For bi-objective problems, recent methods, such as ROBBO [17], provide certified approximation guarantees with worst-case error control. Still, it addresses a complementary setting in bi-objective optimization, where the Pareto set image is a one-dimensional front. It constructs certified upper/lower envelopes of the front (in suitably scaled/rotated coordinates) from finitely many Pareto samples and derives verifiable conditions on the sample set that guarantee a prescribed componentwise tolerance. In contrast, our focus is on $m \geq 3$ objectives, where the Pareto image is a $(m-1)$-dimensional surface and we target the regime where the Pareto image behaves like a smooth $(m-1)$-dimensional manifold, and the desired output is a continuous reconstructed surface with a uniform parametric error certificate, rather than a dominance-coverage guarantee.

2.4. Polyhedral (Sandwiching) Approximations for Convex Vector Optimization

For convex vector optimization there is a rich literature on approximating the upper image (or Pareto set/image) using inner/outer polyhedral approximations. These methods (often related to Benson-type algorithms) provide computable gap-based certificates: if the inner and outer polyhedra are close in a cone-compatible metric, then an ε-accurate approximation is obtained. See, e.g., unbiased/sandwiching/primal–dual schemes, such as those proposed by Klamroth–Tind–Wiecek [18], Rennen–van Dam–den Hertog [19], and Löhne–Rudloff–Ulus [20]. Recent developments in this field include norm-minimization and outer-approximation schemes for convex vector optimization [21,22].

These approaches provide strong guarantees under convexity assumptions and are often the right tool when one seeks a set/image approximation. However, their certificates quantify a polyhedral image gap, not the interpolation error of a smooth parametric map $u \mapsto z(u)$. CATS instead focuses on surface reconstruction and certifies a uniform parametric approximation error for a simplicial interpolant.

2.5. What "Accuracy" Means in Each Paradigm

The phrase "accurate Pareto approximation" refers to different objects and different metrics across communities. Since our main claim concerns oracle calls to reach a certified uniform surface error, we summarize the relevant notions to avoid apples-to-oranges comparisons (see Table 1).

(A) Objective-space triangulations (sample-driven surrogates).

Triangulated surrogates constructed directly in objective space (e.g., [14]) assess quality using indicators tied to sampled outcomes: the coverage, dominance consistency, or geometric regularity of the complex. These measures are useful in practice, but they do

not naturally certify a parametric uniform bound of the form in Equation (7), because the parameter domain is not explicitly controlled.

(B) Sandwiching gaps (convex set/image approximation).

Sandwiching/Benson-type methods maintain inner/outer polyhedral approximations of the upper image and certify accuracy via a distance (gap) between these polyhedra in a cone-compatible metric. This yields a set/image approximation guarantee, but it is not a certificate for parametric interpolation of a smooth manifold.

(C) CATS certificate (uniform parametric surface reconstruction).

CATS assumes a single-valued smooth map $z : \Delta_d \to \mathbb{R}^m$ and reconstructs it by simplicial interpolation on a triangulation $\mathcal{T}$ of the reference simplex. Its local certificate on a cell τ has the classical interpolation form $\eta(\tau) \asymp$ (curvature upper bound on τ) $\cdot$ $\mathrm{diam}(\tau)^2$, which upper bounds the worst-case interpolation error on that cell. Therefore, the stopping condition $\max_{\tau\in\mathcal{T}} \eta(\tau) \le \varepsilon$ implies the certified global bound of Equation (7).

This is precisely the notion relevant when the desired output is a continuous reconstructed Pareto surface.

Table 1. Different deterministic paradigms certify different notions of approximation quality.

Paradigm	Approximated Object	Typical Certificate/Indicator	Typical Assumptions
Objective-space triangulation [14]	Complex/surface in objective space	Sample-driven coverage/consistency indicators	Pareto samples available; geometry-driven
Sandwiching/ Benson-type [19,20]	Upper image/Pareto set (convex)	Gap between inner/outer polyhedra	Convexity (or variants); convex scalar subproblems
CATS (this paper)	Parametric surface $z(u)$ over Δ_d	Uniform interpolation certificate $\max_\tau \eta(\tau)$	C^2 map and curvature bound; no convexity required

In summary, existing deterministic MOO methods provide strong tools for (i) generating point sets from scalarizations and (ii) certifying set/image approximations under convexity or global-optimization assumptions. CATS, instead, given an oracle that maps reference directions to Pareto points, designs the outer-loop sampling and returns a certified triangulated surrogate of the Pareto surface with a uniform parametric reconstruction error guarantee in objective coordinates, together with a rate-optimal oracle-complexity characterization for this reconstruction problem.

This is the lens through which we compare with uniform reference-direction sampling and AWS-inspired refinement in Section 7.

3. Setup: A Preference-to-Pareto Oracle and Surface Reconstruction

This section formalizes the object reconstructed by CATS and the information available to the outer loop. The purpose is to cleanly separate the inner MOO machinery (scalarizations and their solvers) from the outer-loop task studied here: selecting preference vectors and reconstructing the Pareto surface with a uniform certificate.

We consider the deterministic multi-objective optimization problem in (1) where $\mathcal{X} \subset \mathbb{R}^n$ is non-empty and compact and each objective $f_i : \mathcal{X} \to \mathbb{R}$ is continuous. We use the standard partial order induced by $\mathbb{R}^m_+$: for $y, y' \in \mathbb{R}^m$, we write $y \preceq y'$ if $y_i \le y'_i$ for all i. A point $x^\star \in \mathcal{X}$ is Pareto-optimal if there is no $x \in \mathcal{X}$, such that $F(x) \preceq F(x^\star)$ and $F(x) \ne F(x^\star)$. Let $\mathcal{X}^\star$ denote the set of Pareto-optimal solutions and $\mathcal{Z}^\star := F(\mathcal{X}^\star) \subset \mathbb{R}^m$ its image in objective space.

3.1. Reference Simplex and Preference Vectors

Most deterministic outer-loop methods explore trade-offs by varying a preference vector (weights, reference points, directions). We represent preferences by u in a d-dimensional simplex, with $d := m - 1$. We use the standard "truncated" simplex parameterization

$$\Delta_d := \left\{ u \in \mathbb{R}^d : \; u_i \geq 0, \; \sum_{i=1}^{d} u_i \leq 1 \right\}, \qquad u_{d+1} := 1 - \sum_{i=1}^{d} u_i, \tag{8}$$

so that the associated full weight vector $(u_1, \ldots, u_{d+1}) \in \mathbb{R}^m_+$ has a unit sum. Other equivalent parameterizations (e.g., the full simplex in $\mathbb{R}^m$) could be used; this choice is purely notational.

3.2. Scalarization Oracle: What the Inner Solver Returns

The inner step of many deterministic MOO pipelines is as follows: fix a preference vector u and solve a scalarized subproblem to obtain a Pareto outcome. Examples include:

- Weighted sum: minimize $\sum_{i=1}^{m} u_i f_i(x)$, one of the classical scalarization schemes in deterministic multi-objective optimization [8,9];
- NBI/NNC-type methods: solve a scalarized problem that encodes a reference direction or normalized constraint geometry in objective space [10,11];
- Pascoletti–Serafini and related scalarizations: optimize a scalar variable subject to vector constraints parameterized by u [23,24].

Rather than fixing a scalarization, we model the entire inner computation as an oracle

$$\mathcal{O} : \Delta_d \to \mathbb{R}^m, \qquad \mathcal{O}(u) = z(u), \tag{9}$$

where $z(u)$ is a Pareto outcome in objective space.

CATS is an outer-loop algorithm and assumes:

1. Query access: for any $u \in \Delta_d$ chosen by the outer loop, we can evaluate $z(u)$ by running the inner solver once;
2. Determinism/single-valuedness: the oracle returns a deterministic outcome $z(u)$ for each u. In practice, this is ensured by the uniqueness of the scalarized optimum or by an explicit tie-breaking rule (e.g., lexicographic secondary objective, fixed solver seed, smallest-norm solution).

For some problems and scalarizations, the mapping $u \mapsto z(u)$ may fail to be single-valued (see, e.g., [25,26]). This can occur, for example, when the scalarized subproblem admits multiple optimal solutions, when the Pareto image contains flat faces, when non-strict convexity leads to non-unique supported points, or when bifurcations cause discontinuous branch switching as the parameter varies. In such cases, the natural object is a set-valued correspondence rather than a smooth single-valued map; see, e.g., the solution-mapping viewpoint in variational analysis [27]. Since the goal of the present paper is certified reconstruction of a parametric surface $u \mapsto z(u)$, we explicitly restrict attention to the single-valued regime, or to the case where a deterministic tie-breaking rule selects one representative branch.

3.3. Regularity Assumption: Smoothness of the Preference-to-Pareto Map

CATS relies on a classical interpolation principle: affine interpolation of a C^2 map has error controlled by a curvature bound times by the squared cell diameter. Accordingly, we assume twice continuous differentiability of the oracle-induced map.

Definition 1 (Smooth Pareto map)**.** *The oracle induces a smooth Pareto map if* $z : \Delta_d \to \mathbb{R}^m$ *is single-valued and each component* z_k *belongs to* $C^2(\Delta_d)$.

This assumption is natural when the underlying scalarized problems satisfy suitable regularity and second-order sufficient conditions, and the solution varies smoothly with the parameter u away from bifurcations; see, e.g., classical sensitivity and perturbation analyses for parameterized optimization problems [28,29]. If bifurcations, discontinuities, or branch switching occur, differentiability may fail and the single-valued surface model itself may break down. In that regime, the interpolation-based certificate proved in this paper no longer applies as a global guarantee. The algorithm may still be run as a heuristic sampler, but its certified interpretation is restricted to the smooth single-valued case.

3.4. What CATS Reconstructs: A Triangulated Parametric Surface in Objective Space

Let $\mathcal{T}$ be a conforming triangulation of Δ_d into d-simplices and let $\mathcal{V}(\mathcal{T})$ denote its vertex set. After evaluating the oracle at each vertex $v \in \mathcal{V}(\mathcal{T})$, we define the piecewise-affine interpolant $\hat{z} : \Delta_d \to \mathbb{R}^m$ by simplex-wise affine interpolation of the vertex values. Geometrically, $\hat{z}(\Delta_d)$ is a triangulated surface in objective space whose vertices are oracle-returned Pareto outcomes.

We measure error by the uniform (supremum) norm in objective space as

$$\|z - \hat{z}\|_\infty := \sup_{u \in \Delta_d} \|z(u) - \hat{z}(u)\|_\infty, \tag{10}$$

where $\|\cdot\|_\infty$ on $\mathbb{R}^m$ is the maximum absolute component. This criterion controls the deviation uniformly over the entire parameter domain, not only at sampled points.

The outer loop cannot evaluate (10) exactly because z is unknown between sampled vertices. Instead, CATS uses a computable a posteriori local certificate $\eta(\tau)$ for each simplex $\tau \in \mathcal{T}$, such that $\sup_{u \in \tau} \|z(u) - \hat{z}(u)\|_\infty \le \eta(\tau)$. Consequently, if $\max_{\tau \in \mathcal{T}} \eta(\tau) \le \varepsilon$, then $\|z - \hat{z}\|_\infty \le \varepsilon$. This is the stopping condition used by CATS and the basis for comparing oracle complexity across outer-loop strategies.

Each newly added vertex u triggers one oracle call (one scalarized inner solve) to obtain $z(u)$. We therefore measure outer-loop cost primarily by the number of distinct queried vertices, i.e., $|\mathcal{V}(\mathcal{T})|$.

3.5. Summary of the Setup, Scope, Applicability, and Interpretation

Given oracle access to a smooth, single-valued preference-to-Pareto map $z : \Delta_d \to \mathbb{R}^m$, we seek a triangulation $\mathcal{T}$ and a simplicial interpolant $\hat{z}$ such that $\|z - \hat{z}\|_\infty \le \varepsilon$ with a computable certificate, while using as few oracle calls (vertices) as possible.

The contribution of this paper is a certified outer-loop reconstruction method for a specific but meaningful regime of deterministic multi-objective optimization: the regime in which the scalarization oracle induces a single-valued and sufficiently smooth map from preference vectors to Pareto outcomes. This assumption is restrictive, but it is also the natural regime in which a global simplicial interpolant, an a posteriori interpolation certificate, and a rate-optimal oracle-complexity analysis are simultaneously well defined. Accordingly, the paper does not claim a universal reconstruction method for arbitrary Pareto correspondences; rather, it addresses the outer-loop reconstruction problem in the smooth parametric-surface regime considered in this paper.

Three clarifications are important. First, convexity of the Pareto image is not required. Smooth non-convex Pareto surfaces are admissible. What lies outside the present scope are situations in which the preference-to-Pareto relation is genuinely set-valued, discontinuous, disconnected, non-smooth, or branch-switching, because then the object to be reconstructed

is no longer a single smooth parametric surface. Second, the method reconstructs the Pareto image in objective space. Consequently, geometric proximity on $\hat{z}(\Delta_d)$ should not be interpreted as proximity of the corresponding efficient decisions in $\mathcal{X}$. Decision-space geometry is not certified in this paper. Third, the computational study is designed to validate the method within the regime covered by the theory. It does not aim to claim certified effectiveness for non-smooth or set-valued Pareto correspondences, for which a different mathematical framework would be required.

4. The Certificate

This section derives the certificate used by CATS and explains why it is the right adaptivity signal for uniform surface reconstruction. It addresses two questions:

1. What should the outer loop monitor to detect where the Pareto surface is under-resolved?
2. Why does a certificate scale as "(curvature) $\times$ diameter^2" for affine simplicial interpolation?

CATS reconstructs a smooth map $z : \Delta_d \to \mathbb{R}^m$ by affine interpolation over simplices in the reference (parameter) domain. On a simplex τ, the interpolant $\hat{z}$ is the unique affine map that matches z at the vertices of τ. If z were affine on τ, then $\hat{z}$ would be exact on τ; interpolation error arises only from non-zero second derivatives.

A classical principle from approximation theory is that for a C^2 function, the error of affine interpolation on a set of diameter h is $O(h^2)$ and proportional to a bound on the Hessian. Consequently, a region can be difficult to approximate for two distinct reasons:

- the cell is large (large $\operatorname{diam}(\tau)$), even if the curvature is moderate; or
- the curvature is large (large $\|\nabla^2 z\|$), even if the cell is not especially large.

The CATS certificate explicitly combines these two effects and, therefore, targets uniform reconstruction error.

4.1. Interpolation Error on a Simplex

We start with a scalar function $g : \Delta_d \to \mathbb{R}$. Let $\tau = \operatorname{conv}\{v_0, \ldots, v_d\} \subset \mathbb{R}^d$ be a d-simplex and let $\hat{g}$ be the unique affine function interpolating g at the vertices, i.e., $\hat{g}(v_i) = g(v_i)$ for all i. We seek a computable upper bound on $\sup_{u \in \tau} |g(u) - \hat{g}(u)|$ in terms of a curvature bound and the cell size.

Definition 2 (Simplex diameter)**.** *For a simplex $\tau = \operatorname{conv}\{v_0, \ldots, v_d\} \subset \mathbb{R}^d$, define* $\operatorname{diam}(\tau) := \max_{i,j} \|v_i - v_j\|_2$.

Lemma 1 (Affine interpolation error from a Hessian bound)**.** *Let $g \in C^2(\tau)$ on a d-simplex $\tau \subset \mathbb{R}^d$, and let $\hat{g}$ be the affine interpolant of g on τ. Assume that for all $u \in \tau$, $\|\nabla^2 g(u)\|_2 \le M_\tau$, where $\|\cdot\|_2$ is the operator (spectral) norm. Then, for all $u \in \tau$, $|g(u) - \hat{g}(u)| \le \frac{1}{2} M_\tau \operatorname{diam}(\tau)^2$.*

Proof. Fix $u \in \tau$ and let $(\lambda_0, \ldots, \lambda_d)$ be its barycentric coordinates, so $u = \sum_{i=0}^d \lambda_i v_i$, $\lambda_i \ge 0$, and $\sum_{i=0}^d \lambda_i = 1$. Recall that the affine interpolant satisfies $\hat{g}(u) = \sum_{i=0}^d \lambda_i g(v_i)$.

For each vertex v_i, Taylor's theorem with integral remainder around u gives

$$g(v_i) = g(u) + \nabla g(u)^\top (v_i - u) + \int_0^1 (1-t)\,(v_i - u)^\top \nabla^2 g(u + t(v_i - u))(v_i - u)\,dt.$$

Multiply by λ_i and sum over i. Using $\sum_i \lambda_i = 1$ and $\sum_i \lambda_i (v_i - u) = 0$, the constant and linear terms cancel, yielding

$$\hat{g}(u) - g(u) = \sum_{i=0}^d \lambda_i \int_0^1 (1-t)\,(v_i - u)^\top \nabla^2 g(u + t(v_i - u))(v_i - u)\,dt.$$

Taking absolute values and using $\|\nabla^2 g(\cdot)\|_2 \le M_\tau$ gives

$$|\hat{g}(u) - g(u)| \le \sum_{i=0}^{d} \lambda_i \int_0^1 (1-t)\, M_\tau \, \|v_i - u\|_2^2 \, dt = \frac{M_\tau}{2} \sum_{i=0}^{d} \lambda_i \|v_i - u\|_2^2.$$

Finally, $\|v_i - u\|_2 \le \mathrm{diam}(\tau)$ for all i and $\sum_i \lambda_i = 1$, so $|\hat{g}(u) - g(u)| \le \frac{M_\tau}{2}\mathrm{diam}(\tau)^2$. □

Lemma 1 makes the adaptivity signal explicit: affine interpolation error is controlled by (i) a local curvature bound M_τ and (ii) the squared cell diameter $\mathrm{diam}(\tau)^2$. Thus, one reduces error either by refining (shrinking $\mathrm{diam}(\tau)$) or by prioritizing cells with large curvature. CATS does both by ranking cells through $M_\tau\,\mathrm{diam}(\tau)^2$.

4.2. Vector-Valued Map z (Pareto Surface) and a Local Certificate

We now apply Lemma 1 componentwise to the vector-valued Pareto map $z : \Delta_d \to \mathbb{R}^m$. On a simplex τ, let $\hat{z}$ denote the affine interpolant obtained by interpolating each component at the vertices.

Theorem 1 (Local certificate for a Pareto map)**.** *Let $z : \Delta_d \to \mathbb{R}^m$ with components $z_k \in C^2(\tau)$ on a simplex τ. Assume that for each component $k \in \{1, \ldots, m\}$, there exists $M_{\tau,k}$ such that $\|\nabla^2 z_k(u)\|_2 \le M_{\tau,k}, \forall u \in \tau$. Then, $\forall u \in \tau$, $\|z(u) - \hat{z}(u)\|_\infty \le \eta(\tau) := \frac{1}{2}(\max_k M_{\tau,k})\mathrm{diam}(\tau)^2$.*

Proof. For each k, Lemma 1 yields $|z_k(u) - \hat{z}_k(u)| \le \frac{1}{2} M_{\tau,k}\mathrm{diam}(\tau)^2$. Taking the maximum over k and using $\|z(u) - \hat{z}(u)\|_\infty = \max_k |z_k(u) - \hat{z}_k(u)|$ gives the claim. □

The true surface patch over τ is $z(\tau) \subset \mathbb{R}^m$. The reconstructed patch $\hat{z}(\tau)$ is an affine image of τ and hence a flat simplex in $\mathbb{R}^m$. The certificate $\eta(\tau)$ upper-bounds the maximum coordinate-wise deviation between these two patches uniformly over all parameter values in τ.

4.3. Availability and Interpretation of Curvature Bounds

The certificate in Theorem 1 is rigorous only when the quantities $M_{\tau,k}$ are valid upper bounds on the Hessian norms of the oracle-induced map components over the simplex τ. This is an important applicability condition. In analytic test problems, such bounds may be computed exactly or conservatively from closed-form derivatives. In structured deterministic optimization problems, they may also be obtained from sensitivity analysis of the scalarized subproblem, provided that suitable regularity, uniqueness, constraint qualification, and second-order sufficient conditions hold. In other settings, conservative bounds may be available from interval arithmetic, validated numerical differentiation, or known Lipschitz constants for the relevant derivatives.

However, in a general multi-objective optimization pipeline, the map $u \mapsto z(u)$ is usually defined implicitly by the solution of a scalarized optimization problem, and the inner solver may not provide second-order sensitivity information with respect to the preference vector. In simulation-based settings, such as finite-element, CFD, or other black-box engineering models, estimating second derivatives from oracle values may be numerically delicate and may depend strongly on solver tolerances, discretization error, and finite-difference step sizes. Therefore, finite-difference or regression-based estimates of $M_{\tau,k}$ do not by themselves provide a rigorous certificate unless they are accompanied by validated error bounds.

Accordingly, the word "certified" is used in this paper for the regime in which valid local curvature upper bounds are available. When such bounds are unavailable and are replaced by numerical curvature estimates, CATS can still be used as an adaptive sampling strategy, but the resulting quantity $\eta(\tau)$ should be interpreted as a refinement indicator

rather than as a mathematical certificate. In that indicator-based regime, the stopping condition $\max_{\tau\in\mathcal{T}} \eta(\tau) \le \varepsilon$ is a practical rule, not a formal proof of the uniform error bound $\|z - \hat{z}\|_\infty \le \varepsilon$.

4.4. Numerical Stability of the Local Certificate

The local quantity

$$\eta(\tau) = \frac{1}{2}\left(\max_k M_{\tau,k}\right)\operatorname{diam}(\tau)^2 \tag{11}$$

depends on two numerical ingredients: a curvature bound and a geometric diameter. These two quantities have different stability properties and should be treated separately.

If the curvature bounds $M_{\tau,k}$ are available analytically or through validated numerical bounds, then the computation of $\eta(\tau)$ is stable in the usual conservative sense: rounding errors can be absorbed by slightly inflating the quantities entering the certificate. For instance, one may replace $\eta(\tau)$ by the safe value

$$\eta_{\text{safe}}(\tau) = \frac{1}{2}(1+\delta_M)(1+\delta_h)^2\left(\max_k M_{\tau,k}\right)\operatorname{diam}(\tau)^2, \tag{12}$$

where $\delta_M, \delta_h \ge 0$ are small safety factors accounting for floating-point evaluation of the curvature and diameter. The stopping test is then applied to $\eta_{\text{safe}}(\tau)$ rather than to $\eta(\tau)$. This preserves the conservative character of the certificate.

The situation is different when $M_{\tau,k}$ is estimated from oracle values. Finite-difference Hessian estimates are sensitive to solver tolerances, discretization errors, and the finite-difference step size. If the step is too small, roundoff and solver noise may dominate the second-difference calculation; if the step is too large, the estimate may be biased by higher-order terms. This issue is particularly relevant when the oracle is a numerical simulation, such as an FEA, CFD, or multi-physics solver, because the returned values may contain discretization and non-linear-solver errors. In such cases, a finite-difference estimate of $\nabla^2 z_k$ should not be interpreted as a rigorous upper bound unless it is accompanied by an explicit validation margin.

Very small simplices require particular care. Although the factor $\operatorname{diam}(\tau)^2$ decreases under refinement, numerical estimates of curvature may become unreliable when the variation in z over τ is comparable to the numerical noise of the oracle. A practical implementation should therefore ensure that the tolerance of the inner solver is significantly smaller than the target reconstruction tolerance, avoid refinement below a prescribed minimum diameter when curvature estimates are not validated, and use conservative inflation factors whenever a bound is intended to be certified.

In summary, the certificate is numerically robust when its curvature inputs are themselves reliable upper bounds. If these inputs are obtained from unvalidated numerical differentiation, the same formula remains useful for ranking cells and guiding refinement, but the corresponding stopping test should be regarded as indicator-based rather than certified.

4.5. From Local to Global: Controlling the Worst Cell Certifies the Whole Surface

Local certificates imply a global guarantee because the simplices in $\mathcal{T}$ partition the reference simplex.

Theorem 2 (Global certified approximation)**.** *Let $\mathcal{T}$ be a triangulation of Δ_d and let $\hat{z}$ be the piecewise-affine interpolant of z on $\mathcal{T}$. If every simplex $\tau \in \mathcal{T}$ satisfies $\sup_{u\in\tau} \|z(u) - \hat{z}(u)\|_\infty \le \eta(\tau)$, then $\|z - \hat{z}\|_\infty \le \max_{\tau\in\mathcal{T}} \eta(\tau)$. In particular, if $\max_{\tau\in\mathcal{T}} \eta(\tau) \le \varepsilon$, then $\|z - \hat{z}\|_\infty \le \varepsilon$.*

Proof. Any $u \in \Delta_d$ lies in some simplex $\tau \in \mathcal{T}$, so $\|z(u) - \hat{z}(u)\|_\infty \leq \eta(\tau) \leq \max_{\tau'} \eta(\tau')$. Taking the supremum over u yields the result. □

Theorem 2 shows that the maximum local certificate controls the global reconstruction error. Therefore, to reach tolerance ε with as few oracle calls as possible, it is natural to repeatedly refine the simplex with the largest $\eta(\tau)$, which is precisely the CATS selection rule.

The bound in Theorem 1 is worst-case and can be conservative. CATS uses it deliberately: it is a certificate, not merely a heuristic indicator. If curvature bounds are loose, CATS remains correct but may over-refine; this is the price of certification.

4.6. A Further Remark on Multiple Optima, Tie-Breaking, and Bifurcations

The assumption that the oracle is single-valued is essential for the certified reconstruction result. In general non-convex or degenerate multi-objective problems, a fixed preference vector u may lead to a scalarized subproblem with multiple globally optimal solutions, and these solutions may have different images in objective space. In that case, the natural object is a set-valued correspondence

$$u \rightrightarrows Z(u) \subseteq \mathbb{R}^m$$

rather than a single map $u \mapsto z(u)$.

A deterministic implementation may still induce a single-valued map by imposing a fixed tie-breaking rule. Examples include selecting, among all optimal solutions of the scalarized subproblem, the solution with minimum Euclidean norm in decision space, applying a lexicographic secondary objective, choosing the solution closest to the previously selected branch in a continuation procedure, or using any other deterministic selection criterion that is applied consistently for all values of u. With such a rule, the oracle becomes

$$z(u) \in Z(u)$$

where $z(u)$ denotes the selected representative. CATS then reconstructs this selected branch of the Pareto image, not the full set-valued correspondence.

This distinction is important in the presence of bifurcations or branch switching. If the selected branch varies continuously and is C^2 on the region of interest, the interpolation certificate applies to that branch. If, however, the selected representative changes discontinuously, or if two branches merge, split, or exchange optimality as u varies, then the map $u \mapsto z(u)$ may fail to be differentiable or even continuous. In such cases, the smooth parametric-surface model underlying the certificate is violated. The algorithm may still be used as an adaptive sampling heuristic to reveal regions of rapid variation or branch transition, but the stopping condition based on $\eta(\tau)$ should not be interpreted as a rigorous global uniform-error certificate across the bifurcation.

5. The Proposed Algorithm

We now describe Certified Adaptive Triangulation Sampling (CATS) as an outer-loop procedure for Pareto-surface reconstruction. The inner optimization machinery is abstracted as the oracle in Equation (9), which returns one Pareto outcome for each queried preference vector u. The role of CATS is therefore not to solve the scalarized subproblem itself, but to decide which preference vector should be queried next so as to reconstruct the full Pareto surface with a certified uniform accuracy using as few oracle calls as possible.

5.1. Basic Idea

CATS maintains a conforming triangulation $\mathcal{T}$ of the reference simplex Δ_d. Its vertex set $\mathcal{V}(\mathcal{T})$ consists of all preference vectors queried so far. After evaluating the oracle at those vertices, CATS constructs the piecewise-affine interpolant $\hat{z} : \Delta_d \to \mathbb{R}^m$ by simplex-wise barycentric interpolation of the sampled values $\{z(v) : v \in \mathcal{V}(\mathcal{T})\}$.

For each simplex $\tau \in \mathcal{T}$, the local certificate $\eta(\tau)$ from Section 4 provides an upper bound on the worst interpolation error over that cell. Since $\|z - \hat{z}\|_\infty \leq \max_{\tau \in \mathcal{T}} \eta(\tau)$, the global reconstruction problem reduces to controlling the largest local certificate. This naturally suggests a greedy strategy, which operates as follows: repeatedly identify the simplex with the largest certificate, refine it, and query the oracle at any newly created vertices.

Algorithm 1 summarizes this logic at a high level. Algorithm 2 then gives the precise CATS refinement rule.

Algorithm 1 High-level workflow for certified Pareto-surface reconstruction

Require: Reference simplex Δ_d; oracle $\mathcal{O} : u \mapsto z(u)$; initial conforming triangulation $\mathcal{T}_0$; target tolerance $\varepsilon > 0$
Ensure: Conforming triangulation $\mathcal{T}$, sampled values $\{z(v) : v \in \mathcal{V}(\mathcal{T})\}$, and simplicial interpolant $\hat{z}$ satisfying $\|z - \hat{z}\|_\infty \leq \varepsilon$

Initialize $\mathcal{T} \leftarrow \mathcal{T}_0$
Query the oracle at all vertices of $\mathcal{T}$
Construct the piecewise-affine interpolant $\hat{z}$
while the certified error bound exceeds ε **do**
    Compute local certificates on all simplices of $\mathcal{T}$
    Select one simplex for refinement according to the chosen outer-loop rule
    Refine the triangulation conformingly
    Query the oracle at all newly created vertices
    Update $\hat{z}$ and the cell certificates
end while
return $\mathcal{T}$ and $\hat{z}$

Algorithm 2 CATS: Certified Adaptive Triangulation Sampling

Require: Target tolerance $\varepsilon > 0$; dimension $d = m - 1$; initial conforming triangulation $\mathcal{T}_0$ of Δ_d; oracle $\mathcal{O}$; cell-wise curvature bounds $\{M_{\tau,k}\}$
Ensure: Final conforming triangulation $\mathcal{T}$; sampled oracle values at all vertices; piecewise-affine interpolant $\hat{z}$ with certified bound $\|z - \hat{z}\|_\infty \leq \varepsilon$

$\mathcal{T} \leftarrow \mathcal{T}_0$
for all $v \in \mathcal{V}(\mathcal{T})$ **do**
    Evaluate $z(v) \leftarrow \mathcal{O}(v)$
end for
repeat
    for all $\tau \in \mathcal{T}$ **do**
        Compute $\eta(\tau) = \frac{1}{2}(\max_k M_{\tau,k})\mathrm{diam}(\tau)^2$
    end for
    Select $\tau^\star \in \arg\max_{\tau \in \mathcal{T}} \eta(\tau)$
    if $\eta(\tau^\star) \leq \varepsilon$ **then**
        Build/update the piecewise-affine interpolant $\hat{z}$
        return $\mathcal{T}$, $\{z(v)\}_{v \in \mathcal{V}(\mathcal{T})}$, and $\hat{z}$
    end if
    Refine $\tau^\star$ by conforming longest-edge bisection with conformity closure
    for all new vertices v created by the refinement **do**
        Evaluate $z(v) \leftarrow \mathcal{O}(v)$
    end for
until $\eta(\tau^\star) \leq \varepsilon$

5.2. Refinement Primitive and Oracle Accounting

Any refinement strategy used in this setting must satisfy two structural requirements. First, it must preserve conformity, so that the interpolant $\hat{z}$ remains globally continuous on Δ_d. Second, it should avoid arbitrarily degenerate simplices, because both the interpolation estimates and the oracle-complexity analysis rely on mesh regularity.

For this reason, we use conforming longest-edge bisection (LEB), a standard refinement primitive from adaptive finite element methods. Given a marked simplex τ, LEB inserts the midpoint of its longest edge and subdivides τ accordingly. If this creates non-conformity with adjacent simplices, a finite closure step is performed: neighboring simplices are bisected recursively until conformity is restored. Under mild assumptions, repeated LEB refinement generates shape-regular meshes and prevents arbitrarily skinny simplices [30,31]. This is precisely the mesh property used later in the oracle-complexity analysis.

Refining a single marked simplex typically creates one new vertex, namely the midpoint of its longest edge. However, the conformity closure may create additional midpoints on adjacent simplices. Each newly created vertex is queried once through the oracle and therefore contributes one oracle call. Accordingly, one refinement step may incur one or a small bounded number of oracle calls, depending on the closure.

5.3. CATS Refinement Rule

The specific CATS rule is simple: at each iteration, compute the certificate of every current simplex and refine the one with the largest value. The method stops as soon as the largest certificate falls below the prescribed tolerance ε. By Theorem 2, this guarantees $\|z - \hat{z}\|_\infty \leq \varepsilon$.

The rationale for this rule is immediate: since the global certified error is controlled by the largest local certificate, the simplex attaining that maximum is the current bottleneck. Refining it is therefore a natural greedy strategy for reducing an explicit upper bound on the uniform reconstruction error.

The dominant computational cost is the number of oracle calls, i.e., the number of scalarized inner solves. We count one oracle call per distinct vertex at which $\mathcal{O}$ is evaluated. Equivalently, the oracle complexity of the outer loop is $|\mathcal{V}(\mathcal{T})|$ at termination, including any vertices introduced during conformity restoration.

Stopping Criterion and Engineering Tolerances

The stopping criterion in CATS is not based on a prescribed number of sampled points. Instead, it is driven by the desired reconstruction accuracy in objective space. In the simplest case, a single tolerance $\varepsilon > 0$ is specified and the algorithm stops when

$$\max_{\tau \in \mathcal{T}} \eta(\tau) \leq \varepsilon.$$

This means that every component of the reconstructed Pareto surface is uniformly accurate, in the ℓ_∞ sense, up to the prescribed tolerance ε.

In engineering applications, however, different objectives may have different physical units and different acceptable error levels. For example, admissible errors in power loss, wear rate, cost, energy consumption, or service life may not be comparable on the same numerical scale. In such cases, it is more appropriate to prescribe componentwise tolerances

$$\varepsilon_1, \ldots, \varepsilon_m > 0,$$

where ε_k represents the admissible reconstruction error for the kth objective. The stopping rule can then be written in normalized form as

$$\max_{\tau\in\mathcal{T}} \max_{1\le k\le m} \frac{\eta_k(\tau)}{\varepsilon_k} \le 1, \qquad \eta_k(\tau) = \frac{1}{2} M_{\tau,k} \mathrm{diam}(\tau)^2.$$

Equivalently, the algorithm stops only when

$$\eta_k(\tau) \le \varepsilon_k \qquad \text{for every } k = 1, \dots, m \text{ and every } \tau \in \mathcal{T}.$$

This formulation makes the stopping test directly interpretable in terms of application-specific engineering tolerances rather than in terms of the number of sampled reference vectors. The number of oracle calls is then an output of the algorithm, not an input stopping parameter.

5.4. Baselines and Controlled Proxy Comparisons

The computational comparisons in this paper are designed to isolate the effect of the refinement indicator within a common reconstruction framework. Therefore, all methods use the same reference-simplex triangulation, the same conforming longest-edge bisection refinement primitive, and the same stopping rule based on $\max_{\tau\in\mathcal{T}} \eta(\tau) \le \varepsilon$. The compared rules differ only in how they select the next simplex to refine.

The baselines below should be interpreted as controlled proxy rules rather than full reproductions of complete MOO algorithms. In particular, AREA-MAX isolates an objective-space patch-size idea inspired by AWS, whereas HMW-DIAM isolates an objective-space spread idea inspired by triangulated Pareto-front approximation. This controlled design is useful for testing whether objective-space geometric indicators can match the curvature-weighted interpolation certificate used by CATS under identical refinement and stopping conditions. It is not intended as a comprehensive benchmark against the full original implementations.

Baseline 1: UNIFORM reference directions.

UNIFORM performs global refinement of the reference triangulation. At each refinement level, every simplex is bisected once, producing a quasi-uniform mesh over Δ_d. The refinement level is increased until the resulting triangulation satisfies the same certificate threshold.

Baseline 2: AREA-MAX, an AWS-inspired patch-area proxy.

Adaptive Weighted Sum (AWS) [12] is a complete adaptive scalarization method. Its original design is tied to the weighted-sum framework and uses more information than patch area alone, including assumptions and procedures related to the continuity of the weighted-sum solution path. The AREA-MAX rule used in this paper is a controlled proxy that isolates one objective-space idea associated with AWS-type refinement: prioritizing large mapped patches of the current approximation. Within our reference-simplex reconstruction framework, for a simplex $\tau = \mathrm{conv}\{v_0, \dots, v_d\}$, define its mapped patch as

$$\mathrm{conv}\{z(v_0), \dots, z(v_d)\} \subset \mathbb{R}^m.$$

AREA-MAX marks a simplex maximizing the intrinsic d-dimensional volume of this mapped patch,

$$\tau^\star \in \arg\max_{\tau\in\mathcal{T}} \mathrm{vol}_d\Big(\mathrm{conv}\{z(v_i)\}_{i=0}^d\Big),$$

where $\mathrm{vol}_d(\cdot)$ denotes d-dimensional volume in the affine hull of the simplex. The marked simplex is then refined by the same conforming longest-edge bisection procedure used by CATS and UNIFORM.

Thus, AREA-MAX tests whether objective-space patch size alone is an effective refinement signal for the uniform reconstruction criterion considered in this paper. It does not claim to reproduce all features, safeguards, or advantages of the original AWS algorithm.

UNIFORM represents the most common non-adaptive choice in deterministic sampling-based multi-objective optimization. AREA-MAX represents a natural adaptive alternative based on the geometric size of the reconstructed surface patches. CATS differs from both by refining according to the largest certified interpolation error; that is, where cell size and local curvature together indicate that the surface is under-resolved.

Thus, all the above methods share the same reconstruction framework and differ only in the rule used to choose the next simplex for refinement. This isolates the effect of the refinement criterion itself.

Additional Objective-Space Diagnostic: HMW-DIAM

Hartikainen–Miettinen–Wiecek (HMW) [14] proposed a triangulation-based Pareto-front approximation framework that works directly with the geometry of sampled Pareto outcomes in objective space. We include an additional HMW-inspired diagnostic rule, denoted HMW-DIAM, on the certified core tests. This rule should not be interpreted as a full implementation of Hartikainen–Miettinen–Wiecek. Rather, it isolates a simple objective-space coverage principle: refine the current simplex whose sampled image has the largest diameter.

For a simplex $\tau = \mathrm{conv}\{v_0, \ldots, v_d\}$, define

$$D_z(\tau) = \max_{0 \le i < j \le d} \|z(v_i) - z(v_j)\|_2. \tag{13}$$

The HMW-DIAM rule marks

$$\tau^\star \in \arg\max_{\tau \in \mathcal{T}} D_z(\tau) \tag{14}$$

and refines $\tau^\star$ using the same conforming longest-edge bisection procedure and stopping criterion as the other methods. This diagnostic comparison indicates whether an objective-space spread indicator can match a curvature-weighted interpolation certificate under identical refinement and stopping conditions.

6. Oracle Complexity: Upper and Lower Bounds

This section quantifies the fundamental question behind outer-loop design, i.e., how many oracle calls are necessary (and sufficient) to reconstruct a d-dimensional Pareto surface to uniform tolerance ε using a simplicial, piecewise-affine surrogate with a certificate?

Recall that $d = m - 1$ is the dimension of the reference simplex Δ_d and hence of the reconstructed parameter domain. The central message is that, under C^2 smoothness, certified uniform reconstruction by simplicial affine interpolation has the rate $N(\varepsilon) = \Theta(\varepsilon^{-d/2})$, and this exponent is unavoidable for certified reconstructions based on point queries.

Affine interpolation on a simplex τ has error bounded by $\eta(\tau) \lesssim (\text{curvature}) \cdot \mathrm{diam}(\tau)^2$ (Section 4). Under a uniform curvature bound, achieving $\eta(\tau) \le \varepsilon$ requires $\mathrm{diam}(\tau) \lesssim \sqrt{\varepsilon}$. On a quasi-uniform mesh in d dimensions with mesh size h, the number of simplices (and vertices) scales like h^{-d}. Substituting $h \asymp \sqrt{\varepsilon}$ yields the exponent $d/2$:

$$N(\varepsilon) \asymp (\sqrt{\varepsilon})^{-d} = \varepsilon^{-d/2}.$$

The results below formalize both the sufficiency (upper bound) and necessity (lower bound) of this scaling.

6.1. Upper Bound: $O(\varepsilon^{-d/2})$ Oracle Calls (Sufficiency)

We first formalize a standard mesh-quality condition.

Definition 3 (Shape regularity)**.** *A family of d-simplices is shape-regular if there exists $\rho > 0$ such that for every simplex τ in the family, $\frac{\mathrm{inrad}(\tau)}{\mathrm{circrad}(\tau)} \geq \rho$, where $\mathrm{inrad}(\tau)$ and $\mathrm{circrad}(\tau)$ are the inradius and circumradius, respectively.*

Shape regularity controls simplex degeneracy and implies that simplex volume is comparable to $\mathrm{diam}(\tau)^d$ up to constants.

Lemma 2 (Volume–diameter equivalence under shape regularity)**.** *Fix dimension d and shape-regularity constant ρ. Then, there exist constants $c_1, c_2 > 0$ (depending only on d and ρ), such that every shape-regular d-simplex τ satisfies $c_1 \,\mathrm{diam}(\tau)^d \leq \mathrm{vol}_d(\tau) \leq c_2 \,\mathrm{diam}(\tau)^d$.*

Proof. The upper bound follows since τ is contained in a ball of radius $\mathrm{diam}(\tau)$, so $\mathrm{vol}_d(\tau) \leq c_2 \mathrm{diam}(\tau)^d$ for a constant $c_2 = c_2(d)$. For the lower bound, shape regularity implies $\mathrm{inrad}(\tau) \geq \rho \,\mathrm{circrad}(\tau)$ and $\mathrm{circrad}(\tau) \geq c \,\mathrm{diam}(\tau)$ for a constant $c = c(d) > 0$. Hence, τ contains a d-ball of radius at least $\alpha \mathrm{diam}(\tau)$ for $\alpha = \alpha(d, \rho) > 0$, implying $\mathrm{vol}_d(\tau) \geq c_1 \mathrm{diam}(\tau)^d$. □

Lemma 3 (Vertices vs. simplices)**.** *Let $\mathcal{T}$ be a conforming triangulation of Δ_d into d-simplices, and let $\mathcal{V}(\mathcal{T})$ be its vertex set. Then, $|\mathcal{V}(\mathcal{T})| \leq (d+1)\,|\mathcal{T}|$.*

Proof. Each simplex has exactly $d+1$ vertices, so counting vertex–simplex incidences yields $(d+1)|\mathcal{T}| = \sum_{v \in \mathcal{V}(\mathcal{T})} \deg(v) \geq \sum_{v \in \mathcal{V}(\mathcal{T})} 1 = |\mathcal{V}(\mathcal{T})|$. □

The next result is an existence (sufficiency) bound: it shows that the certified tolerance ε can be achieved with $O(\varepsilon^{-d/2})$ oracle calls, for example by a uniform simplex lattice.

Theorem 3 (Oracle complexity upper bound (existence/sufficiency))**.** *Assume that $z : \Delta_d \to \mathbb{R}^m$ is C^2 and has uniformly bounded curvature: for each component $k \in \{1, \ldots, m\}$ and every $u \in \Delta_d$, $\|\nabla^2 z_k(u)\|_2 \leq M$ for some constant $M > 0$. Then, for every $\varepsilon \in (0, 1)$, there exists a conforming, shape-regular triangulation $\mathcal{T}_\varepsilon$ of Δ_d and a piecewise-affine interpolant $\hat{z}$ on $\mathcal{T}_\varepsilon$ such that $\|z - \hat{z}\|_\infty \leq \varepsilon$ and $|\mathcal{V}(\mathcal{T}_\varepsilon)| \leq C\left(\frac{M}{\varepsilon}\right)^{d/2}$, where C depends only on d (and the reference simplex geometry).*

Proof. By Theorem 1 and the uniform curvature bound, on any simplex τ, we have $\sup_{u \in \tau} \|z(u) - \hat{z}(u)\|_\infty \leq \eta(\tau) \leq \frac{1}{2} M \,\mathrm{diam}(\tau)^2$. Thus, it suffices to ensure $\mathrm{diam}(\tau) \leq h := \sqrt{2\varepsilon/M}$ for every simplex.

Consider a standard uniform simplex-lattice triangulation of Δ_d with mesh size proportional to h (e.g., evenly spaced barycentric grids). Such meshes are conforming and shape-regular, and satisfy $\max_{\tau \in \mathcal{T}_\varepsilon} \mathrm{diam}(\tau) \leq c\,h$ for a constant $c = c(d)$. Hence, choosing a fine-enough mesh ensures $\eta(\tau) \leq \varepsilon$ on every simplex and therefore $\|z - \hat{z}\|_\infty \leq \varepsilon$.

Moreover, the number of simplices and vertices in a uniform lattice triangulation scales as $|\mathcal{T}_\varepsilon| = O(h^{-d})$ and $|\mathcal{V}(\mathcal{T}_\varepsilon)| = O(h^{-d})$. Substituting $h = \sqrt{2\varepsilon/M}$ yields $|\mathcal{V}(\mathcal{T}_\varepsilon)| = O\left(h^{-d}\right) = O\left(\left(\frac{M}{\varepsilon}\right)^{d/2}\right)$, giving the claim. □

Theorem 3 is a sufficiency bound: it shows that ε-accurate certified reconstruction is achievable with $O(\varepsilon^{-d/2})$ oracle calls (e.g., by uniform sampling). An arbitrary triangulation satisfying $\max_\tau \eta(\tau) \leq \varepsilon$ may contain many more elements if it over-refines; the theorem bounds the order of oracle calls needed by an optimal (or near-optimal) outer-loop design.

6.2. Lower Bound: $\Omega(\varepsilon^{-d/2})$

We now show that the exponent $d/2$ cannot be improved in a worst-case, information-theoretic sense for certified simplicial piecewise-affine reconstructions based on point evaluations. We state the lower bound for scalar functions; it implies the same lower bound for vector maps by embedding a scalar hard instance into one component.

Theorem 4 (Minimax lower bound for certified simplicial interpolation)**.** *Fix $d \geq 1$. Let $\mathcal{G}$ be the class of scalar functions $g : \Delta_d \to \mathbb{R}$ such that $g \in C^2(\Delta_d)$ and $\|\nabla^2 g(u)\|_2 \leq 1, \forall u \in \Delta_d$. Consider any (possibly adaptive) algorithm that may query point values $g(u)$ at points $u \in \Delta_d$ and must output: (i) a conforming triangulation $\mathcal{T}$ of Δ_d whose vertex set equals the queried points, and (ii) the nodal piecewise-affine interpolant $\hat{g}$ on $\mathcal{T}$ (so $\hat{g}(v) = g(v)$ for all $v \in \mathcal{V}(\mathcal{T})$), together with (iii) a certificate guaranteeing $\sup_{u\in\Delta_d} |g(u) - \hat{g}(u)| \leq \varepsilon, \forall g \in \mathcal{G}$. Then, there exists a constant $c > 0$ (depending only on d), such that, in the worst case over $\mathcal{G}$, the algorithm requires at least $N(\varepsilon) \geq c\,\varepsilon^{-d/2}$ queries.*

Proof. Step 1: packing. Let $r := \gamma\sqrt{\varepsilon}$ where $\gamma > 0$ is a fixed constant chosen below. For ε small enough, we can place a maximal family of disjoint Euclidean balls $\{B(u_i, r)\}_{i=1}^{K}$ contained in the interior of Δ_d. A standard volumetric packing argument yields $K \geq c_0\, r^{-d} = c_0\, \gamma^{-d}\, \varepsilon^{-d/2}$, for some constant $c_0 > 0$ depending only on d and the geometry of Δ_d.

Step 2: a smooth bump with controlled Hessian. Let $\psi : [0,\infty) \to [0,1]$ be a fixed C^2 function with $\psi(t) = 0$ for $t \geq 1$, $\psi(0) = 1$, and bounded derivatives $\sup_{t\geq 0} |\psi'(t)| + \sup_{t\geq 0} |\psi''(t)| \leq C_\psi$. For each center u_i, define $\phi_i(u) := a\,\psi\left(\frac{\|u-u_i\|_2^2}{r^2}\right)$. Then ϕ_i is C^2, supported in $B(u_i, r)$, and satisfies $\phi_i(u_i) = a$. A direct chain-rule calculation gives $\|\nabla^2\phi_i(u)\|_2 \leq C\,\frac{a}{r^2}$, for a constant $C = C(d, C_\psi)$. Choose the amplitude $a := 2\varepsilon$ and choose γ large enough so that $C\,\frac{a}{r^2} = C\,\frac{2\varepsilon}{\gamma^2\varepsilon} = \frac{2C}{\gamma^2} \leq 1$; e.g., take $\gamma := \sqrt{2C}$. Then, $\phi_i \in \mathcal{G}$.

Step 3: indistinguishability from point queries. Let the algorithm make N queries at points $q_1, \ldots, q_N$. If $N < K$, then at least one ball $B(u_i, r)$ contains no query point. For that index i, consider $g_0(u) \equiv 0$ and $g_1(u) := \phi_i(u)$. Since g_1 is supported in $B(u_i, r)$ and no query lies in that ball, we have $g_0(q_j) = g_1(q_j) = 0$ for all j. Thus, the algorithm observes identical data under g_0 and g_1.

Step 4: contradiction with certified interpolation. By construction, the triangulation $\mathcal{T}$ uses only queried points as vertices, and $\hat{g}$ is the nodal interpolant. Hence, all vertex values are zero under both g_0 and g_1, so the interpolant is identical in both cases. In particular, every simplex of $\mathcal{T}$ has all-zero nodal data under g_1, implying $\hat{g} \equiv 0$ everywhere. Therefore, $\sup_{u\in\Delta_d} |g_1(u) - \hat{g}(u)| \geq |g_1(u_i) - \hat{g}(u_i)| = 2\varepsilon > \varepsilon$, contradicting the claimed certificate for $g_1 \in \mathcal{G}$. Thus, any such algorithm must have $N \geq K \geq c\,\varepsilon^{-d/2}$ for a constant $c > 0$ depending only on d. □

The lower bound applies to $z : \Delta_d \to \mathbb{R}^m$ as well: embed a hard scalar instance as one component (e.g., $z_1 = g$) and set the remaining components identically zero.

Theorems 3 and 4 show that, for certified simplicial piecewise-affine reconstructions under C^2 smoothness, the oracle complexity exponent $d/2$ is tight. Improvements from CATS over uniform sampling are therefore improvements in constant factors (by focusing refinement where curvature is high), not in the information-theoretic exponent.

7. Computational Results

This section evaluates CATS as an outer-loop strategy for Pareto-surface reconstruction. The objective is not simply to generate non-dominated points, but to reconstruct the entire Pareto surface as a continuous triangulated manifold. In the rigorously certified core study,

the comparison is made under a uniform accuracy guarantee; in the broader supplemental studies, the same local interpolation indicator is used as a common practical stopping rule across methods. Accordingly, the main performance measure is oracle complexity, i.e., the number of oracle calls required to reach the prescribed stopping threshold.

To make the comparison meaningful, all methods reconstruct the same object—a piecewise-affine interpolant $\hat{z}$ on a triangulation $\mathcal{T}$ of Δ_2—and use the same refinement primitive, namely conforming longest-edge bisection. In the certified core study of Sections 7.3–7.5, all methods are compared under the stopping rule $\max_{\tau\in\mathcal{T}} \eta(\tau) \leq \varepsilon$. In the additional benchmark studies reported later in this section, the same quantity $\eta(\tau)$ is used as a common practical stopping indicator across methods. In all cases, the reported oracle calls equal the number of distinct queried vertices, i.e., the number of scalarized inner solves one would run in a real MOO pipeline (including any additional vertices created to restore conformity).

The reason for using a controlled oracle in these experiments is as follows. In a full MOO application, the oracle $\mathcal{O}(u)$ would solve a scalarized subproblem and return a Pareto point. This couples the outer-loop question studied here (where to sample?) with solver-dependent effects (tolerances, non-uniqueness, convergence issues). To isolate the contribution of CATS as a reconstruction-driven sampling rule, we use test problems where the oracle map is known analytically and smooth. This lets us (i) compute curvature bounds exactly, (ii) focus evaluation on outer-loop sampling efficiency, and (iii) produce faithful visualizations of true vs. reconstructed surfaces.

These experiments should therefore be interpreted as validation within the regular regime covered by the theory. They are not intended as evidence for discontinuous, branch-switching, or set-valued Pareto correspondences, for which the present certification framework does not apply.

All experiments were run in Python 3.11.2. Numerical computations used NumPy 1.24.0 and SciPy 1.14.1. All plots were generated with Matplotlib 3.7.5, and data aggregation used Pandas 2.2.3.

7.1. Scope of the Computational Validation

The experiments below are designed to validate the outer-loop reconstruction mechanism studied in this paper. They should not be interpreted as a comprehensive empirical comparison of complete multi-objective optimization solvers. The reason is that CATS operates on top of a scalarized inner solver, while the theory concerns the sampling and interpolation of the oracle-induced map $u \mapsto z(u)$. Analytic smooth test surfaces allow us to isolate this outer-loop question, compute or bound the curvature quantities entering the certificate, and evaluate all methods under the same stopping rule. Consequently, the reported gains support the claim that curvature-aware adaptive sampling can reduce the number of oracle calls in the smooth certified reconstruction regime. They do not imply universal superiority for arbitrary non-smooth, noisy, high-dimensional, or set-valued multi-objective optimization problems.

7.2. A Remark on Oracle-Query Counts Versus Full MOO Computational Cost

The computational study does not evaluate the wall-clock performance of complete multi-objective optimization solvers. In a full deterministic MOO pipeline, each query selected by the outer loop would require solving a scalarized optimization subproblem. Therefore, the total computational cost of a complete implementation would have the form

$$T_{\text{total}} = \sum_{u\in\mathcal{V}(\mathcal{T})} T_{\text{solve}}(u) + T_{\text{outer}}(\mathcal{T}), \tag{15}$$

where $T_{\text{solve}}(u)$ is the cost of solving the scalarized problem associated with the preference vector u, and $T_{\text{outer}}(\mathcal{T})$ is the cost of maintaining the triangulation, computing refinement indicators, updating certificates, and enforcing conformity.

In the experiments reported here, the oracle map $u \mapsto z(u)$ is available analytically. Consequently, the reported quantities are oracle-query counts, i.e., the number of distinct sampled preference vectors, not wall-clock runtimes of complete MOO solvers. This controlled setting is deliberate: it isolates the sampling and reconstruction question from solver-dependent effects such as convergence tolerances, local minima, non-uniqueness, infeasibility of scalarized subproblems, and numerical noise.

The query-count comparison is meaningful when scalarized solves dominate the outer-loop overhead and have comparable cost across preference vectors. In that case, reducing $|\mathcal{V}(\mathcal{T})|$ is expected to reduce the dominant part of the total computational burden. However, if the inner-solve costs $T_{\text{solve}}(u)$ vary substantially across the reference simplex, or if the solver does not return a smooth and reliable branch of Pareto outcomes, then oracle-query counts alone are not sufficient to predict wall-clock performance. A full runtime comparison with specific scalarization solvers is outside the scope of this paper and is left for future work.

7.3. Two Deterministic Tri-Objective Test Problems

We consider decision variables $x = (x_1, x_2) \in \mathbb{R}^2$ constrained to the simplex $\mathcal{X} = \{x_1 \geq 0,\ x_2 \geq 0,\ x_1 + x_2 \leq 1\}$. We define three objectives to minimize:

$$f_1(x) = x_1, \qquad f_2(x) = x_2, \qquad f_3(x) = g(x_1, x_2). \tag{16}$$

The first two objectives prefer small (x_1, x_2), while f_3 is constructed to induce trade-offs and strongly non-uniform curvature.

To isolate reconstruction, we identify preferences with the decision variables $u = (u_1, u_2) \in \Delta_2, x(u) = (u_1, u_2), u_3 := 1 - u_1 - u_2$. Then, the oracle map becomes

$$z(u) = (f_1(x(u)), f_2(x(u)), f_3(x(u))) = \big(u_1, u_2, g(u_1, u_2)\big), \tag{17}$$

so the (image) surface is the graph of g over the reference simplex.

For both problems below, g is strictly decreasing in each of u_1 and u_2 on Δ_2. Hence, if $u_1 \leq u_1'$ and $u_2 \leq u_2'$, then $g(u) \geq g(u')$, so no point can improve (f_1, f_2) without worsening f_3 (and vice versa). Therefore, all points in $z(\Delta_2)$ are mutually non-dominated, and the full image is the Pareto surface.

We consider the following test problems:

$$\text{P1)}\ \ g(u_1, u_2) = u_3^4 + 0.05\Big(e^{-5u_1} + e^{-5u_2}\Big). \tag{18}$$

$$\text{P2)}\ \ g(u_1, u_2) = u_3^4 + 0.03\Big(e^{-7u_1} + e^{-7u_2} + e^{-7(u_1+u_2)}\Big). \tag{19}$$

Both P1 and P2 have highly non-uniform curvature: the exponential terms create sharp features near the edges $u_1 \approx 0$ and $u_2 \approx 0$ (and near the vertex $(u_1, u_2) \approx (0, 0)$ in P2), while much of the interior is comparatively flat. This stresses the main failure mode of uniform sampling: it oversamples flat regions and undersamples the sharp regions that govern uniform reconstruction error.

7.4. How Certificates Are Computed in These Experiments

Only the third component $z_3(u) = g(u_1, u_2)$ is non-linear; $z_1(u) = u_1$ and $z_2(u) = u_2$ are affine and therefore interpolated exactly by $\hat{z}$ on any triangulation. Consequently,

$\|z(u) - \hat{z}(u)\|_\infty = |g(u) - \hat{g}(u)|$, and the reconstruction certificate reduces to controlling the scalar interpolation error for g.

For a triangle (2-simplex) $\tau \subset \Delta_2$, the theory in Section 4 requires an upper bound on the spectral norm of the Hessian $\sup_{u\in\tau} \|\nabla^2 g(u)\|_2 \le M_\tau$. In these controlled experiments, we compute a conservative bound as follows:

$$\tilde{M}_\tau = \max_{v\in \mathrm{verts}(\tau)} \|\nabla^2 g(v)\|_\infty, \qquad M_\tau = \sqrt{2}\,\tilde{M}_\tau, \tag{20}$$

and then use $\eta(\tau) = \frac{1}{2} M_\tau \,\mathrm{diam}(\tau)^2$. The conversion factor is valid because for 2×2 matrices, $\|A\|_2 \le \sqrt{2}\,\|A\|_\infty$, so M_τ upper-bounds $\sup_{u\in\tau} \|\nabla^2 g(u)\|_2$ whenever $\tilde{M}_\tau$ upper-bounds $\sup_{u\in\tau} \|\nabla^2 g(u)\|_\infty$.

For P1 and P2, each entry of $\nabla^2 g(u)$ is a sum of terms of the form $c\,u_3^2$ and $c\,e^{-b\ell(u)}$, where $\ell(u)$ is affine over τ. Both u_3^2 and $e^{-b\ell(u)}$ achieve their maxima over a simplex at a vertex (since u_3 and ℓ are affine and these functions are monotone in their arguments on Δ_2). Thus, for these test problems, $\sup_{u\in\tau} \|\nabla^2 g(u)\|_\infty$ is attained at a vertex, and (20) yields a valid cell-wise curvature bound.

7.5. Pareto-Surface Reconstructions and Analysis on the Results

All geometric interpretations in this section concern the Pareto image in objective space. No claim is made here about distances or continuity properties in the corresponding decision-variable space.

Figures 1 and 2 visualize the key output of each outer loop (shown for P1; P2 is similar): a triangulated surface $\hat{z}(\Delta_2)$.

In Figure 1, the transparent surface is the ground truth $z(\Delta_2)$. The wireframe is the reconstructed surface $\hat{z}(\Delta_2)$. Deviations indicate regions where the local certificate remains large and refinement is still needed.

In Figure 2, plots show where the algorithm spent oracle calls in the reference domain. Since $z(u) = (u_1, u_2, g(u))$, the reference coordinates are also the (f_1, f_2) coordinates. CATS concentrates triangles near edges where curvature is high, whereas UNIFORM distributes triangles evenly, and AREA-MAX follows objective-space patch size, which is correlated with but not equivalent to curvature-driven interpolation error.

Tables 2 and 3 report the number of oracle calls required to certify $\max_\tau \eta(\tau) \le \varepsilon$. Because all methods stop under the same certified criterion, the ratios UNIF/CATS and AREA/CATS quantify the outer-loop efficiency gain of curvature-aware refinement.

Table 2. P1: oracle calls required to certify $\max_\tau \eta(\tau) \le \varepsilon$.

ε	CATS	AREA-MAX	UNIFORM	UNIF/CATS	AREA/CATS
0.05	111	132	300	2.70	1.19
0.02	242	360	703	2.90	1.49
0.01	431	687	1378	3.20	1.59
0.005	815	1325	2701	3.31	1.63
0.002	1895	2820	6555	3.46	1.49

Table 3. P2: oracle calls required to certify $\max_\tau \eta(\tau) \le \varepsilon$.

ε	CATS	AREA-MAX	UNIFORM	UNIF/CATS	AREA/CATS
0.05	115	150	325	2.83	1.30
0.02	249	362	780	3.13	1.45

Table 3. *Cont.*

ε	CATS	AREA-MAX	UNIFORM	UNIF/CATS	AREA/CATS
0.01	439	689	1540	3.51	1.57
0.005	838	1341	3003	3.58	1.60
0.002	1936	2907	7381	3.81	1.50

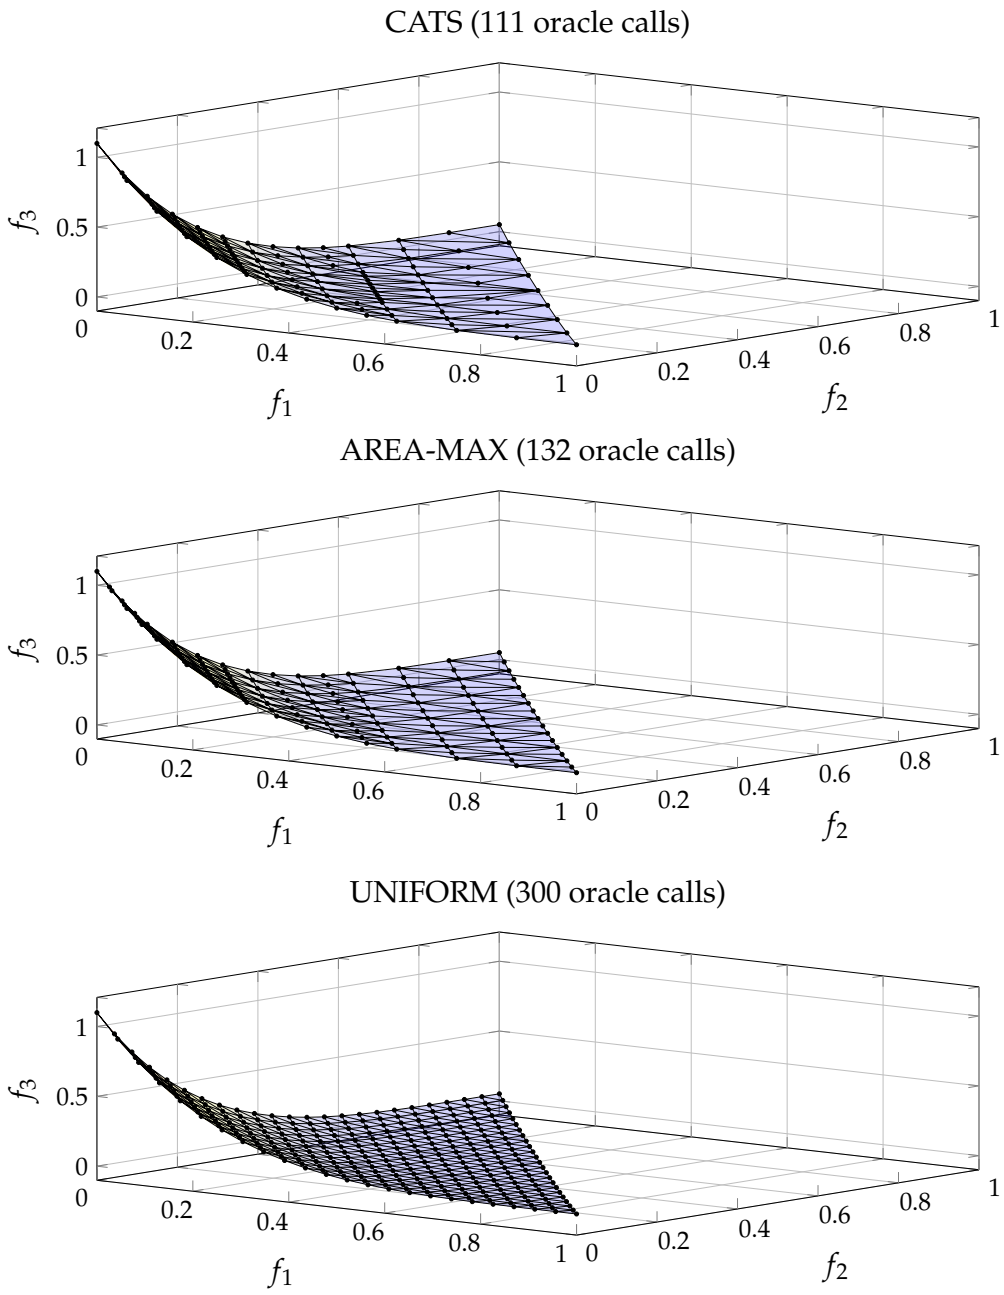


Figure 1. P1 at tolerance $\varepsilon = 0.05$: true Pareto surface (transparent) and reconstructed triangulated surfaces (wireframes) from CATS, AREA-MAX (AWS-inspired), and UNIFORM reference directions.

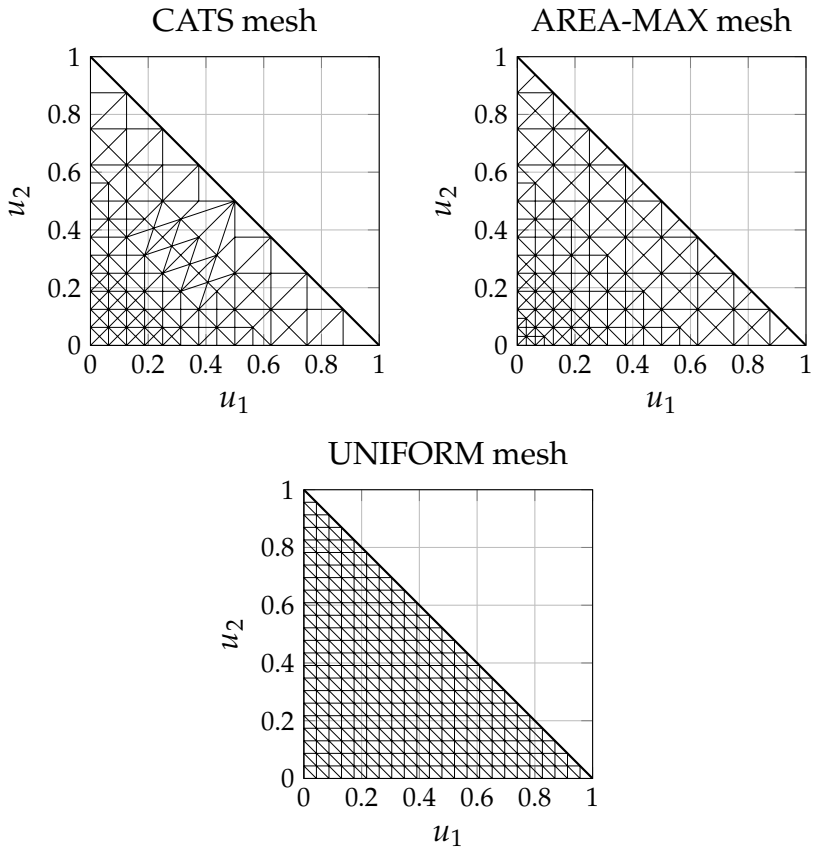


Figure 2. Same setting as Figure 1, shown in the reference domain. CATS refines near the high-curvature regions that dominate the interpolation certificate, which explains the oracle-call savings relative to uniform refinement.

Figure 3 visualizes the same data on log–log axes. Since $d = 2$ here, the theory predicts $N(\varepsilon) = \Theta(\varepsilon^{-1})$, and the curves are consistent with an approximately linear trend in $1/\varepsilon$ (up to method-dependent constants).

Figure 4 shows how the maximum local certificate decreases as oracle calls accumulate. This plot reveals whether an outer loop is efficiently attacking the current bottleneck. CATS reduces the dominant simplex error directly, whereas AREA-MAX may spend iterations refining large mapped patches that are already sufficiently flat from the interpolation viewpoint.

Across both test problems and across tolerances spanning more than an order of magnitude, CATS consistently reduces oracle calls compared to

- UNIFORM reference-direction sampling (often by factors around 3× in our tests), and
- AREA-MAX (often by factors around 1.2×–1.6×),

while providing the same uniform certified accuracy guarantee. The mesh visualizations explain these gains: CATS concentrates resolution where curvature makes the surface hard to approximate, which is exactly what governs uniform interpolation error.

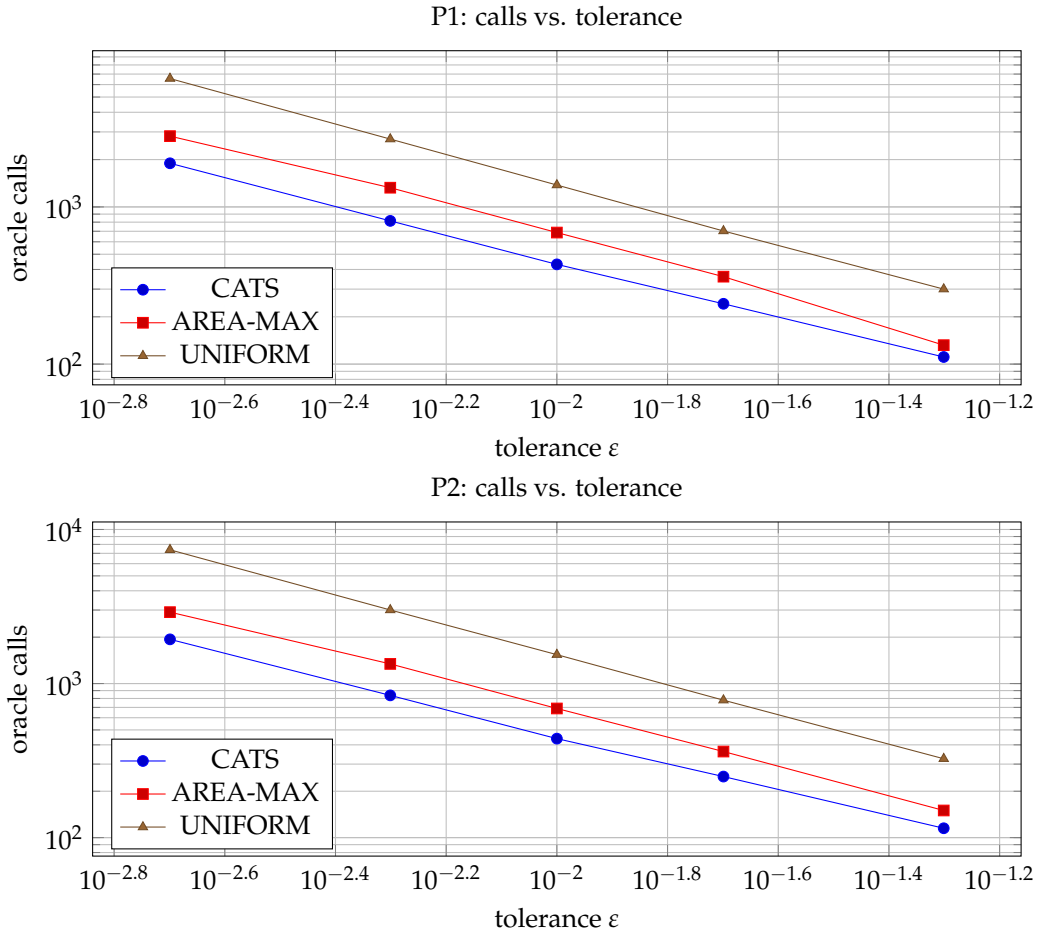


Figure 3. Certified oracle complexity on P1 and P2. CATS consistently reduces oracle calls compared to uniform reference directions and patch-area refinement.

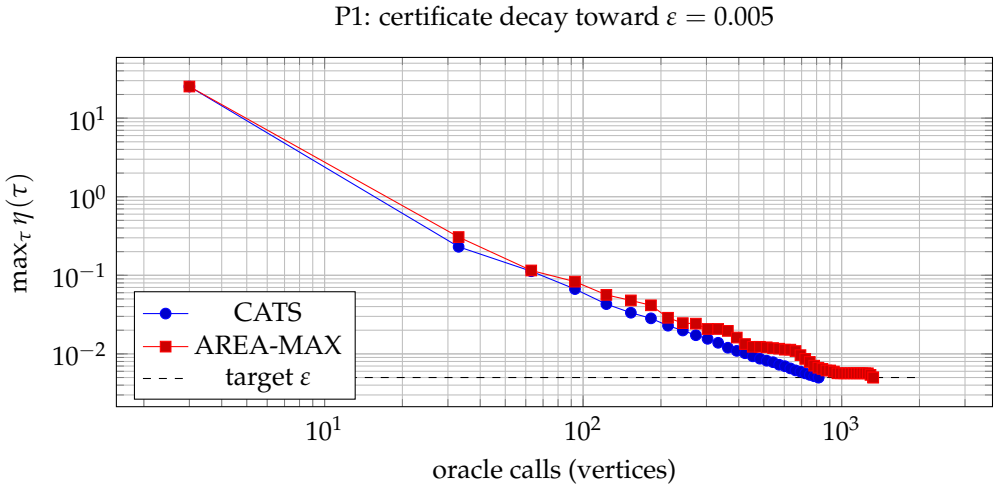


Figure 4. Certificate decay on P1. Both methods terminate when $\max_\tau \eta(\tau) \leq \varepsilon$, but CATS reaches that regime with fewer oracle calls.

7.6. Focused Comparison with an Objective-Space Triangulation Indicator

In order to provide a further comparison with an adaptive objective-space triangulation method, we also tested the HMW-DIAM rule introduced in Section 5.4 on the two rigorously certified core problems P1 and P2. This additional comparison is deliberately restricted to the certified core tests, where the curvature bounds are analytic and the stopping condition $\max_\tau \eta(\tau) \leq \varepsilon$ has the rigorous interpretation proved in Section 4.

Table 4 reports the oracle calls obtained by HMW-DIAM together with the corresponding CATS values. The results show that objective-space diameter is not a reliable surrogate for certified interpolation error. Although HMW-DIAM refines large gaps between sampled Pareto outcomes, it does not explicitly account for the local curvature of the map $u \mapsto z(u)$. Consequently, it may refine large but almost affine patches while under-prioritizing smaller high-curvature regions. CATS remains more efficient because it directly targets the local interpolation quantity that controls the certified uniform reconstruction bound.

Table 4. Focused objective-space triangulation comparison on the certified core tests. HMW-DIAM is an HMW-inspired objective-space diameter indicator, included as an additional robustness check against objective-space triangulation criteria.

Problem	ε	CATS	HMW-DIAM	UNIFORM	HMW/CATS
P1	0.05	111	276	300	2.49
P1	0.02	242	563	703	2.33
P1	0.01	431	1569	1378	3.64
P1	0.005	815	2188	2701	2.68
P1	0.002	1895	6147	6555	3.24
P2	0.05	115	344	325	2.99
P2	0.02	249	566	780	2.27
P2	0.01	439	1559	1540	3.55
P2	0.005	838	3069	3003	3.66
P2	0.002	1936	6133	7381	3.17

The table also shows that HMW-DIAM does not uniformly improve over UNIFORM. This is expected: objective-space diameter measures the spread of sampled outcomes, whereas the stopping criterion is governed by a curvature-weighted interpolation bound. Hence, a large objective-space patch may still be nearly affine, while a smaller patch may dominate the certified reconstruction error because of high local curvature.

7.7. Additional Benchmark Study on Smooth Tri-Objective Surfaces

To strengthen the empirical assessment, we added a second computational campaign on four additional smooth tri-objective parametric surface-reconstruction benchmarks. The purpose of these tests is twofold. First, they broaden the evidence beyond the two original examples. Second, they probe specific geometric difficulties—edge layers, localized hotspots, and narrow anisotropic ridges—for which the choice of the next reference direction is expected to matter.

All additional benchmarks are defined on the reference simplex $\Delta_2 = \{(u_1, u_2) \in \mathbb{R}^2 : u_1 \geq 0,\ u_2 \geq 0,\ u_1 + u_2 \leq 1\}$, $u_3 := 1 - u_1 - u_2$, and generate smooth tri-objective surfaces in objective space of graph form $z(u) = \big(u_1, u_2, g(u_1, u_2)\big)$. As in the original experiments, the first two components are affine, while the third component controls the surface geometry. The four additional test functions are

$$g_{\text{EdgeTriple}}(u) = u_3^4 + 0.02\Big(e^{-10u_1} + e^{-10u_2} + e^{-10u_3}\Big), \tag{21}$$

$$g_{\text{RidgeLayer}}(u) = u_3^4 + 0.04e^{-18u_1} + 0.024e^{-14u_2} + 0.09e^{-120(u_1+u_2-0.55)^2}, \tag{22}$$

$$g_{\text{CornerLayer}}(u) = u_3^4 + 0.025\Big(e^{-10u_1} + e^{-6u_2} + e^{-9(u_1+u_2)}\Big) + 0.12e^{-45(u_1^2+u_2^2)}, \tag{23}$$

$$g_{\text{LocalGaussian}}(u) = u_3^4 + 0.03e^{-14u_1} + 0.02e^{-8u_2} + 0.06e^{-\left(80(u_1-0.18)^2+120(u_2-0.08)^2\right)}. \tag{24}$$

The EdgeTriple surface contains three boundary layers and serves as a smooth multi-edge stress test. CornerLayer concentrates curvature near the simplex corner and introduces directional anisotropy.

A specific worst-case-type scenario for Pareto-surface reconstruction occurs when the surface is almost flat over most of the reference simplex but contains a small localized region of high curvature. In this case, a non-adaptive uniform lattice must refine the entire domain until the small difficult region is resolved, thereby spending many oracle calls in regions that are already accurately approximated. This scenario is also challenging for refinement rules based only on mapped patch size, because the region that dominates the interpolation error may occupy a small portion of the domain and may not correspond to the largest objective-space patch.

The benchmark $g_{\mathrm{LocalGaussian}}$ was included precisely to test this behavior. Its Gaussian term creates a localized high-curvature hotspot near $(u_1, u_2) = (0.18, 0.08)$, while the rest of the surface remains comparatively smoother. Thus, this test directly addresses the case in which only a small portion of the Pareto surface controls the global reconstruction error. Anticipating what will be discussed in the next paragraph, the results reported in Tables 5 and 6 show that CATS identifies this localized bottleneck and reaches the stopping threshold with substantially fewer oracle calls than UNIFORM and AREA-MAX. Specifically, on $g_{\mathrm{LocalGaussian}}$, CATS requires 83, 126, and 174 oracle calls for $\varepsilon = 0.08$, 0.05, and 0.03, respectively, compared with 289, 561, and 561 calls for UNIFORM and 108, 189, and 241 calls for AREA-MAX. This corresponds to savings of 3.22×–4.45× relative to UNIFORM and 1.30×–1.50× relative to AREA-MAX.

Table 5. Additional benchmark study: oracle calls required to satisfy the common stopping threshold $\max_\tau \eta(\tau) \le \varepsilon$. The ratios UNIF/CATS and AREA/CATS quantify the extra sampling effort required by the competing outer loops relative to CATS. Entries marked with † indicate that AREA-MAX hit the preset refinement budget before reaching the target stopping threshold.

Benchmark	ε	CATS	AREA-MAX	UNIFORM	UNIF/CATS	AREA/CATS
CornerLayer	0.08	79	111	153	1.94	1.41
CornerLayer	0.05	116	158	289	2.49	1.36
CornerLayer	0.03	186	260	561	3.02	1.40
EdgeTriple	0.08	87	103	289	3.32	1.18
EdgeTriple	0.05	130	132	561	4.32	1.02
EdgeTriple	0.03	199	210	561	2.82	1.06
LocalGaussian	0.08	83	108	289	3.48	1.30
LocalGaussian	0.05	126	189	561	4.45	1.50
LocalGaussian	0.03	174	241	561	3.22	1.39
RidgeLayer	0.10	167	304 †	289	1.73	–
RidgeLayer	0.08	189	304 †	289	1.53	–
RidgeLayer	0.05	305	304 †	561	1.84	–

Table 6. Additional benchmark study: ex post validation error E_{grid} on a dense simplex grid after termination. These values are not used for stopping; they are reported only to verify that the ranking induced by the common stopping indicator is consistent with the true interpolation error. Entries marked with † correspond to runs where AREA-MAX hit the preset budget before reaching the target stopping threshold.

Benchmark	ε	CATS	AREA-MAX	UNIFORM
CornerLayer	0.08	0.0187	0.0136	0.0187
CornerLayer	0.05	0.0075	0.0075	0.0060

Table 6. *Cont.*

Benchmark	ε	CATS	AREA-MAX	UNIFORM
CornerLayer	0.03	0.0054	0.0035	0.0035
EdgeTriple	0.08	0.0154	0.0106	0.0062
EdgeTriple	0.05	0.0072	0.0072	0.0025
EdgeTriple	0.03	0.0052	0.0052	0.0025
LocalGaussian	0.08	0.0098	0.0098	0.0082
LocalGaussian	0.05	0.0075	0.0071	0.0028
LocalGaussian	0.03	0.0049	0.0034	0.0028
RidgeLayer	0.10	0.0164	0.0162 †	0.0091
RidgeLayer	0.08	0.0132	0.0162 †	0.0091
RidgeLayer	0.05	0.0082	0.0162 †	0.0075

The $g_{\text{RidgeLayer}}$ benchmark provides an even more anisotropic version of the same stress test, because its difficult region is a narrow diagonal ridge rather than a compact hotspot. In that case, the limitation of patch-area refinement is particularly visible: AREA-MAX reaches the preset refinement budget before satisfying the stopping threshold, whereas CATS continues to reduce the largest local interpolation quantity by targeting the ridge-dominated error. These two tests clarify what happens when the surface is highly curved only in a small portion of the reference domain.

Experimental Protocol

The same three outer loops are compared as in the main study: CATS, AREA-MAX, and UNIFORM. All methods use the same refinement primitive, namely conforming longest-edge bisection, and are compared under the same stopping rule based on the quantity $\eta(\tau) = \frac{1}{2}\left(\max_k M_{\tau,k}\right)\text{diam}(\tau)^2$. For each leaf triangle τ, $M_{\tau,k}$ is estimated from the closed-form Hessian of the benchmark surface by taking the maximum componentwise spectral norm at the three triangle vertices and at the centroid. This yields a reproducible cell-wise curvature surrogate used to rank cells and define a common stopping threshold across methods. Each method is stopped when $\max_{\tau\in\mathcal{T}} \eta(\tau) \le \varepsilon$. In addition, after termination we compute an ex post validation error on a dense 30×30 simplex grid, $E_{\text{grid}} := \max_{u\in\mathcal{U}_{30}} \|z(u) - \hat{z}(u)\|_\infty$, where $\mathcal{U}_{30}$ denotes the validation grid. Thus, oracle calls measure the outer-loop cost under a common practical stopping rule, while E_{grid} provides an independent numerical check of the reconstructed surface accuracy.

For the especially difficult RidgeLayer test, AREA-MAX was additionally given a safeguard cap of 300 refinement rounds. Entries marked with † in the tables indicate that this budget was reached before the stopping threshold was met.

Table 5 confirms the same qualitative message as the original study, but now across a broader collection of surface geometries. CATS is consistently the most sample-efficient method among those that successfully reach the target. Across all successful comparisons, the UNIFORM/CATS ratio ranges from 1.53× to 4.45×, while the AREA/CATS ratio ranges from 1.02× to 1.50×. The gains over UNIFORM are largest on EdgeTriple and LocalGaussian, where a large portion of the simplex is relatively flat and global refinement wastes oracle calls. The gains over AREA-MAX are most visible on CornerLayer and LocalGaussian, where the difficult regions are localized and highly anisotropic.

The RidgeLayer benchmark is particularly revealing. Here, the hard region is a thin diagonal ridge superimposed on an otherwise smoother structure. Because the mapped patch area is only an indirect proxy for interpolation difficulty, AREA-MAX keeps spending refinement effort in triangles that are geometrically large but not necessarily responsible

for the stopping-indicator bottleneck. As a consequence, it fails to meet the requested threshold within the imposed budget, whereas CATS reaches the target with 167, 189, and 305 oracle calls for $\varepsilon = 0.10$, 0.08, and 0.05, respectively. UNIFORM eventually reaches the same tolerances, but only after substantially more global refinement.

Table 6 helps interpret the indicator-based results. In several cases, UNIFORM attains a smaller ex post grid error than CATS or AREA-MAX. This does not contradict the main conclusion. Rather, it reflects the fact that UNIFORM refines globally and can only stop at discrete lattice levels; hence, it often overshoots the requested threshold and returns a mesh that is much denser than necessary. By contrast, CATS typically stops much closer to the prescribed stopping threshold and therefore achieves the same target with far fewer oracle calls. From the perspective of the outer-loop reconstruction problem studied in this paper, this is precisely the desired behavior.

Figure 5 visualizes the scaling of oracle calls with the requested tolerance. The broad trend is consistent with the theory: all methods exhibit increasing cost as ε decreases, but the constant factor depends strongly on how effectively the outer loop identifies the local bottleneck. The curves show that CATS preserves its advantage over a range of tolerances rather than only at a single operating point. This is especially clear on EdgeTriple and LocalGaussian, where the gap to UNIFORM widens rapidly as the target becomes more stringent.

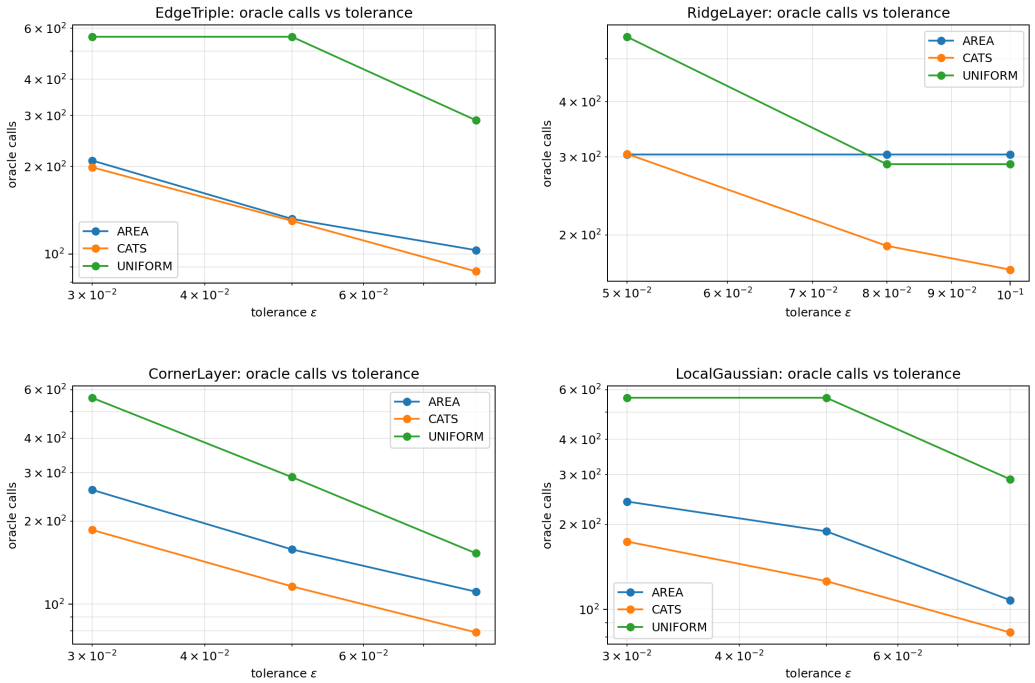


Figure 5. Additional benchmark study: oracle calls versus tolerance on the four new smooth tri-objective surfaces. On all benchmarks, CATS requires fewer oracle calls than UNIFORM; on the more anisotropic and localized cases, it also improves consistently over AREA-MAX.

Figure 6 illustrates the mechanism behind the gain on the hardest benchmark. CATS directly attacks the cell with the largest local interpolation indicator, so the maximum stopping quantity decreases steadily. AREA-MAX, in contrast, follows mapped patch size, which becomes misaligned with the true interpolation bottleneck once the narrow ridge dominates the error. This explains why the performance gap is modest on more benign surfaces, but becomes decisive when the difficult geometry is highly localized.

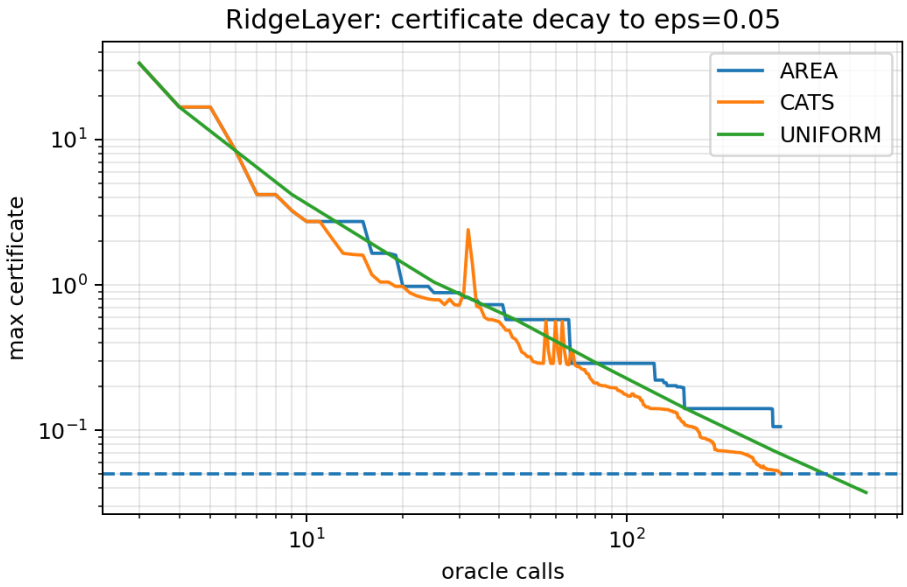


Figure 6. Decay of the maximum local stopping indicator on the challenging RidgeLayer benchmark for target threshold $\varepsilon = 0.05$ (the dashed line). CATS decreases the dominant local quantity substantially faster than AREA-MAX, which explains why AREA-MAX reaches the refinement budget without satisfying the stopping rule.

Overall, the additional benchmark campaign strengthens the empirical evidence in three ways. First, it confirms on four further smooth tri-objective surfaces that CATS is consistently more sample-efficient than global uniform refinement. Second, it shows that the advantage over AREA-MAX persists beyond the original examples and becomes larger on anisotropic or ridge-dominated geometries. Third, the combination of closed-form benchmarks, explicit indicator computation, and dense-grid ex post validation makes the study fully reproducible and closely aligned with the reconstruction objective of the paper.

7.8. Compact Standard-Family Benchmark on DTLZ-Style Tri-Objective Surfaces

To complement the analytical benchmarks used in the previous experiments, we propose a further experiment based on two DTLZ-style tri-objective surfaces inspired by the scalable benchmark suite of Deb et al. [32]. The first benchmark is DTLZ2-inspired, in the sense that it reproduces the spherical Pareto-front geometry of DTLZ2. The second is DTLZ4-inspired, in that it starts from the same spherical geometry but introduces a power-type reparameterization, mirroring DTLZ4's role as a modified DTLZ2 problem designed to test an algorithm's ability to maintain a good distribution of solutions on a biased front.

The purpose is to show that the same qualitative message also appears on a more standard benchmark family. We considered the following two parametric surfaces over the reference simplex Δ_2:

1. DTLZ2-simplex (near-isotropic control). Let $w(u) = (u_1,\ u_2,\ 1 - u_1 - u_2) \in \mathbb{R}^3_+, u \in \Delta_2$, and define $z_{\mathrm{DTLZ2}}(u) = \frac{w(u)}{\|w(u)\|_2}$. This yields a smooth spherical Pareto surface with relatively uniform curvature over the simplex.
2. DTLZ4-inspired-a4 (anisotropic standard-family variant). Using the same weight vector $w(u)$, define $v_i(u) = w_i(u)^\alpha, \alpha = 4$, and $z_{\mathrm{DTLZ4}}(u) = \frac{v(u)}{\|v(u)\|_2}$. This surface remains smooth but becomes significantly more anisotropic, with interpolation difficulty concentrated near the simplex corners.

The DTLZ2-simplex case acts as a control experiment: because the surface is nearly isotropic, a quasi-uniform allocation of points is already close to appropriate, and therefore UNIFORM is expected to be competitive. By contrast, the DTLZ4-inspired surface is designed precisely to create non-uniform approximation difficulty, which is the regime where a curvature-aware adaptive rule should be advantageous.

In both cases, all methods use the same triangulation framework and the same stopping rule based on the quantity $\eta(\tau)$, exactly as in the additional smooth benchmark study above. The comparison is therefore again based on the number of oracle calls required to reach the same stopping threshold.

On the near-isotropic DTLZ2-simplex control, the three methods perform similarly, which is consistent with the geometry of that surface and should be viewed as a useful sanity check rather than a limitation.

On the anisotropic DTLZ4-inspired-a4 surface, however, CATS clearly outperformed both competing strategies. At tolerance $\varepsilon = 0.12$, CATS required only 40 oracle calls, compared with 55 for UNIFORM and 126 for AREA-MAX (see Table 7). At $\varepsilon = 0.08$, CATS required 69 oracle calls, compared with 91 for UNIFORM and 176 for AREA-MAX (see Table 7). Thus, on this standard-family anisotropic surface, the advantage of adaptive indicator-based refinement is already substantial.

Table 7. DTLZ4-inspired-a4: oracle calls required to satisfy the stopping threshold.

ε	CATS	AREA-MAX	UNIFORM	UNIF/CATS	AREA/CATS
0.12	40	126	55	1.38	3.15
0.08	69	176	91	1.32	2.55

Figure 7 reports the corresponding oracle-call curves, while Figure 8 shows the speedup ratios with CATS as the denominator. Even in this additional experiment, the conclusion is consistent with the main computational section: when approximation difficulty is non-uniform, refining the simplex with the largest interpolation indicator is more effective than either global uniform refinement or patch-area refinement.

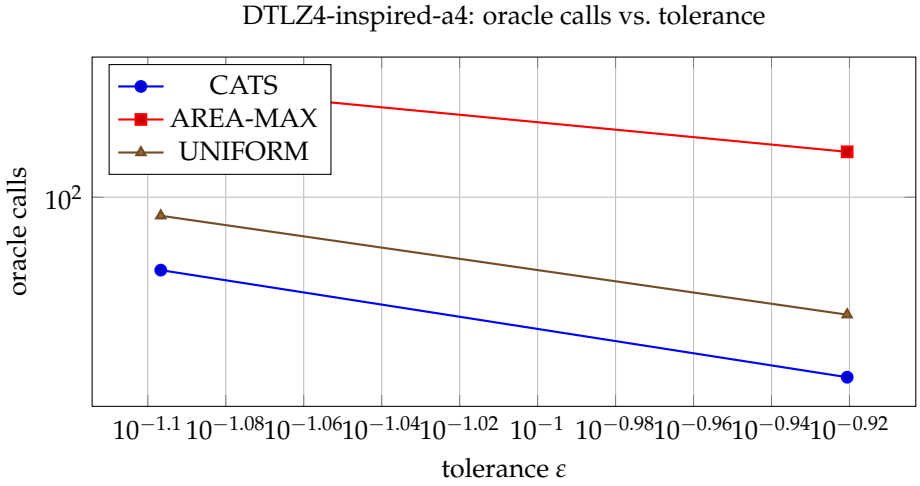


Figure 7. DTLZ4-inspired-a4: number of oracle calls required by CATS, AREA-MAX, and UNIFORM to reach the same stopping threshold.

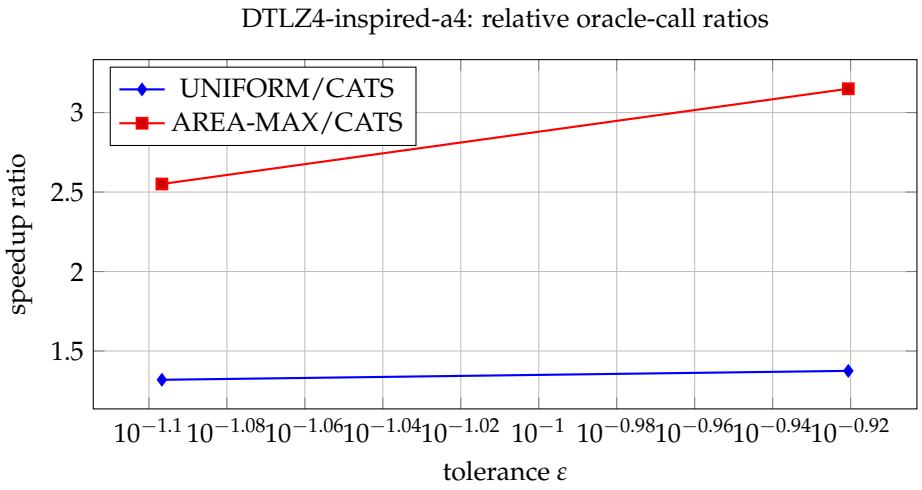
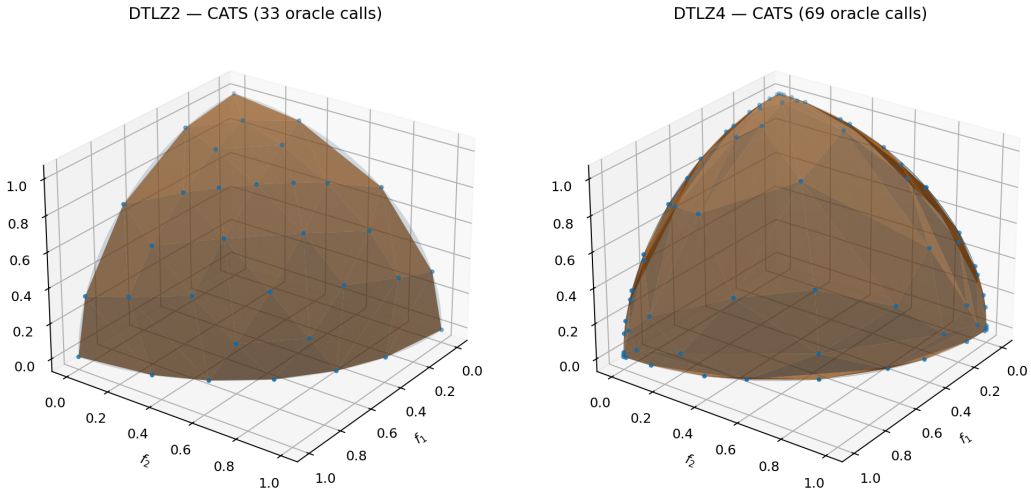


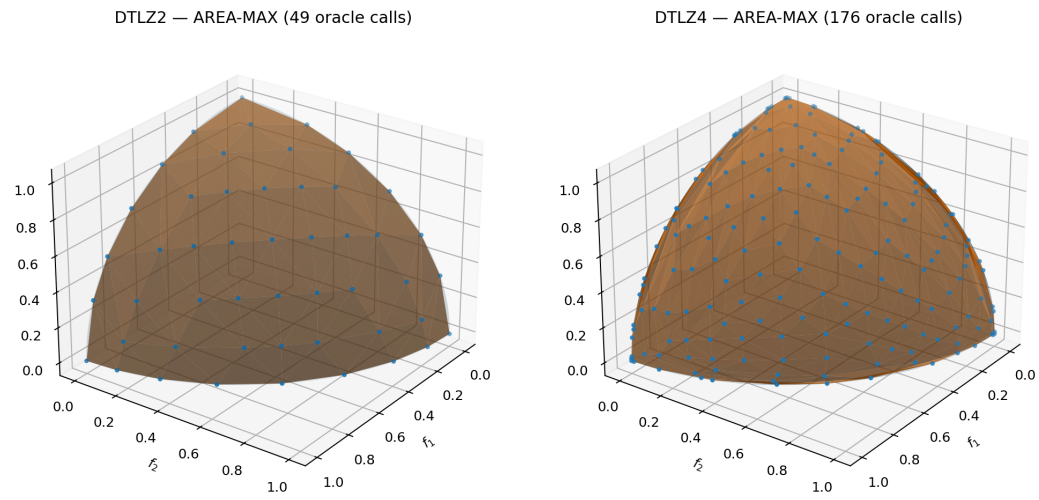
Figure 8. DTLZ4-inspired-a4: relative oracle-call ratios. Values above 1 indicate that the competing method required more oracle calls than CATS.

Figure 9 complements the quantitative results by showing the actual reconstructed triangulated surfaces for the three competing outer-loop strategies. The comparison is particularly informative because it juxtaposes a nearly isotropic case (DTLZ2), where uniform sampling is already well matched to the surface geometry, with a more anisotropic case (DTLZ4), where the benefit of curvature-aware refinement becomes visually evident. In each panel, the transparent surface represents the ground-truth surface, while the darker triangulated mesh is the reconstructed simplicial surrogate. The DTLZ2 panels illustrate a near-isotropic control case, for which all three methods produce qualitatively similar reconstructions. The DTLZ4 panels illustrate an anisotropic case, for which the benefit of refining according to the local interpolation indicator becomes visually clearer: CATS concentrates samples where the surface bends most strongly and achieves a comparably accurate reconstruction with fewer oracle calls than AREA-MAX and UNIFORM.

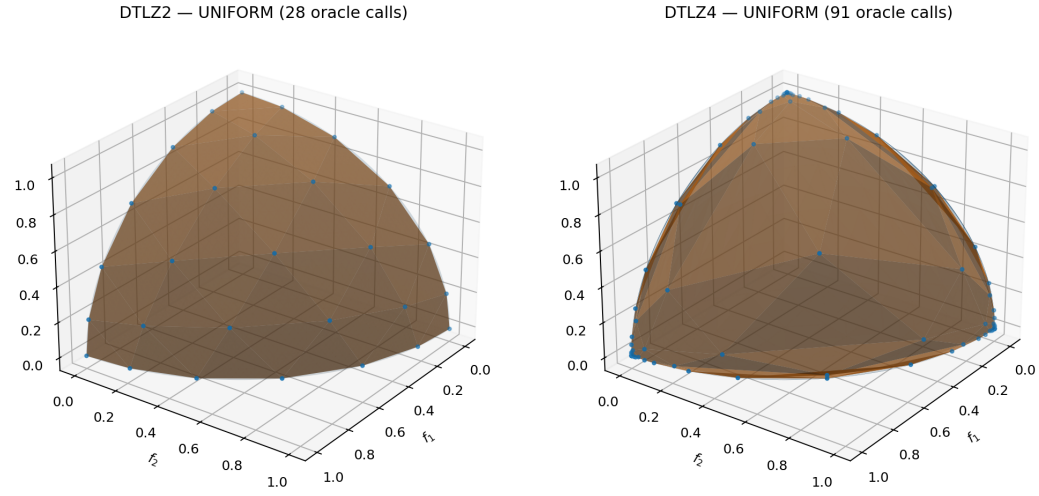


(**a**) DTLZ2 — CATS (**b**) DTLZ4 — CATS

Figure 9. *Cont.*



(**c**) DTLZ2 — AREA-MAX (**d**) DTLZ4 — AREA-MAX



(**e**) DTLZ2 — UNIFORM (**f**) DTLZ4 — UNIFORM

Figure 9. Triangulated reconstructions of two DTLZ-style tri-objective surfaces obtained with CATS, AREA-MAX, and UNIFORM. The DTLZ2 panels are shown at stopping threshold $\varepsilon = 0.03$, whereas the DTLZ4 panels are shown at stopping threshold $\varepsilon = 0.08$.

Overall, this supports a deliberately nuanced conclusion. CATS is not expected to dominate on every smooth benchmark surface: on nearly isotropic geometries, a uniform allocation of samples is already close to optimal. However, as soon as the approximation difficulty becomes anisotropic—as in the DTLZ4-inspired case—the benefit of refining according to the local interpolation indicator becomes clearly visible. This is precisely the regime targeted by the proposed method and is consistent with the main experiments reported in Section 7.

8. Conclusions

This paper studied the outer-loop reconstruction problem that arises in deterministic multi-objective optimization when Pareto outcomes are generated by repeatedly solving scalarized subproblems. For problems with $m \geq 3$ objectives, the sampled Pareto outcomes lie on an $(m-1)$-dimensional surface in objective space. In many applications, a finite cloud of non-dominated points is not sufficient: one needs a continuous representation of the Pareto surface that can support visualization, trade-off exploration, sensitivity analysis, and interactive decision making.

We proposed Certified Adaptive Triangulation Sampling (CATS), an adaptive outer-loop method for reconstructing a Pareto surface from an oracle mapping reference vectors to Pareto outcomes. CATS builds a conforming triangulation of the reference simplex and constructs a simplicial piecewise-affine surrogate of the Pareto surface. The refinement criterion is based on a Hessian-driven interpolation quantity of the form

$$\eta(\tau) = \frac{1}{2}\left(\max_{k} M_{\tau,k}\right)\mathrm{diam}(\tau)^2,$$

which reflects the classical error scale of affine interpolation for C^2 maps. By refining the simplex with the largest local quantity, CATS concentrates oracle calls in the regions where the surface is hardest to approximate.

The certified interpretation of the method applies in the smooth single-valued regime, namely when the preference-to-Pareto map is C^2 and reliable local upper bounds on its Hessian norms are available. Under these assumptions, we proved that the stopping condition

$$\max_{\tau \in \mathcal{T}} \eta(\tau) \leq \varepsilon$$

implies the global uniform reconstruction guarantee

$$\|z - \hat{z}\|_\infty \leq \varepsilon.$$

We also established an oracle-complexity upper bound of order $O(\varepsilon^{-d/2})$, where $d = m - 1$, and a matching minimax lower bound $\Omega(\varepsilon^{-d/2})$ for certified simplicial piecewise-affine reconstructions based on point queries. Thus, the exponent $d/2$ is information-theoretically tight for the reconstruction model considered in this paper. The advantage of adaptive sampling is therefore not a change in the asymptotic exponent, but a reduction in the constants by allocating oracle calls where they are most needed.

The computational results confirm the initial expectation motivating the method. On the rigorously certified core tests, CATS consistently required fewer oracle calls than both uniform reference-direction sampling and an AWS-inspired patch-area refinement rule. The reference-domain meshes and reconstructed surfaces show the mechanism behind these savings: CATS refines near localized high-curvature regions, whereas uniform sampling spends many oracle calls in nearly flat parts of the surface. The additional smooth benchmark surfaces and compact DTLZ-style tests further support this conclusion. In particular, the largest gains occur on anisotropic or localized geometries, where the interpolation difficulty is highly non-uniform. On nearly isotropic surfaces, the advantage of CATS is naturally smaller, because uniform sampling is already reasonably aligned with the geometry of the surface. This behavior is consistent with the theoretical interpretation of the method.

It is important to emphasize the scope of the proposed certificate. CATS provides a rigorous uniform-error guarantee only when valid local curvature bounds for the oracle-induced map are available. If these quantities are replaced by finite-difference, regression-based, or simulation-based estimates that are not themselves validated, then the same refinement rule remains meaningful as a practical adaptive indicator, but the resulting stopping test should not be interpreted as a formal certificate. Similarly, the theory assumes that the oracle induces a single-valued smooth map. Problems involving discontinuities, bifurcations, branch switching, non-smooth fronts, or genuinely set-valued Pareto correspondences require additional mathematical tools beyond the smooth interpolation framework developed here.

From an engineering implementation viewpoint, CATS is naturally positioned as an outer-loop sampling layer that can be coupled with existing deterministic design-

optimization workflows. In a CAD/CAE environment, each oracle call may correspond to a scalarized non-linear optimization problem, a finite-element simulation, a CFD computation, a multi-body model evaluation, or a tribological/lubrication model. Since such evaluations can be computationally expensive, reducing the number of required reference-vector queries while maintaining a prescribed reconstruction accuracy is practically relevant. In this sense, CATS is not intended to replace commercial CAD/CAE solvers or application-specific optimization modules; rather, it can guide which scalarized design evaluations should be performed next in order to reconstruct the Pareto surface efficiently.

This perspective is particularly relevant for mechanical and tribological design problems, where several physical quantities may interact non-linearly and conflict with one another. For example, in transmission systems, lubrication optimization, wear reduction, thermal behavior, efficiency, and service life may all contribute to the design trade-off. A reliable reconstruction of the Pareto surface can help engineers identify sensitive regions, compare design alternatives, and select robust compromises. The same principle applies to other simulation-driven engineering settings in which each objective evaluation is expensive and where a continuous description of the trade-off surface is more informative than a sparse set of non-dominated points.

Several directions remain open for future research. First, it would be important to develop certified or semi-certified curvature-estimation procedures when closed-form Hessian bounds are unavailable. This could involve validated sensitivity analysis of scalarized optimization problems, interval arithmetic, conservative Lipschitz estimates, or surrogate models equipped with rigorous error bounds. Second, future work should extend the framework to noisy, stochastic, or simulation-based oracles, where numerical discretization and solver tolerances must be incorporated into the reconstruction guarantee. Third, the treatment of non-smooth, discontinuous, or set-valued Pareto geometries remains an important challenge. In such cases, the appropriate notion of reconstruction may require tools from non-smooth analysis, set-valued approximation, or probabilistic certification. Finally, the integration of CATS into CAD/CAE and tribological design platforms represents a promising applied direction, especially for lubrication optimization, mechanical transmission design, and other engineering problems where adaptive Pareto-surface reconstruction could reduce computational effort while improving decision support.

Funding: This research received no external funding.

Data Availability Statement: The datasets generated and analyzed during the current study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Bukvić, M.; Vencl, A.; Milojević, S.; Skulić, A.; Gajević, S.; Stojanović, B. The Influence of Carbon Nanotube Additives on the Efficiency and Tribological Characteristics of Worm Gear Units. *Lubricants* **2025**, *13*, 327. [CrossRef]
2. Ma, Z.; Zhu, A.; Feng, H.; Zhou, C.; Liu, H.; Sheng, D.; Guo, H.; Su, C. Research on Path and Energy Consumption Control of Six-Axis Manipulator Based on Multi-Objective Optimization Method. *Int. Sci. Tech. Econ. Res.* **2024**, *2*, 96–113.
3. Wu, X.; Zhang, Y.; Lai, K.W.; Yang, M.Z.; Yang, G.; Wang, H. A Novel Centralized Federated Deep Fuzzy Neural Network with Multi-Objectives Neural Architecture Search for Epistatic Detection. *IEEE Trans. Fuzzy Syst.* **2025**, *33*, 94–107. [CrossRef]
4. Li, W.; Zhang, T.; Wang, R.; Huang, S.; Liang, J. Multimodal multi-objective optimization: Comparative study of the state-of-the-art. *Swarm Evol. Comput.* **2023**, *77*, 101253. [CrossRef]
5. Pereira, J.L.J.; Oliver, G.A.; Francisco, M.B.; Cunha, S.S., Jr.; Gomes, G.F. A review of multi-objective optimization: Methods and algorithms in mechanical engineering problems. *Arch. Comput. Methods Eng.* **2022**, *29*, 2285–2308. [CrossRef]
6. Sharma, S.; Kumar, V. A Comprehensive Review on Multi-objective Optimization Techniques: Past, Present and Future: S. Sharma, V. Kumar. *Arch. Comput. Methods Eng.* **2022**, *29*, 5605–5633. [CrossRef]

7. Wang, Z.; Pei, Y.; Li, J. A survey on search strategy of evolutionary multi-objective optimization algorithms. *Appl. Sci.* **2023**, *13*, 4643. [CrossRef]
8. Ehrgott, M. *Multicriteria Optimization*; Springer: Berlin/Heidelberg, Germany, 2005. [CrossRef]
9. Miettinen, K. *Nonlinear Multiobjective Optimization*; Springer: Berlin/Heidelberg, Germany, 1999. [CrossRef]
10. Das, I.; Dennis, J.E. Normal-Boundary Intersection: A New Method for Generating the Pareto Surface in Nonlinear Multicriteria Optimization Problems. *SIAM J. Optim.* **1998**, *8*, 631–657. [CrossRef]
11. Messac, A.; Ismail-Yahaya, A.; Mattson, C.A. The Normalized Normal Constraint Method for Generating the Pareto Frontier. *Struct. Multidiscip. Optim.* **2003**, *25*, 86–98. [CrossRef]
12. Kim, I.Y.; de Weck, O.L. Adaptive weighted sum method for multiobjective optimization: A new method for Pareto front generation. *Struct. Multidiscip. Optim.* **2006**, *31*, 105–116. [CrossRef]
13. Hoseinpoor, N.; Ghaznavi, M. A flexible objective-constraint approach and a new algorithm for constructing the Pareto front of multiobjective optimization problems. *Acta Math. Sci.* **2024**, *44*, 702–720. [CrossRef]
14. Hartikainen, M.; Miettinen, K.; Wiecek, M.M. Constructing a Pareto front approximation for decision making. *Math. Methods Oper. Res.* **2011**, *73*, 209–234. [CrossRef]
15. Evtushenko, Y.G.; Posypkin, M.A. Nonuniform covering method as applied to multicriteria optimization problems with guaranteed accuracy. *Comput. Math. Math. Phys.* **2013**, *53*, 144–157. [CrossRef]
16. Wu, W.t.; Yang, X.m. Reference-point-based branch and bound algorithm for multiobjective optimization. *J. Glob. Optim.* **2024**, *88*, 927–945. [CrossRef]
17. Boffadossi, R.; Leonesio, M.; Fagiano, L. ROBBO: An Efficient Method for Pareto Front Estimation with Guaranteed Accuracy. *arXiv* **2025**, arXiv:2506.18004. [CrossRef]
18. Klamroth, K.; Tind, J.; Wiecek, M.M. Unbiased approximation in multicriteria optimization. *Math. Methods Oper. Res.* **2003**, *56*, 413–437. [CrossRef]
19. Rennen, G.; van Dam, E.R.; den Hertog, D. Enhancement of sandwich algorithms for approximating higher-dimensional convex Pareto sets. *Inf. J. Comput.* **2011**, *23*, 493–517. [CrossRef]
20. Löhne, A.; Rudloff, B.; Ulus, F. Primal and dual approximation algorithms for convex vector optimization problems. *J. Glob. Optim.* **2014**, *60*, 713–736. [CrossRef]
21. Ararat, Ç.; Ulus, F.; Umer, M. A Norm Minimization-Based Convex Vector Optimization Algorithm. *J. Optim. Theory Appl.* **2022**, *194*, 681–712. [CrossRef]
22. Keskin, İ.N.; Ulus, F. Outer approximation algorithms for convex vector optimization problems. *Optim. Methods Softw.* **2023**, *38*, 723–755. [CrossRef]
23. Pascoletti, A.; Serafini, P. Scalarizing vector optimization problems. *J. Optim. Theory Appl.* **1984**, *42*, 499–524. [CrossRef]
24. Zhou, Z.; Huang, F.; Ansari, Q.H. Pascoletti-Serafini Scalarizations for Approximate Quasi Efficient Solutions for Multiobjective Optimization Problems. *J. Optim. Theory Appl.* **2025**, *207*, 61. [CrossRef]
25. Das, M.; Nahak, C.; Biswal, M.P. Treatment of set-valued robustness via separation and scalarization. *J. Optim. Theory Appl.* **2024**, *201*, 843–865. [CrossRef]
26. Linh, H.M.; Xuan Dai, L.; Tam, T.N.; Chi Thang, N. Stability of solutions to set-valued equilibrium problems via scalarization. *Optimization* **2025**, *78*, 2003–2024. [CrossRef]
27. Dontchev, A.L.; Rockafellar, R.T. *Implicit Functions and Solution Mappings: A View from Variational Analysis*, 2nd ed.; Springer: Berlin/Heidelberg, Germany, 2014.
28. Bonnans, J.F.; Shapiro, A. *Perturbation Analysis of Optimization Problems*; Springer: Berlin/Heidelberg, Germany, 2000.
29. Fiacco, A.V. *Introduction to Sensitivity and Stability Analysis in Nonlinear Programming*; Academic Press: New York, NY, USA, 1983.
30. Maubach, J.M. Local Bisection Refinement for N-Simplicial Grids Generated by Reflection. *SIAM J. Sci. Comput.* **1995**, *16*, 210–227. [CrossRef]
31. Perdomo, F.; Plaza, A. Properties of triangulations obtained by the longest-edge bisection. *Cent. Eur. J. Math.* **2014**, *12*, 1796–1810. [CrossRef]
32. Deb, K.; Thiele, L.; Laumanns, M.; Zitzler, E. Scalable Test Problems for Evolutionary Multiobjective Optimization. In *Evolutionary Multiobjective Optimization: Theoretical Advances and Applications*; Goldberg, A.J., Ed.; Springer: Berlin/Heidelberg, Germany, 2005; pp. 105–145.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.