# Picking a Representative Set in Multiobjective Optimization: Axioms, Algorithms, and Experiments

**Niclas Boehmer**[1] , **Maximilian T. Wittmann**[1] ,
[1]Hasso-Plattner-Institut, Potsdam, Germany
{niclas.boehmer, maximilian.wittmann}@hpi.de,

## Abstract

Many real-world decision-making problems involve optimizing multiple objectives simultaneously, rendering the selection of the most preferred solution a non-trivial problem: All Pareto optimal solutions are viable candidates, and it is typically up to a decision maker to select one for implementation based on their subjective preferences. To reduce the cognitive load on the decision maker, previous work has introduced the Pareto pruning problem, where the goal is to compute a fixed-size subset of Pareto optimal solutions that best represent the full set, as evaluated by a given quality measure. Reframing Pareto pruning as a multiwinner voting problem, we conduct an axiomatic analysis of existing quality measures, uncovering several unintuitive behaviors. Motivated by these findings, we introduce a new measure, *directed coverage*. We also analyze the computational complexity of optimizing various quality measures, identifying previously unknown boundaries between tractable and intractable cases depending on the number and structure of the objectives. Finally, we present an experimental evaluation, demonstrating that the choice of quality measure has a decisive impact on the characteristics of the selected set of solutions and that our proposed measure performs competitively or even favorably across a range of settings.

## 1 Introduction

A core challenge in multiobjective optimization is the absence of a single, objectively best solution. Instead, all solutions on the Pareto front, i.e., solutions that are not (weakly) outperformed by another solution in every objective, are viable options for implementation, with each of them reflecting a different tradeoff among the objectives. To resolve this, the multiobjective literature assumes the presence of a decision maker (DM) who selects a final solution to be implemented based on their subjective preferences. However, in practice, the Pareto front is often very large, making it cognitively infeasible for the DM to process all solutions effectively. This motivates the study of the *Pareto pruning* problem: Compute a fixed-size subset of Pareto optimal solutions that represents the overall structure and available tradeoffs of the full Pareto front well [Vaz *et al.*, 2015; Petchrompo *et al.*, 2022; Sayin, 2000]. Two widely discussed measures that evaluate the quality of such a subset are *uniformity*, which maximizes the minimum distance between any two selected solutions, and *coverage*, which minimizes the maximum distance from any non-selected solution to its nearest selected neighbor [Sayin, 2000].

Despite the wide range of studied measures, systematic comparisons of their formal properties and a comprehensive analysis of their computational complexity remain largely absent from the literature. For our analysis, we approach the Pareto pruning problem through the lens of social choice theory. One of the most actively studied problems in social choice is multiwinner voting, where the goal is to select a subset of $k$ candidates based on the preferences of voters over a given set of candidates. We observe a natural analogy: solutions correspond to candidates, and each objective constitutes a voter evaluating solutions according to their performance. This connection lends itself to an axiomatic analysis of quality measures, a core method in social choice theory. The common distinction between ordinal, cardinal, and approval preferences in social choice further motivates an analysis of the Pareto pruning problem under analogous restrictions on the objectives.

### 1.1 Contributions

We present a systematic study of quality measures for Pareto pruning in multiobjective problems, taking a holistic perspective by integrating axiomatic, algorithmic, and experimental analyses. Our goal is to contribute formal structure and arguments to the discussion around quality measures that have traditionally remained quite fragmented. Our analysis connects multiobjective optimization, social choice, and computational geometry. Besides considering $\ell_1$-variants of the widely studied uniformity and coverage measures, we introduce a novel measure, *directed coverage*.

In Section 2, we propose five axioms capturing whether solutions reflecting distinct tradeoffs are guaranteed to be selected (standout consistency and outlier consistency), and how the selected subset responds to changes in solutions' performance or the addition of new solutions (extremism monotonicity, monotonicity, and $\varepsilon$-split-proofness). Motivated by shortcomings of uniformity and coverage revealed in our ax-

iomatic analysis, we introduce *directed coverage*, a measure that, like coverage, aims to cover all solutions well, but evaluates how well a solution $s$ covers another solution $a$ not by their $\ell_1$-distance but by the summed extent to which $a$ outperforms $s$, i.e., $\|a - s\|_+ = \sum_i \max(a_i - s_i, 0)$.

In Section 3, we show that selecting a subset that optimizes uniformity or coverage is NP-hard for three objectives, thereby identifying the precise boundary of tractability left open by Vaz *et al.* [2015]. As a byproduct, we fill a surprising gap in the computational geometry literature by proving NP-hardness of DISCRETE $k$-CENTER for the $\ell_1$-distance in two dimensions. We also show that approval objectives render all three pruning problems tractable for any constant number of objectives.

In Section 4, we observe experimentally that the three measures yield distinctly different slates, and that directed coverage introduces a new perspective, resulting in the selection of more efficient solutions.

The full version of this work, code, and experimental results are publicly available.[1]

## 2 Quality Measures & Axioms

We study the Pareto pruning problem, where given a set of Pareto optimal alternatives $A \subseteq \mathbb{R}^d$ and an integer $k$, we select a size-$k$ slate $S \subseteq A$ to present to a DM. Three natural desiderata are **diversity** ($S$ should be redundancy-free), **representativity** (each non-selected alternative should be close to some selected one), and **efficiency** ($S$ should contain high-quality alternatives). *Uniformity* $I_U(S) = \min_{x,y \in S} \|x - y\|_1$ (maximized) is inspired by diversity. *Coverage* $I_C(S, A) = \max_{a \in A} \min_{s \in S} \|a - s\|_1$ (minimized) captures representativity. Our new measure *directed coverage* $I_{DC}(S, A) = \max_{a \in A} \min_{s \in S} \|a - s\|_+$ (minimized) quantifies the total efficiency loss of presenting $s$ rather than $a$. Unlike coverage, it does not penalize $s$ for objectives in which $s$ outperforms $a$, making it inherently sensitive to solution quality.

We conduct an axiomatic analysis of these three measures (Table 1). Both uniformity and coverage violate *monotonicity*: improving a solution can reduce its Manhattan distance to others, causing it to be dropped from the slate. Directed coverage satisfies monotonicity, as improving a solution $x$ can only increase $\|x - z\|_+$ for any $z$. Only uniformity satisfies *$\varepsilon$-split proofness* (one should never select two arbitrarily close alternatives). Directed coverage uniquely satisfies *standout consistency*: a solution whose weakest aggregate lead over any other exceeds any other's strongest lead against it is always selected. Coverage uniquely satisfies *outlier consistency*: an alternative further away from every other alternative than any of the remaining solutions are from each other is always selected for $k \geq 2$.

## 3 Algorithmic Analysis

We conduct a thorough algorithmic analysis of computing the optimal pruning under the three quality measures (see Table 2). While Vaz *et al.* [2015] established that Pareto pruning for uniformity and coverage is solvable in polynomial time for two objectives, the complexity for more than two objectives had been left open. Building on results from computational geometry, we prove NP-hardness for uniformity and coverage for three objectives, thereby identifying the precise boundary of tractability. Along the way, we present a proof of NP-hardness for the classic DISCRETE $k$-CENTER problem for the $\ell_1$-distance in two dimensions, which has surprisingly been missing from the computational geometry literature. We further extend the algorithmic analysis to directed coverage and explore variants for ordinal (ranking-based) and approval (binary) objectives. While we cannot observe a difference in complexity when moving from cardinal to ordinal objectives, approval objectives render all three pruning problems solvable in polynomial time for any constant number of objectives.

[1]Full version: arxiv.org/abs/2511.10716, Code & Experimental Results: github.com/maxitw/picking_representative_moo

| | Monotonicity | $\varepsilon$-Split Proofness | Extremism Monoton. | Standout Consist. | Outlier Consist. |
|---|---|---|---|---|---|
| Uniformity | ✗ | ✓ | ✓ | ✗ | ✗ |
| Coverage | ✗ | ✗ | ✗ | ✗ | ✓ |
| Dir. Cov. | ✓ | ✗ | ✗ | ✓ | ✗ |

Table 1: Axiomatic results. ✓ = satisfied, ✗ = violated.

| Measure | #Obj. | Cardinal | Ordinal | Approval |
|---|---|---|---|---|
| Uniformity | $d = 2$ | P† | P† | P† |
| | fixed $d \geq 3$ | NP-hard | ? | P |
| | unbounded $d$ | NP-hard | NP-hard | NP-hard |
| Coverage | $d = 2$ | P† | P† | P† |
| | fixed $d \geq 3$ | NP-hard | ? | P |
| | unbounded $d$ | NP-hard | NP-hard | NP-hard |
| Dir. Cov. | $d = 2$ | P | P | P |
| | fixed $d \geq 3$ | NP-hard | ? | P |
| | unbounded $d$ | NP-hard | NP-hard | NP-hard |

Table 2: Computational results. †: Vaz *et al.* [2015].

## 4 Experiments

We evaluate slates from all three solution methods on benchmark datasets ZDT [Zitzler *et al.*, 2000] (six instances, two objectives), DTLZ [Deb *et al.*, 2002] (seven instances, three objectives), and PGMORL [Xu *et al.*, 2020] (six instances, two objectives, simulated agents). We compute slates via ILP using Gurobi and evaluate them under five measures: uniformity, coverage, directed coverage, hypervolume, and average summed objective value.

The three solution methods yield substantially different slates. Uniformity and coverage behave more similarly to each other than either does to directed coverage: both produce slates that are typically over 50% worse under the directed coverage measure than the dedicated method. With respect to the efficiency-oriented measures (hypervolume and average summed objective), directed coverage consistently outperforms the other two methods on ZDT and PGMORL, confirming that it tends to select more efficient solutions—as also suggested by its axiomatic properties.

## References


K. Deb, L. Thiele, M. Laumanns, and E. Zitzler. Scalable multi-objective optimization test problems. In *Proceedings of the 2002 Congress on Evolutionary Computation (CEC '02)*, volume 1, pages 825–830, 2002.

Sanyapong Petchrompo, David W. Coit, Alexandra Brintrup, Anupong Wannakrairot, and Ajith Kumar Parlikad. A review of pareto pruning methods for multi-objective optimization. *Comput. Ind. Eng.*, 167:108022, 2022.

Serpil Sayin. Measuring the quality of discrete representations of efficient sets in multiple objective mathematical programming. *Math. Program.*, 87(3):543–560, 2000.

Daniel Vaz, Luís Paquete, Carlos M. Fonseca, Kathrin Klamroth, and Michael Stiglmayr. Representation of the nondominated set in biobjective discrete optimization. *Comput. Oper. Res.*, 63:172–186, 2015.

Jie Xu, Yunsheng Tian, Pingchuan Ma, Daniela Rus, Shinjiro Sueda, and Wojciech Matusik. Prediction-guided multiobjective reinforcement learning for continuous robot control. In *Proceedings of the 37th International Conference on Machine Learning (ICML '20)*, pages 10607–10616. PMLR, 2020.

Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. *Evol. Comput.*, 8(2):173–195, June 2000.