

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

Niclas Boehmer¹, Maximilian T. Wittmann¹

¹Hasso Plattner Institute, University of Potsdam, Germany
niclas.boehmer@hpi.de, 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.

Code —

https://github.com/maxitw/picking_representative_moo

Extended version — <https://arxiv.org/abs/2511.10716>

1 Introduction

Many real-world decision-making problems in domains such as systems design, engineering, operations research, and healthcare are inherently multiobjective (Stewart et al. 2008; Marler and Arora 2004; Eriskin, Karatas, and Zheng 2024). As a result, multiobjective optimization has become a central research area (Branke et al. 2008; Ehrgott 2005), and multiobjective variants of many classical algorithmic techniques, including reinforcement learning (Hayes et al. 2022), integer programming (Sylva and Crema 2007), scheduling (Guo et al. 2013), and flows (Eusébio, Figueira, and Ehrgott 2014), have been intensively studied.

A key 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 preferences. Herein, a canonical approach is to first compute the Pareto front and then present it to the DM for selection (see the full version for a discussion of alternative approaches and additional background on multiobjective optimization). However, in practice, the Pareto front is often very large, making it cognitively infeasible for the DM to process all solutions and compare them effectively. This motivates the study of the *Pareto pruning* problem (also known as the representation 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; Taboada and Coit 2007; Petchrompo, Wannakrairot, and Parlikad 2022; Sayin 2000; Zio and Bazzo 2011; Wang et al. 2020; Taboada et al. 2007; Eusébio, Figueira, and Ehrgott 2014). A wide range of quality measures have been proposed to evaluate the selected subset (Li and Yao 2019; Faulkenberg and Wiecek 2010), each inducing a different solution method by selecting the subset optimizing the measure. Two widely used measures are *uniformity*, which aims to maximize 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. Existing comparative work (e.g., Li and Yao (2019) and Faulkenberg and Wiecek (2010)) typically groups quality measures into different categories based on soft criteria or compares them experimentally. Prior algorithmic work has mostly considered the case of two objectives (Vaz et al. 2015) or has focused on heuristic and evolutionary methods (Taboada and Coit 2008; Petchrompo et al. 2022).

For our analysis, we approach the Pareto pruning problem through the lens of social choice theory, also known as the theory of collective decision making. 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 to the multiobjective optimization problem correspond to candidates, and each objective constitutes a voter, who evaluates the solutions according to their performance in the respective objective. This connection lends itself to an axiomatic analysis of quality measures, a core method in social choice theory, to enable a structured comparison between them. Second, the common distinction between ordinal, cardinal, and approval preferences in social choice motivates an analysis of the Pareto pruning problem under analogous assumptions about the structure of the objectives, as ordinal and approval objectives are generally easier to elicit, especially if objectives correspond to different human evaluators.

This work also offers a new perspective on multiwinner voting that complements the three classical paradigms of *individual excellence*, *proportional representation*, and *diversity* (Faliszewski et al. 2017). The goal of Pareto pruning is distinct from these three in that the focus lies on “satisfying” the candidates, i.e., solutions, and not the voters, i.e., objectives. In Pareto pruning, voters are merely used to assess the similarity between two candidates, i.e., two candidates are considered close if they are evaluated similarly by all voters. Uniformity then seeks a set of mutually dissimilar candidates, while coverage aims to ensure that every non-selected candidate is close to at least one selected candidate. To the best of our knowledge, the only prior work in the voting space adopting a somewhat similar perspective is that of Delemazure et al. (2024), who, in the spirit of uniformity, consider the problem of selecting two distant candidates.

1.1 Our 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 to the discussion around quality measures that have traditionally remained fragmented and disconnected across different approaches. Our analysis connects three traditionally distinct areas: multiobjective optimization, social choice, and computational geometry. Besides considering ℓ_1 -variants of the widely studied uniformity and coverage measures, we also introduce a novel measure, directed coverage. Proofs and additional experimental results can be found in the full version.

In Section 3, we discuss different desiderata for Pareto pruning and initiate the axiomatic analysis of quality measures. 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 ε -split proofness). An overview of which measures satisfy which axioms is provided in Table 1. Motivated by the shortcomings of uniformity and coverage revealed in our axiomatic analysis, we introduce *directed coverage*, a measure that, like coverage, aims to cover all solutions well, but, unlike coverage, evaluates how well a solution a covers another solution b not by their ℓ_1 -distance but by the summed extent to which b outperforms a . Unlike uniformity and coverage, directed coverage guarantees, for instance, that

a selected solution will continue to get selected in case it improves its performance.

In Section 4, we conduct a thorough algorithmic analysis of Pareto pruning under the three considered quality measures (see Table 2). Vaz et al. (2015) established that Pareto pruning for uniformity and coverage is solvable in polynomial time for two objectives, but left the complexity for more than two and even an arbitrary number of objectives open.¹ 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 for the NP-hardness of the classic DISCRETE k -CENTER problem for the ℓ_1 -distance in two dimensions, which has surprisingly been missing from the literature. We extend the algorithmic analysis to our new *directed coverage* measure and explore variants of the problem for different practical restrictions on the type of information provided by each objective, i.e., besides cardinal (score-based) objectives, we also explore ordinal (ranking-based) and approval (binary) ones. While we are unable to observe a difference in complexity when moving from cardinal to ordinal objectives, we find that approval objectives render all three pruning problems solvable in polynomial time for any constant number of objectives.

In Section 5, we conduct an experimental analysis of our three considered measures, observing that each yields distinctly different results, and that optimizing for directed coverage introduces a new perspective, resulting in the selection of slightly more efficient solutions.

2 Preliminaries

For some $n \in \mathbb{N}$, let $[n] := \{1, \dots, n\}$. In a d -dimensional multiobjective optimization problem, a finite set X of alternatives is evaluated by $d \in \mathbb{N}$ objective functions $f_i : X \rightarrow \mathbb{R}$ for $i \in [d]$, where $f_i(x) < f_i(y)$ for two alternatives x and y means that y *outperforms* x under the i -th objective. The overarching goal of a multiobjective optimization problem is to maximize all objectives simultaneously, that is, to analyze $\max_{x \in X} (f_1(x), \dots, f_d(x))$. We write $f : X \rightarrow \mathbb{R}^d$, $f(x) = (f_1(x), \dots, f_d(x))$ for the function f aggregating all objectives into the *objective space* \mathbb{R}^d . For two alternatives $x, y \in X$, we say that x is *dominated* by y if $f_i(x) \leq f_i(y)$ for all $i \in [d]$ and there exists $j \in [d]$ with $f_j(x) < f_j(y)$. In addition, we say x is *Pareto dominated* if there exists some $y \in X$ such that x is dominated by y . Otherwise, we call x *Pareto optimal*.

For $i \in [d]$, an objective f_i is called an *approval objective* if $f_i(x) \in \{0, 1\}$ for all $x \in X$, i.e., each alternative is either approved or disapproved by the objective, and an *ordinal objective* if f_i is a bijection from X to $[|X|]$, i.e., f_i arranges all alternatives from X in a strict ranking. We refer to the general, unrestricted case as a *cardinal objective*.

Pareto dominated alternatives are of little importance to a DM, since there is a strictly better option available. Accordingly, we will only implicitly assume the existence of

¹This focus on few objectives reflects that many classical multiobjective problems involve only two to four objectives (Marler and Arora 2004; Branke et al. 2008; Ehrgott 2005).

X and instead operate directly on the set of Pareto optimal alternatives, i.e., we “preprocess” our instances to only include Pareto optimal alternatives. Similarly, we will only implicitly assume the existence of f_i and instead treat each alternative as a point in \mathbb{R}^d with its i -th component denoting its value according to f_i . Formally, as input to our problem, we receive the *set of Pareto optimal alternatives* $A = \{f(x) \mid x \in X \wedge x \text{ is Pareto optimal}\} \subseteq \mathbb{R}^d$, to which we will refer as *alternatives* for short. Our goal is to “inform” the DM about A by selecting k alternatives from A for some given $k \in \mathbb{N}$. We call a subset $S \subseteq A$ with $|S| = k$ a *slate*.

To measure the similarity between two alternatives $x, y \in A$, we use the Manhattan norm $\|x - y\| = \sum_{i=1}^d |x_i - y_i|$. Intuitively, two alternatives that are close to each other present similar tradeoff decisions. We further introduce a “directed” variant of the Manhattan norm $\|x - y\|_+ = \sum_{i=1}^d \max(x_i - y_i, 0)$. Note that $\|\cdot\|_+$ is not a metric, as it is not symmetric, i.e., we generally have $\|x - y\|_+ \neq \|y - x\|_+$.

We will use different *measures* to evaluate the quality of a slate $S \subseteq A$. We refer to a generic measure as I , where it will always be clear from context whether lower or higher values of I are preferable. For some set of alternatives A and integer k , we let $\mathcal{S}(I, A, k)$ be the set of slates which are optimal according to measure I , i.e., subsets $S \subseteq A$ with $|S| = k$ that maximize (resp. minimize) the value of I .

3 Pareto Pruning:

Problem, Quality Measures, and Axioms

We present a general formulation of the Pareto pruning problem in Section 3.1, the three quality measures we examine in Section 3.2, and our axiomatic analysis in Section 3.3.

3.1 Problem Setting and Desiderata

We study the Pareto pruning problem, where given a set of alternatives A and an integer k , we want to select a size- k slate $S \subseteq A$ (to be presented to a DM). Three natural desiderata for the selected slate S , regularly discussed in the literature under potentially different names (Petchrompo et al. 2022; Branke et al. 2008; Li and Yao 2019), are:

Diversity S should be “redundancy-free”, i.e., no two selected alternatives should be similar to each other.²

Representativity S should represent every alternative in A , i.e., each non-selected alternative from A should be close to one from S .

Efficiency S should contain “high-quality” alternatives, i.e., alternatives which score well across objectives.³

Which of these three desiderata is most important or appropriate depends on the context and the demands of the DM, making it hard to argue for or against each of them in general.

²Note that in the literature, the term “diversity” sometimes also refers to what we call representativity. Our notion of diversity is also distinct from the notion of diversity of Faliszewski et al. (2017) from the multiwinner voting literature.

³We use the term “efficiency” as an umbrella term to refer to notions explicitly capturing solution quality. This differs from parts of the multiobjective literature, where efficiency is used as a synonym for Pareto optimality (see, e.g., (Ehrgott 2005)).

3.2 Quality Measures⁴

We focus on two of the arguably most popular quality measures for Pareto pruning: uniformity and coverage (Sayin 2000; Petchrompo et al. 2022; Li and Yao 2019). Inspired by the desideratum of diversity, the *uniformity* of a slate S is $I_U(S) = \min_{x, y \in S, x \neq y} \|x - y\| = \min_{x, y \in S, x \neq y} \sum_{i=1}^d |x_i - y_i|$. UNIFORMITY PARETO PRUNING is the problem of finding a slate S , i.e., a size- k subset of A , maximizing uniformity $\max_{S \subseteq A, |S|=k} I_U(S)$.

Inspired by the idea of representativity, the *coverage* of a slate S with respect to a set of alternatives A is $I_C(S, A) = \max_{a \in A} \min_{s \in S} \|a - s\| = \max_{a \in A} \min_{s \in S} \sum_{i=1}^d |a_i - s_i|$. Note that a lower coverage value is better, since it signals that every point in A is close to a point in S . COVERAGE PARETO PRUNING is the problem of finding a slate with a minimum coverage value $\min_{S \subseteq A, |S|=k} I_C(S, A)$.⁵

A New Quality Measure: Directed Coverage Our new measure *directed coverage* is inspired by the coverage measure, but aims to correct some of its flaws that surface in our axiomatic analysis. The difference between the two is best illustrated by means of the following example. Consider $a = (1, 0)$ and $b = (0, \varepsilon)$ for some small $\varepsilon > 0$. Asked to present one alternative to the decision maker, which alternative should we choose? Coverage alone provides no guidance on which alternative is preferable, yet there is a strong case that one should select option a , since it significantly outperforms b under objective one and is almost as good as b under objective two. This is because coverage is based on the symmetric Manhattan distance, making it irrelevant whether we take an efficient alternative to cover a less-efficient one or the other way around. Directed coverage fixes this issue: When quantifying how suitable an alternative s is to cover an alternative a , we do not take into account the distance between the two with respect to objectives in which s outperforms a , as s covers a in these objectives “perfectly” in any case. Instead, we purely focus on and sum over the objectives in which a outperforms s , i.e., $\|a - s\|_+$, as this quantifies the total efficiency loss we suffer by presenting s rather than a to the decision maker. Formally, we define the *directed coverage* of a slate $S \subseteq A$ as $I_{DC}(S, A) = \max_{a \in A} \min_{s \in S} \|a - s\|_+ = \max_{a \in A} \min_{s \in S} \sum_{i=1}^d \max(a_i - s_i, 0)$. DIRECTED COVERAGE PARETO PRUNING is the problem of finding a slate minimizing directed coverage: $\min_{S \subseteq A, |S|=k} I_{DC}(S, A)$.

To illustrate the different selections made by the three measures, we refer to the full version, where we show their behavior on instances from our experiments.

3.3 Axiomatic Analysis

While numerous quality measures have been proposed in the literature (Li and Yao 2019; Faulkenberg and Wiecek 2010),

⁴Technically speaking, our quality measures can also be viewed as objectives we optimize. However, to distinguish them from the objectives present in multiobjective optimization problems, we exclusively refer to them as measures.

⁵Uniformity and coverage are connected. In the full version we show that the optimal coverage value with k points and the optimal uniformity value with $k + 1$ points differ by a factor of at most 2.

	Mono- tonicity	ε -Split Proofness	Extremism Monotonicity	Standout Consistency	Outlier Consistency
Uni.	✗	✓	✓	✗	✗
Cov.	✗	✗	✗	✗	✓
Dir. Cov.	✓	✗	✗	✓	✗

Table 1: Overview of axiomatic results. ✓ indicates that the measure fulfills the axiom. ✗ means that it violates it.

there is a lack of theoretical comparisons between them. In this section, we conduct an axiomatic analysis of the three measures introduced above, aiming to provide formal arguments for and against each measure. This approach allows us to move beyond intuitive arguments for and against different measures on disconnected grounds and instead evaluate measures based on explicitly stated criteria.

We consider two types of axioms. The first type concerns how optimal slates change in response to modifications of the underlying instance. The second set examines whether certain “extreme” alternatives are guaranteed to be included in an optimal slate. Our axioms serve two main purposes: (i) to identify measures that exhibit unintuitive or unreasonable behavior, and (ii) to identify how measures align with the three desiderata introduced in Section 3.1. An overview of which measures satisfy which axioms is provided in Table 1. Formal statements and proofs are given in the full version.

We begin with the axiom of *monotonicity*, which intuitively demands that improving an alternative x with respect to one or more objectives should not result in x being kicked out from the selected slate. Such behavior would be counterintuitive, as it implies that strictly improving an alternative’s performance can make it less likely for the DM to be presented with the alternative.

Axiom 1 (Monotonicity). *A measure I satisfies monotonicity if, for any set of alternatives A , $k \in \mathbb{N}$, and $S \in \mathcal{S}(I, A, k)$ with $x \in S$, the following holds: If $y \in \mathbb{R}^d$ dominates x , then there exists an optimal slate $S' \in \mathcal{S}(I, (A \setminus \{x\}) \cup \{y\}, k)$ with $y \in S'$.*

Both uniformity and coverage violate monotonicity. One reason for this is that improving an alternative can reduce its Manhattan distance to other alternatives, thereby diminishing its appeal to diversity (as it decreases the quality measure) or coverage (as the alternative becomes easier to cover). In contrast, directed coverage avoids this issue: if x strictly improves, then for any other alternative z , $\|x - z\|_+$ can only increase, while $\|z - x\|_+$ can only decrease, implying that x is not better covered by z than before. As a result, directed coverage satisfies monotonicity.

The second type of instance modification we consider is splitting an alternative into two alternatives. A popular variant of this idea, known as clone-robustness, requires that adding a perfect duplicate of an alternative should not affect the selected slate (up to potentially replacing the alternative with the duplicate). All three of our measures trivially satisfy clone-robustness, as selecting two identical alternatives

is never optimal. To obtain a more meaningful distinction between measures, we consider a stronger axiom, which we call *ε -split proofness*. It requires that no alternative x can be replaced by two arbitrarily close alternatives y_ε and z_ε so that both y_ε and z_ε get selected. Additionally, we demand that if either y_ε or z_ε is selected in the modified instance, replacing them with x should still yield an optimal slate in the original instance. This ensures that arbitrarily small perturbations cannot cause any changes to the slate.

Axiom 2 (ε -split proofness). *A measure I satisfies ε -split proofness if, for any set of alternatives A and $k \in \mathbb{N}$, there exists some $\varepsilon > 0$ such that for all $x \in A$ and $y_\varepsilon, z_\varepsilon \in \mathbb{R}^d$ with $\|x - y_\varepsilon\| < \varepsilon$ and $\|x - z_\varepsilon\| < \varepsilon$, the following holds: If $S_\varepsilon \in \mathcal{S}(I, (A \setminus \{x\}) \cup \{y_\varepsilon, z_\varepsilon\}, k)$, then (i) $S_\varepsilon \subseteq A$ and $S_\varepsilon \in \mathcal{S}(I, A, k)$ or (ii) $S_\varepsilon \setminus \{y_\varepsilon, z_\varepsilon\} \cup \{x\} \in \mathcal{S}(I, A, k)$.*

Notably, the axiom implies that a measure never selects two alternatives that are arbitrarily close to one another, a property particularly desirable from the perspective of the diversity desideratum. Among the measures we consider, only uniformity satisfies ε -split proofness. Both coverage and directed coverage violate the axiom, as it can be beneficial for these measures to select two alternatives arbitrarily close to each other if they cover different halves of the space.

While monotonicity and ε -split proofness can be considered broadly desirable, the desirability of the remaining axioms is more subjective, as each of them captures some form of alignment with one of the three desiderata introduced above. We begin with a variant of monotonicity tailored to the diversity desideratum, which we call *extremism monotonicity*. This axiom requires that if a selected alternative is the most extreme according to some objective, then pushing it even further away from the other alternatives in this objective should not result in its exclusion from the slate.

Axiom 3 (Extremism monotonicity). *A measure I satisfies extremism monotonicity if for any set of alternatives A , $k \in \mathbb{N}$, $t > 0$, and $S \in \mathcal{S}(I, A, k)$ with $x \in S$, the following holds: If for some objective $i \in [d]$, we have $x_i = \max_{a \in A} a_i$ (resp. $x_i = \min_{a \in A} a_i$), then there exists an optimal slate $S' \in \mathcal{S}(I, (A \setminus \{x\}) \cup \{x'\}, k)$ with $x' \in S'$, where $x'_i := x_i + t$ (resp. $x'_i := x_i - t$) and $x'_j := x_j$ for all $j \in [d] \setminus \{i\}$.*

This axiom formalizes the intuition that alternatives corresponding to particularly distinct tradeoff decisions should remain part of the slate when they become more distinct. As expected, uniformity satisfies extremism monotonicity, while both coverage and directed coverage violate it.

Our next axiom is inspired by the notion of Condorcet-consistency. Translated to our setting, Condorcet-consistency says that an alternative outperforming each of the others in a majority of objectives is always selected if it exists. We introduce a cardinal, weighted variant we call a *standout alternative*. To formalize this, we interpret $\|x - y\|_+$ as the “lead” of alternative x over alternative y , as it captures the total amount by which x outperforms y across all objectives in which x outperforms y . An alternative is a standout alternative if its weakest lead against any other alternative exceeds the strongest lead any other alternative has against it:

	#Objectives	Cardinal	Ordinal	Approval
Uni. / Cov.	$d = 2$	P^\dagger	P^\dagger	P^\dagger
	fixed $d \geq 3$	NP-h [Th. 8]	?	P [Pr. 11]
	unbounded d	NP-h [Th. 8]	NP-h [Pr. 10]	NP-h [Pr. 12]
Dir. Cov.	$d = 2$	P [Pr. 7]	P [Pr. 7]	P [Pr. 7]
	fixed $d \geq 4$	NP-h [Pr. 9]	?	P [Pr. 11]
	unbounded d	NP-h [Pr. 9]	NP-h [Pr. 10]	NP-h [Pr. 12]

Table 2: Summary of computational results.⁷ Results marked with \dagger are by Vaz et al. (2015).

Axiom 4 (Standout consistency). *An alternative $x \in A$ is a standout alternative if $\min_{a \in A \setminus \{x\}} \|x - a\|_+ > \max_{a \in A \setminus \{x\}} \|a - x\|_+$. A measure I is standout consistent if for any set of alternatives A containing a standout alternative $x \in A$, we have $x \in S$ for each optimal slate $S \in \mathcal{S}(I, A, k)$ and $k \geq 1$.*

From the perspective of efficiency, standout alternatives are highly desirable, as they are significantly better than all other alternatives in aggregate. Among the measures we consider, only directed coverage satisfies standout consistency. Uniformity and coverage, in contrast, do not satisfy this axiom. When faced with the decision of which of two alternatives to pick, they disregard which one is more efficient.

We conclude with the concept of an *outlier alternative*, an alternative that is further away from every other alternative than any two non-outlier alternatives are from each other:

Axiom 5 (Outlier consistency). *An alternative $x \in A$ is an outlier alternative if $\min_{a \in A \setminus \{x\}} \|x - a\| > \max_{y, z \in A \setminus \{x\}} \|y - z\|$. A measure I is outlier consistent if for any $k \geq 2$ and any set of alternatives A containing an outlier alternative $x \in A$, we have $x \in S$ for each optimal slate $S \in \mathcal{S}(I, A, k)$.*

From the perspective of representativity, an outlier should be selected, as it lies too far from all other alternatives to be adequately “covered” by any of them. Among the measures we consider, only coverage satisfies outlier consistency, while both uniformity and directed coverage do not.

4 Algorithmic Analysis

We present our algorithmic analysis (see Table 2). We start by discussing some related problems from computational geometry (Section 4.1), before we analyze the complexity of Pareto pruning for cardinal (Section 4.2), ordinal (Section 4.3), and approval (Section 4.4) objectives.

4.1 Connections to Computational Geometry

UNIFORMITY PARETO PRUNING and COVERAGE PARETO PRUNING are special cases of geometric variants of two well-known computational problems on graphs: the DISCRETE k -CENTER problem (Hakimi 1964) and the p -DISPERSION problem (Erkut 1990). Given a set of points B , a metric $d : B \times B \rightarrow \mathbb{R}_{\geq 0}$, and an integer k , DISCRETE k -CENTER (resp. p -DISPERSION) asks for a size- k subset $S \subseteq B$ minimizing $\max_{a \in B} \min_{s \in S} d(s, a)$ (resp. maximizing $\min_{x, y \in S, x \neq y} d(x, y)$). Note that in case d is the

Manhattan distance, these problems only differ from COVERAGE PARETO PRUNING (resp. UNIFORMITY PARETO PRUNING) in that B and S can contain Pareto dominated points. Wang and Kuo (1988) studied the geometric variant of p -DISPERSION when d is the Euclidean distance, establishing NP-hardness in \mathbb{R}^2 . Considering the case when d is the Euclidean or Manhattan distance, Megiddo and Supowit (1984) showed NP-hardness in \mathbb{R}^2 for a continuous version of DISCRETE k -CENTER, where the selected points are not restricted to be from B , but one can select any subset $S \subseteq \mathbb{R}^d$ of k points. In the literature, it is commonly assumed that DISCRETE k -CENTER in \mathbb{R}^2 is NP-hard as well. However, we were unable to track down a readily available proof.⁶ To fill this gap and to use the results in our later analysis, we provide a proof for the Manhattan distance in two dimensions following the key ideas from Megiddo and Supowit (1984):

Theorem 6. DISCRETE k -CENTER for the Manhattan distance is NP-hard, even in two dimensions.

4.2 Cardinal Objectives

When we restrict ourselves to Pareto optimal points in two dimensions, DISCRETE k -CENTER and p -DISPERSION become tractable: Vaz et al. (2015) have presented polynomial-time algorithms for UNIFORMITY PARETO PRUNING and COVERAGE PARETO PRUNING for the case of two objectives by exploiting that a set of Pareto optimal alternatives $A \subseteq \mathbb{R}^2$ can be embedded into \mathbb{R} in a way that maintains the Manhattan distance between alternatives. A dynamic programming approach for the embedded problem in \mathbb{R} yields a polynomial-time algorithm. This general approach can also be adapted to show an analogous result for directed coverage:

Proposition 7. For at most two objectives, DIRECTED COVERAGE PARETO PRUNING can be solved in $O(|A|k + |A| \log |A|)$.

Vaz et al. (2015) state in their conclusion: “[Pareto pruning] for more than two objectives may become an intractable task”. In fact, we were unable to find an NP-hardness result for Pareto pruning, even for an arbitrary number of objectives. We complement their tractability results with an NP-hardness for UNIFORMITY / COVERAGE PARETO PRUNING for three objectives. We establish this result by adapting NP-hardness proofs for DISCRETE k -CENTER and p -DISPERSION for the Manhattan distance in \mathbb{R}^2 . The general idea is that it is possible to construct a hyperplane $H \subseteq \mathbb{R}^3$ in which there is no pair of points $x, y \in H$, such that x dominates y . Embedding the constructions from these hardness proofs into such a hyperplane H then allows us to derive hardness results for UNIFORMITY PARETO PRUNING and COVERAGE PARETO PRUNING for three objectives.

Theorem 8. UNIFORMITY / COVERAGE PARETO PRUNING are NP-hard, even for three objectives.

For DIRECTED COVERAGE PARETO PRUNING, we are able to show hardness for four objectives. To derive this

⁶For example, Agarwal and Sharir (1998) cite the works of Megiddo and Supowit (1984), and Fowler, Paterson, and Tanimoto (1981) as a reference, yet both sources only contain a proof for the continuous version.

result, consider the map $f : \mathbb{R}^2 \rightarrow \mathbb{R}^4, (x_1, x_2) \mapsto (x_1, -x_1, x_2, -x_2)$. This map fulfills two key properties: For arbitrary $x, y \in \mathbb{R}^2$ it holds that $f(x)$ does not dominate $f(y)$ and vice versa, and second, $\|x - y\| = \|f(x) - f(y)\|_+$. Therefore, we can apply f to any instance of DISCRETE k -CENTER in \mathbb{R}^2 to get an equivalent instance of DIRECTED COVERAGE PARETO PRUNING in \mathbb{R}^4 . This establishes:

Proposition 9. DIRECTED COVERAGE PARETO PRUNING is NP-hard, even for four objectives.⁷

4.3 Ordinal Objectives

For the special case of ordinal objectives, the polynomial-time algorithm for two objectives still applies. However, complementing this result with a hardness result for a fixed number of ordinal objectives turns out to be surprisingly difficult and remains an open problem: The restriction of having to map bijectively to $[|A|]$ is not strong enough to provide clear properties that an algorithm can exploit, yet seems too restrictive to allow us to nicely control the distances $\|x - y\|$ or $\|x - y\|_+$ within a larger set of points.

Nevertheless, we show that all three problems are NP-hard for an unbounded number of objectives. For coverage and uniformity, the proof builds upon the hardness proofs for DISCRETE k -CENTER and p -DISPERSION in dimension two. For directed coverage, we employ a reduction from EXACT COVER BY 3-SETS.

Proposition 10. UNIFORMITY / COVERAGE / DIRECTED COVERAGE PARETO PRUNING are NP-hard, even if all objectives are ordinal objectives.

4.4 Approval Objectives

For approval objectives, our problems become easier from a computational perspective. We establish polynomial-time solvability for every fixed number of objectives $d \in \mathbb{N}$: For this, we call two alternatives *equivalent* if they are evaluated the same under every objective. Observe that there can be at most 2^d pairwise non-equivalent alternatives. As it is never optimal for any of our measures to pick two equivalent alternatives, it suffices to brute force over all at most $\binom{2^d}{k} \leq 2^{2^d}$ size- k subsets of pairwise non-equivalent alternatives:

Proposition 11. For any fixed $d \in \mathbb{N}$, UNIFORMITY / COVERAGE / DIRECTED COVERAGE PARETO PRUNING are solvable in polynomial time for d approval objectives.

We complement this result with an NP-hardness result for an unbounded number of approval objectives. To show this result, we draw inspiration from the classic (non-metric) hardness proofs for DISCRETE k -CENTER and p -DISPERSION on graphs (Hakimi 1964; Erkut 1990). The idea is that given a graph $G = (V, E)$, we construct an alternative $a_v \in A$ for every $v \in V$ and add objectives such that the distance between a_v and a_w is small if $\{a_v, a_w\} \in E$ and large otherwise. Hardness is then a straightforward reduction from INDEPENDENT SET for uniformity, and DOMINATING SET for coverage and directed coverage.

⁷We strengthened this result post-submission to show hardness for three objectives. A proof can be found in the full version.

Proposition 12. UNIFORMITY / COVERAGE / DIRECTED COVERAGE PARETO PRUNING are NP-hard, even if all objectives are approval objectives.

5 Experiments

We conduct an experimental evaluation of the slates returned by the three solution methods we consider. In this section, we use the terms uniformity, coverage, and directed coverage to refer both to the underlying quality measures (I_U , I_C , and I_{DC} , respectively), which we use to evaluate slates, and the respective solution methods that optimize for one of them. To distinguish, we use typewriter font when referring to the solution method, i.e., the slate obtained by solving the corresponding optimization problem (e.g., we write `Uniformity` to refer to UNIFORMITY PARETO PRUNING).

Setup We consider three different datasets. Datasets ZDT (Zitzler, Deb, and Thiele 2000) containing six instances with two objectives and DTLZ (Deb et al. 2002) containing seven instances with three objectives are widely used for the evaluation of multiobjective evolutionary algorithms. For a more realistic example, we consider the dataset PGMORL containing six instances, where the alternatives correspond to simulated agents performing a simple task evaluated under two objectives. Xu et al. (2020) created these benchmark instances to evaluate their multiobjective evolutionary algorithm PGMORL. We compute all slates⁸ via integer linear programming (ILP) formulations, solved using Gurobi. For feasibility reasons, for the six instances from these datasets in which the Pareto front contains more than 200 alternatives, we delete all but 200 randomly sampled alternatives from the instance. We consider three different values of k , i.e., $k = 5\% \cdot |A|$, $k = 10\% \cdot |A|$, and $k = 25\% \cdot |A|$.

Results We evaluate each computed slate S using four quality measures: uniformity I_U , coverage I_C , directed coverage I_{DC} , and hypervolume⁹. To enable a meaningful comparison across solution methods and aggregation across instances, we normalize all scores within each instance by dividing by the score of the best-performing slate under the respective measure. For example, when evaluating uniformity I_U , we divide the uniformity score of each slate by the maximum uniformity achieved across all methods, which is by definition `Uniformity`, for that instance. Table 3 reports the normalized values, averaged over all instances in each dataset. For measures marked with (\uparrow), higher values indicate better performance; for those marked with (\downarrow), lower values are preferred.

Analysis We discuss some patterns observed in Table 3. While the choice of k does influence methods' performance, no clear trend emerges. Therefore, we focus on observations that hold across all three considered values of k . First, we

⁸The Pareto fronts for ZDT and DTLZ are taken from the pymoo library (Blank and Deb 2020). The Pareto fronts for PGMORL have been calculated and made public by Xu et al. (2020).

⁹For a set of alternatives A , and a reference point r , the hypervolume of $S \subseteq A$ is the volume of $C = \{x \in \mathbb{R}^d \mid x \text{ dominates } r \text{ and } x \text{ is dominated by some } a \in S\}$. Hypervolume is seen to capture both efficiency and diversity of alternatives.

Method	Uniformity (\uparrow)			Coverage (\downarrow)			Directed Coverage (\downarrow)			Hypervolume (\uparrow)		
	k as fraction of $ A $	5%	10%	25%	5%	10%	25%	5%	10%	25%	5%	10%
Dataset ZDT												
Uniformity	100.0%	100.0%	100.0%	121.8%	108.3%	117.4%	195.3%	187.1%	150.4%	90.8%	97.9%	99.9%
Coverage	81.0%	83.0%	74.7%	100.0%	100.0%	100.0%	201.0%	204.1%	182.3%	97.3%	98.9%	99.9%
Directed Coverage	69.3%	68.4%	41.3%	165.7%	212.9%	364.6%	100.0%	100.0%	100.0%	99.3%	99.8%	99.9%
Dataset DTLZ												
Uniformity	100.0%	100.0%	100.0%	131.2%	125.1%	123.2%	200.0%	158.1%	158.5%	92.6%	98.6%	99.2%
Coverage	70.3%	72.8%	66.7%	100.0%	100.0%	100.0%	178.2%	188.7%	185.6%	99.9%	97.9%	95.9%
Directed Coverage	72.2%	58.9%	59.1%	155.8%	188.7%	246.6%	100.0%	100.0%	100.0%	96.4%	99.6%	99.1%
Dataset PGMORL												
Uniformity	100.0%	100.0%	100.0%	123.0%	132.0%	144.8%	187.5%	234.6%	254.7%	94.8%	97.9%	99.5%
Coverage	79.9%	60.8%	56.4%	100.0%	100.0%	100.0%	169.7%	182.3%	227.1%	98.3%	98.7%	99.4%
Directed Coverage	51.1%	38.5%	45.9%	280.4%	347.4%	482.3%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%

Table 3: Comparison of three methods for Pareto pruning. We report average values of four measures, each normalized by the best solution at the instance level. (\uparrow) indicates that higher values are better, and (\downarrow) indicates that lower values are better.

observe substantial relative differences between the three solution methods in terms of their performance under the uniformity, coverage, and directed coverage measures. This underscores that the choice of method can have significant practical implications. We observe that `Coverage` consistently outperforms `Directed Coverage` with respect to uniformity, and `Uniformity` outperforms `Directed Coverage` with respect to coverage. This suggests that, despite differences in their formal definitions, `Uniformity` and `Coverage` exhibit more similar behavior to each other than either does to `Directed Coverage`. In contrast, when evaluating performance under the directed coverage measure, no consistent trend emerges as to whether `Uniformity` or `Coverage` performs better. Both return states that, from the perspective of directed coverage, are typically more than 50% worse than those produced by the dedicated `Directed Coverage` method. This illustrates that if one cares about the directed coverage measure, using one of the two more established approaches is insufficient.

When evaluating performance with respect to hypervolume, which is more efficiency-focused, the differences between the solution methods are less pronounced. On `ZDT` and `PGMORL`, `Directed Coverage` consistently outperforms `Coverage`, which in turn outperforms `Uniformity`. For `DTLZ`, which method performs better depends on the choice of k . While the differences are smaller than for the other measures, these results still provide evidence that `Directed Coverage` tends to select more efficient solutions. This is also intuitive: by design, `Directed Coverage` avoids selecting alternatives that are only marginally better in some objectives while being worse in all others in comparison to other alternatives. At the instance level, we further observe that, unlike the other two methods, `Directed Coverage` tends to avoid selecting large numbers of alternatives from regions populated by less-efficient alternatives. See the full version for some more tangible examples.

6 Discussion

We presented a systematic study of quality measures for Pareto pruning, including the first axiomatic analysis and a comprehensive complexity investigation. We hope that our work enables more principled arguments for and against different measures in multiobjective optimization and contributes to a clearer understanding of their tractability. Motivated by the shortcomings revealed in our axiomatic analysis, we proposed the new measure of directed coverage, which performs competitively or even favorably in our experiments.

There are several promising directions for future work. First, it would be valuable to complement our axiomatic analysis with characterization and impossibility results, and design axioms tailored to ordinal or approval objectives (in particular, ε -split proofness and extremism monotonicity do not translate to these settings). Second, our algorithmic analysis leaves open whether Pareto pruning remains hard for ordinal objectives with a fixed number of objectives. Third, extending our analysis to further quality measures would be worthwhile. Lastly, it would be intriguing to further explore the connection between Pareto pruning and previous work in social choice, particularly the paradigms of proportional representation and diversity in multiwinner voting (Faliszewski et al. 2017). While in our work, we interpreted solutions as candidates and objectives as voters, it would also be fruitful to explore a social choice modeling in which solutions serve as both candidates and voters, ranking other solutions by similarity. This would embed the problem in recent work on centroid clustering in the social choice literature (Micha and Shah 2020; Kellerhals and Peters 2024). It would be interesting to analyze whether existing Pareto pruning methods satisfy solution concepts from this setting, and conversely, whether algorithms from that literature can offer meaningful guarantees or performances for Pareto pruning.

References

Agarwal, P. K.; and Sharir, M. 1998. Efficient algorithms for geometric optimization. *ACM Comput. Surv.*, 30(4): 412–458.

- Blank, J.; and Deb, K. 2020. Pymoo: Multi-Objective Optimization in Python. *IEEE Access*, 8: 89497–89509.
- Branke, J.; Deb, K.; Miettinen, K.; and Slowinski, R., eds. 2008. *Multiobjective Optimization, Interactive and Evolutionary Approaches*, volume 5252 of *Lecture Notes in Computer Science*. Springer.
- Deb, K.; Thiele, L.; Laumanns, M.; and Zitzler, E. 2002. Scalable multi-objective optimization test problems. In *Proceedings of the 2002 Congress on Evolutionary Computation (CEC '02)*, volume 1, 825–830.
- Delemazure, T.; Janeczko, L.; Kaczmarczyk, A.; and Szufa, S. 2024. Selecting the Most Conflicting Pair of Candidates. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI '24)*, 2766–2773. ijcai.org.
- Ehrgott, M. 2005. *Multicriteria optimization*, volume 491. Springer Science & Business Media.
- Eriskin, L.; Karatas, M.; and Zheng, Y. 2024. A robust multi-objective model for healthcare resource management and location planning during pandemics. *Ann. Oper. Res.*, 335(3): 1471–1518.
- Erkut, E. 1990. The discrete p -dispersion problem. *Eur. J. Oper. Res.*, 46(1): 48–60.
- Eusébio, A.; Figueira, J. R.; and Ehrgott, M. 2014. On finding representative non-dominated points for bi-objective integer network flow problems. *Comput. Oper. Res.*, 48: 1–10.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2017. Multiwinner voting: A new challenge for social choice theory. *Trends in computational social choice*, 74(2017): 27–47.
- Faulkenberg, S. L.; and Wiecek, M. M. 2010. On the quality of discrete representations in multiple objective programming. *Optim. Eng.*, 11(3): 423–440.
- Fowler, R. J.; Paterson, M. S.; and Tanimoto, S. L. 1981. Optimal packing and covering in the plane are NP-complete. *Inf. Process. Lett.*, 12(3): 133–137.
- Guo, Z. X.; Wong, W. K.; Li, Z.; and Ren, P. 2013. Modeling and Pareto optimization of multi-objective order scheduling problems in production planning. *Comput. Ind. Eng.*, 64(4): 972–986.
- Hakimi, S. L. 1964. Optimum locations of switching centers and the absolute centers and medians of a graph. *Oper. Res.*, 12(3): 450–459.
- Hayes, C. F.; Radulescu, R.; Bargiacchi, E.; Källström, J.; Macfarlane, M.; Reymond, M.; Verstraeten, T.; Zintgraf, L. M.; Dazeley, R.; Heintz, F.; Howley, E.; Irissappane, A. A.; Mannion, P.; Nowé, A.; de Oliveira Ramos, G.; Restelli, M.; Vamplew, P.; and Roijers, D. M. 2022. A practical guide to multi-objective reinforcement learning and planning. *Auton. Agents Multi Agent Syst.*, 36(1): 26.
- Kellerhals, L.; and Peters, J. 2024. Proportional Fairness in Clustering: A Social Choice Perspective. In *Proceedings of the Thirty-Eighth Annual Conference on Neural Information Processing Systems (NeurIPS '24)*.
- Li, M.; and Yao, X. 2019. Quality Evaluation of Solution Sets in Multiobjective Optimisation: A Survey. *ACM Comput. Surv.*, 52(2): 26:1–26:38.
- Marler, R. T.; and Arora, J. S. 2004. Survey of multi-objective optimization methods for engineering. *Struct. Multidiscip. Optim.*, 26: 369–395.
- Megiddo, N.; and Supowit, K. J. 1984. On the Complexity of Some Common Geometric Location Problems. *SIAM J. Comput.*, 13(1): 182–196.
- Micha, E.; and Shah, N. 2020. Proportionally Fair Clustering Revisited. In *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP '20)*, 85:1–85:16. Schloss Dagstuhl.
- Petchrompo, S.; Coit, D. W.; Brintrup, A.; Wannakrairot, A.; and Parlikad, A. K. 2022. A review of Pareto pruning methods for multi-objective optimization. *Comput. Ind. Eng.*, 167: 108022.
- Petchrompo, S.; Wannakrairot, A.; and Parlikad, A. K. 2022. Pruning Pareto optimal solutions for multi-objective portfolio asset management. *Eur. J. Oper. Res.*, 297(1): 203–220.
- Sayin, S. 2000. Measuring the quality of discrete representations of efficient sets in multiple objective mathematical programming. *Math. Program.*, 87(3): 543–560.
- Stewart, T. J.; Bandte, O.; Braun, H.; Chakraborti, N.; Ehrgott, M.; Göbelt, M.; Jin, Y.; Nakayama, H.; Poles, S.; and Stefano, D. D. 2008. Real-World Applications of Multiobjective Optimization. In *Multiobjective Optimization, Interactive and Evolutionary Approaches*, volume 5252 of *Lecture Notes in Computer Science*, 285–327. Springer.
- Sylva, J.; and Crema, A. 2007. A method for finding well-dispersed subsets of non-dominated vectors for multiple objective mixed integer linear programs. *Eur. J. Oper. Res.*, 180(3): 1011–1027.
- Taboada, H. A.; Baheranwala, F.; Coit, D. W.; and Wattanapongsakorn, N. 2007. Practical solutions for multi-objective optimization: An application to system reliability design problems. *Reliab. Eng. Syst. Saf.*, 92(3): 314–322.
- Taboada, H. A.; and Coit, D. W. 2007. Data clustering of solutions for multiple objective system reliability optimization problems. *Qual. Technol. Quant. Manag.*, 4(2): 191–210.
- Taboada, H. A.; and Coit, D. W. 2008. Multi-objective scheduling problems: Determination of pruned Pareto sets. *Iie Transactions*, 40(5): 552–564.
- Vaz, D.; Paquete, L.; Fonseca, C. M.; Klamroth, K.; and Stiglmayr, M. 2015. Representation of the non-dominated set in biobjective discrete optimization. *Comput. Oper. Res.*, 63: 172–186.
- Wang, D. W.; and Kuo, Y.-S. 1988. A study on two geometric location problems. *Inf. Process. Lett.*, 28(6): 281–286.
- Wang, W.; Lin, M.; Fu, Y.; Luo, X.; and Chen, H. 2020. Multi-objective optimization of reliability-redundancy allocation problem for multi-type production systems considering redundancy strategies. *Reliab. Eng. Syst. Saf.*, 193: 106681.
- Xu, J.; Tian, Y.; Ma, P.; Rus, D.; Sueda, S.; and Matusik, W. 2020. Prediction-Guided Multi-Objective Reinforcement Learning for Continuous Robot Control. In *Proceedings of the 37th International Conference on Machine Learning (ICML '20)*, 10607–10616. PMLR.

Zio, E.; and Bazzo, R. 2011. A clustering procedure for reducing the number of representative solutions in the Pareto Front of multiobjective optimization problems. *Eur. J. Oper. Res.*, 210(3): 624–634.

Zitzler, E.; Deb, K.; and Thiele, L. 2000. Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. *Evol. Comput.*, 8(2): 173–195.