



Approximating multiobjective optimization problems: How exact can you be?

Cristina Bazgan¹ · Arne Herzel^{2,3} · Stefan Ruzika² · Clemens Thielen^{3,4}  · Daniel Vanderpooten¹

Received: 28 February 2023 / Revised: 25 July 2023 / Accepted: 8 September 2023 /

Published online: 27 September 2023

© The Author(s) 2023

Abstract

It is well known that, under very weak assumptions, multiobjective optimization problems admit $(1 + \varepsilon, \dots, 1 + \varepsilon)$ -approximation sets (also called ε -Pareto sets) of polynomial cardinality (in the size of the instance and in $\frac{1}{\varepsilon}$). While an approximation guarantee of $1 + \varepsilon$ for any $\varepsilon > 0$ is the best one can expect for singleobjective problems (apart from solving the problem to optimality), even better approximation guarantees than $(1 + \varepsilon, \dots, 1 + \varepsilon)$ can be considered in the multiobjective case since the approximation might be exact in some of the objectives. Hence, in this paper, we consider *partially exact approximation sets* that require to approximate each feasible solution exactly, i.e., with an approximation guarantee of 1, in some of the objectives while still obtaining a guarantee of $1 + \varepsilon$ in all others. We characterize the types of polynomial-cardinality, partially exact approximation sets that are guaran-

✉ Clemens Thielen
clemens.thielen@tum.de

Cristina Bazgan
bazgan@lamsade.dauphine.fr

Arne Herzel
arne.herzel@hswt.de

Stefan Ruzika
ruzika@mathematik.uni-kl.de

Daniel Vanderpooten
daniel.vanderpooten@lamsade.dauphine.fr

¹ CNRS, LAMSADE, Université Paris-Dauphine, PSL Research University, 75016 Paris, France

² Department of Mathematics, University of Kaiserslautern-Landau, Paul-Ehrlich-Str. 14, 67663 Kaiserslautern, Germany

³ TUM Campus Straubing for Biotechnology and Sustainability, Weihenstephan-Triesdorf University of Applied Sciences, Am Essigberg 3, 94315 Straubing, Germany

⁴ Department of Mathematics, School of Computation, Information and Technology, Technical University of Munich, Boltzmannstr. 3, 85748 Garching bei München, Germany

teed to exist for general multiobjective optimization problems. Moreover, we study minimum-cardinality partially exact approximation sets concerning (weak) efficiency of the contained solutions and relate their cardinalities to the minimum cardinality of a $(1 + \varepsilon, \dots, 1 + \varepsilon)$ -approximation set.

Keywords Multiobjective optimization · Approximation · Efficient set · Intractability

1 Introduction

Many real-world optimization problems require taking into account multiple conflicting objective functions. Since solutions that optimize all objectives simultaneously do usually not exist, the solutions of interest are the so-called *efficient* (or *Pareto optimal*) solutions. A solution is called efficient if any other solution that is better in some objective is worse in at least one other objective. The images of the efficient solutions in the objective space are called *nondominated points*. A main goal in multiobjective optimization is to determine the set of all nondominated points (the *nondominated set*) and, for each of them, one corresponding efficient solution. One major difficulty, however, is *intractability*, that is, the fact that the nondominated set (and the efficient set) may be exponentially large for discrete problems (see, e.g., Ehrgott 1998), and typically infinite for continuous problems.

This provides a strong motivation to consider *approximations* of multiobjective optimization problems. Approximations for general multiobjective problems have been investigated since the 1990s (Safer 1992; Safer and Orlin 1995). Their existence and cardinality for general multiobjective problems have been systematically investigated in the seminal work of Papadimitriou and Yannakakis (2000). They show that, for any instance of a multiobjective optimization problem with a constant number of positive-valued, polynomial-time computable objective functions and any $\varepsilon > 0$, there exists a $(1 + \varepsilon, \dots, 1 + \varepsilon)$ -approximation set (also called a $(1 + \varepsilon)$ -approximation set or an ε -Pareto set) whose cardinality is *fully* polynomial, i.e., polynomial in the encoding length of the instance and $\frac{1}{\varepsilon}$. Additionally, they prove that a $(1 + \varepsilon)$ -approximation set can be computed in (fully) polynomial time for every $\varepsilon > 0$ if and only if a certain auxiliary problem GAP_δ (the *gap problem*) can be solved in (fully) polynomial time for every $\delta > 0$.

It has been observed, however, that $(1 + \varepsilon)$ -approximation sets are far from unique and different $(1 + \varepsilon)$ -approximation sets can differ significantly in their cardinality. Consequently, further work has focused on the computation of $(1 + \varepsilon)$ -approximation sets whose cardinality can be bounded relative to the cardinality of a smallest (minimum-cardinality) $(1 + \varepsilon)$ -approximation set for the given instance (Vassilvitskii and Yannakakis 2005; Koltun and Papadimitriou 2007; Diakonikolas and Yannakakis 2009; Bazgan et al. 2015).

While an approximation guarantee of $1 + \varepsilon$ for any $\varepsilon > 0$ is the best solution quality one can expect for singleobjective problems (apart from solving the problem to optimality), even better approximation guarantees than $(1 + \varepsilon, \dots, 1 + \varepsilon)$ can be considered in the multiobjective case since the approximation might be *partially exact*, i.e., exact in some of the objectives (and $(1 + \varepsilon)$ -approximate in the others). In fact,

it has recently been shown in Herzel et al. (2021b) that, under similar assumptions as in Papadimitriou and Yannakakis (2000), there exist fully-polynomial-cardinality *one-exact* $(1 + \varepsilon)$ -approximation sets that achieve an approximation guarantee of 1 in one of the objective functions (without loss of generality the first one) and $1 + \varepsilon$ in the others.

In this paper, we systematically investigate which types of partially exact approximation sets of polynomial cardinality are guaranteed to exist for general multiobjective problems. In particular, we show that, for any constant number p of objective functions and any $k \leq \lceil \frac{p}{2} \rceil$, polynomial-cardinality $(1 + \varepsilon)$ -approximation sets exist that approximate each feasible solution exactly in k of the objectives if we allow the objectives in which the approximation is exact to differ between different feasible solutions. We call these *quasi- k -exact* $(1 + \varepsilon)$ -approximation sets. Their existence contrasts the fact that the efficient and nondominated sets are already intractable for many biobjective problems, which implies that polynomial-cardinality $(1 + \varepsilon)$ -approximation sets that are exact in the same objectives for all solutions can in general be exact in at most one objective (Herzel et al. 2021b). We obtain our general result by providing a general existence proof that generalizes the existence proofs for $(1 + \varepsilon)$ -approximation sets and one-exact $(1 + \varepsilon)$ -approximation sets provided in Papadimitriou and Yannakakis (2000) and Herzel et al. (2021b), respectively.

Moreover, we investigate the question of (weak) efficiency of the solutions contained in minimum-cardinality partially exact $(1 + \varepsilon)$ -approximation sets of different types and relate their cardinalities to the minimum cardinality of a $(1 + \varepsilon)$ -approximation set for a given instance. In particular, we show that, for every instance, the minimum cardinality of a quasi-1-exact $(1 + \varepsilon)$ -approximation set equals the minimum cardinality of an ordinary $(1 + \varepsilon)$ -approximation set.¹ This contrasts the result from Herzel et al. (2021b) stating that, for any positive integer n , there exist instances for which the minimum cardinality of a one-exact $(1 + \varepsilon)$ -approximation set is more than n times larger than the minimum cardinality of an ordinary $(1 + \varepsilon)$ -approximation set. By generalizing the corresponding proof from Herzel et al. (2021b), however, we show that, for any positive integer n , there also exist instances for which the minimum cardinality of a quasi-2-exact $(1 + \varepsilon)$ -approximation set is more than n times larger than the minimum cardinality of an ordinary $(1 + \varepsilon)$ -approximation set.

Finally, we present several results concerning the polynomial-time computability of partially exact approximation sets. In particular, we show that the (fully) polynomial-time solvability of the gap problem does not suffice for the computation of any type of partially exact approximation set in general.

While we focus on results that are applicable to general multiobjective optimization problems under weak assumptions, there also exists a large body of work on approximations for specific multiobjective problems. For an up-to-date survey on both general and problem-specific approximation results, we refer to Herzel et al. (2021a).

¹ We sometimes use the term *ordinary* $(1 + \varepsilon)$ -approximation set to refer to $(1 + \varepsilon)$ -approximation sets if we want to explicitly distinguish them from the more demanding partially exact $(1 + \varepsilon)$ -approximation sets.

2 Preliminaries

Multiobjective optimization problems can contain objective functions that are to be minimized or maximized (or even a combination of both). However, we only consider the case of minimization problems in this article, i.e., all objectives are to be minimized. This is without loss of generality here since all our arguments can be straightforwardly adapted to the case where some or all objective functions are to be maximized.

Definition 2.1 (*Multiobjective optimization problem*) A multiobjective optimization problem Π is given by a set of instances. Each instance I consists of a (finite or infinite) set X^I of feasible solutions and a vector $f^I = (f_1^I, \dots, f_p^I)$ of p objective functions $f_i^I : X^I \rightarrow \mathbb{Q}$ for $i = 1, \dots, p$ to be minimized. The feasible set X^I might not be given explicitly.

Here, the number p of objective functions in a multiobjective optimization problem Π is assumed to be constant. Moreover, as is common in the context of approximation of multiobjective optimization problems, we assume that the objective functions take only positive, rational values and are polynomial-time computable (cf. Papadimitriou and Yannakakis 2000; Vassilvitskii and Yannakakis 2005; Diakonikolas and Yannakakis 2009; Herzel et al. 2021b). Additionally, we use the following standard assumption:²

Assumption 1 For any multiobjective optimization problem Π , there exists a polynomial \mathcal{P} such that, for any instance I of Π , there exists a polynomial-time computable value $M^I \leq \mathcal{P}(\text{enc}(I))$ such that $\text{enc}(f_i^I(x)) \leq M^I$ for any $x \in X^I$ and any $i \in \{1, \dots, p\}$. Here, $\text{enc}(I)$ denotes the encoding length of the instance I and $\text{enc}(f_i^I(x))$ denotes the encoding length of the value $f_i^I(x) \in \mathbb{Q}_{>0}$ in binary. This, in particular, implies that, for any instance I and any objective function value $f_i^I(x)$, we have $2^{-M^I} \leq f_i^I(x) \leq 2^{M^I}$. Also, any two values $f_i^I(x)$ and $f_i^I(x')$ differ by at least 2^{-2M^I} if they are not equal.

In the following, we sometimes blur the distinction between the problem Π and a concrete instance $I = (X^I, f^I)$ and we drop the superscript I indicating the dependence on the instance in X^I , f^I , M^I , etc. whenever the instance is clear from the context.

For multiobjective optimization problems, the solutions of interest are the so-called *efficient solutions* for which any other solution that is better in some objective is worse in at least one other objective:

² Given the previously-stated assumptions on the objective functions, this general assumption also used in Vassilvitskii and Yannakakis (2005), Diakonikolas and Yannakakis (2009), and Herzel et al. (2021b) holds for large classes of problems—in particular, for combinatorial problems. Together with the previously-stated assumptions, it guarantees the existence of all versions of approximation sets considered here and in Papadimitriou and Yannakakis (2000), Vassilvitskii and Yannakakis (2005), Diakonikolas and Yannakakis (2009), and Herzel et al. (2021b). We note, however, that specific types of approximation sets can also exist for some classes of problems for which this assumption does not formally hold (e.g., for multiobjective linear programming problems, see Papadimitriou and Yannakakis 2000).

Definition 2.2 For an instance I of a multiobjective optimization problem, a solution $x \in X$ *dominates* another solution $x' \in X$ (denoted as $x \preceq x'$) if $f_i(x) \leq f_i(x')$ for $i = 1, \dots, p$ and $f_i(x) < f_i(x')$ for at least one i . Moreover, a solution $x \in X$ *strictly dominates* another solution $x' \in X$ if $f_i(x) < f_i(x')$ for $i = 1, \dots, p$. A solution $x \in X$ is called (weakly) *efficient* if it is not (strictly) dominated by any other solution $x' \in X$. In this case, we call the corresponding image $f(x) \in f(X) \subseteq \mathbb{Q}^p$ a (weakly) *nondominated point*. The set $X_E \subseteq X$ of all efficient solutions is called the *efficient set* (or *Pareto set*) and the set $Y_N := f(X_E)$ of nondominated points is called the *nondominated set*. Similarly, the set $X_{WE} \subseteq X$ of all weakly efficient solutions is called the *weakly efficient set* and the set $Y_{WN} := f(X_{WE})$ of weakly nondominated points is called the *weakly nondominated set*.

One way to circumvent the problem that the nondominated set often has cardinality exponential in the input size (see, e.g., Ehrgott 1998) is the concept of approximation. Here, instead of requiring to select at least one corresponding efficient solution for each nondominated point, the requirement is relaxed so that each image point is only required to be dominated “approximately” by some image of a selected solution. This idea is formalized as follows:

Definition 2.3 Let (X, f) be an instance of a multiobjective optimization problem and let $\alpha = (\alpha_1, \dots, \alpha_p) \in \mathbb{R}^p$ with $\alpha_i \geq 1$ for all $i \in \{1, \dots, p\}$. We say that a feasible solution $x \in X$ α -*approximates* another feasible solution $x' \in X$ if $f_i(x) \leq \alpha_i \cdot f_i(x')$ for all $i \in \{1, \dots, p\}$.

When fixing a desired vector α of approximation guarantees (or a set of possible desired vectors), approximation as in Definition 2.3 can equivalently be expressed as a relation $R \subseteq X \times X$, where $(x, x') \in R$ (also denoted as xRx') if and only if x' is α -approximated by x for some desired vector α . In the following, we consider general approximate dominance relations that are *monotonic* in the following sense:

Definition 2.4 Let (X, f) be an instance of a multiobjective optimization problem. A relation $R \subseteq X \times X$ is called *monotonic* if $f_i(x) \leq f_i(x')$ for all $i \in \{1, \dots, p\}$ implies that xRx' .

Note that any monotonic relation R is, in particular, reflexive, i.e., xRx for all $x \in X$. We are particularly interested in the following monotonic approximate dominance relations, which give rise to ordinary $(1 + \varepsilon)$ -approximation sets as introduced in Papadimitriou and Yannakakis (2000) and the different types of partially exact approximation sets considered here:

Definition 2.5 Let (X, f) be an instance of a multiobjective optimization problem and let $\varepsilon > 0$. We define the following relations on $X \times X$:

- $(1 + \varepsilon)$ -*dominance*: $x \preceq_\varepsilon x' :\Leftrightarrow x$ α -approximates x' , where $\alpha_i = 1 + \varepsilon$ for all $i \in \{1, \dots, p\}$,
- *one-exact* $(1 + \varepsilon)$ -*dominance*: $x \preceq_\varepsilon^1 x' :\Leftrightarrow x$ α -approximates x' , where $\alpha_1 = 1$ and $\alpha_i = 1 + \varepsilon$ for all $i \geq 2$,
- *two-exact* $(1 + \varepsilon)$ -*dominance*: $x \preceq_\varepsilon^{1,2} x' :\Leftrightarrow x$ α -approximates x' , where $\alpha_1 = \alpha_2 = 1$ and $\alpha_i = 1 + \varepsilon$ for all $i \geq 3$,

- *quasi- k -exact* $(1 + \varepsilon)$ -dominance: $x \preceq_{\varepsilon}^{(k)} x' : \Leftrightarrow x$ α -approximates x' for some α where k components α_i are equal to 1 and the other $p - k$ are equal to $1 + \varepsilon$,
- *one-exact, quasi- k -exact* $(1 + \varepsilon)$ -dominance: $x \preceq_{\varepsilon}^{1,(k)} x' : \Leftrightarrow x$ α -approximates x' for some α where $\alpha_1 = 1$, $k - 1$ of the other components α_i are equal to 1, and the remaining $p - k$ are equal to $1 + \varepsilon$.³

The relations \preceq_{ε}^1 , $\preceq_{\varepsilon}^{1,2}$, $\preceq_{\varepsilon}^{(k)}$, and $\preceq_{\varepsilon}^{1,(k)}$ will also be referred to as *partially exact* (approximate dominance) relations in the following.

Given an approximate dominance relation, approximation sets for multiobjective optimization problems are defined as follows:

Definition 2.6 Let (X, f) be an instance of a multiobjective optimization problem and let R be a relation on $X \times X$. A set $P \subseteq X$ of feasible solutions is called an *R -approximation set* for (X, f) if, for any feasible solution $x' \in X$, there exists a solution $x \in P$ such that $x R x'$ (i.e., x *R -dominates* x').

For $\varepsilon > 0$, R -approximation sets for the relations \preceq_{ε} , \preceq_{ε}^1 , $\preceq_{\varepsilon}^{1,2}$, $\preceq_{\varepsilon}^{(k)}$, and $\preceq_{\varepsilon}^{1,(k)}$ are referred to as (*ordinary*) $(1 + \varepsilon)$ -approximation sets, *one-exact* $(1 + \varepsilon)$ -approximation sets, *two-exact* $(1 + \varepsilon)$ -approximation sets, *quasi- k -exact* $(1 + \varepsilon)$ -approximation sets, and *one-exact, quasi- k -exact* $(1 + \varepsilon)$ -approximation sets, respectively.

In the following, we will often use that, since monotonic relations R are reflexive, R -approximation sets for monotonic relations correspond to *R -dominating sets* for X , i.e., to subsets $P \subseteq X$ of feasible solutions such that any feasible solution $x' \notin P$ is R -dominated by some solution $x \in P$. This correspondence, in particular, holds for the monotonic relations \preceq_{ε} , \preceq_{ε}^1 , $\preceq_{\varepsilon}^{1,2}$, $\preceq_{\varepsilon}^{(k)}$, and $\preceq_{\varepsilon}^{1,(k)}$.

We are interested in approximation sets whose cardinality is bounded by a polynomial in the encoding length of the given instance (X, f) . For the relations \preceq_{ε} , \preceq_{ε}^1 , $\preceq_{\varepsilon}^{1,2}$, $\preceq_{\varepsilon}^{(k)}$, and $\preceq_{\varepsilon}^{1,(k)}$ that depend on a given value $\varepsilon > 0$, we are specifically interested in approximation sets whose cardinality is *fully polynomial*, i.e., polynomial in the encoding length of the instance and in $\frac{1}{\varepsilon}$.

3 Existence of approximation sets

We start by investigating the existence of R -approximations of (fully) polynomial cardinality under weak assumptions on the approximate dominance relation R . For the $(1 + \varepsilon)$ -dominance relation $R = \preceq_{\varepsilon}$, this is established in the seminal work of Papadimitriou and Yannakakis (2000), who show the existence of $(1 + \varepsilon)$ -approximations of fully polynomial cardinality $\mathcal{O}((\frac{M}{\varepsilon})^{p-1})$.

We now generalize the existence proof from Papadimitriou and Yannakakis (2000) to more demanding approximate dominance relations R that refine \preceq_{ε} , which, as shown afterwards in Theorem 3.2, include the partially exact approximate dominance relations \preceq_{ε}^1 and $\preceq_{\varepsilon}^{(k)}$.

³ Note that this means that $\preceq_{\varepsilon}^{1,(k)} = \preceq_{\varepsilon}^1 \cap \preceq_{\varepsilon}^{(k)}$ as subsets of $X \times X$.

Theorem 3.1 *Let (X, f) be an instance of a p -objective optimization problem and let $\varepsilon > 0$. Consider a monotonic relation R' defined on $X \times X$ and let $R := \preceq_\varepsilon \cap R'$ (so R -approximation sets correspond to $(1 + \varepsilon)$ -approximation sets that are simultaneously R' -approximation sets). If there exists a constant c (independent of the instance size and $\frac{1}{\varepsilon}$) such that every nonempty subset $X' \subseteq X$ admits an R' -dominating set $P_{X'} \subseteq X'$ of cardinality at most c , then there exists an R -approximation set of cardinality $\mathcal{O}\left(\left(\frac{M}{\varepsilon}\right)^{p-1}\right)$.*

Proof Consider the hyperrectangle $[2^{-M}, 2^M] \times \dots \times [2^{-M}, 2^M]$, in which all images of feasible solutions are contained. We create a grid by subdividing this hyperrectangle into smaller hyperrectangles such that, in each dimension, the ratio of the larger to the smaller coordinate is $1 + \varepsilon$ (see Fig. 1 on Page 9 for an illustration). Thus, the total number of subdivisions in each dimension is $\lceil \log_{1+\varepsilon} \left(\frac{2^M}{2^{-M}}\right) \rceil = \lceil \log_{1+\varepsilon} (2^{2M}) \rceil$. By the assumption on R' , any subset X' of feasible solutions admits an R' -dominating set $P_{X'} \subseteq X'$ of cardinality at most c . Moreover, any two solutions x and x' with images in the same hyperrectangle of the grid satisfy $x \preceq_\varepsilon x'$ by construction of the grid. Thus, if X' is the preimage of a nonempty hyperrectangle of the grid (i.e., X' is the set of all solutions $x \in X$ with images in this hyperrectangle), then the R' -dominating set $P_{X'}$ is actually an R -dominating set for X' . Consequently, forming the union of the sets $P_{X'}$ over all subsets X' that are preimages of nonempty hyperrectangles of the grid yields an R -approximation set. Moreover, since R is monotonic, it suffices to consider only weakly nondominated nonempty hyperrectangles.

To bound the cardinality of this R -approximation set, we note that, for each hyperrectangle $[l_1, u_1] \times \dots \times [l_p, u_p]$ of the grid, we have $u_i = (1 + \varepsilon) \cdot l_i$ for all $i \in \{1, \dots, p\}$, so the vectors $l = (l_1, \dots, l_p)$ and $u = (u_1, \dots, u_p)$ of lower and upper bounds of the hyperrectangle lie on a straight line through the origin. Calling the set of all hyperrectangles of the grid whose bounds lie on the same straight line through the origin a *diagonal*, we see that (1) at most one hyperrectangle on each diagonal is weakly nondominated and nonempty, and (2) the number of diagonals is in $\mathcal{O}(p \cdot \lceil \log_{1+\varepsilon} (2^{2M}) \rceil^{p-1}) = \mathcal{O}\left(\left(\frac{M}{\varepsilon}\right)^{p-1}\right)$. Hence, since the cardinality of the constructed R -approximation set is at most c times the number of weakly nondominated nonempty hyperrectangles and c is a constant, its cardinality is also in $\mathcal{O}\left(\left(\frac{M}{\varepsilon}\right)^{p-1}\right)$. \square

We remark that the proof of Theorem 3.1 also yields an R -approximation set of (fully) polynomial cardinality if the cardinality of the R' -dominating set $P_{X'}$ for each subset $X' \subseteq X$ is only required to be polynomial in the instance size (and in $\frac{1}{\varepsilon}$) instead of constant. In this case, the cardinality of the obtained R -approximation set would be upper bounded by $\mathcal{O}\left(\left(\frac{M}{\varepsilon}\right)^{p-1}\right)$ times the worst-case cardinality of $P_{X'}$ taken over all preimages X' of weakly nondominated nonempty hyperrectangles of the grid.

The following theorem shows that the general result in Theorem 3.1, in particular, yields existence results for $(1 + \varepsilon)$ -approximation sets, one-exact $(1 + \varepsilon)$ -approximation sets, and quasi- k -exact $(1 + \varepsilon)$ -approximation sets of fully polynomial cardinality. Here the existence results for $(1 + \varepsilon)$ -approximation sets and for one-exact $(1 + \varepsilon)$ -approximation sets, which are already known from Papadimitriou and Yannakakis (2000) and Herzel et al. (2021b), respectively, follow immediately from the theorem.

The result concerning quasi- k -exact $(1 + \varepsilon)$ -approximation sets, however, has not been shown in the literature before and requires a more involved proof, which is illustrated in Fig. 1.

Theorem 3.2 *For any instance (X, f) of a p -objective optimization problem and any $\varepsilon > 0$, there exists an R -approximation set of cardinality $\mathcal{O}((\frac{M}{\varepsilon})^{p-1})$*

- (1) for $R = \preceq_{\varepsilon}^1$ ($(1 + \varepsilon)$ -approximation set),
- (2) for $R = \preceq_{\varepsilon}^1$ (one-exact $(1 + \varepsilon)$ -approximation set),
- (3) for $R = \preceq_{\varepsilon}^{(k)}$ for any $k \leq \lceil \frac{p}{2} \rceil$ (quasi- k -exact $(1 + \varepsilon)$ -approximation set).

Proof (1) Choosing $R' := X \times X$ in Theorem 3.1 yields $R = \preceq_{\varepsilon}$. Moreover, R' is clearly monotonic and, for any nonempty subset $X' \subseteq X$, any solution $x \in X'$ constitutes an R' -dominating set, so the result follows from Theorem 3.1.

(2) Choosing $R' := \{(x, x') \in X \times X : f_1(x) \leq f_1(x')\}$ in Theorem 3.1 yields $R = \preceq_{\varepsilon}^1$. Again, R' is clearly monotonic. Moreover, given any nonempty subset $X' \subseteq X$, any solution $x \in X'$ such that $f_1(x) = \min_{x' \in X'} f_1(x')$ constitutes an R' -dominating set (note that, by Assumption 1, the number of possible images is finite, so the minimum is guaranteed to exist).

(3) Given $k \leq \lceil \frac{p}{2} \rceil$, choosing

$$R' := \{(x, x') \in X \times X : f_j(x) \leq f_j(x') \text{ for at least } k \text{ objectives } f_j\},$$

in Theorem 3.1 yields $R = \preceq_{\varepsilon}^{(k)}$. Since R' is clearly monotonic, it remains to show that there exists an R' -dominating set of constant cardinality for any nonempty subset $X' \subseteq X$. To show this, we use a result from Alon et al. (2006) about dominating sets in k -majority tournaments. A k -majority tournament is a directed graph on a finite vertex set V where, given $2k - 1$ linear orders of V , there exists a directed arc from u to v if and only if u is ranked higher than v in at least k of the orders. The result from Alon et al. (2006) states that any k -majority tournament admits a dominating set of cardinality $\mathcal{O}(k \log k)$.

In our situation, given some nonempty subset $X' \subseteq X$, each objective f_j defines a reflexive, transitive, and strongly connected order on $f(X')$ by ranking $f(x)$ higher than $f(x')$ if and only if $f_j(x) \leq f_j(x')$. This order can be made antisymmetric (i.e., turned into a linear order) by breaking ties within any subset of elements that have the same f_j -value via an arbitrary linear order on this subset. Since $k \leq \lceil \frac{p}{2} \rceil$, this construction yields $p \geq 2k - 1$ linear orders on $f(X')$. Moreover, by Assumption 1, the number of possible images the objective space is finite, so $f(X')$ is a finite set. Consequently, the directed graph on $f(X')$ in which there exists a directed arc from $f(x)$ to $f(x')$ if and only if $f(x)$ is ranked higher than $f(x')$ in at least k of the orders is finite and contains a k -majority tournament on $f(X')$ (since, in case that $p > 2k - 1$, the additional $p - 2k + 1$ linear orders only lead to additional arcs in the graph). Consequently, by the result of Alon et al. (2006), the graph contains a dominating set of constant cardinality $\mathcal{O}(k \log k)$. By construction of the graph, choosing one corresponding solution x for each image in this dominating set yields an R' -dominating set for X' . □




Fig. 1 Existence of quasi- k -exact $(1 + \epsilon)$ -approximation sets of fully polynomial cardinality. For each weakly nondominated nonempty hyperrectangle in the objective space as in the proof of Theorem 3.1, a dominating set for the supergraph of a k -majority tournament defined on the images in the hyperrectangle is selected to obtain a quasi- k -exact $(1 + \epsilon)$ -approximation set. There exists at most one weakly nondominated nonempty hyperrectangle in each diagonal of the grid. A quasi- k -exact $(1 + \epsilon)$ -approximation set is given by preimages of the bold points. Weakly nondominated nonempty hyperrectangles are indicated by shaded boxes. One diagonal is indicated by the bold lines. The constructed supergraph of a k -majority tournament is illustrated for one hyperrectangle

In the remainder of this section, we show that the existence results for one-exact $(1 + \epsilon)$ -approximation sets and quasi- k -exact $(1 + \epsilon)$ -approximation sets with $k = \lceil \frac{p}{2} \rceil$ from Theorem 3.2 are tight in the sense that polynomial-cardinality $(1 + \epsilon)$ -approximation sets with additional exact components do not exist in general. This is made precise for the case of one-exact $(1 + \epsilon)$ -approximation sets in the following theorem.

Theorem 3.3 *There exist instances of p -objective combinatorial optimization problems, including the multiobjective versions of shortest path, minimum spanning tree, assignment, minimum s - t -cut, knapsack, and traveling salesman, that do not admit a polynomial-cardinality one-exact, quasi-2-exact $(1 + \epsilon)$ -approximation set for any $\epsilon > 0$.*

Proof The biobjective versions of all of these problems are known to be intractable, i.e., there exist instances for which the cardinality of the nondominated set is exponential in the instance size. From such an intractable biobjective instance, construct a p -objective instance where objective f_1 corresponds to the first objective function of the biobjective instance and the other $p - 1$ objective functions correspond to the second objective function of the biobjective instance. Then, at least one solution corresponding to each nondominated point of the biobjective instance is required in any one-exact, quasi-2-exact $(1 + \epsilon)$ -approximation set of the p -objective instance. \square

The above theorem shows that, for many important problems, one-exact $(1 + \varepsilon)$ -approximation sets are the best polynomial-cardinality partially exact approximation sets one can hope for when requiring to be exact in one fixed objective. In fact, we are not aware of many (non-trivial) problems that admit polynomial-cardinality one-exact, quasi-2-exact $(1 + \varepsilon)$ -approximation sets. One notable exception is the biobjective minimum cut problem, for which it is known that the cardinality of the efficient set is polynomial in the instance size (Aissi et al. 2015). Since the efficient set corresponds to a two-exact $(1 + \varepsilon)$ -approximation set in the biobjective case, this in particular shows the existence of a polynomial-cardinality one-exact, quasi-2-exact $(1 + \varepsilon)$ -approximation set.

We now consider quasi- k -exact $(1 + \varepsilon)$ -approximation sets, that is, approximation sets that approximate each feasible solution exactly in k objectives (and with a factor of $(1 + \varepsilon)$ -approximate in all other objectives), but allow the objectives in which the approximation is exact to differ depending on the approximated solution. Here, a similar argument as in the proof of Theorem 3.3 shows that, for many multiobjective combinatorial optimization problems, polynomial-cardinality quasi- k -exact $(1 + \varepsilon)$ -approximation sets for $k > \lceil \frac{p}{2} \rceil$ are not guaranteed to exist for all instances.

Theorem 3.4 *There exist instances of many p -objective combinatorial optimization problems, including the multiobjective versions of shortest path, minimum spanning tree, assignment, minimum s - t -cut, knapsack, and traveling salesman, that do not admit a polynomial-cardinality quasi- k -exact $(1 + \varepsilon)$ -approximation set for any $k > \lceil \frac{p}{2} \rceil$ and any $\varepsilon > 0$.*

Proof As noted in the proof of Theorem 3.3, the biobjective versions of all of these problems are known to be intractable. From such an intractable biobjective instance, construct a p -objective instance where the first $\lceil \frac{p}{2} \rceil$ objectives corresponds to the first objective of the biobjective instance and the other $\lfloor \frac{p}{2} \rfloor$ objectives correspond to the second objective function of the biobjective instance. Then, for $k > \lceil \frac{p}{2} \rceil$, at least one solution corresponding to each nondominated point of the biobjective instance is required in any quasi- k -exact $(1 + \varepsilon)$ -approximation set of the p -objective instance. \square

Figure 2 summarizes the known tractability and intractability results for different types of approximation sets.

4 Minimum-cardinality approximation sets

Theorem 3.2 implies that an asymptotic upper bound for the minimum cardinality of an ordinary, a one-exact, and a quasi- k -exact $(1 + \varepsilon)$ -approximation set with $k \leq \lceil \frac{p}{2} \rceil$ is in $\mathcal{O}\left(\left(\frac{M}{\varepsilon}\right)^{p-1}\right)$ and it is easy to see that this bound is tight, i.e., there exist instances where even a minimum-cardinality ordinary $(1 + \varepsilon)$ -approximation set consists of $\Theta\left(\left(\frac{M}{\varepsilon}\right)^{p-1}\right)$ solutions. For a specific instance, however, as pointed out in Vassilvitskii and Yannakakis (2005) and Diakonikolas and Yannakakis (2009), there might still exist $(1 + \varepsilon)$ -approximation sets of vastly different cardinalities. This motivates to study

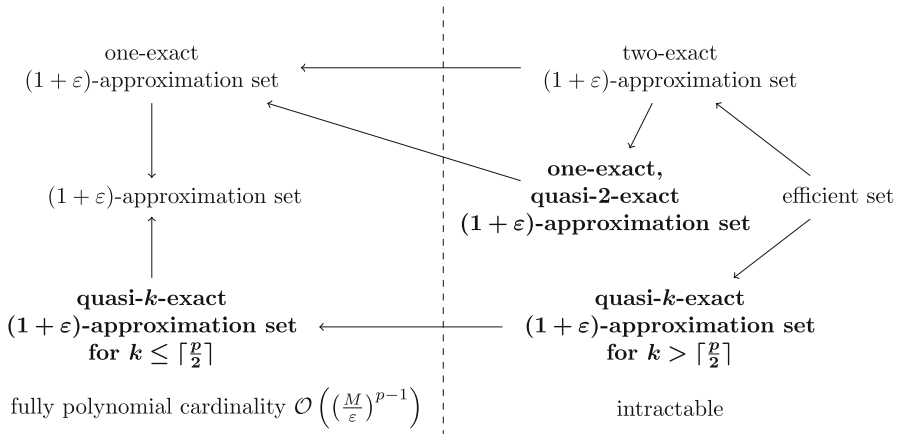


Fig. 2 Tractability and intractability results for different types of approximation sets. An arrow indicates that one type of approximation set fulfills the requirements of the other. The dashed line marks the boundary between tractability and intractability in general p -objective optimization problems. New results obtained in this paper are shown in bold. The result on $(1 + \varepsilon)$ -approximation sets is shown in Papadimitriou and Yannakakis (2000), and the results on one-exact and two-exact $(1 + \varepsilon)$ -approximation sets are shown in Herzel et al. (2021b)

properties of $(1 + \varepsilon)$ -approximation sets of minimum cardinality. Moreover, it raises the question whether the minimum possible cardinality can increase for a specific instance when considering partially exact $(1 + \varepsilon)$ -approximation sets. As shown in Herzel et al. (2021b), this is indeed the case when considering one-exact $(1 + \varepsilon)$ -approximation sets: for any positive integer n , there exist instances for which the minimum cardinality of a one-exact $(1 + \varepsilon)$ -approximation set is more than n times larger than the minimum cardinality of an ordinary $(1 + \varepsilon)$ -approximation set.

Hence, in this section, we systematically study minimum-cardinality ordinary and partially exact $(1 + \varepsilon)$ -approximation sets. In Sect. 4.1, we first investigate the (weak) efficiency of solutions in minimum-cardinality approximation sets and show that minimum-cardinality ordinary $(1 + \varepsilon)$ -approximation sets and quasi-1-exact $(1 + \varepsilon)$ -approximation sets can consist of dominated solutions only. For minimum-cardinality one-exact $(1 + \varepsilon)$ -approximation sets, however, we show that one contained solution must be weakly efficient, but not more than one in general. Afterwards, in Sect. 4.2, we relate the minimum cardinalities of different types of partially exact approximation sets to the minimum cardinality of an ordinary $(1 + \varepsilon)$ -approximation set. Here, we show that requiring the $(1 + \varepsilon)$ -approximation set to be quasi-1-exact never changes the minimum cardinality for any instance. In contrast, for any positive integer n , there exist instances for which requiring to be one-exact or quasi-2-exact increases the minimum cardinality by more than a factor of n .

4.1 Efficiency of minimum-cardinality approximation sets

In this subsection, we investigate the question of (weak) efficiency of solutions in minimum-cardinality ordinary and partially exact $(1 + \varepsilon)$ -approximation sets.

Since any dominated solution in any type of partially exact or ordinary $(1 + \varepsilon)$ -approximation set can always be replaced by an efficient solution dominating it without impairing the obtained approximation guarantee, it follows that there always exist minimum-cardinality $(1 + \varepsilon)$ -approximation sets of each type that consist of efficient solutions only. In contrast, it is well known that, even in the biobjective case, there are ordinary $(1 + \varepsilon)$ -approximation sets of minimum cardinality that consist of strictly dominated solutions only. The following proposition generalizes this result to quasi-1-exact $(1 + \varepsilon)$ -approximation sets.

Proposition 4.1 *For any $\varepsilon > 0$, there is an instance of a biobjective optimization problem that admits a quasi-1-exact $(1 + \varepsilon)$ -approximation set of minimum cardinality consisting of strictly dominated solutions only.*

Proof The proof is illustrated in Fig. 3. Consider a biobjective optimization problem instance consisting of six feasible solutions $x^1, x^2, x^3, x^4, x^5, x^6$ with

$$\begin{array}{ll} f_1(x^1) = 1, & f_2(x^1) = (1 + \varepsilon)^2, \\ f_1(x^2) = 1 + \frac{\varepsilon}{2}, & f_2(x^2) = (1 + \varepsilon) \cdot \left(1 + \frac{\varepsilon}{4}\right), \\ f_1(x^3) = (1 + \varepsilon) \cdot \left(1 + \frac{\varepsilon}{4}\right), & f_2(x^3) = 1 + \frac{\varepsilon}{2}, \\ f_1(x^4) = (1 + \varepsilon)^2, & f_2(x^4) = 1, \\ f_1(x^5) = 1 + \varepsilon, & f_2(x^5) = (1 + \varepsilon) \cdot \left(1 + \frac{\varepsilon}{2}\right), \\ f_1(x^6) = (1 + \varepsilon) \cdot \left(1 + \frac{\varepsilon}{2}\right), & f_2(x^6) = 1 + \varepsilon. \end{array}$$

First, observe that there is no solution that \leq_ε -dominates the five other solutions, which means that any (quasi-1-exact) $(1 + \varepsilon)$ -approximation set consists of at least two solutions. However, even though x^5 and x^6 are strictly dominated by x^2 and x^3 , respectively, the set $\{x^5, x^6\}$ is a quasi-1-exact $(1 + \varepsilon)$ -approximation set: x^1 is $(1 + \varepsilon, 1)$ -approximated by x^5 , x^2 is $(1 + \varepsilon, 1)$ -approximated by x^6 , x^3 is $(1, 1 + \varepsilon)$ -approximated by x^5 , and x^4 is $(1, 1 + \varepsilon)$ -approximated by x^6 . \square

For one-exact $(1 + \varepsilon)$ -approximation sets, the definition implies that each such set must contain at least one weakly efficient solution, namely a solution with minimum f_1 -value. The following result, however, shows that, even in the biobjective case, there are minimum-cardinality one-exact $(1 + \varepsilon)$ -approximation sets in which all other solutions are strictly dominated.

Proposition 4.2 *For any $\varepsilon > 0$ and any positive integer $n \in \mathbb{N}_+$, there exists an instance of a biobjective optimization problem that admits a one-exact $(1 + \varepsilon)$ -approximation set of minimum cardinality $n + 1$ in which all but one of the solutions are strictly dominated.*

Proof Given $\varepsilon > 0$ and $n \in \mathbb{N}_+$, let $\delta > 0$ so that $(1 + \delta)^{2n} = 1 + \varepsilon$ and consider the following instance of a biobjective optimization problem: The feasible set is given as




Fig. 3 Illustration of the instance constructed in the proof of Proposition 4.1. The set $\{x^5, x^6\}$ is a minimum-cardinality quasi-1-exact $(1 + \varepsilon)$ -approximation set. The regions that are approximated by x^5 and x^6 are indicated by a solid and a dashed line, respectively

$X := \{x^0, \dots, x^n, \bar{x}^1, \dots, \bar{x}^n, \tilde{x}^0, \dots, \tilde{x}^n\}$ and

$$\begin{aligned}
 f_1(x^0) &:= 1, & f_2(x^0) &:= (1 + \varepsilon)^n, \\
 f_1(x^i) &:= 3i + 1, & f_2(x^i) &:= (1 + \varepsilon)^{n-i} \cdot (1 + \delta)^i, \\
 f_1(\bar{x}^i) &:= 3i, & f_2(\bar{x}^i) &:= (1 + \varepsilon)^{n-i} \cdot (1 + \delta)^{i-1}, \\
 f_1(\tilde{x}^i) &:= 3i + 2, & f_2(\tilde{x}^i) &:= (1 + \varepsilon)^{n-i} \cdot \frac{1}{(1 + \delta)^i},
 \end{aligned}$$

which implies that

$$f_2(\bar{x}^i) = \frac{1}{1 + \delta} \cdot f_2(x^i) = \frac{1}{1 + \varepsilon} \cdot f_2(x^{i-1}), \text{ and} \tag{1}$$

$$f_2(\tilde{x}^i) = \frac{1}{(1 + \delta)^{2i}} \cdot f_2(x^i). \tag{2}$$

We now show that $\{x^0, x^1, \dots, x^n\}$ is a minimum-cardinality one-exact $(1 + \varepsilon)$ -approximation set, even though x^0 is the only solution in this set that is not strictly

dominated. To this end, first note that

$$\begin{aligned}
 f_1(x^0) &< f_1(\bar{x}^1) < f_1(x^1) < f_1(\tilde{x}^1) \\
 &< f_1(\bar{x}^2) < f_1(x^2) < f_1(\tilde{x}^2) \\
 &< \dots \\
 &< f_1(\bar{x}^n) < f_1(x^n) < f_1(\tilde{x}^n).
 \end{aligned}
 \tag{3}$$

Thus, x^0 is not \preceq_{ε}^1 -dominated by any other solution, so x^0 must be contained in every one-exact $(1 + \varepsilon)$ -approximation set. Moreover, the following statements hold:

- \bar{x}^i strictly dominates x^i for $i \geq 1$,
- $x^i \preceq_{\varepsilon}^1 \tilde{x}^i$ for $i \geq 1$ (by (2) and (3)),
- $x^i \preceq_{\varepsilon}^1 \bar{x}^{i+1}$ (by (1) and (3)), and, therefore, also $x^i \preceq_{\varepsilon}^1 x^{i+1}$ for $0 \leq i \leq n - 1$,
- Neither x^i , nor \bar{x}^i , nor $\tilde{x}^i \preceq_{\varepsilon}^1$ -dominates \tilde{x}^{i+1} for $0 \leq i \leq n - 1$ since

$$\begin{aligned}
 f_2(\tilde{x}^{i+1}) &= \frac{1}{(1 + \varepsilon) \cdot (1 + \delta)} \cdot f_2(\tilde{x}^i) \\
 &< \frac{1}{1 + \varepsilon} \cdot f_2(\tilde{x}^i) < \frac{1}{1 + \varepsilon} \cdot f_2(\bar{x}^i) < \frac{1}{1 + \varepsilon} \cdot f_2(x^i).
 \end{aligned}$$

- Neither x^i , nor \bar{x}^i , nor $\tilde{x}^i \preceq_{\varepsilon}^1$ -dominates \tilde{x}^j for $0 \leq j < i \leq n$ (by (3))
- Neither x^i , nor \bar{x}^i , nor $\tilde{x}^i \preceq_{\varepsilon}^1$ -dominates \tilde{x}^j for $i + 2 \leq j \leq n$ since

$$f_2(\tilde{x}^j) < f_2(\tilde{x}^{i+1}) < \frac{1}{1 + \varepsilon} \cdot f_2(\tilde{x}^i) < \frac{1}{1 + \varepsilon} \cdot f_2(\bar{x}^i) < \frac{1}{1 + \varepsilon} \cdot f_2(x^i).$$

In total, this implies that $\{x^0, x^1, \dots, x^n\}$ is a one-exact $(1 + \varepsilon)$ -approximation set, even though x^0 is the only solution in this set that is not strictly dominated. Moreover, for each $i \in \{1, \dots, n\}$, in order to \preceq_{ε}^1 -dominate \tilde{x}^i , at least one of the three solutions x^i , \bar{x}^i , and \tilde{x}^i must be contained in any one-exact $(1 + \varepsilon)$ -approximation set. Consequently, any one-exact $(1 + \varepsilon)$ -approximation set has cardinality at least $n + 1$. □

The final result of this subsection, which will be used when analyzing the minimum cardinality of quasi-1-exact $(1 + \varepsilon)$ -approximation sets in the following subsection, states that any $(1 + \varepsilon)$ -approximation set consisting only of weakly efficient solutions must actually be quasi-1-exact.

Proposition 4.3 *Any $(1 + \varepsilon)$ -approximation set in which all solutions are weakly efficient is a quasi-1-exact $(1 + \varepsilon)$ -approximation set.*

Proof Let P be a $(1 + \varepsilon)$ -approximation set in which all solutions are weakly efficient. Let $x' \in X$ be an arbitrary feasible solution and let $x \in P$ be a solution that \preceq_{ε} -dominates x' . If x' is not $\preceq_{\varepsilon}^{(1)}$ -dominated by x , i.e., we do not have $f_i(x) \leq f_i(x')$ for any $i \in \{1, \dots, p\}$, we immediately obtain that x is strictly dominated by x' , which contradicts the weak efficiency of x . □

4.2 Relations between minimum cardinalities

We now relate the minimum cardinalities of different types of partially exact approximation sets to the minimum cardinality of an ordinary $(1 + \varepsilon)$ -approximation set.

Concerning one-exact $(1 + \varepsilon)$ -approximation sets, it is shown in Herzel et al. (2021b) that, for any positive integer n , there exist biobjective problem instances for which the minimum cardinality of a such a set is more than n times larger than the minimum cardinality of an ordinary $(1 + \varepsilon)$ -approximation set. In contrast to this, we now use Proposition 4.3 to show that requiring a quasi-1-exact $(1 + \varepsilon)$ -approximation set instead of only an ordinary $(1 + \varepsilon)$ -approximation set never changes the minimum cardinality for any instance.

Theorem 4.4 *For any instance of a multiobjective optimization problem and any $\varepsilon > 0$, there exists a minimum-cardinality ordinary $(1 + \varepsilon)$ -approximation set that is quasi-1-exact. In particular, the minimum cardinalities of an ordinary $(1 + \varepsilon)$ -approximation set and of a quasi-1-exact $(1 + \varepsilon)$ -approximation set coincide for every instance.*

Proof Let P^* be a smallest ordinary $(1 + \varepsilon)$ -approximation set. Turn P^* into a quasi-1-exact $(1 + \varepsilon)$ -approximation set P' of the same cardinality that consists of weakly efficient solutions only by replacing any strictly dominated solution $x \in P^*$ by a weakly efficient solution that strictly dominates it (this does not impair the approximation guarantee). By Proposition 4.3, the set P' is then indeed a quasi-1-exact $(1 + \varepsilon)$ -approximation set. □

Contrasting the result on the minimum cardinality of quasi-1-exact $(1 + \varepsilon)$ -approximation sets shown in Theorem 4.4, we now show that, for any positive integer n , there exist instances for which the minimum cardinality of a quasi-2-exact $(1 + \varepsilon)$ -approximation set is more than n times larger than the minimum cardinality of an ordinary $(1 + \varepsilon)$ -approximation set—even in the case of three objectives ($p = 3$). The proof of the following theorem generalizes the construction used in the proof of Theorem 2 in Herzel et al. (2021b).

Theorem 4.5 *For any $\varepsilon > 0$ and any positive integer $n \in \mathbb{N}_+$, there exist instances of 3-objective optimization problems such that $|P^{(2),*}| > n \cdot |P^*|$, where $P^{(2),*}$ denotes a minimum-cardinality quasi-2-exact $(1 + \varepsilon)$ -approximation set and P^* denotes a minimum-cardinality ordinary $(1 + \varepsilon)$ -approximation set.*

Proof Given $\varepsilon > 0$ and $n \in \mathbb{N}_+$, we construct an instance of a 3-objective optimization problem with $|P^{(2),*}| = n + 1$ and $|P^*| = 1$ as follows: The feasible set is given as $X := \{x^0, \dots, x^n\}$ and, for $j = 0, \dots, n$, we define $f(x^j) = (f_1(x^j), f_2(x^j), f_3(x^j))$ via

$$f(x^j) := \left(1 + \frac{n-j}{n} \cdot \varepsilon, 1 + \frac{n-j}{n} \cdot \varepsilon, (1 + \varepsilon)^{2j+1} \right).$$

Note that this implies that

$$f_i(x^0) > f_i(x^1) > \dots > f_i(x^n) \quad \text{for } i = 1, 2,$$

but

$$f_3(x^j) = \frac{1}{(1+\varepsilon)^2} \cdot f_3(x^{j+1}) < \frac{1}{1+\varepsilon} \cdot f_3(x^{j+1})$$

for $j = 0, \dots, n-1$. Consequently, no solution $\preceq_{\varepsilon}^{(2)}$ -dominates any other solution, which implies that $P^{(2),*} = X$ with cardinality $n+1$. The solution x^0 with $f(x^0) = (1+\varepsilon, 1+\varepsilon, 1+\varepsilon)$, however, \preceq_{ε} -dominates all other solutions, so $\{x^0\}$ is a $(1+\varepsilon)$ -approximation set. \square

5 Computation of approximation sets

After studying the existence of (fully) polynomial-cardinality approximation sets and the impact of using partially exact approximate dominance relations on the minimum cardinality of approximation sets, we now consider the computability of such sets.

In the literature, several auxiliary problems have been defined for the computation of (fully) polynomial-cardinality approximation sets (Papadimitriou and Yannakakis 2000; Vassilvitskii and Yannakakis 2005; Diakonikolas and Yannakakis 2009; Herzel et al. 2021b), which are also used in the following. The first such problem, whose (fully) polynomial-time solvability has been shown to characterize the (fully) polynomial-time compatibility of $(1+\varepsilon)$ -approximation sets in Papadimitriou and Yannakakis (2000) is the *gap problem*:

Definition 5.1 (*Gap Problem*) Given an instance (X, f) of a p -objective optimization problem, a point $b \in \mathbb{Q}^p$, and some $\delta > 0$, the problem $\text{GAP}_{\delta}(b)$ is the following: Either return a feasible solution $x \in X$ with $f_i(x) \leq b_i$ for $i = 1, \dots, p$, or answer correctly that there does not exist any feasible solution x' such that $f_i(x') \leq \frac{b_i}{1+\delta}$ for all $i = 1, \dots, p$.

The following auxiliary problem, which scalarizes the multiobjective problem via budget constraints on all but one objective function, is widely used both in practice and in the theoretical literature on multiobjective optimization:

Definition 5.2 (*Budget-Constrained Problem*) Given an instance (X, f) of a multiobjective optimization problem and bounds $b_i > 0$, $i = 2, \dots, p$, for all objective functions except the first one, the problem $\text{CONSTRAINED}(b_2, \dots, b_p)$ is the following: Either answer that there does not exist a feasible solution $x' \in X$ with $f_i(x') \leq b_i$ for $i = 1, \dots, p$, or return a feasible solution that minimizes f_1 among all such solutions, i.e., return $x \in X$ with

$$f_1(x) = \text{OPT}_1(b_2, \dots, b_p) := \min_{x' \in X} \{f_1(x') : f_i(x') \leq b_i \text{ for } i = 2, \dots, p\}, \text{ and} \\ f_i(x) \leq b_i, \quad i = 2, \dots, p.$$

Note that CONSTRAINED , which is also often referred to as the ε -constrained problem in the literature, is hard to solve even for the biobjective version of most relevant

problems such as shortest path. A way to circumvent the hardness of CONSTRAINED is to consider solutions that violate the given bounds slightly, while requiring an objective value that is at least as good as the objective value of any solution that respects the bounds (Diakonikolas and Yannakakis 2009):

Definition 5.3 (DUALRESTRICT) Given an instance (X, f) of a multiobjective optimization problem, bounds $b_i > 0, i = 2, \dots, p$, for all objective functions except the first one, and some $\delta > 0$, the problem $\text{DUALRESTRICT}_\delta(b_2, \dots, b_p)$ is the following: Either answer that there does not exist a feasible solution $x' \in X$ with $f_i(x') \leq b_i$ for $i = 1, \dots, p$, or return $x \in X$ with

$$\begin{aligned} f_1(x) &\leq \text{OPT}_1(b_2, \dots, b_p) \\ f_i(x) &\leq (1 + \delta) \cdot b_i, \quad i = 2, \dots, p. \end{aligned}$$

Note that CONSTRAINED can be viewed as the limit case where $\delta = 0$ in DUALRESTRICT.

The auxiliary problems CONSTRAINED and DUALRESTRICT can also be defined such that, instead of the first one, some other objective is to be optimized subject to budgets on the remaining objectives. In the following, we sometimes indicate which objective is optimized by an upper index. For example, DUALRESTRICT^i denotes the DUALRESTRICT problem with a bound on all objectives but the i -th one.

As shown in Papadimitriou and Yannakakis (2000), a $(1 + \varepsilon)$ -approximation set can be computed in (fully) polynomial time for every $\varepsilon > 0$ if and only if the auxiliary problem $\text{GAP}_\delta(b)$ can be solved in (fully) polynomial time for every vector b and every $\delta > 0$. Hence, a natural question is whether (fully) polynomial-time solvability of $\text{GAP}_\delta(b)$ is also sufficient for the (fully) polynomial-time computability of partially exact $(1 + \varepsilon)$ -approximation sets. The following theorem gives a negative answer to this question even for biobjective problems and the least-demanding type of partially exact $(1 + \varepsilon)$ -approximation sets. The proof is partly similar to the proof of Theorem 1 in Vassilvitskii and Yannakakis (2005), where it is shown that polynomial-time algorithms only generating solutions by solving $\text{GAP}_\delta(b)$ cannot approximate the minimum cardinality of an ordinary $(1 + \varepsilon)$ -approximation set to a factor better than 3 in the biobjective case (such algorithms are called *generic* in Vassilvitskii and Yannakakis 2005).

Theorem 5.4 *There exists no polynomial-time algorithm that returns a quasi-1-exact $(1 + \varepsilon)$ -approximation set for every biobjective problem instance and every $\varepsilon > 0$, and generates feasible solutions only by solving $\text{GAP}_\delta(b)$ for points b of polynomial encoding length and polynomial values of $\frac{1}{\delta}$.*

Proof Given some (rational) $\varepsilon > 0$, consider two instances I_1, I_2 with feasible sets $X^{I_1} = \{x^1\}$ and $X^{I_2} = \{x^1, x^2\}$ and the objective function values given by

$$f(x^1) = \left(1 + \frac{1}{l(\varepsilon)}, 1 + \frac{1}{l(\varepsilon)}\right), \quad f(x^2) = (1, 1),$$

where $l(\varepsilon)$ is a positive integer that is exponential in the input size and in $\frac{1}{\varepsilon}$. We show that no algorithm that generates feasible solutions only by solving $\text{GAP}_\delta(b)$ for values $\delta \geq \frac{1}{l(\varepsilon)}$ can distinguish between the two instances (i.e., detect whether x^2 is part of the feasible set). Since any smaller value of δ would be exponential in the input size and in $\frac{1}{\varepsilon}$ by definition of $l(\varepsilon)$, and any quasi-1-exact $(1 + \varepsilon)$ -approximation set for instance I_2 must include x^2 , this will show the claim.

So consider a call of $\text{GAP}_\delta(b)$ for some point b and some $\delta \geq \frac{1}{l(\varepsilon)}$. We distinguish two cases in order to show that $\text{GAP}_\delta(b)$ can always return either x^1 or NO as a correct answer for both instance I_1 and instance I_2 . If $b_i \geq f_i(x^1)$ for $i = 1, 2$, then $\text{GAP}_\delta(b)$ can return x^1 for both instances. Otherwise, we have $b_j < f_j(x^1) = 1 + \frac{1}{l(\varepsilon)}$ for some $j \in \{1, 2\}$. Using that $\delta \geq \frac{1}{l(\varepsilon)}$, this implies that

$$f_j(x^2) = 1 \geq \frac{1 + \frac{1}{l(\varepsilon)}}{1 + \delta} > \frac{b_j}{1 + \delta}.$$

Consequently, for both instances, there exists no solution with j -th objective value less than or equal to $\frac{b_j}{1 + \delta}$, so $\text{GAP}_\delta(b)$ can return NO. This shows that an algorithm that generates feasible solutions only by solving $\text{GAP}_\delta(b)$ for values of $\delta \geq \frac{1}{l(\varepsilon)}$ cannot detect whether x^2 is part of the feasible set as claimed. \square

Theorem 5.4 shows that harder-to-solve auxiliary problems than the gap problem must be used for computing partially exact $(1 + \varepsilon)$ -approximation sets in polynomial time. We now use Proposition 4.3 to show that, for biobjective problems, several known algorithms for computing (small-cardinality) $(1 + \varepsilon)$ -approximation sets based on such auxiliary problems actually yield quasi-1-exact $(1 + \varepsilon)$ -approximation sets.

Theorem 5.5 (1) *For biobjective optimization problems for which CONSTRAINED^1 and CONSTRAINED^2 are solvable in (fully) polynomial time, a quasi-1-exact $(1 + \varepsilon)$ -approximation set with minimum cardinality can be computed in (fully) polynomial time for every instance.*

(2) *For biobjective optimization problems for which $\text{DUALRESTRICT}_\delta^1$ or $\text{DUALRESTRICT}_\delta^2$ is solvable in (fully) polynomial time, a quasi-1-exact $(1 + \varepsilon)$ -approximation set with cardinality at most twice the cardinality of a smallest quasi-1-exact $(1 + \varepsilon)$ -approximation set can be computed in (fully) polynomial time for every instance.*

Proof (1) If CONSTRAINED^1 and CONSTRAINED^2 are solvable in (fully) polynomial time, an ordinary $(1 + \varepsilon)$ -approximation set with minimum cardinality can be computed in (fully) polynomial time for every instance by a greedy procedure as shown in Diakonikolas and Yannakakis (2009). Since this set contains only weakly efficient solutions (Bazgan et al. 2017), it is actually a quasi-1-exact $(1 + \varepsilon)$ -approximation set of minimum cardinality by Proposition 4.3.

(2) If $\text{DUALRESTRICT}_\delta^1$ or $\text{DUALRESTRICT}_\delta^2$ is solvable in (fully) polynomial time, an ordinary $(1 + \varepsilon)$ -approximation set with cardinality at most twice the cardinality of a smallest ordinary $(1 + \varepsilon)$ -approximation set (which equals the minimum cardinality of a quasi-1-exact $(1 + \varepsilon)$ -approximation set by Theorem 4.4) can be computed

in (fully) polynomial time for every instance as shown in Diakonikolas and Yannakakis (2009) and Bazgan et al. (2017). Moreover, the $(1 + \varepsilon)$ -approximation sets returned by the algorithms in Diakonikolas and Yannakakis (2009) and Bazgan et al. (2017) contain only weakly efficient solutions (Bazgan et al. 2017). Consequently, by Proposition 4.3, these sets are actually quasi-1-exact $(1 + \varepsilon)$ -approximation sets. \square

Examples of biobjective optimization problems for which $\text{DUALRESTRICT}_\delta^1$ or $\text{DUALRESTRICT}_\delta^2$ is solvable in polynomial time include, e.g., biobjective shortest path and biobjective spanning tree (Diakonikolas and Yannakakis 2009). We remark that, as shown in Herzel et al. (2021b), the same results as in Theorem 5.5 can be achieved even for one-exact $(1 + \varepsilon)$ -approximation sets, but modified algorithms targeted specifically at the computation of one-exact $(1 + \varepsilon)$ -approximation sets are necessary in this case. Moreover, it has been shown in Herzel et al. (2021b) that, for general p -objective problems, one-exact $(1 + \varepsilon)$ -approximation sets can be computed in (fully) polynomial time if and only if $\text{DUALRESTRICT}_\delta^1(b_2, \dots, b_p)$ can be solved for any choice of the bounds b_2, \dots, b_p and any $\delta > 0$ in (fully) polynomial time.

We finish this section by considering the case where all feasible solutions are given explicitly in the input. For this case, we now present two related approaches for the polynomial-time computation of R -approximation sets for approximate dominance relations R as in Theorem 3.1 under the assumption that checking whether a given solution $x \in X$ R' -dominates another given solution $x' \in X$ (for R' as in the theorem) is possible in polynomial time. This class of relations, in particular, includes quasi- k -exact $(1 + \varepsilon)$ -dominance for any $k \leq \lceil \frac{p}{2} \rceil$ as well as one-exact $(1 + \varepsilon)$ -dominance and ordinary $(1 + \varepsilon)$ -dominance.

The first approach is based directly on the proof of Theorem 3.1. Given that checking whether a given solution $x \in X$ R' -dominates another given solution $x' \in X$ (for R' as in the theorem) is possible in polynomial time, we can generate the directed graph corresponding to the approximate dominance relation R' in each weakly nondominated nonempty hyperrectangle considered in the proof in polynomial time and compute a dominating set in each of these graphs. While computing a minimum-cardinality dominating set is NP-hard in general, a simple greedy algorithm can be used to compute a dominating set that is larger by at most a factor $1 + \log|V|$, where V denotes the vertex set (Johnson 1974). This yields a polynomial-cardinality set in our situation since $|V|$ is bounded by $|X|$ in all cases, which is polynomial in the input size since X is given explicitly in the input.

A second related approach, which additionally yields a bound on the cardinality of the obtained R -approximation set, consists of constructing the directed graph on X corresponding to the approximate dominance relation R (which is possible in polynomial time given that R' -dominance of solutions can be checked in polynomial time) and to compute a dominating set of cardinality at most $1 + \log|X|$ times the minimum cardinality of dominating set in this graph using the algorithm from Johnson (1974). The resulting set then corresponds to an R -approximation set of cardinality at most $1 + \log|X|$ times the minimum cardinality of an R -approximation set.

6 Conclusion

In this paper, we explore the borderline between tractability and intractability when approximating multiobjective optimization problems using more demanding approximate dominance relations than $(1 + \varepsilon)$ -dominance. We show that, under very weak assumptions, there always exist $(1 + \varepsilon)$ -approximation sets of fully polynomial cardinality such that every feasible solution is, in fact, 1-approximated with respect to at least half of the objectives functions if we allow these objective functions to differ between different feasible solutions. If a polynomial-cardinality approximation set is required to be exact in one specified objective function, however, exactness cannot be ensured in a second objective function in general, even if this second objective function is allowed to differ depending on the approximated solution. This leads to the “frontier of intractability” being located in between one-exact $(1 + \varepsilon)$ -approximation sets and one-exact, quasi-2-exact $(1 + \varepsilon)$ -approximation sets and in between quasi- $\lceil \frac{p}{2} \rceil$ -exact $(1 + \varepsilon)$ -approximation sets and quasi- $(\lceil \frac{p}{2} \rceil + 1)$ -exact $(1 + \varepsilon)$ -approximation sets.

While first positive and negative results concerning the polynomial-time computability of quasi- k -exact $(1 + \varepsilon)$ -approximation sets have been obtained in Sect. 5, an important question for future research is to obtain further computability results that extend beyond the biobjective case or the situation where the feasible set is given explicitly in the input. Specifically, since both the polynomial-time computability of ordinary $(1 + \varepsilon)$ -approximation sets and of one-exact $(1 + \varepsilon)$ -approximation sets have been characterized by the polynomial-time solvability of certain auxiliary problems (GAP and DUALRESTRICT, respectively), it would be interesting to investigate whether a similar characterization via an auxiliary problem can also be obtained for quasi- k -exact $(1 + \varepsilon)$ -approximation sets for $k \leq \lceil \frac{p}{2} \rceil$. As our results shown, such an auxiliary problem would have to be strictly harder to solve than GAP.

Funding. Open Access funding enabled and organized by Projekt DEAL. This work was supported by the bilateral cooperation project “Approximation methods for multiobjective optimization problems” funded by the German Academic Exchange Service (DAAD, Project-ID 57388848) and by Campus France, PHC PROCOPE 2018 (Project no. 40407WF) as well as by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—Project number 398572517.

Data Availability Not applicable

Declarations

Competing Interests The authors have no relevant financial or non-financial interests to disclose.

Author Contributions All authors contributed to the study conception and design and the writing of the manuscript. All authors read and approved the final manuscript.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aissi H, Mahjoub AR, McCormick ST, Queyranne M (2015) Strongly polynomial bounds for multiobjective and parametric global minimum cuts in graphs and hypergraphs. *Math Program* 154(1–2):3–28
- Alon N, Brightwell GR, Kierstead HA, Kostochka AV, Winkler P (2006) Dominating sets in k -majority tournaments. *J Comb Theory Ser B* 96(3):374–387
- Bazgan C, Jamain F, Vanderpooten D (2015) Approximate Pareto sets of minimal size for multi-objective optimization problems. *Oper Res Lett* 43(1):1–6
- Bazgan C, Jamain F, Vanderpooten D (2017) Discrete representation of the non-dominated set for multi-objective optimization problems using kernels. *Eur J Oper Res* 260(3):814–827
- Diakonikolas I, Yannakakis M (2009) Small approximate Pareto sets for biobjective shortest paths and other problems. *SIAM J Comput* 39(4):1340–1371
- Ehrgott M (1998) Hard to say it's easy—four reasons why combinatorial multiobjective programmes are hard. In: *Proceedings of the 14th international conference on multiple criteria decision making (MCDM)*. Lecture notes in economics and mathematical systems, vol 487, pp 69–80
- Herzel A, Ruzika S, Thielen C (2021a) Approximation methods for multiobjective optimization problems: a survey. *INFORMS J Comput* 33(4):1284–1299
- Herzel A, Bazgan C, Ruzika S, Thielen C, Vanderpooten D (2021b) One-exact approximate pareto sets. *J Glob Optim* 80(1):87–115
- Johnson DS (1974) Approximation algorithms for combinatorial problems. *J Comput Syst Sci* 9(3):256–278
- Koltun V, Papadimitriou C (2007) Approximately dominating representatives. *Theor Comput Sci* 371(3):148–154
- Papadimitriou C, Yannakakis M (1991) Optimization, approximation, and complexity classes. *J Comput Syst Sci* 43(3):425–440
- Papadimitriou C, Yannakakis M (2000) On the approximability of trade-offs and optimal access of web sources. In: *Proceedings of the 41st annual IEEE symposium on the foundations of computer science (FOCS)*, pp 86–92
- Safer HM (1992) Fast approximation schemes for multi-criteria combinatorial optimization. PhD thesis, Massachusetts Institute of Technology
- Safer HM, Orlin JB (1995) Fast approximation schemes for multi-criteria combinatorial optimization. Working paper 3756-95, Sloan School of Management, Massachusetts Institute of Technology
- Vassilvitskii S, Yannakakis M (2005) Efficiently computing succinct trade-off curves. *Theor Comput Sci* 348(2–3):334–356

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.