Convergence of stochastic search algorithms to finite size pareto set approximations — **Source link** ⧉

Oliver Schütze, Marco Laumanns, Carlos A. Coello Coello, Michael Dellnitz ...+1 more authors

Institutions: CINVESTAV, ETH Zurich, University of Paderborn, French Institute for Research in Computer Science and Automation

Related papers:

- Combining convergence and diversity in evolutionary multiobjective optimization

- A fast and elitist multiobjective genetic algorithm: NSGA-II

- Computing gap free pareto front approximations with stochastic search algorithms

- Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach

- SMS-EMOA : Multiobjective selection based on dominated hypervolume

Share this paper: 🟦 🟦 in ✉

View more about this paper here: https://typeset.io/papers/convergence-of-stochastic-search-algorithms-to-finite-size-12cik9p401

Convergence of Stochastic Search Algorithms to Finite Size Pareto Set Approximations

Oliver Schuetze, Marco Laumanns, Carlos Coello Coello, Michael Dellnitz,

El-Ghazali Talbi

Convergence of Stochastic Search Algorithms to Finite Size Pareto Set Approximations

Oliver Schütze[1], Marco Laumanns[2], Carlos A. Coello Coello[3],
Michael Dellnitz[4], and El-ghazali Talbi[1]

[1] INRIA Futurs, LIFL, CNRS Bât M3, Cité Scientifique
e-mail: {schuetze,talbi}@lifl.fr

[2] Institute for Operations Research, ETH Zurich, 8092 Zurich, Switzerland
e-mail: laumanns@ifor.math.ethz.ch

[3] CINVESTAV-IPN, Electrical Engineering Department
e-mail: ccoello@cs.cinvestav.mx

[4] University of Paderborn
e-mail: dellnitz@upb.de

N° 6063

November 2006

*R apport
de recherche*

Convergence of Stochastic Search Algorithms to Finite Size Pareto Set Approximations

Oliver Schütze[1], Marco Laumanns[2], Carlos A. Coello Coello[3],
Michael Dellnitz[4], and El-ghazali Talbi[1]

[1] INRIA Futurs, LIFL, CNRS Bât M3, Cité Scientifique
e-mail: {schuetze,talbi}@lifl.fr

[2] Institute for Operations Research, ETH Zurich, 8092 Zurich, Switzerland
e-mail: laumanns@ifor.math.ethz.ch

[3] CINVESTAV-IPN, Electrical Engineering Department
e-mail: ccoello@cs.cinvestav.mx

[4] University of Paderborn
e-mail: dellnitz@upb.de

Abstract: In this work we study the convergence of generic stochastic search algorithms toward the Pareto set of continuous multi-objective optimization problems. The focus is on obtaining a finite approximation that should capture the entire solution set in a suitable sense, which will be defined using the concept of ϵ-dominance. Under mild assumptions about the process to generate new candidate solutions, the limit approximation set will be determined entirely by the archiving strategy. We investigate two different archiving strategies which lead to a different limit behavior of the algorithms, yielding bounds on the obtained approximation quality as well as on the cardinality of the resulting Pareto set approximation. Finally, we demonstrate the potential for a possible hybridization of a given stochastic search algorithm with a particular local search strategy – multi-objective continuation methods – by showing that the concept of ϵ-dominance can be integrated into this approach in a suitable way.

Key-words: multi-objective optimization, convergence, ϵ-dominance, stochastic search algorithms.) performances

1   Introduction

A common goal in multi-objective optimization is to identify the set of Pareto-optimal solutions (the efficient set) and its image in objective space, the Pareto front (the efficient frontier). Except for special cases, where the Pareto set is finite or representable by a finite collection of line segments (such as in multi-objective linear programming), it is in general not practicable to determine the entire Pareto set. Instead, a suitable approximation concept is needed.

Various approximation concepts based on ϵ-efficiency are given in [4]. As most of them deal with infinite sets, they are not practical for our purpose of producing and maintaining a representative subset of finite size. Using discrete ϵ-approximations of the Pareto set was suggested simultaneously by [1], [7], and [9]. The general idea is that each Pareto-optimal point is approximately dominated by some point of the approximation set.

Despite the existence of suitable approximation concepts, investigations on the *convergence* of particular algorithms towards such approximation sets, that is, their ability to obtain a suitable Pareto set approximation in the limit, have remained rare. Several studies, such as [2, 8], consider only the convergence to the entire Pareto set, or to a certain subset without considering the approximation quality.

In [6] the issue of convergence towards a finite-size Pareto set approximation was finally addressed for a general class of iterative search algorithms. Two archiving algorithms were proposed that provably maintain a finite-size approximation of all points ever generated during the search process. This led to the claim that these archiving strategies will ensure convergence to a Pareto set approximation of given quality for any iterative search algorithm that fulfills certain mild assumptions about the process to generate new search points. While this claim holds trivially in the case of discrete (or discretized) search spaces, its extension to the continuous case is not straightforward. Consideration of *discretized* models, however, can lead to problems when, e.g., using memetic strategies (metaheuristic search algorithms mixed with local search strategies which itself use step size control).

The goal of this paper is to establish convergence results with respect to finite Pareto set approximations for stochastic multi-objective optimization algorithms working in continuous domains. We start by considering the first archiving strategy from [6] and prove convergence with probability one to an ϵ-approximate Pareto set in the limit. Then we propose a new archiving strategy that additionally ensures that all elements of the limit set are Pareto-optimal points itself. For both strategies we give bounds on the approximation quality and on the cardinality of the limit solution set.

2   Background

In the following we consider continuous unconstrained multi-objective optimization problems

$$\min_{x \in \mathbb{R}^n} \{F(x)\}, \tag{MOP}$$

where the function F is defined as the vector of the objective functions

$$F : \mathbb{R}^n \to \mathbb{R}^k, \qquad F(x) = (f_1(x), \ldots, f_k(x)),$$

and where each $f_i : \mathbb{R}^n \to \mathbb{R}$ is continuous.

Definition 2.1   *(a) Let $v, w \in \mathbb{R}^k$. Then the vector v is less than w ($v <_p w$), if $v_i < w_i$ for all $i \in \{1, \ldots, k\}$. The relation $\leq_p$ is defined analogously.*

(b) *A vector $y \in \mathbb{R}^n$ is dominated by a vector $x \in \mathbb{R}^n$ (in short: $x \prec y$) with respect to (MOP) if $F(x) \leq_p F(y)$ and $F(x) \neq F(y)$ (i.e. there exists a $j \in \{1, \ldots, k\}$ such that $f_j(x) < f_j(y)$), else y is called non-dominated by x.*

(c) *A point $x \in \mathbb{R}^n$ is called Pareto optimal or a Pareto point if there is no $y \in \mathbb{R}^n$ which dominates x.*

(d) *A point $x \in \mathbb{R}^n$ is weakly Pareto optimal if there does not exist another point $y \in \mathbb{R}^n$ such that $F(y) <_p F(x)$.*

In the following we will define a weaker concept of dominance, so-called (absolute) ϵ-dominance, which will be used for our further studies.

Definition 2.2 *Let $\epsilon = (\epsilon_1, \ldots, \epsilon_k) \in \mathbb{R}_+^k$ and $x, y \in \mathbb{R}^n$. x is said to ϵ-dominate y (in short: $x \prec_\epsilon y$) with respect to (MOP) if*

(i) *$f_i(x) - \epsilon_i \leq f_i(y) \quad \forall i = 1, \ldots, k$, and*

(ii) *$f_j(x) - \epsilon_j < f_j(y) \quad$ for at least one $j \in \{1, \ldots, k\}$.*

We have to emphasize that the ϵ-dominance relation – unlike the 'classical' one defined above – is not transitive, i.e., if $x \prec_\epsilon y$ and $y \prec_\epsilon z$ it does *not* follow that $x \prec_\epsilon z$, but it follows that $x \prec_{2\epsilon} z$. This fact will be used in later considerations as well as the following: if $x \prec y$ and $y \prec_\epsilon z$ it follows that $x \prec_\epsilon z$.

Definition 2.3 *Let $\epsilon \in \mathbb{R}_+^k$.*

(a) *A set $F_\epsilon \subset \mathbb{R}^n$ is called an ϵ-approximate Pareto set of (MOP) if every point $x \in \mathbb{R}^n$ is ϵ-dominated by at least one $f \in F_\epsilon$, i.e.*

$$\forall x \in \mathbb{R}^n \ : \ \exists f \in F_\epsilon : \quad f \prec_\epsilon x$$

(b) *A set $F_\epsilon^* \subset \mathbb{R}^n$ is called an ϵ-Pareto set if F_ϵ^* is an ϵ-approximate Pareto set and if every point $f \in F_\epsilon^*$ is a Pareto point of (MOP).*

Algorithm 1 Generic Stochastic Search Algorithm

1: $P_0 \subset Q$ drawn at random
2: $A_0 = ArchiveUpdate(P_0, \emptyset)$
3: **for** $j = 0, 1, 2, \ldots$ **do**
4:    $P_{j+1} = Generate(P_j)$
5:    $A_{j+1} = ArchiveUpdate(P_{j+1}, A_j)$
6: **end for**

Further, let $B_\delta(x_0) := \{x \in \mathbb{R}^n \ : \ \|x - x_0\| < \delta\}$ be the open ball with center $x_0 \in \mathbb{R}^n$ and radius $\delta \in \mathbb{R}_+$.

Algorithm 1 gives a framework of a generic stochastic multi-objective optimization algorithm, which will be considered in this work. Theorem 2.4 states a convergence result which is closely related to the present work, but which leads in general to unbounded archive sizes.

Theorem 2.4 *[11] Let an MOP $F : \mathbb{R}^n \to \mathbb{R}^k$ be given, where F is continuous, let $Q \subset \mathbb{R}^n$ be compact. Further, let there be no weak Pareto point in $Q \backslash P_Q$ (where P_Q denotes the set of Pareto points of $F\big|_Q$), and*

$$\forall x \in Q \text{ and } \forall \delta > 0 : \qquad P\left(\exists l \in \mathbb{N} \ : \ P_l \cap B_\delta(x) \cap Q \neq \emptyset\right) = 1 \tag{1}$$

Then an application of Algorithm 1, where all non-dominated points are kept, i.e.,

$$ArchiveUpdate(P, A) := \{x \in P \cup A : y \nprec x \ \forall y \in P \cup A\},$$

generates a sequence of archives $\{A_i\}_{i \in \mathbb{N}}$, such that

$$\lim_{i \to \infty} d(F(P_Q), F(A_i)) = 0 \quad \text{with probability one,}$$

where $d(\cdot, \cdot)$ denotes the Hausdorff distance.

An archiving scheme to maintain an archive of finite size was recently proposed in [3]. New archive members are required to have a distance of at least a prescribed value of ϵ from all current archive members, unless they dominate (and hence replace) a current archive member. The subsequent proof of convergence (in probability) is based on the claim that Pareto points that lie within an ϵ-neighborhood of an archive member in iteration i also do so in the next iteration $i + 1$, since archive members are only deleted when substituted by a dominating alternative. The following simple example shows, however, that this claim does not hold for any L_p norm as the chosen distance metric. Consider the problem

$$\max_{x \in \mathbb{R}^2} F(x) = x$$

subject to the constraints

$$x_2 \leq 1 - \sqrt{1 - (x_1 - 1)^2},$$
$$x_1, x_2 \in [0, 1],$$

so that the objective functions are the projections to the i-th coordinate and the constraint defines the Pareto set with extreme points $(1, 0)$ and $(0, 1)$. Let $A_i = \{(0.5, 0), (0, 0.5)\}$ be the archive at iteration i, which is an ϵ-approximation[1] with $\epsilon = 0.5$: For all points in the Pareto set, in particular for the point $(0.25, 0.25)$, the distance to either archive member is not larger than 0.5. Now let the new points $(1, 0)$ and $(0, 1)$ be generated, which dominate, and hence replace, both archive members. The new archive $\{(1, 0), (0, 1)\}$ is only a 0.75-approximation, as the distance of point $(0.25, 0.25)$ to either archive member is 0.75 assuming the maximum norm and certainly greater than 0.5 in all reasonable distance metrics. At least for obtaining an ϵ-Pareto set in the objective space, this problem can be overcome by using the ϵ-dominance instead of a distance metric for defining the approximation quality as well as for the archive updating strategy, as proposed in [6].

Further, the next example shows that we can run into trouble when using an elitist archiving strategy as proposed in [3] in case F is not injective: for a given $\epsilon > 0$ let f_ϵ be as shown in Figure 1. That is, let f_ϵ have two isolated global minima m_1 and m_2 with $m_1 < m_2$ and with $d(m_1, m_2) > \epsilon$. Define $F := (f_\epsilon, f_\epsilon + C)$, where $C \in \mathbb{R}$ is a constant. If the domain is e.g. chosen as $A := [m_1 - \epsilon, m_2 + \epsilon]$, the Pareto set of the resulting MOP is given by $\mathcal{P} = \{m_1, m_2\}$. However, since the probability to find a point $p_2 \in A$ which has the same objective values $F(p_2) = F(p_1)$ of a given point $p_1 \in A$ is zero in the underlying setting, it follows that the set of nondominated points of a given population consists with probability one of one single point. Thus, an ϵ-approximation can in general not be obtained when only nondominated points are stored in the archive.

3 The Algorithms

In the following we investigate two different strategies for the archiving of the solutions found by the algorithm leading to different limit behaviors of the sequence of archives (under certain additional conditions).

First, we assume that the entries of $\epsilon \in \mathbb{R}_+^k$ are 'small', and thus that it is sufficient to obtain an ϵ-approximate Pareto set. For this, we consider the archiving strategy proposed in [6], here given as Algorithm 2. It computes the subsequent archive A of a given archive A_0, a population P, and an $\epsilon \in \mathbb{R}_+^k$. Using this strategy, the sequence of archives has a limit behavior described in Theorem 3.2. To show this, we need first the following obvious but crucial property of the archiving strategy.

---

[1]See Appendix for the definition. Note that this concept of ϵ-efficiency is not – as the ϵ-dominance – defined in the objective space but instead in parameter space.

Figure 1: Example of a function f_ϵ with two isolated global minima $m_1 = 0$ and $m_2 = 2\epsilon$.

Algorithm 2 $A := ArchiveUpdate1_\epsilon \ (P, A_0)$

1:   $A := A_0$
2: **for all** $p \in P$ **do**
3:     **if** $\exists a \in A : \ a \prec_{\epsilon/3} p$ **then**
4:        CONTINUE $\triangleright$ *do not execute lines* $6 - 11$
5:     **end if**
6:     **for all** $a \in A$ **do**
7:        **if** $p \prec a$ **then**
8:           $A := A \backslash \{a\}$
9:        **end if**
10:    **end for**
11:    $A := A \cup \{p\}$
12: **end for**

Lemma 3.1 *Let $A_0, P \subset \mathbb{R}^n$ be finite sets, $\epsilon \in \mathbb{R}_+^k$, and $A := ArchiveUpdate1_\epsilon (P, A_0)$. Then the following holds:*

$$\forall x \in P \cup A_0 : \quad \exists a \in A : a \prec_{\epsilon/3} x.$$

Proof: Roughly speaking, the statement follows since points a are only discarded from the archive if in turn another point p with $p \prec a$ is inserted (this 'replacement' is given in lines 7, 8 and 11 in Algorithm 2). To be more precise, let $P = \{p_1, p_2, \ldots, p_l\}, l \in \mathbb{N}$. Without loss of generality we assume that all points p_i are considered in this ordering (i.e., in the for-loop in line 2 of Algorithm 2). There are two cases we have to distinguish.
Case A: $x \in A_0$. Define $p_0' := x$ and

$$p_i' := \left\{ \begin{array}{ll} p_i & \text{if } p_i \text{ 'replaces' } p_{i-1}' \\ p_{i-1}' & \text{else} \end{array} \right. , \quad i = 1, \ldots, l.$$

It holds that $p_l' \in A$ and either $p_l' = x$ or $p_l' \prec x$ (due to the transitivity of $\prec$). In both cases it is $p_l' \prec_{\epsilon/3} x$.
Case B: $x \in P$. Let $x = p_j, j \in \{1, \ldots, l\}$. After the j-th iteration of the outer for-loop in Algorithm 2 there exists an element $a_j \in A$ with $a_j \prec_{\epsilon/3} p_j$ (line 3 resp. line 11 of Algorithm 2). Define $p_j' := a_j$ and $p_i', \ i = j+1, \ldots, l$, as above. It follows that $p_l' \in A$ and $p_l' \prec_{\epsilon/3} x$ as claimed.

Theorem 3.2 *Let an MOP $F : \mathbb{R}^n \to \mathbb{R}^k$ be given, where F is continuous, let $Q \subset \mathbb{R}^n$ be a compact set and $\epsilon \in \mathbb{R}_+^k$. Further let*

$$\forall x \in Q \text{ and } \forall \delta > 0 : \qquad P\left(\exists l \in \mathbb{N} \ : \ P_l \cap B_\delta(x) \cap Q \neq \emptyset\right) = 1 \tag{2}$$

Then an application of Algorithm 1, where $ArchiveUpdate1_\epsilon$ is used to update the archive, leads to a sequence of archives $A_l, l \in \mathbb{N}$, where the following holds:

(a) *There exists with probability one a $l_0 \in \mathbb{N}$ such that A_l is an ϵ-approximate Pareto set for all $l \geq l_0$.*

(b) *Assume there exists a $l_0 \in \mathbb{N}$ such that A_{l_0} is an ϵ-approximate Pareto set. Then*

$$A_l = A_{l_0}, \quad \forall l \geq l_0.$$

Proof:

(a) Since Q is compact and F is continuous it follows that $F\big|_Q$ is uniformly continuous. Hence for $\epsilon/3 \in \mathbb{R}_+^k$ there exists a $\delta > 0$ such that

$$x \prec_{\epsilon/3} y \quad \forall x, y \in Q \text{ with } \|x - y\| < \delta. \tag{3}$$

Define

$$G := \bigcup_{p \in P_Q} B_\delta(p)$$

G is an open cover of P_Q. Since P_Q is compact it follows – due to the theorem of Heine-Borel – that there exists a finite subcover

$$\mathcal{S} := \bigcup_{i=1}^{s} B_\delta(p_i) \supset P_Q, \quad p_i \in P_Q, i = 1, \ldots, s.$$

By (2) it follows that there exist with probability one s numbers $l_1, \ldots, l_s \in \mathbb{N}$ such that each $B_\delta(p_i) \cap Q, i = 1, \ldots, s$, gets 'visited' by *Generate* after l_i iteration steps. That is, $P_{l_i}, i = 1, \ldots, s$, contains with probability one a point $b_i \in B_\delta(p_i) \cap Q$, and thus, A_{l_i} contains with probability one a vector d_i with $d_i \prec_{\epsilon/3} b_i$. By construction of $ArchiveUpdate1_\epsilon$ there exists for all $l \geq l_i$ with probability one a $d_i^l \in A_l$ such that $d_i^l \prec_{\epsilon/3} b_i$ (see Lemma 1). Set $l_0 := \max\{l_1, \ldots, l_s\}$.

Now let $x \in Q$. For x there exists a $p \in P_Q$ such that $F(p) \leq_p F(x)$ and since $\mathcal{S}$ is a cover of P_Q there exists an $i \in \{1, \ldots, s\}$ with $p \in B_\delta(p_i)$. Let l_0, b_i, and d_i^l be as described above and let $l \geq l_0$. Since b_i and p are inside $B_\delta(p_i)$ it follows by (3) that $b_i \prec_{\epsilon/3} p_i$ and $p_i \prec_{\epsilon/3} p$. Hence we have with probability one:

$$d_i^l \prec_{\epsilon/3} b_i \prec_{\epsilon/3} p_i \prec_{\epsilon/3} p, \quad l \geq l_i.$$

Thus, we have that $d_i^l \prec_\epsilon x$, $l \geq l_0$, with probability one as desired.

(b) This follows immediately by the construction of $ArchiveUpdate1_\epsilon$ (to be more precise, by lines $3 - 5$ of Algorithm 2).

Remarks 3.3 *(a) Assumption (2) is the crucial part to obtain the convergence. For general ϵ and general F it is certainly not possible to postulate less. Given a fixed $\epsilon \in \mathbb{R}_+^k$ it would in principle be sufficient to require condition (2) only for the δ which is given in the proof above as well as for finitely many vectors $x \in Q$. However, this is nearly impossible to check in practise.*

(b) Here we have used the absolute ϵ-dominance. If $0 \notin f_i(P_Q), i = 1, \ldots, k$, alternatively the relative ϵ-dominance as in [6] can be used yielding similar results.

(c) We have restricted the domain to a compact subset of the $\mathbb{R}^n$. The following (academic) example shows that we can run into trouble if Q is not compact: consider the MOP

$$F : \mathbb{R}_+ \to \mathbb{R}^2$$
$$F(x) = (-x, -\frac{1}{x})$$

In this case, the Pareto set is given by $P = \mathbb{R}_+$. Since $F(P)$ is not bounded below it can not be represented by a finite archive using ϵ-dominance. However, this changes if $Q = [a, b]$, $a < b$, $a, b > 0$ is chosen as the domain.

Next, we assume that the entries of ϵ are relatively large. This can be the case when the decision maker prefers to obtain few, widespread solutions of the MOP, or in order to be able to 'capture' the entire Pareto set with a limited archive, in particular when considering more than two objectives. Hence, convergence of the entries of the sequence of archives toward the Pareto set is desired. For this, we propose to use the archiving strategy which is described in Algorithm 3. In the following we will discuss the limit behavior of this approach.

Algorithm 3 $A := ArchiveUpdate2_\epsilon (P, A_0)$

1: $A := A_0$
2: **for all** $p \in P$ **do**
3:      **if** $\nexists a \in A : a \prec_{\epsilon/3} p$ **then**
4:          $A := A \cup \{p\}$
5:      **end if**
6:      **for all** $a \in A$ **do**
7:          **if** $p \prec a$ **then**
8:              $A := A \cup \{p\}\backslash\{a\}$
9:          **end if**
10:      **end for**
11: **end for**

Lemma 3.4 *Let $A_0, P \subset \mathbb{R}^n$ be finite sets, $\epsilon \in \mathbb{R}_+^k$, and $A := ArchiveUpdate2_\epsilon (P, A_0)$. Then the following holds:*

$$\forall x \in P \cup A_0 : \quad \exists a \in A : a \prec_{\epsilon/3} x.$$

Proof: Analogue to the proof of Lemma 3.1.

Theorem 3.5 *Let (MOP) be given and $Q \subset \mathbb{R}^n$ be compact, and let there be no weak Pareto points in $Q\backslash P_Q$. Further, let F be injective and*

$$\forall x \in Q \text{ and } \forall \delta > 0 : \qquad P (\exists l \in \mathbb{N} : P_l \cap B_\delta(x) \cap Q \neq \emptyset) = 1 \qquad (4)$$

Then an application of Algorithm 1, where $ArchiveUpdate2_\epsilon$ is used to update the archive, leads to a sequence of archives $A_l, l \in \mathbb{N}$, where the following holds:

(a) *There exists with probability one a $l_0 \in \mathbb{N}$ such that A_l is an ϵ-approximate Pareto set for all $l \geq l_0$.*

(b) *There exists with probability one a $l_1 \in \mathbb{N}$ such that*

$$|A_{l+1}| = |A_l|, \quad \forall l \geq l_1.$$

(c) The limit archive

$$A_\infty := \lim_{l \to \infty} A_l$$

is an ϵ-Pareto set with probability one.

Proof:

(a) Analogue to the proof of Theorem 3.2 (a).

(b) By (a) it follows that there exists with probability one a $l_0 \in \mathbb{N}$ such that A_{l_0} is an ϵ-approximate Pareto set. Assume that this number l_0 is given. $|A_{l_0}|$ is certainly finite. Further let $l \geq l_0$. By construction of *ArchiveUpdate2$_\epsilon$* the archive A_l is also an ϵ-approximate Pareto set. That is, further points are only inserted to the archive if in turn at least one dominated solution is deleted (line 8 of Algorithm 3). Thus it holds that

$$|A_{l+1}| \leq |A_l| \quad \forall l \geq l_0.$$

Since on the other hand $|A_l| \geq 1 \ \forall l \in \mathbb{N}$, the sequence $\{|A_l|\}_{l \in \mathbb{N}}$ of the magnitudes of the archives is bounded below and monotonically decreasing and converges thus to an element $N_A \in \mathbb{N}$. Further, since $|A_l| \in \mathbb{N}, l \in \mathbb{N}$, there exists a $l_1 \in \mathbb{N}$ such that $|A_l| = N_A, \ \forall l \geq l_1$.

(c) By (b) it follows that there exists with probability one a $l_1 \in \mathbb{N}$ such that $|A_{l+1}| = |A_l|, \ \forall l \geq l_1$. Assume that this number l_1 is given. Consider an element $a_0 \in A_l$ with $l \geq l_1$. If $a_0 \in P_Q$ it follows that $a_0 \in A_{l+m}, \ \forall m \in \mathbb{N}$, and thus also $a_0 \in A_\infty$. Assume that $a_0 \notin P_Q$. Define

$$M : Q \to \mathbb{R}$$
$$M(x) := \max_{p \in P_Q} \min_{i=1,\dots,k} (f_i(x) - f_i(p)) \tag{5}$$

Under the assumptions made above it holds that

$$M(x) \geq 0 \ \forall x \in Q \quad \text{and} \quad M(x) = 0 \Leftrightarrow x \in P_Q.$$

Let $p_0 \in P_Q$ be the argument of the maximum of $M(a_0)$. Since $a_0 \notin P_Q$ and a_0 is no weak Pareto point it follows that $M(a_0) > 0$ and $F(p_0) <_p F(a_0)$. Since F is continuous there exists a neigborhood U_{p_0} of p_0 such that

$$F(y) <_p F(p_0) + \frac{M(a_0)}{2} \cdot (1, \dots, 1) \quad \forall y \in U_{p_0},$$

and thus, that $F(y) <_p F(a_0), \ \forall y \in U_{p_0}$. By (4) it follows that *Generate* generates with probability one after finitely many steps a point $b \in U_{p_0} \cap Q$. Now there are two cases: (1) b is added to the archive (in this case set $a_1 := b$), and (2), a_0 has already

been replaced by an element $\tilde{a} \in \mathbb{R}^n$ such that b and $\tilde{a}$ are mutually non-dominating (in this case set $a_1 := \tilde{a}$). In both cases there exists a $j \in \{1, \ldots, k\}$ such that

$$f_j(a_1) < f_j(p_0) + \frac{M(a_0)}{2}.$$

Proceeding in an analogous way we obtain a sequence $\{a_i\}_{i \in \mathbb{N}}$ of dominating points. Since the sequence $\{F(a_i)\}_{i \in \mathbb{N}}$ is below bounded and F is injective it follows that $a_i \to a^* \in Q$ for $i \to \infty$.

It remains to show that $a^* \in P_Q$. For this, assume that $a^* \notin P_Q$. Define p^* as the argument of the maximum of $M(a^*)$. Since $a^* \notin P_Q$ and a^* is no weak Pareto point it follows that $F(p^*) <_p F(a^*)$ and $M(a^*) > 0$. Proceeding further as above we obtain a point a^{**} and an element $j \in \{1, \ldots, k\}$ such that

$$f_j(a^{**}) < f_j(p^*) + \frac{M(a^*)}{2} \leq f_j(p^*) + \frac{f_j(a^*) - f_j(p^*)}{2}$$
$$= \frac{f_j(p^*) + f_j(a^*)}{2} < f_j(a^*)$$

This is a contradiction to the assumption of the convergence of the sequence, and thus it must be that $a^* \in P_Q \cap A_\infty$. Since $a_0 \in A_l, l \geq l_1$, was chosen arbitrarily it follows that A_∞ is a ϵ-Pareto set and the proof is complete.

4 Bounds on the Archive Sizes

In the following we give bounds on the magnitude of the limit archives A_∞ with respect to $\epsilon \in \mathbb{R}_+^k$ and the chosen archiving strategy.

For this, we have to introduce some notations: denote by m_i and M_i the minimal resp. maximal value of objective f_i, $i = 1, \ldots, k$, inside Q (these values exist since F is continuous and Q is compact). Further, we need k-dimensional boxes, which can be represented by a center $c \in \mathbb{R}^k$ and a radius $r \in \mathbb{R}_+^k$:

$$B = B(c, r) = \{x \in \mathbb{R}^k : |x_i - c_i| \leq r_i \ \forall i = 1, \ldots, k\}.$$

In the following we assume that $|P_0| = 1$, and thus also $|A_0| = 1$. The lower bound of $|A_\infty|$ for both archiving strategies is obviously given by 1. For this, consider e.g. $f_1 = f_2 = \ldots = f_k$ to be a convex function which takes its (unique) minimum inside Q. The upper bounds for the different archiving strategies are derived separately in the following.

Theorem 4.1 *Let $m_i = \min_{x \in Q} f_i(x)$ and $M_i = \max_{x \in Q} f_i(x), 1 \leq i \leq k$, and $|A_0| = 1$. Then, when using $ArchiveUpdate1_\epsilon$, the archive size maintained in Algorithm 1 for all $l \in \mathbb{N}$ is bounded as*

$$|A_l| \leq \left\lceil \frac{1}{\epsilon_m} \sum_{\substack{i_1, \ldots, i_{k-1} = 1 \\ i_1 > \ldots > k_{k-1}}}^{k} \prod_{j=1}^{k-1} (M_{i_j} - m_{i_j}) \right\rceil, \tag{6}$$

where $\epsilon_m := \min\limits_{i=1,\ldots,k} \frac{\epsilon_i}{3}$.

Proof: Consider a sequence $p_1, p_2, \ldots$ of points which are all accepted by $ArchiveUpdate1_\epsilon$ in this order (i.e., starting with $A_0 = \{p_1\}$). Consider the i-th step and let $A_i = \{a_1, \ldots, a_l\}$ with $l \le i$. Define $B_j := B(F(a_j) - \epsilon/6, \epsilon/6)$, $j = 1, \ldots, l$. Using inductive arguments we see that (a) all elements in A_i are mutually non-dominating, and that (b) the interiors of all the boxes B_j, $j = 1, \ldots, l$, are mutually non-intersecting. Since the points a_j are the upper right corners of the boxes B_j and since the interiors of these boxes are mutually non-intersecting the minimal distance between two points a_{j_1} and $a_{j_2}, j_1 \neq j_2$, is given by ϵ_m (see Figure 2). Thus we are able to bound the number of entries in the archives if we can bound the number of such boxes which can be placed in the image space.

Let us first consider a bi-objective model (i.e., $k = 2$), since in this case the proof is geometrically descriptive and already captures the basic idea. Since all points a_j are mutually non-dominating, the images of these points are all located on a (virtual) continuously differentiable curve

$$
\begin{aligned}
c : [m_1, M_1] &\to \mathbb{R}^2 \\
u &\mapsto (u, f(u))
\end{aligned}
\tag{7}
$$

where $f : [m_1, M_1] \to [m_2, M_2]$ is a strictly monotonically decreasing (but not necessarily surjective) function. The length of this curve can be bounded as follows:

$$
\begin{aligned}
L(c) &= \int_{m_1}^{M_1} \|c'(u)\| du = \int_{m_1}^{M_1} \sqrt{|1|^2 + |f'(u)|^2} \, du \\
&\le \int_{m_1}^{M_1} 1 \, du + \int_{m_1}^{M_1} |f'(u)| du = \int_{m_1}^{M_1} 1 \, du - \int_{m_1}^{M_1} f'(u) du \\
&\le (M_1 - m_1) + (M_2 - m_2)
\end{aligned}
\tag{8}
$$

Thus, for $k = 2$ we see that $|A_i| \le \left\lceil \frac{(M_1 - m_1) + (M_2 - m_2)}{\epsilon_m} \right\rceil$, $i \in \mathbb{N}$, as claimed above.

Now we turn our attention to the general case, i.e. let $k \ge 2$ be given. Define

$$
\begin{aligned}
K &:= [m_1, M_1] \times \ldots [m_{k-1}, M_{k-1}], \\
K_{(i)} &:= [m_1, M_1] \times \ldots \times [m_{i-1}, M_{i-1}] \times [m_{i+1}, M_{i+1}] \times \ldots \times [m_{k-1}, M_{k-1}], \text{ and} \\
u_{(i)} &:= (u_1, \ldots, u_{i-1}, u_{i+1}, \ldots, u_{k-1}), \; i = 1, \ldots, k-1.
\end{aligned}
\tag{9}
$$

In analogy to the bi-objective case, the images of the elements of the archives are located in the graph of a map Φ which is characterized as follows:

$$
\begin{aligned}
\Phi : K &\to \mathbb{R}^k \\
\Phi(u_i, \ldots, u_{k-1}) &= (u_1, \ldots, u_{k-1}, f(u_1, \ldots, u_{k-1})),
\end{aligned}
\tag{10}
$$

where $f : K \to [m_k, M_k]$ is a sufficiently smooth function satisfying the monotonicity conditions $\frac{\partial f}{\partial u_i} u < 0$, $\forall u \in K$ and $\forall i = 1, \ldots, k-1$. Then, the $(k-1)$-dimensional volume of Φ with parameter range K can be bounded as follows:

$$Vol(\Phi) = \int_K \sqrt{||\nabla f||^2 + 1}\, du = \int_K \sqrt{\left(\frac{\partial f}{\partial u_1}\right)^2 + \ldots + \left(\frac{\partial f}{\partial u_{k-1}}\right)^2 + 1}\, du$$

$$\leq \int_K \left|\frac{\partial f}{\partial u_1}\right| du + \ldots + \int_K \left|\frac{\partial f}{\partial u_{k-1}}\right| du + \int_K 1\, du$$

$$= \sum_{i=1}^{k-1} \left(\int_{K_{(i)}} \left(\int_{m_i}^{M_i} \left|\frac{\partial f}{\partial u_i}\right| du_i \right) du_{(i)} \right) + \int_K 1\, du \tag{11}$$

$$= \sum_{i=1}^{k-1} \left(\int_{K_{(i)}} \left(-\int_{m_i}^{M_i} \frac{\partial f}{\partial u_i} du_i \right) du_{(i)} \right) + \int_K 1\, du$$

$$\leq \sum_{\substack{i_1,\ldots,i_{k-1}=1 \\ i_1>\ldots>k_{k-1}}}^{k} \prod_{j=1}^{k-1} (M_{i_j} - m_{i_j})$$

This bound of the volume leads directly to the bound of the cardinality of the archives as stated above which concludes the proof.


Figure 2: The entries a_i of each archive lie on a (virtual) curve c. Since the boxes B_i (with upper right corners $F(a_i)$) are mutually non-intersecting, it follows that the minimal distance of two entries is given by ϵ_m.

Remarks 4.2 *(a) Since the considertations on the 'dominating map' (10) hold also for the Pareto front, the obtained bounds on the archive size are tight.*

(b) As described above, ϵ_m is the minimal distance between the images of two distinct points in A_∞. Further, for every point y in the Pareto front there exists a $a \in A_\infty$

with $d(F(a), y) \leq \Delta$, where $d(\cdot, \cdot)$ denotes the maximum norm and $\Delta := \max_i(M_i - m_i)$ (e.g., when ϵ is too large or the Pareto front is 'flat'). Thus, following [10], the set $F(A_\infty)$ can be viewed as an ϵ_m-uniform d_Δ-representation[2] of the Pareto front (see Appendix for the definition). The huge value of Δ – in fact, the largest possible value – may be unsatisfying for certain applications, and thus it could be interesting to search for archiving strategies which generate such representations with prescribed (smaller) values of Δ.

Theorem 4.3 *Let $m_i = \min_{x \in Q} f_i(x)$ and $M_i = \max_{x \in Q} f_i(x), 1 \leq i \leq k$, and $|A_0| = 1$. Then, when using ArchiveUpdate2$_\epsilon$, the archive size maintained in Algorithm 1 is bounded for all $l \in \mathbb{N}$ as*

$$|A_l| \leq \prod_{i=1}^{k} \left\lceil 3 \frac{M_i - m_i}{\epsilon_i} \right\rceil. \tag{12}$$

Proof: We can consider the process of including solutions into the archive over time as a process for constructing a directed graph G. Starting with an empty graph, we add a new node for each new solution p that is added to the archive A in line 4 or line 8 of the algorithm. If p is added in line 8 (meaning the condition in line 7 is true), we add arcs (p, a) from p to each solution a that is discarded in line 8 due to $p \prec a$. Let $V_t := \bigcup_{1 \leq j \leq t} A_j$ be the union of all archives up to iteration t and $V_t' \subseteq V_t$ the subset of those archive members that have been added in line 4. Thus, the node set of G_t after iteration t is V_t, and G_t itself is a forest whose roots are the current archive members A_t and whose leafs are the elemets of V_t'. Obviously, the number of roots must be smaller than the number of leafs, so $|A_t| \leq |V_t'|$.

To bound $|V_t'|$, the number of elements that ever entered the archive in line 4, we again consider the boxes $B_v := B(F(v) - \epsilon/6, \epsilon/6)$ for all $v \in V_t'$. Due to line 3, a solution p generated in iteration $t' \leq t$ cannot be accepted in line 4 if $F(p)$ lies inside the box B_v of any previously accepted element of $v \in V_t'$, otherwise $a \prec_{\epsilon/3} p$ for some current archive member $a \in A_t$ as there exists $a \in A_t$ with $F(a) \leq F(v)$ and $v \prec_{\epsilon/3} p$. If p was accepted in line 4, then $F(p)$ cannot lie inside the box B_v of any subsequently accepted element of $v \in V_t'$ neither, as this would entail $p \prec v$. Hence, the interiors of the boxes B_v must be mutually non-intersecting. The maximum number of non-intersecting boxes with side length $\epsilon/3$ and centers c with $m_i \leq c_i \leq M_i$ is $\prod_{i=1}^{k} \lceil 3(M_i - m_i)/\epsilon_i \rceil$, thus the claimed bound on the archive size follows.

5 Outlook: Hybridizing with Multi-Objective Continuation Methods

In order to increase the overall computational performance, it is often desired to combine the (global) stochastic search algorithm with a local search strategy. In this section, we

---

[2]When changing the definition slightly.

want to show that in the underlying context a hybridization with multi-objective continuation methods (e.g., [5], [12]) could be advantageous since the concept of ϵ-dominance can be directly integrated into them.

In the following we construct a step size strategy for our purpose and show numerical results on two (easy) MOPs, indicating the possible benefit of such a hybridization.

The basic idea of multi-objective continuation methods is, roughly speaking, to move *along* the set of (local) Pareto points. To be more precise, in the course of the algorithm one is faced with the following setting: given a point $x_0 \in P|_Q$, an $\epsilon \in \mathbb{R}_+^k$, and a *search direction* $v \in \mathbb{R}^n$ with $\|v\| = 1$, the task is to find a step size $h \in \mathbb{R}_+$ such that for the next guess $y_0 = x_0 + hv$ it holds

$$\|F(x_0) - F(y_0)\|_\infty = \Theta\epsilon_m, \tag{13}$$

where $\Theta \in (0,1)$ is a safety factor. In case F is Lipschitz continuous there exists an $L \geq 0$ such that

$$\|F(x) - F(y)\| \leq L\|x - y\|, \quad \forall x, y \in Q. \tag{14}$$

The Lipschitz constant around x_0 can be estimated by

$$L_{x_0} := \|DF(x_0)\|_\infty = \max_{i=1,\dots,k} \|\nabla f_i(x_0)\|_1.$$

Combining (13) and (14), using $\|x_0 - y_0\| = h$, and assuming that h is sufficiently small, we obtain the following estimation

$$h \approx \frac{\Theta m_\epsilon}{L_{x_0}} \tag{15}$$

Note that this estimation only holds for small values of ϵ_m since in the other case h will be too large, and thus L_{x_0} can not serve as a suitable Lipschitz estimation.

5.1 Example 1

In order to understand the possible impact of the discussion made above on the continuation methods, we first apply the step size control on an academic example (see also [11]):

$$F : \mathbb{R}^2 \to \mathbb{R}^2$$
$$F(x) = \begin{pmatrix} (x_1 - 1)^4 + (x_2 - 1)^4 \\ (x_1 + 1)^2 + (x_2 + 1)^2 \end{pmatrix} \tag{16}$$

The Pareto set of MOP (16) is given by

$$\mathcal{P} = \left\{ \lambda \begin{pmatrix} -1 \\ -1 \end{pmatrix} + (1-\lambda) \begin{pmatrix} 1 \\ 1 \end{pmatrix} : \lambda \in [0,1] \right\}.$$

Figure 3 shows two different discretizations of $\mathcal{P}$. In Figure 3 (a) the Pareto set is approximated by points $x_i, i = 1, \ldots, N$, which are placed equidistant in parameter space:

$$x_i = \begin{pmatrix} -1 \\ -1 \end{pmatrix} + \frac{2i}{N} \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix}.$$

Next, the Pareto set was discretized using the adaptive step size control which is proposed above:

$$x_0 = \begin{pmatrix} -1 \\ -1 \end{pmatrix}, \quad x_{i+1} = x_i + h_i \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix},$$

where h_i is taken from (15) and $v_i = (1/\sqrt{2}, 1/\sqrt{2})^T$ was chosen as the search direction. Figure 3 (b) shows the discretization points x_i for $\epsilon_m = 1$ and $\Theta = 0.99$ yielding a satisfying distribution of the solutions on the Pareto front according to the value of ϵ_m.

5.2 Example 2

Next we consider the following MOP:

$$f_1, f_2 : \mathbb{R}^n \to \mathbb{R}$$
$$f_i(x) = \sum_{\substack{j=1 \\ j \neq i}}^{n} (x_j - a_j^i)^2 + (x_i - a_i^i)^4, \tag{17}$$

where

$$a^1 = (1, 1, 1, 1, \ldots) \quad \in \mathbb{R}^n$$
$$a^2 = (-1, -1, -1, -1, \ldots) \quad \in \mathbb{R}^n$$

In Figures 4 and 5 some numerical results are presented, where we have used the continuation method proposed in [12]. To be more precise, we have applied the step size control on the distance between the current solution and the *predictor*, since this point mainly determines the distance of two solutions.

Figure 4 shows the result for $n = 3$ and $\epsilon_m = 2$. In total, 23 solutions were obtained. This fits quite well with the bound in Theorem 4.1, which is given by

$$|A_i| \leq \left\lceil \frac{(25 - 0) + (25 - 0)}{2} \right\rceil = 25.$$

Note that the points have not been stored according to one of the archiving strategies proposed above. In that case, many solutions – depending on the insertion ordering – would have been discarded.

6 Conclusion and Future Work

We have proposed generic stochastic search algorithms for obtaining ϵ-approximate Pareto sets as well as ϵ-Pareto sets of a continuous multi-objective optimization problem in the

(a) fixed step size

(b) adaptive step size

Figure 3: Discretizations of the Pareto set of MOP (16) with (a) fixed step size and (b) adaptive step size control.

limit. We have presented a convergence result for these algorithms, and have given bounds on the cardinality of the corresponding archives.

For future work, there are a lot of interesting topics which can be addressed to advance the present work. One could e.g. consider the speed of the convergence, in particular when the methods presented above are hybridized with local search strategies. Further, we intend to

Figure 4: Result of the continuation method with step size control on MOP (17) for $n = 3$ in parameter space (left) and image space (right).

Figure 5: Result of the continuation method with step size control on MOP (17) for $n = 20$ in image space: all solutions (left) and zoom (right).

apply this theoretical framework in search for the development of fast and reliable multi-objective optimization algorithms.

7 Appendix

In the following we state some definitions which are used in Theorem 2.4 and Remark 4.2 (b).

Definition 7.1 *Let $u \in \mathbb{R}^n$ and $A, B \subset \mathbb{R}^n$. The semi-distance $dist(\cdot, \cdot)$ and the* Hausdorff *distance $d(\cdot, \cdot)$ are defined as follows:*

(a) $dist(u, A) := \inf\limits_{v \in A} \|u - v\|$

(b) $dist(B, A) := \sup\limits_{u \in B} dist(u, A)$

(c) $d(A, B) := \max \{dist(A, B), dist(B, A)\}$

Definition 7.2 *Let $\epsilon > 0$ and let $D \subset Z$ be a discrete set. D is called a d_ϵ-representation of Z if for any $z \in Z$, there exists $y \in D$ such that $d(z, y) \leq \epsilon$.*

Definition 7.3 *Let $Z \subset \mathbb{R}^n$ be any set and let D be a d_ϵ-representation of Z. Then D is called a δ-uniform d_ϵ-representation if*

$$\min\limits_{x, y \in D, x \neq y} d(x, y) \geq \delta.$$

References

[1] Y. G. Evtushenko and M. A. Potapov. Methods of numerical solution of multicriterion problem. *Soviet mathematics – doklady*, 34:420–423, 1987.

[2] T. Hanne. On the convergence of multiobjective evolutionary algorithms. *European Journal Of Operational Research*, 117(3):553–564, 1999.

[3] T. Hanne. A multiobjective evolutionary algorithm for approximating the efficient set. *European Journal Of Operational Research*, 176, 2007.

[4] S. Helbig and D. Pateva. On several concepts for ϵ-efficiency. *OR Spektrum*, 16(3):179–186, 1994.

[5] C. Hillermeier. *Nonlinear Multiobjective Optimization - A Generalized Homotopy Approach*. Birkhäuser, 2001.

[6] M. Laumanns, L. Thiele, K. Deb, and E. Zitzler. Combining convergence and diversity in evolutionary multiobjective optimization. *Evolutionary Computation*, 10(3):263–282, 2002.

[7] H. Reuter. An approximation method for the efficiency set of multiobjective programming problems. *Optimization*, 21:905–911, 1990.

[8] G. Rudolph and A. Agapie. On a multi-objective evolutionary algorithm and its convergence to the Pareto set. In *Congress on Evolutionary Computation (CEC2000)*, pages 1010–1016, 2000.

[9] G. Ruhe and B. Fruhwirt. ϵ-optimality for bicriteria programs and its application to minimum cost flows. *Computing*, 44:21–34, 1990.

[10] S. Sayin. Measuring the quality of discrete representations of efficient sets in multiple objective mathematical programming. *Mathematical Programming*, 87:543–560, 2000.

[11] O. Schütze. *Set Oriented Methods for Global Optimization*. PhD thesis, University of Paderborn, 2004. <http://ubdata.uni-paderborn.de/ediss/17/2004/schuetze/>.

[12] O. Schütze, A. Dell'Aere, and M. Dellnitz. On continuation methods for the numerical treatment of multi-objective optimization problems. In Jürgen Branke, Kalyanmoy Deb, Kaisa Miettinen, and Ralph E. Steuer, editors, *Practical Approaches to Multi-Objective Optimization*, number 04461 in Dagstuhl Seminar Proceedings. Internationales Begegnungs- und Forschungszentrum (IBFI), Schloss Dagstuhl, Germany, 2005. <http://drops.dagstuhl.de/opus/volltexte/2005/349>.