

Seeding the initial population of multi-objective evolutionary algorithms: A computational study



Tobias Friedrich^a, Markus Wagner^{b,*}

^a Hasso Plattner Institute, Potsdam, Germany

^b University of Adelaide, Australia

ARTICLE INFO

Article history:

Received 25 November 2014

Received in revised form 21 April 2015

Accepted 22 April 2015

Available online 30 April 2015

Keywords:

Multi-objective optimization

Approximation

Comparative study

Limited evaluations

ABSTRACT

Most experimental studies initialize the population of evolutionary algorithms with random genotypes. In practice, however, optimizers are typically seeded with good candidate solutions either previously known or created according to some problem-specific method. This *seeding* has been studied extensively for single-objective problems. For multi-objective problems, however, very little literature is available on the approaches to seeding and their individual benefits and disadvantages. In this article, we are trying to narrow this gap via a comprehensive computational study on common real-valued test functions. We investigate the effect of two seeding techniques for five algorithms on 48 optimization problems with 2, 3, 4, 6, and 8 objectives. We observe that some functions (e.g., DTLZ4 and the LZ family) benefit significantly from seeding, while others (e.g., WFG) profit less. The advantage of seeding also depends on the examined algorithm.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

In many real-world applications trade-offs between conflicting objectives play a crucial role. As an example, consider engineering a bridge, where one objective might be costs to build and another durability of the bridge. For such problems, we need specialized optimizers that determine the Pareto front of mutually non-dominated solutions. There are several established multi-objective evolutionary algorithms (MOEA) and many comparisons on various test functions. However, most of them start with random initial solutions.

If prior knowledge exists or can be generated at a low computational cost, good initial estimates may generate better solutions with faster convergence. These good initial estimates are often referred to as *seeds*, and the method of using good initial estimates is referred to as *seeding*. These botanical terms are used to express the possibility that good solutions for the environment can develop from these starting points. In practice, a good initial seeding can make problem solving approaches competitive that would otherwise be inferior.

For single-objective evolutionary algorithms, methods such as seeding have been studied for about two decades; see, e.g., [17,20,23,26,30,41] for studies and examples (see [27] for a recent

categorization). For example, the effects of seeding for the Traveling Salesman Problem (TSP) and the job-shop scheduling problem (JSSP) were investigated in [32]. The algorithms were seeded with known good solutions in the initial population, and it was found that the results were significantly improved on the TSP but not on the JSSP. To investigate the influence of seeding on the optimization, a varying percentage of seeding was used, ranging from 25 to 75%. Interestingly, it was also pointed out that a 100% seed is not necessarily very successful on either problems [28]. This is one of the very few reports that shows seeding can in some cases be beneficial to an optimization process, but not necessarily always is. In [21] a seeding technique for dynamic environments was investigated. There, the population was seeded when a change in the objective landscape arrived, aiming at a faster convergence to the new global optimum. Again, some of the investigated seeding approaches were more successful than others.

One of the very few studies that can be found on seeding techniques for MOEAs is the one performed by Hernandez-Diaz et al. [22]. There, seeds were created using gradient-based information. These were then fed into the algorithm called Non-Dominated Sorting Genetic Algorithm II (NSGA-II, [10]) and the quality was assessed on the benchmark family ZDT ([44], named after the authors Zitzler, Deb, and Thiele). The results indicate that the proposed approach can produce a significant reduction in the computational cost of the approach.

In general, seeding is not well documented for multi-objective problems, even for real-world problems. If seeding is done, then

* Corresponding author. Tel.: +61 8 8313 5405; fax: +61 8 8313 4366.

E-mail address: markus.wagner@adelaide.edu.au (M. Wagner).

typically the approach is outlined and used with the comment that it worked in “preliminary experiments”—the reader is left in the dark on the design process behind the used seeding approach. This is quite striking as one expects that humans can construct a few solutions by hand, even if they do not represent the ranges of the objectives well. The least that one should be able to do is to reuse existing designs, and to modify these iteratively towards extremes. Nevertheless, even this manual seeding is rarely reported.

In this paper, we are going to investigate the effects of two structurally different seeding techniques for five algorithms on 48 multi-objective optimization (MOO) problems.

1.1. Seeding

As seeding we use the weighted-sum method, where the trade-off preferences are specified by non-negative weights for each objective. Solutions to these weighted-sums of objectives can be found with an arbitrary classical single-objective evolutionary algorithm. In our experiments we use the algorithm Covariance Matrix Adaptation Evolution Strategy (CMA-ES, [18]). Details of the two studied weighting schemes are presented in Section 2.1.

1.2. Quality measure

There are different ways to measure the quality of the solutions. A recently very popular measure is the hypervolume indicator, which measures the volume of the objective space dominated by the set of solutions relative to a reference point [43]. Its disadvantage is its high computational complexity [4,3] and the arbitrary choice of the reference point. We instead consider the mathematically well founded *approximation constant*. In fact, it is known that the worst-case approximation obtained by optimal hypervolume distributions is asymptotically equivalent to the best worst-case additive approximation constant achievable by all sets of the same size [6]. For a rigorous definition, see Section 2. This notion of multi-objective approximation was introduced by several authors [19,15,31,35,36] in the 80s and its theoretical properties have been extensively studied [9,12,33,34,37].

1.3. Algorithms

We use the jMetal framework [13] and its implementation of NSGA-II [10], Strength Pareto Evolutionary Algorithm (SPEA2, [45]), S-Metric Selection Evolutionary Multi-Objective Algorithm (SMS-EMOA, [14]), and Indicator Based Evolutionary Algorithm (IBEA, [42]). Additionally to these more classical MOEAs, we also study Approximation Guided Evolution (AGE, [7]), which aims at directly minimizing the approximation constant and has shown to perform very well for larger dimensions [38–40]. For each of these algorithms we compare their regular behavior after a certain number of iterations with their performance when initialized with a certain seeding.

1.4. Benchmark families

We compare the aforementioned algorithms on four common families of benchmark functions. These are DTLZ ([11], named after the authors Deb, Thiele, Laumanns and Zitzler), LZ09 ([29], named after the authors Li and Zhang), WFG ([24], named after the authors’ research group Walking Fish Group) and ZDT [44]. While the last three families only contain two- and three-dimensional problems, DTLZ can be scaled to an arbitrary number of dimensions.

2. Preliminaries

We consider minimization problems with d objective functions, where $d \geq 2$ holds. Each objective function $f_i : S \mapsto \mathbb{R}$, $1 \leq i \leq d$, maps from the considered search space S into the real values. In order to simplify the presentation we only work with the dominance relation on the objective space and mention that this relation transfers to the corresponding elements of S .

For two points $x = (x_1, \dots, x_d)$ and $y = (y_1, \dots, y_d)$, with $x, y \in \mathbb{R}^d$ we define the following dominance relation:

$$x \preceq y \quad :\Leftrightarrow \quad x_i \leq y_i \text{ for all } 1 \leq i \leq d,$$

$$x \prec y \quad :\Leftrightarrow \quad x \preceq y \text{ and } x \neq y.$$

We assess the seeding schemes and algorithms by their achieved *additive approximation* of the (known) Pareto front. We use the following definition.

Definition 1. For finite sets $S, T \subset \mathbb{R}^d$, the additive approximation of T with respect to S is defined as

$$\alpha(S, T) := \max_{s \in S} \min_{t \in T} \max_{1 \leq i \leq d} (s_i - t_i).$$

We measure the approximation constant with respect to the known Pareto front of the test functions. The better an algorithm approximates a Pareto front, the smaller the additive approximation value is. Perfect approximation is achieved if the additive approximation constant becomes 0. However, the approximation constant achievable for a (finite) population with respect to a continuous Pareto front (consisting of an infinite number of points) is always strictly larger than 0. It depends on the fitness function what is the smallest possible approximation constant achievable with a population of bounded size.

2.1. Seeding

For the task of computing the seeds, we employ an evolutionary strategy (ES), because it “self-adapts” the extent to which it perturbs decision variables when generating new solutions based on previous ones. CMA-ES [18] self-adapts the covariance matrix of a multivariate normal distribution. This normal distribution is then used to sample from the multidimensional search space where each variate is a search variable. The co-variance matrix allows the algorithm to respect the correlations between the variables making it a powerful evolutionary search algorithm.

To compute a seed, a (2,4)-CMA-ES minimizes $\sum_{i=1}^d a_i f_i(x)$, where the $f_i(x)$ are the objective values of the solution x . In preliminary testing, we noticed that larger population values for CMA-ES tended to result in seeds with better objective values. This came at the cost of significantly increased evaluation budgets, as the learning of the correlations takes longer. Our choice does not necessarily represent the optimal choice across all 48 benchmark functions, however, it is our take on striking a balance between (1) investing evaluations in the seeding and (2) investing evaluations in the regular multi-objective optimization. Note that large computational budgets for the seeding have the potential to put the unseeded approaches at a disadvantage, if the final performance assessment is not done carefully.

The number of seeds, the coefficients used, and the budget of evaluations is determined by the seeding approaches, which we will describe in the following.

CCORNERSANDCENTRE: A total of 10,000 evaluations is equally distributed over the generation of $d+1$ seeds. The rest of the

population is generated randomly. For the i -th seed, $1 \leq i \leq d$, the coefficients a_j ($1 \leq j \leq d$) are set in the following way:

$$a_i = \begin{cases} 10 & \text{if } i = j, \\ 1 & \text{otherwise.} \end{cases}$$

Thus, we prevent the seeding mechanism from treating the optimization problem in a purely single-objective way by entirely neglecting any trade-off relationships between the objectives.¹ Lastly, the $(d + 1)$ -th weight vector uses equal weights of 1 per objective. This way, we aim at getting a seed that is relatively central with respect to the others.

LINEARCOMBINATIONS: Here a total of 100 seeds is generated, where each seed is the result of running CMA-ES for 1000 evaluations. The coefficients of the linear combinations are integer values and we construct them in the following way. First, we consider all “permutations” of coefficients with $a_i = 1$ for one coefficient and $a_{j \neq i} = 0$ for all others. Then, we consider all permutations where two coefficients have the value 1, then those where three coefficients have the value 1, and so on. When all such permutations that are based on $\{0, 1\}$ are considered, we consider all permutations based on $\{0, 1, 2\}$, then based on $\{0, 1, 3\}$, then based on $\{0, 2, 3\}$, and so on.

Consequently, we achieve a better distribution of points in the objective space. This comes, however, at the increased initial computational cost. Furthermore, the budget per seed is lower than in the CORNERSANDCENTRE approach, which typically results in less optimized seeds.

NOSEED: All solutions of the initial population are generated randomly. This is the approach that is typically used for the generation of the initial population.

In the classification of population initialization techniques for evolutionary algorithms by Kazimipour et al. [27], our approaches CORNERSANDCENTRE and LINEARCOMBINATIONS fall into the following categories:

- Randomness: stochastic, as they rely on the stochastic hill-climber CMA-ES,
- Compositionality: composite multi-step, as the initial population is generated by several individual methods,
- Generality: generic, as the coefficients can easily be adjusted for other problems.

2.2. Evolutionary multi-objective optimization algorithms

In the following, we outline the five optimization algorithms for which we will later-on investigate the benefits of seeding the initial populations.

Many approaches try to produce good approximations of the true Pareto front by incorporating different preferences. For example, the environmental selection in NSGA-II [10] first ranks the individuals using non-dominated sorting. Then, in order to distinguish individuals with the same rank, the crowding distance metric is used, which prefers individuals from less crowded sections of the objective space. The metric value for each solution is computed by adding the edge lengths of the cuboids in which the solutions reside, bounded by the nearest neighbors.

SPEA2 [45] works similarly. The raw fitness of the individuals according to Pareto dominance relations between them is calculated, and then a density measure to break the ties is used. The individuals that reside close together in the objective space are less likely to enter the archive of best solutions.

In contrast to these two algorithms, IBEA [42] is a general framework, which uses no explicit diversity preserving mechanism. The fitness of individuals is determined solely based on the value of a predefined indicator. Typically, implementations of IBEA come with the epsilon indicator or the hypervolume indicator, where the latter measures the volume of the dominated portion of the objective space.

SMS-EMOA [14] is a frequently used IBEA, which uses the hypervolume indicator directly in the search process. It is a steady-state algorithm that uses non-dominated sorting as a ranking criterion, and the hypervolume as the selection criterion to discard the individual that contributes the least hypervolume to the worst-ranked front. While SMS-EMOA often outperforms its competition, its runtime unfortunately increases exponentially with the number of objectives. Nevertheless, with the use of fast approximation algorithms (e.g., [2,5,25]), this algorithm can be applied to solve problems with many objectives as well.

Recently, approximation-guided evolution (AGE) [7] has been introduced, which allows to incorporate a formal notion (such as Definition 1) of approximation into a multi-objective algorithm. This approach is motivated by studies in theoretical computer science studying multiplicative and additive approximations for given multi-objective optimization problems [8,9,12,37]. As the algorithm cannot have complete knowledge about the true Pareto front, it uses the best knowledge obtained so far during the optimization process. It stores an archive A consisting of the non-dominated objectives vectors found so far. Its aim is to minimize the additive approximation $\alpha(A, P)$ of the population P with respect to the archive A . The experimental results presented in [7] show that given a fixed time budget it outperforms current state-of-the-art algorithms in terms of the desired additive approximation, as well as the covered hypervolume on standard benchmark functions.

3. Experimental setup

We use the jMetal framework [13], and our code for the seeding as well as all used seeds are available online.² As test problems we used the benchmark families DTLZ [11], ZDT [44], LZ09 [29], and WFG [24]. We used the functions DTLZ 1–4, each with 30 function variables and with $d \in \{2, 4, 6, 8\}$ objective values/dimensions.

In order to investigate the benefits of seeding even in the long run, we limit the calculations of the algorithms to a maximum of 10^6 fitness evaluations and a maximum computation time of four hours per run. Note that the time restriction had to be used as the runtime of some algorithms increases exponentially with respect to the size of the objective space.

AGE uses random parent selection; in all other algorithms parents are selected via a binary tournament. As variation operators, the polynomial mutation and the simulated binary crossover [1] were applied, which are both used widely in MOEAs [10,16,45]. The distribution parameters associated with the operators were $\eta_m = 20.0$ and $\eta_c = 20.0$. The crossover operator is biased towards the creation of offspring that are close to the parents, and was applied with $p_c = 0.9$. The mutation operator has a specialized explorative effect for MOO problems, and was applied with $p_m = 1/(\text{number of decision variables})$. Population size was set to $\mu = 100$ and $\lambda = 100$. Each setup was repeated 100 times. Note that these parameter settings are the default settings in the jMetal framework, and they can often be found in the literature, which makes a cross-comparison easier. To the best of our knowledge, this parameter setting does not favor any particular algorithm or put one at a disadvantage, even

¹ If the ranges of the objective values differ significantly, then the coefficients should be adjusted accordingly.

² <http://cs.adelaide.edu.au/markus/publications.html>.

though individual algorithms can have differing optimal settings for individual problems.

In a real-world scenario, if an algorithm is run several times (e.g. because of restarts), the seeding might be only calculated once. In this case, it might make sense to compare the unseeded and seeded variant of an algorithm with the same number of fitness evaluations. However, we observed the expected outcome that in

this case seeding is almost always beneficial. We therefore consider a more difficult scenario where the optimization is only run once and the number of fitness function evaluations used for the seeding is deduced from the number of fitness evaluations available for the MOEA.

As pointed out earlier, we assess the seeding schemes and algorithms using the additive approximation of the Pareto front.

Table 1

Summary of our results for the improvement through LINEARCOMBINATIONS seeding. We compare the default strategy NoSEED with 10^6 fitness function evaluations and LINEARCOMBINATIONS seeding, which uses 10^5 fitness function evaluations, plus $9 \cdot 10^5$ fitness function evaluations. The table shows the ratio of the median approximation constant of NoSEED (50 runs) divided by the median approximation constant of LINEARCOMBINATIONS (100 runs). Values >1.00 indicate where LINEARCOMBINATIONS achieves a better additive approximation, as the default strategy's outcome is in the dividend. To facilitate qualitative observations, we show only two decimal place. " $>$ " marks statistically significant improvements, " $<$ " marks statistically significant worsenings, " $=$ " marks statistically insignificant findings. In case a MOEA needed more than 4 h time, the approximation constant after 4 h is used. Dashes indicate scenarios where not even the first iteration of the algorithm was completed within the allotted 4 h.

Function	AGE (↔)	IBEA (↔)	NSGA-II (↔)	SMS-EMOA (↔)	SPEA2 (↔)
DTLZ1 2D	0.99=	1.04=	1.00=	1.00=	1.02=
DTLZ1 4D	0.95<	0.98=	0.97=	1.39>	1.11>
DTLZ1 6D	0.41<	1.01=	1.00=	–	1.02>
DTLZ1 8D	1.04=	0.97=	1.00=	–	0.99=
DTLZ2 2D	0.96<	1.01=	1.00=	0.99=	1.00=
DTLZ2 4D	1.08>	1.00=	1.00=	1.01=	1.01=
DTLZ2 6D	0.95=	1.00=	1.00=	–	1.00=
DTLZ2 8D	0.85<	1.00=	0.97=	–	0.99=
DTLZ3 2D	0.85<	1.00=	1.00=	0.99=	1.02=
DTLZ3 4D	0.97<	0.50=	1.00=	0.97=	1.03=
DTLZ3 6D	0.09<	1.00=	0.99=	–	1.00=
DTLZ3 8D	0.86<	1.00=	1.00=	–	1.02=
DTLZ4 2D	0.98<	1.15>	1.00>	1.01>	1.06>
DTLZ4 4D	1.10>	3.79>	1.00>	5.22>	2.89>
DTLZ4 6D	1.07>	2.25>	1.00=	1.95>	0.98<
DTLZ4 8D	1.05>	1.79>	1.02=	–	1.00=
LZ09 F1	3.65>	1.57>	1.05>	1.03>	0.98=
LZ09 F2	1.47>	0.97<	0.74<	1.47>	1.49>
LZ09 F3	1.03>	1.05>	0.99=	1.02>	1.00=
LZ09 F4	6.35>	4.15>	0.94<	6.03>	6.49>
LZ09 F5	1.06>	1.08>	0.98=	1.10>	1.02=
LZ09 F6	1.19>	1.00>	1.00=	1.09>	1.01=
LZ09 F7	1.61>	1.93>	0.86<	1.77>	1.53>
LZ09 F8	1.00=	1.44>	1.06=	1.16>	0.98=
LZ09 F9	2.17>	1.31>	0.81<	2.22>	2.26>
ZDT1	0.98<	1.03=	1.00=	1.00=	0.99=
ZDT2	1.01=	1.01=	1.00=	1.01>	1.03=
ZDT3	1.02=	1.01=	1.00=	0.99<	1.03=
ZDT4	1.02=	0.98=	1.00=	1.01=	1.02=
ZDT6	1.02=	1.05>	1.00>	0.99=	1.02=
WFG1 2D	2.97=	1.00=	1.00=	1.05=	0.99=
WFG1 3D	1.01=	0.99=	1.01>	1.01=	1.00=
WFG2 2D	0.01<	0.05<	0.50<	0.01<	0.02<
WFG2 3D	1.00=	0.95=	1.00=	1.00=	1.02=
WFG3 2D	1.00=	1.00<	1.00=	1.00=	1.00=
WFG3 3D	0.99=	0.91<	1.00=	1.01>	0.95=
WFG4 2D	1.00=	1.00=	1.00=	1.00=	1.01=
WFG4 3D	0.99=	0.99=	1.00=	1.00=	1.03>
WFG5 2D	1.00=	0.99>	1.00>	1.00=	1.00=
WFG5 3D	0.98=	0.99=	1.00>	1.00=	0.99=
WFG6 2D	0.06<	0.15<	0.73<	0.06<	0.09<
WFG6 3D	0.38<	0.57<	0.86<	0.41<	0.62<
WFG7 2D	1.01=	0.90<	1.00=	0.99<	1.00=
WFG7 3D	1.00=	1.00=	1.00=	1.00=	1.00=
WFG8 2D	0.54<	0.57<	0.96<	1.00<	0.53<
WFG8 3D	0.99<	0.83<	1.21>	0.99<	0.96<
WFG9 2D	1.02=	0.86<	1.00=	1.01=	1.00=
WFG9 3D	1.01>	1.01=	1.00=	1.02>	1.00=

Table 2

Summary of our results for the improvement through CORNERSANDCENTRE seeding. We compare the default strategy NoSEED with 10^6 fitness function evaluations and CORNERSANDCENTRE seeding, which uses 10^4 fitness function evaluations, plus $9 \cdot 10^5$ fitness function evaluations. The table shows the ratio of the median approximation constant of NoSEED (50 runs) divided by the median approximation constant of CORNERSANDCENTRE (100 runs). Values >1.00 indicate where CORNERSANDCENTRE achieves a better additive approximation, as the default strategy's outcome is in the dividend. To facilitate qualitative observations, we show only two decimal place. " $>$ " marks statistically significant improvements, " $<$ " marks statistically significant worsenings, " $=$ " marks statistically insignificant findings. In case a MOEA needed more than 4 h time, the approximation constant after 4 h is used. Dashes indicate scenarios where not even the first iteration of the algorithm was completed within the allotted 4 h.

Function	AGE (↔)	IBEA (↔)	NSGA-II (↔)	SMS-EMOA (↔)	SPEA2 (↔)
DTLZ1 2D	1.00=	0.99=	1.01=	1.00=	1.01=
DTLZ1 4D	0.99<	1.00=	1.00=	0.78<	0.93=
DTLZ1 6D	1.57>	0.99=	1.00=	–	0.98=
DTLZ1 8D	1.02>	0.98=	1.01=	–	0.99=
DTLZ2 2D	1.00=	1.00=	1.00=	1.00=	1.02=
DTLZ2 4D	1.06>	1.01=	1.00=	1.00=	1.01=
DTLZ2 6D	1.05>	0.99<	1.00=	–	0.99=
DTLZ2 8D	0.99=	1.00=	1.04=	–	1.01>
DTLZ3 2D	1.00=	1.00=	1.00=	1.00=	1.01=
DTLZ3 4D	1.00=	1.00=	1.00=	0.38<	0.91=
DTLZ3 6D	1.29>	1.00=	1.00=	–	0.98<
DTLZ3 8D	1.10>	1.00=	1.00=	–	1.03>
DTLZ4 2D	0.98=	1.10>	1.00>	1.05>	1.03>
DTLZ4 4D	4.55>	0.99=	1.00=	5.14>	1.01=
DTLZ4 6D	0.62<	0.95<	1.00=	1.12>	1.05=
DTLZ4 8D	0.97=	0.97<	1.04>	–	1.00=
LZ09 F1	3.01>	1.26>	1.06>	1.02=	0.97=
LZ09 F2	0.90<	1.03>	0.90<	1.08>	1.06>
LZ09 F3	1.07>	0.98=	0.92<	1.02=	0.97=
LZ09 F4	2.85>	2.43>	1.05>	2.95>	3.50>
LZ09 F5	1.03=	0.90<	1.02=	0.86<	0.91<
LZ09 F6	0.90<	1.00>	1.00=	0.89<	0.28<
LZ09 F7	1.03>	0.87>	0.96=	1.26>	1.19>
LZ09 F8	1.06>	1.71>	0.75<	1.35>	1.10>
LZ09 F9	1.76>	1.53>	0.82<	1.69>	1.68>
ZDT1	1.00=	0.98=	1.00=	1.00=	0.98<
ZDT2	1.01=	1.02=	1.00=	1.01=	1.00=
ZDT3	1.01=	0.99=	1.00=	1.00=	0.94=
ZDT4	0.99=	0.98=	1.00=	1.01=	0.97<
ZDT6	1.00=	1.03=	1.00>	1.01=	0.99=
WFG1 2D	1.00=	1.04=	0.87=	2.15>	1.02>
WFG1 3D	1.00=	0.99=	1.00=	1.01=	1.02>
WFG2 2D	0.98=	1.00=	1.00=	1.03=	0.99=
WFG2 3D	1.00=	0.94=	1.00=	1.00=	0.98=
WFG3 2D	1.00=	1.00=	1.00=	1.00=	1.00=
WFG3 3D	1.01=	0.99=	1.00=	1.00=	0.99=
WFG4 2D	1.00=	1.00=	1.00=	1.01=	0.96<
WFG4 3D	1.00=	0.97<	1.00=	1.00=	0.99=
WFG5 2D	1.00=	1.02>	1.00>	1.00=	0.99=
WFG5 3D	0.99=	1.01=	1.00>	1.00=	0.98=
WFG6 2D	0.91<	0.99=	1.00=	0.86<	0.97=
WFG6 3D	1.01=	1.00=	1.00=	1.01=	1.00=
WFG7 2D	1.01=	1.00=	1.00=	0.99<	1.00=
WFG7 3D	1.00=	1.02=	1.00=	1.00=	1.03=
WFG8 2D	1.71=	0.69<	0.96>	1.00=	1.79=
WFG8 3D	0.99=	1.00=	1.21>	0.99=	1.26>
WFG9 2D	0.99=	1.05>	1.00=	0.99=	1.00=
WFG9 3D	1.00=	1.02=	1.00=	1.00=	1.01=

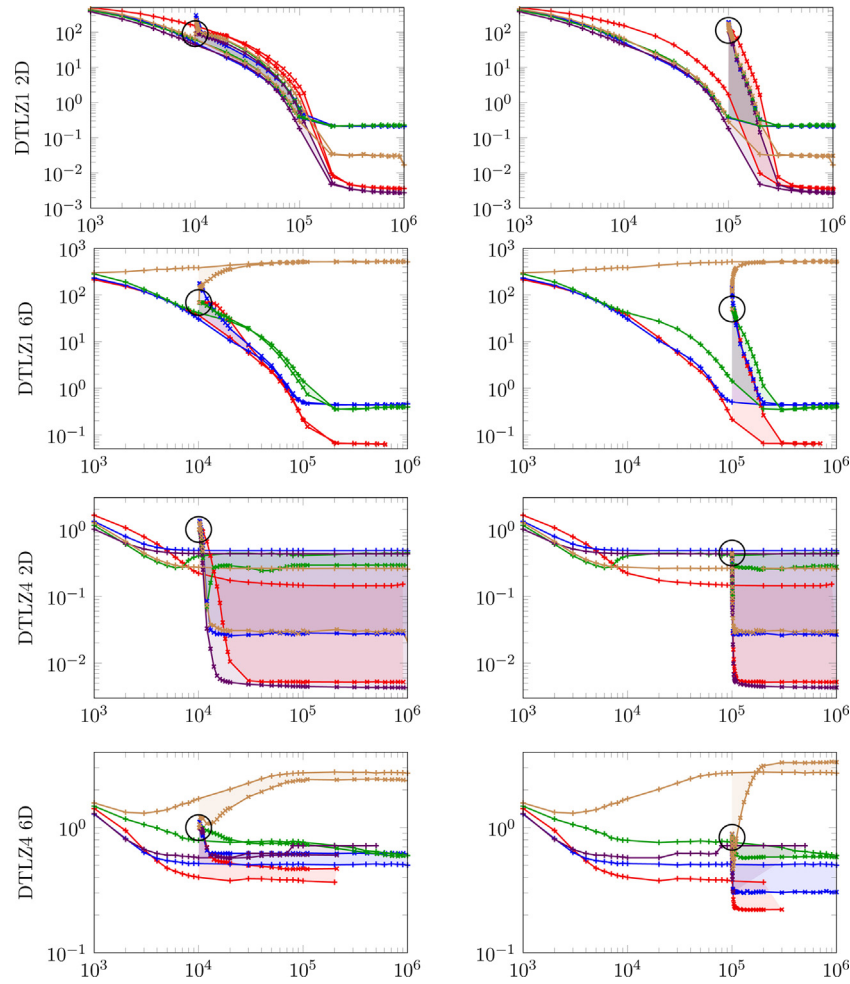


Fig. 1. Comparison of seeding with CORNERSANDCENTRE (left column) and LINEARCOMBINATIONS (right column) on DTLZ1 and DTLZ4 for two and six dimensions. The approximation constant of the Pareto front (y-axis) is shown as a function of the number of fitness function evaluations (x-axis) for the seeded (x) and unseeded (+) versions of AGE (—), IBEA (—), NSGA-II (—), SMS-EMOA (—), and SPEA2 (—). The figures show the average of 100 repetitions each. Smaller approximation constants indicate a better approximation of the front. The plots for the seeded versions are shifted by the number of iterations required by the CORNERSANDCENTRE seeding (10^4 iterations) and the LINEARCOMBINATIONS seeding (10^5 iterations); circles indicate the approximation of the initial seeding. The shaded areas illustrate the difference between seeding and no seeding for a specific algorithm. Plots end prematurely if the time limit of four hours was reached.

However, as it is difficult to compute the exact achieved approximation constant of a known Pareto front, we approximate it. For the quality assessment on the LZ, WFG and ZDT functions, we compute the achieved additive approximations with respect to the Pareto fronts given in the jMetal package. For the DTLZ functions, we draw one million points of the front uniformly at random, and then compute the additive approximation achieved for this set.

We also measure the hypervolume for all experiments. As the behaviors of the five algorithms differ significantly, there is no single reference point that allows for a meaningful comparison of all functions. However, we observe the same qualitative comparison with the hypervolume as we do with the additive approximation. Therefore, we omit all hypervolume values in this paper, because the additive approximation constant gives a much better way to compare the results for these benchmark functions, where the Pareto fronts are known in advance.

In addition to calculating the average ratio of the achieved approximation constant with and without seeding, we also perform a non-parametric test on the significance of the observed behavior. For this, we compare the final approximation of the 100 runs without seeding and the 100 runs with seeding using the Wilcoxon–Mann–Whitney two-sample rank-sum test at the 95% confidence level.

4. Experimental results

Our results are summarized in Tables 1 and 2. They compare the approximation constant achieved with CORNERSANDCENTRE seeding (Table 1) and LINEARCOMBINATIONS seeding (Table 2) with the same number of iterations without seeding. As the seeding itself requires a number of fitness function evaluations (10^4 for CORNERSANDCENTRE and 10^5 for LINEARCOMBINATIONS), we allocate the seeded algorithms fewer fitness function evaluations. This makes it harder for the seeded algorithms to outperform its unseeded variant, as discussed above.

Figs. 1 and 2 show some representative charts how the approximation constant behaves over the runtime of the algorithms. First note that the approximation constant is mostly monotonically decreasing. As a smaller approximation constant corresponds to a better approximation of the Pareto front, this means that most algorithms achieve a better approximation over time. Exceptions are SPEA2 (—), which is unable to handle the six dimensional variants of DTLZ, and NSGA-II (—), which sometimes gets worse after a certain time. For most problems and algorithms, the total maximal number of fitness function evaluations (10^6) was enough such that the algorithms have converged.




Fig. 2. Comparison of seeding with CORNERSANDCENTRE (left column) and LINEARCOMBINATIONS (right column) on ZDT1, ZDT2, LZ09 F1 and LZ09 F2. The approximation constant of the Pareto front (y-axis) is shown as a function of the number of fitness function evaluations (x-axis) for the seeded (\times) and unseeded ($+$) versions of AGE (\rightarrow), IBEA (\leftarrow), NSGA-II (\leftarrow), SMS-EMOA (\leftarrow), and SPEA2 (\leftarrow). The figures show the average of 100 repetitions each. Smaller approximation constants indicate a better approximation of the front. The plots for the seeded versions are shifted by the number of iterations required by the CORNERSANDCENTRE seeding (10^4 iterations) and the LINEARCOMBINATIONS seeding (10^5 iterations); circles indicate the approximation of the initial seeding. The shaded areas illustrate the difference between seeding and no seeding for a specific algorithm. Plots end prematurely if the time limit of four hours was reached.

The small black circles in Figs. 1 and 2 indicate the average approximation constant of the initial seeding after the number of fitness function evaluations needed to calculate it. Note that because of their specific selection schemes, some algorithms like SPEA2 (\leftarrow) and NSGA-II (\leftarrow) sometimes increase the approximation constant after the initial seeding (e.g. on LZ09 F1). Another surprising effect can be observed on LZ09 F2. There, seeding is disadvantageous to NSGA-II (\leftarrow), while it is advantageous to SPEA2 (\leftarrow).

For all considered test problems (not only the ones shown in Figs. 1 and 2), either AGE (\rightarrow) or SMS-EMOA (\leftarrow) reach the best approximation constant. However, for test problems with more than two or three dimensions (cf. Fig. 1), SMS-EMOA (\leftarrow) fails due to the high computational cost of calculating the hypervolume. On some problems AGE (\rightarrow) does not finish all iterations within 10^6 steps, but still achieves the best approximation constant (e.g. DTLZ1 6D and DTLZ4 6D).

For all algorithms, both seedings are beneficial on some test functions. However, the generally more performant algorithms (AGE and SMS-EMOA) typically gain the most from both seedings. On some functions, these algorithms not only achieve a better approximation faster with seeding, but it seems that all best

approximations can only be achieved with seeding (e.g. DTLZ4). On DTLZ4 2D the gap between the approximation constant achieved by SMS-EMOA (\leftarrow) with and without seeding is about two orders of magnitude, which is the difference between a very good approximation of the Pareto front and basically no approximation of the Pareto front.

Tables 1 and 2 give a numerical comparison assuming that about 10% of the fitness evaluations are used for seeding. The shown numbers are the ratios of the median approximation constant without seeding to the median approximation constant with seeding. Values >1.00 indicate where seeding is beneficial in the median. We additionally show statistical significance based on the Wilcoxon–Mann–Whitney two-sample rank-sum test at the 95% confidence level. “ $>$ ” marks statistically significant improvements, “ $<$ ” marks statistically significant worsenings, “ $=$ ” marks statistically insignificant findings. In order to avoid making a Familywise error, we have split our 100 baseline results of NoSEED in two independent subsets and compared our results of CORNERSANDCENTRE and LINEARCOMBINATIONS seeding in Tables 1 and 2 with two independent sets of 50 runs of NoSEED for each seeding. The ratios in Table 1 correspond to the approximation constants after 10^5 function evaluations in the left column of Figs. 1 and 2. The ratios in

Table 2 correspond to the approximation constants after 10^6 function evaluations in the right column of Figs. 1 and 2.

Counting only the statistically significant results over all functions and seedings, the tables show that the majority profits from the seeding. The algorithms which benefit the most are SPEA2 (—) with $20\times">"$ and $12\times"<"$, AGE (—) with $26\times">"$ and $20\times"<"$, and IBEA (—) with $24\times">"$ and $16\times"<"$. There are significant differences depending on the test function. The LZ09 benchmark family profits the most: Summing up the significant results for all algorithms, there are $55\times">"$ and $15\times"<"$. Also for DTLZ4 there are $21\times">"$ and $5\times"<"$. The worst performance of the seeding is achieved on the rather difficult WFG functions: While the CORNERSANDCENTRE seeding achieves over all algorithms $10\times">"$ and $6\times"<"$, the LINEARCOMBINATIONS seeding only achieves $9\times">"$ and $29\times"<"$.

We can do a similar analysis to assess the benefits of the investigated seeding approaches. We observe that over all algorithms the CORNERSANDCENTRE seeding yields in total $52\times">"$ and $27\times"<"$, which is a bit better than the LINEARCOMBINATIONS seeding which yields in total $60\times">"$ and $46\times"<"$. In order to answer the question whether this is statistically significant, we calculate the average rank of all runs with and without seeding, for each of the 48 functions, and each of the 5 algorithms. With this combined data from all runs, functions and algorithms, the Wilcoxon–Mann–Whitney two-sample rank-sum test shows significance at the 95% confidence level that both seedings improve upon no seeding.

5. Conclusions

Seeding can result in a significant reduction of the computational cost and the number of fitness function evaluations needed. We observe that there is an advantage on many common real-valued fitness functions even if computing an initial seeding reduces the number of fitness function evaluations available for the MOEA. For some functions we observe a dramatic improvement in quality and needed runtime (e.g., DTLZ4 and the LZ09 family).

For practitioners, our results show that it can be worthwhile to apply some form of seeding (especially when evaluations are expensive), but also to investigate different MOEAs as well, as they have proven to benefit differently from seeding. While we observed that seeding can be very beneficial, our experiments could not reveal why this is the case for a particular combination of seeding, algorithm, and function landscape. To answer this, many parts have to be studied: the mappings that the benchmark functions create from the search spaces into the objective spaces, the connectedness between different local Pareto fronts, the adequacy of using CMA-ES in the seeding procedure, and much more. As a next step towards this goal, we propose to investigate seeding for combinatorial optimization problems.

Acknowledgements

The research leading to these results has received funding from the Australian Research Council (ARC) under grant agreement DP140103400 and from the European Union Seventh Framework Programme (FP7/2007–2013) under grant agreement no 618091 (SAGE).

References

- [1] R.B. Agrawal, K. Deb, Simulated Binary Crossover for Continuous Search Space. Technical report, 1994.
- [2] J. Bader, K. Deb, E. Zitzler, Faster hypervolume-based search using Monte Carlo sampling, in: Multiple Criteria Decision Making for Sustainable Energy and Transportation Systems (MCDM '10), Vol. 636 of Lecture Notes in Economics and Mathematical Systems, Springer, 2010, pp. 313–326.
- [3] K. Bringmann, T. Friedrich, Parameterized average-case complexity of the hypervolume indicator, in: 15th Annual Conference on Genetic and Evolutionary Computation Conference (GECCO '13, ACM Press, 2013, pp. 575–582.
- [4] K. Bringmann, T. Friedrich, Approximating the volume of unions and intersections of high-dimensional geometric objects, *Comput. Geom.: Theory Appl.* 43 (2010) 601–610.
- [5] K. Bringmann, T. Friedrich, Approximating the least hypervolume contributor: NP-hard in general, but fast in practice, *Theor. Comput. Sci.* 425 (2012) 104–116.
- [6] K. Bringmann, T. Friedrich, Approximation quality of the hypervolume indicator, *Artif. Intell.* 195 (2013) 265–290.
- [7] K. Bringmann, T. Friedrich, F. Neumann, M. Wagner, Approximation-guided evolutionary multi-objective optimization, in: Proc. 22nd International Joint Conference on Artificial Intelligence (IJCAI '11), Barcelona, Spain, 2011, pp. 1198–1203, IJCAI/AAAI.
- [8] T.C.E. Cheng, A. Janiak, M.Y. Kovalyov, Bicriterion single machine scheduling with resource dependent processing times, *SIAM J. Optim.* 8 (1998) 617–630.
- [9] C. Daskalakis, I. Diakonikolas, M. Yannakakis, How good is the Chord algorithm? in: 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '10), 2010, pp. 978–991.
- [10] K. Deb, A. Pratap, S. Agrawal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II *IEEE Trans. Evolut. Comput.* 6 (2002) 182–197.
- [11] K. Deb, L. Thiele, M. Laumanns, E. Zitzler, Scalable test problems for evolutionary multiobjective optimization, in: *Evolutionary Multiobjective Optimization, Advanced Information and Knowledge Processing*, 2005, pp. 105–145.
- [12] I. Diakonikolas, M. Yannakakis, Small approximate Pareto sets for biobjective shortest paths and other problems, *SIAM J. Comput.* 39 (2009) 1340–1371.
- [13] J.J. Durillo, A.J. Nebro, E. Alba, The jMetal framework for multi-objective optimization: design and architecture, in: *IEEE Congress on Evolutionary Computation (CEC '10)*, 2010, pp. 4138–4325.
- [14] M.T.M. Emmerich, N. Beume, B. Naujoks, An EMO algorithm using the hypervolume measure as selection criterion, in: *3rd International Conference on Evolutionary Multi-Criterion Optimization (EMO '05)*, Springer, 2005, pp. 62–76.
- [15] Y.G. Evtushenko, M. Potapov, Methods of numerical solution of multicriterion problem, in: *Soviet mathematics – doklady*, Vol. 34, 1987, pp. 420–423.
- [16] M. Gong, L. Jiao, H. Du, L. Bo, Multiobjective immune algorithm with nondominated neighbor-based selection, *Evolut. Comput.* 16 (2008) 225–255.
- [17] J.J. Grefenstette, Incorporating problem specific knowledge into genetic algorithms, *Genet. Algorithms Simul. Annealing* 4 (1987) 42–60.
- [18] N. Hansen, The CMA evolution strategy: a comparing review, in: *Towards a New Evolutionary Computation. Advances in Estimation of Distribution Algorithms*, Springer, 2006, pp. 75–102.
- [19] P. Hansen, Bicriterion path problems, in: *Multiple Criteria Decision Making: Theory and Applications*, Vol. 177 of Lecture Notes in Economics and Mathematical Systems, 1980, pp. 109–127.
- [20] G.R. Harik, D.E. Goldberg, Linkage learning through probabilistic expression, *Comput. Methods Appl. Mech. Eng.* 186 (2000) 295–310.
- [21] I. Hatzakis, D. Wallace, Dynamic multi-objective optimization with evolutionary algorithms: a forward-looking approach, in: *Proc. 8th Annual Conference on Genetic and Evolutionary Computation Conference (GECCO '06)*, ACM Press, 2006, pp. 1201–1208.
- [22] A. Hernandez-Diaz, C.A. Coello Coello, F. Perez, R. Caballero, J. Molina, L. Santana-Quintero, Seeding the initial population of a multi-objective evolutionary algorithm using gradient-based information, in: *Proc. Congress on Evolutionary Computation (CEC '08)*, IEEE Press, 2008, pp. 1617–1624.
- [23] E. Hopper, B. Turton, An empirical investigation of meta-heuristic and heuristic algorithms for a 2D packing problem, *Eur. J. Oper. Res.* 128 (2001) 34–57.
- [24] S. Huband, L. Barone, R.L. While, P. Hingston, A scalable multi-objective test problem toolkit, in: *3rd International Conference on Evolutionary Multi-Criterion Optimization (EMO '05)*, Vol. 3410 of LNCS, Springer, 2005, pp. 280–295.
- [25] H. Ishibuchi, N. Tsukamoto, Y. Sakane, Y. Nojima, Indicator-based evolutionary algorithm with hypervolume approximation by achievement scalarizing functions, in: *12th Annual Conference on Genetic and Evolutionary Computation Conference (GECCO '10)*, ACM Press, 2010, pp. 527–534.
- [26] B. Kazimipour, X. Li, A. Qin, Effects of population initialization on differential evolution for large scale optimization, in: *Proc. Congress on Evolutionary Computation (CEC '14)*, 2014, pp. 2404–2411.
- [27] B. Kazimipour, X. Li, A. Qin, A review of population initialization techniques for evolutionary algorithms, in: *Proc. Congress on Evolutionary Computation (CEC '14)*, 2014, pp. 2585–2592.
- [28] E. Keedwell, S.-T. Khu, A hybrid genetic algorithm for the design of water distribution networks, *Eng. Appl. Artif. Intell.* 18 (2005) 461–472.
- [29] H. Li, Q. Zhang, Multiobjective optimization problems with complicated Pareto sets, MOEA/D and NSGA-II, *IEEE Trans. Evolut. Comput.* 13 (2009) 284–302.
- [30] C.-F. Liaw, A hybrid genetic algorithm for the open shop scheduling problem, *Eur. J. Oper. Res.* 124 (2000) 28–42.
- [31] P. Lordin, ϵ -solutions in vector minimization problems, *J. Optim. Theory Appl.* 43 (1984) 265–276, <http://dx.doi.org/10.1007/BF00936165>
- [32] S. Oman, P. Cunningham, Using case retrieval to seed genetic algorithms, *Int. J. Comput. Intell. Appl.* 1 (2001) 71–82.

- [33] C.H. Papadimitriou, M. Yannakakis, On the approximability of trade-offs and optimal access of web sources, in: 41st Annual Symposium on Foundations of Computer Science (FOCS '00), IEEE Press, 2000, pp. 86–92.
- [34] C.H. Papadimitriou, M. Yannakakis, Multiobjective query optimization, in: 20th ACM Symposium on Principles of Database Systems (PODS '01), 2001, pp. 52–59.
- [35] H. Reuter, An approximation method for the efficiency set of multiobjective programming problems, *Optimization* 21 (1990) 905–911.
- [36] G. Ruhe, B. Fruhwirth, ϵ -optimality for bicriteria programs and its application to minimum cost flows, *Computing* 44 (1990) 21–34.
- [37] S. Vassilvitskii, M. Yannakakis, Efficiently computing succinct trade-off curves, *Theor. Comput. Sci.* 348 (2005) 334–356.
- [38] M. Wagner, T. Friedrich, Efficient parent selection for approximation-guided evolutionary multi-objective optimization, in: Proc. IEEE Congress on Evolutionary Computation (CEC '13), IEEE, 2013, pp. 1846–1853.
- [39] M. Wagner, F. Neumann, A fast approximation-guided evolutionary multi-objective algorithm, in: Proc. 15th Annual Conference on Genetic and Evolutionary Computation Conference (GECCO '13), ACM, 2013, pp. 687–694.
- [40] M. Wagner, K. Bringmann, T. Friedrich, F. Neumann, Efficient optimization of many objectives by approximation-guided evolution, *Eur. J. Oper. Res.* 243 (2015) 465–479.
- [41] M. Yang, X. Zhang, X. Li, X. Wu, A hybrid genetic algorithm for the fitting of models to electrochemical impedance data, *J. Electroanal. Chem.* 519 (2002) 1–8.
- [42] E. Zitzler, S. Künzli, Indicator-based selection in multiobjective search, in: 8th International Conference on Parallel Problem Solving from Nature (PPSN VIII), Vol. 3242 of LNCS, Springer, 2004, pp. 832–842.
- [43] E. Zitzler, L. Thiele, Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach, *IEEE Trans. Evolut. Comput.* 3 (1999) 257–271.
- [44] E. Zitzler, K. Deb, L. Thiele, Comparison of multiobjective evolutionary algorithms: empirical results, *Evolut. Comput.* 8 (2000) 173–195.
- [45] E. Zitzler, M. Laumanns, L. Thiele, SPEA2. Improving the strength Pareto evolutionary algorithm for multiobjective optimization, in: *Evolutionary Methods for Design, Optimisation and Control with Application to Industrial Problems (EUROGEN 2001)*, 2002, pp. 95–100.