



Universiteit
Leiden
The Netherlands

Uncrowded hypervolume-based multiobjective optimization with gene-pool optimal mixing

Maree, S.C.; Alderliesten, T.; Bosman, P.A.N.

Citation

Maree, S. C., Alderliesten, T., & Bosman, P. A. N. (2022). Uncrowded hypervolume-based multiobjective optimization with gene-pool optimal mixing. *Evolutionary Computation*, 30(3), 329-353. doi:10.1162/evco_a_00303

Version: Publisher's Version
License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)
Downloaded from: <https://hdl.handle.net/1887/3566466>

Note: To cite this publication please use the final published version (if applicable).

Uncrowded Hypervolume-Based Multiobjective Optimization with Gene-Pool Optimal Mixing

S.C. Maree

Amsterdam UMC, University of Amsterdam, The Netherlands
Centrum Wiskunde & Informatica, Amsterdam, The Netherlands

s.c.maree@cw.nl

T. Alderliesten

Leiden University Medical Center, The Netherlands

t.alderliesten@lumc.nl

P.A.N. Bosman

Centrum Wiskunde & Informatica, Amsterdam, The Netherlands

peter.bosman@cw.nl

https://doi.org/10.1162/evco_a_00303

Abstract

Domination-based multiobjective (MO) evolutionary algorithms (EAs) are today arguably the most frequently used type of MOEA. These methods, however, stagnate when the majority of the population becomes nondominated, preventing further convergence to the Pareto set. Hypervolume-based MO optimization has shown promising results to overcome this. Direct use of the hypervolume, however, results in no selection pressure for dominated solutions. The recently introduced Sofomore framework overcomes this by solving multiple interleaved single-objective dynamic problems that iteratively improve a single approximation set, based on the uncrowded hypervolume improvement (UHVI). It thereby however loses many advantages of population-based MO optimization, such as handling multimodality. Here, we reformulate the UHVI as a quality measure for approximation sets, called the uncrowded hypervolume (UHV), which can be used to directly solve MO optimization problems with a single-objective optimizer. We use the state-of-the-art gene-pool optimal mixing evolutionary algorithm (GOMEA) that is capable of efficiently exploiting the intrinsically available grey-box properties of this problem. The resulting algorithm, UHV-GOMEA, is compared with Sofomore equipped with GOMEA, and the domination-based MO-GOMEA. In doing so, we investigate in which scenarios either domination-based or hypervolume-based methods are preferred. Finally, we construct a simple hybrid approach that combines MO-GOMEA with UHV-GOMEA and outperforms both.

Keywords

Multiobjective optimization, black-box optimization, hypervolume indicator.

1 Introduction

A multiobjective (MO) optimization problem is given by a to-be-minimized objective function $\mathbf{f} : \mathcal{X} \rightarrow \mathbb{R}^m$, with $\mathbf{f} = [f_1, \dots, f_m]$ being an m -dimensional vector function and $\mathcal{X} \subseteq \mathbb{R}^n$ the n -dimensional decision space. MO problems do not naturally imply a complete ordering of solutions $\mathbf{x} \in \mathcal{X}$ based on their objective values. When the objectives contradict each other, no single solution minimizes all objectives simultaneously. The optimum of an MO problem can then be defined in terms of *Pareto optimality* (Knowles et al., 2006). A solution is said to *dominate* another solution when it is strictly better in one or more objectives, and is not worse in any of the other objectives. A solution

Manuscript received: 10 April 2020; revised: 3 May 2021, 2 August 2021, 9 November 2021; accepted: 28 November 2021.

is *Pareto optimal* when there exist no solutions that dominates it. The *Pareto set* is the set of all Pareto optimal solutions, and its image under f is known as the *Pareto front*. In practice, the aim of MO optimization is to provide a decision maker with a set of nondominated solutions, known as an *approximation set* $A \subset \mathcal{X}$, whose image under f , the *approximation front*, approximates the Pareto front. The decision maker then selects a preferred solution from this set. Real-valued MO problems generally have an infinite number of Pareto-optimal solutions. As it is impossible to obtain all of these, MO optimization algorithms generally attempt to find an approximation set containing a manageable number of nondominated solutions that forms a good representation of the entire front. This results in an inherent trade-off in the two-sided optimization goal of MO optimization, since it is desirable to obtain a diverse approximation set as well as an approximation set that is close to the Pareto set (Bosman and Thierens, 2003).

MO evolutionary algorithms (MOEAs) have shown to be very successful for (black-box) MO optimization in practice (Deb, 2001). These algorithms maintain a population of solutions, and generally equip a domination-based fitness selection scheme (Deb, Pratap et al., 2002; Zitzler et al., 2001; Bouter, Luong, Alderliesten et al., 2017), such as in the nondominated sorting genetic algorithm (NSGA-II) (Deb, Pratap et al., 2002), arguably the best-known MOEA. By equipping a domination-based selection scheme, MOEAs are capable of quickly obtaining a population of only nondominated solutions. When this happens, there is no further selection pressure toward the Pareto set. MOEAs then typically aim to improve diversity in the population, for example, using the crowding distance in NSGA-II. This implies, however, that proximity to the Pareto set is no longer systematically improved, which could even lead to cyclic behavior, with the MOEA improving diversity but worsening proximity (Berghammer et al., 2012). Consequently, even though the approximation sets obtained with domination-based MOEAs are often sufficient for practical use, these algorithms do not result in approximation sets converging to the Pareto set, which is a limitation we aim to address in this work.

Indicator-based MO optimization, especially based on the hypervolume measure (Zitzler and Thiele, 1999), has shown promising results to overcome stagnation that occurs with domination-based selection schemes (Beume et al., 2007; Igel et al., 2007). Practically used indicators in indicator-based optimization are the R2 indicator (Hansen and Jaszkiewicz, 1998), the epsilon indicator (Zitzler and Künzli, 2004), and the hypervolume measure (Zitzler et al., 2003). The hypervolume measure is particularly interesting as it is the only known measure that is strictly monotonic with respect to Pareto-dominance (Knowles, 2002; Fleischer, 2003). This means that the approximation set with optimal hypervolume is a subset of the Pareto set. A limitation of the hypervolume measure is that dominated solutions have no contribution, in the sense that there is no selection pressure for dominated solutions toward a nondominated region in the decision space (Emmerich et al., 2007). Therefore, in the first usages of the hypervolume measure in the optimization process, it was only partially used to guide the search (Emmerich et al., 2005; Zitzler et al., 2007; Igel et al., 2007; Nicolini, 2005; Mostaghim et al., 2007). Specifically, in the \mathcal{S} -metric selection evolutionary multiobjective optimization algorithm (SMS-EMOA) (Beume et al., 2007), indicators such as the hypervolume measure, are therefore used as a secondary fitness after nondominated sorting. A recently published technique to overcome this limitation of the hypervolume, without relying on domination-based properties, is to assign a fitness value to dominated solutions based on their distance to the boundary of the dominated area,

that is, the boundary surface of the hypervolume (Touré et al., 2019). The resulting uncrowded hypervolume improvement (UHVI) is not a set-based measure, and cannot be used directly to optimize approximation sets. The Sofomore framework (Touré et al., 2019) was therefore formulated, in which multiple optimizers are interleaved that each solve a dynamic single-objective optimization problem. In this framework, a single approximation set is iteratively optimized, but by doing so, it loses some of the advantages of population-based evolutionary algorithms, such as the ability to escape local optima.

The decomposition approach of the Sofomore framework shares some resemblance with the MOEA/D framework (Zhang and Li, 2007), in which the multiobjective problem is decomposed into a number of single-objective problems by the use of a scalarization function. Each of these single-objective problems corresponds to a direction in objective space, along which optimization is performed. The directions that are considered by MOEA/D are static, and a proper scalarization function should thus be chosen for the problem at hand. In contrast, in the Sofomore framework, and the framework we will introduce in this work, the directions along which single-objective optimization is performed are dynamic, and are dependent on each other.

In this work, we show how and when hypervolume-based MO optimization can be used to replace or supplement domination-based MOEAs. To this end, we formulate the uncrowded hypervolume (UHV) measure, which can be directly used to achieve population-based hypervolume-driven MO optimization using a single-objective problem formulation. We use the gene-pool optimal mixing evolutionary algorithm (GOMEA) (Bouter, Alderliesten et al., 2017; Bouter, Luong, Alderliesten et al., 2017) as baseline algorithm to compare domination-based and hypervolume-based MO optimization. This algorithm has both a single-objective (Bouter, Alderliesten et al., 2017) and MO (Bouter, Luong, Alderliesten et al., 2017) version, with recently published excellent results. In both versions, essentially, the same variation operators are used, which provides a more fair and modern comparison than when e.g., NSGA-II and the Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES) (Hansen and Ostermeier, 2001) were to be used. Furthermore, in this article we will show that the single-objective GOMEA can be used to efficiently solve the UHV problem formulation by exploiting grey-box properties that are intrinsically present in the UHV, while still assuming the MO problem itself to be black box.

The remainder of this article is organized as follows. In Section 2, we discuss the hypervolume and related measures. In Section 3, we introduce an indicator-based MO optimization problem formulation based on the UHV, and how we can efficiently solve this problem with GOMEA, resulting in UHV-GOMEA. In Section 4.1, we furthermore incorporate GOMEA in the Sofomore framework (Sofomore-GOMEA), and in Section 4.2, we discuss MO-GOMEA. In Section 5, we empirically compare these three algorithms and discuss which of the approaches is preferred in which scenario, and how to combine them into a simple hybrid approach, on both simple benchmark problems and the commonly used WGF test suite (Huband et al., 2005). We discuss the overall results in Section 6 and conclude in Section 7.

2 Preliminaries: The Hypervolume Measure

A solution $\mathbf{x} \in \mathcal{X}$ is said to *weakly dominate* another solution $\mathbf{y} \in \mathcal{X}$, written as $\mathbf{x} \leq \mathbf{y}$, if and only if $f_i(\mathbf{x}) \leq f_i(\mathbf{y})$ for all $i \in \{1, \dots, m\}$. When the relation $f_i(\mathbf{x}) < f_i(\mathbf{y})$ is furthermore strict for at least one i , we say that \mathbf{x} *dominates* \mathbf{y} , written as $\mathbf{x} < \mathbf{y}$ or, with a slight

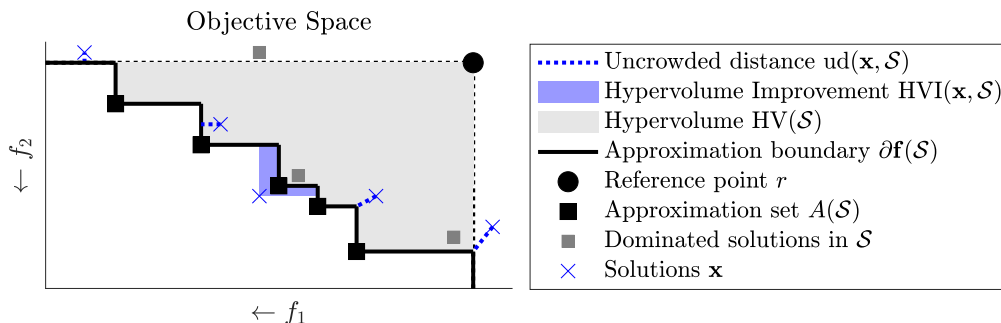


Figure 1: Illustration of the hypervolume measure for a bi-objective minimization problem, together with the uncrowded distance and hypervolume improvement of some example solutions x with respect to solution set S .

abuse of notation, as $f(x) < f(y)$. A solution that is not dominated by any other solution in \mathcal{X} is called *Pareto optimal*. The *Pareto set* A^* can then be formulated as $A^* = \{x \in \mathcal{X} : \nexists y \in \mathcal{X} : y < x\} \subset \mathcal{X}$, while the *Pareto front* is given by $\{f(x) : x \in A^*\} \subset \mathbb{R}^m$.

Let $\wp(\mathcal{X})$ be the powerset of \mathcal{X} , i.e., the set of all solution sets $S \subseteq \mathcal{X}$. The hypervolume measure $HV : \wp(\mathcal{X}) \rightarrow \mathbb{R}$ (Zitzler and Thiele, 1999; Zitzler et al., 2003; Auger et al., 2009a) of a solution set S with respect to a reference point $r \in \mathbb{R}^m$ measures the volume dominated by all $x \in S$, and bounded by r , as is illustrated in Figure 1. Let $A : \wp(\mathcal{X}) \rightarrow \wp(\mathcal{X})$ be the approximation set of S , $A(S) = \{x \in S : f(x) < r, \nexists y \in S : y < x\}$, that is, the largest subset of S that contains only nondominated solutions within the region defined by r . Additionally, let the hypervolume improvement $HVI : \mathcal{X} \times \wp(\mathcal{X}) \rightarrow \mathbb{R}$ of a solution x with respect to a solution set S be defined as the increase in hypervolume when x is added to S , that is, $HVI(x, S) = HV(S \cup \{x\}) - HV(S)$, as shown in Figure 1.

The difference between domination-based improvements and hypervolume-based improvements is illustrated in Figure 2 for the bi-sphere problem with hypervolume reference point $r = (11, 11)$. This problem is composed of two single-objective sphere problems, $f_{\text{sphere}}(x) = \sum_{i=1}^n x_i^2$, of which one is translated, $f_{\text{bi-sphere}}(x) = [f_{\text{sphere}}(x); f_{\text{sphere}}(x - e_1)]$, where e_i is the i^{th} unit vector, with all zeros except a one in the i^{th} position. Figure 2 shows that $HVI(x, S) \geq 0 \Leftrightarrow \{\nexists y \in S : y < x\}$ holds. Indeed, both approaches have the same improvement region (in blue). However, when comparing the *improvement region for the middle solution* (in orange), that is, when that solution is replaced in S , we see that its hypervolume improvement region is larger, essentially making it easier to find improvements. Additionally, the hypervolume improvement region consists of two disconnected subsets, indicating that it takes diversity into account, in contrast to domination-based improvements.

When x is dominated by a solution in S , its hypervolume improvement $HVI(x, S)$ is zero. The uncrowded hypervolume improvement (UHVI) (Touré et al., 2019) was recently introduced to overcome this. Let $\partial f(S)$ be the *approximation boundary*, that is, the boundary between the dominated and nondominated region in objective space, bounded by the reference point r , as illustrated in Figure 1. Let $ud(x, S)$ be the *uncrowded distance*, which measures the shortest Euclidean distance between x and $\partial f(S) = \partial f(A(S))$, when x is dominated by any solution in S or outside the region defined by r . Else, we set $ud(x, S) = 0$. It is called the uncrowded distance as the shortest distance

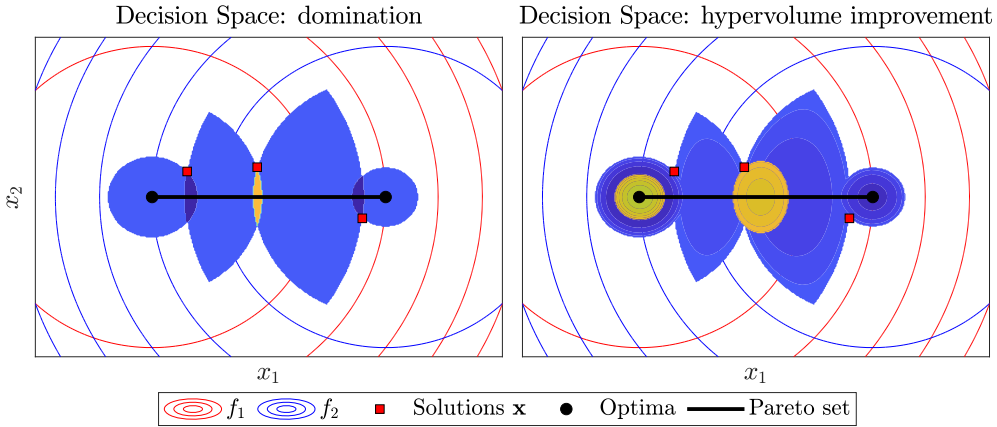


Figure 2: Domination region (left) compared with the hypervolume improvement region (right) for the bi-sphere problem. Additionally, in both subfigures, the orange region is the improvement region when the middle solution is being replaced. Solutions in the light-blue region are nondominated with respect to the current solutions. The dark-blue regions dominate one of the current solutions. In the right subfigure, solutions in the blue region improve the hypervolume if added to the solution set, and darker blue represents larger improvements.

to $\partial f(S)$ is obtained for a point on $\partial f(S)$ that is not in S itself. The UHVI can then be defined as,

$$\text{UHVI}(x, S) = \text{HVI}(x, S) - \text{ud}(x, S).$$

It is interesting to note that the uncrowded distance is similar to the distance measure d^+ that was used to construct a weak Pareto compliance version of the inverted generational distance indicator, called the IGD^+ (Ishibuchi et al., 2015). In the IGD^+ , a reference set Z of nondominated solutions is used (i.e., a subset of the known Pareto set), and the distance of reference solutions $z \in Z$ toward the approximation boundary $\partial f(S)$ is computed. There, the solutions z are nondominated with respect to S . Here, however, we are particularly interested in the other case, where solutions are dominated with respect to S .

3 The Uncrowded Hypervolume

The UHVI is a measure of solution quality with respect to a solution set S , but not a measure of quality for S itself. We therefore introduce the *uncrowded hypervolume* (UHV) in this work as the hypervolume of S penalized by all uncrowded distances,

$$\text{UHV}(S) = \text{HV}(S) - \frac{1}{|S|} \sum_{x \in S} \text{ud}(x, S)^m. \tag{1}$$

The uncrowded distances are taken to the power m such that they have the same unit as the hypervolume. As improving a nondominated solution in S could increase both its hypervolume and the uncrowded distances, the factor $1/|S|$ is added to guarantee that an improvement in hypervolume is not negatively influenced by the increase in uncrowded distances. The uncrowded hypervolume is a strictly monotonic indicator on the space of approximation sets (i.e., sets containing only nondominated solutions), as it is equal to the hypervolume for those sets. However, on the entire space of solution sets, strict monotonicity does not hold.

3.1 UHV-Maximization

In an indicator-based multiobjective optimization problem (IBMOP), a quality indicator is used to assign a quality value to a solution set. The underlying idea is that a single-objective optimizer then can be used to explicitly search for a solution set \mathcal{S} that maximizes this indicator (Zitzler and Künzli, 2004; Beume et al., 2007). Let $I_f : \wp(\mathcal{X}) \rightarrow \mathbb{R}$ be such an indicator, with respect to the multiobjective problem given by \mathbf{f} . To be able to search the space of solution sets, we parameterize solution sets by considering solutions sets of a predetermined size $|\mathcal{S}_p| = p \geq 1$. Let ϕ be a vector of concatenated decision variables of the solutions in \mathcal{S}_p , i.e., $\phi = [\mathbf{x}_1 \cdots \mathbf{x}_p] \in \mathbb{R}^{p \cdot n}$. Inversely, let $S(\phi) = \{\mathbf{x}_1, \dots, \mathbf{x}_p\}$ be the operator that transforms ϕ into a solution set. IBMOPs are then formulated as,

$$\begin{aligned} & \text{maximize} && g_{\mathbf{f},p} : (\mathcal{X})^p \rightarrow \mathbb{R}, \\ & \text{with} && g_{\mathbf{f},p}(\phi) = I_f(S(\phi)), \\ & && \mathbf{f} : \mathcal{X} \subseteq \mathbb{R}^n \rightarrow \mathbb{R}^m, p \geq 1 \\ & && \phi = [\mathbf{x}_1 \cdots \mathbf{x}_p] \in (\mathcal{X})^p \subseteq \mathbb{R}^{p \cdot n}. \end{aligned} \quad (2)$$

We will refer to a solution ϕ of the IBMOP with objective function g as a g -solution, while solutions of the original MO optimization problem given by \mathbf{f} are called MO-solutions from now on.

This IBMOP formulation fully specifies the MO optimization problem, as well as its optimum, that is, the resulting distribution of solutions along the front. In case of the (uncrowded) hypervolume as quality indicator, this optimal distribution is known as the optimal μ -distribution (Auger et al., 2009b). If and only if the front is linear, the solutions in the optimal distribution are equally spaced along the front. In general, the density of solutions is proportional to the negative slope of the front (Auger et al., 2009b).

This clear and unique definition of the optimum of the IBMOP allows us to discuss convergence to optimality, in contrast to the general aim of MO optimization that translates to a trade-off between proximity and diversity (Bosman and Thierens, 2003). Furthermore, since the hypervolume measure is strictly monotonic with respect to Pareto dominance, optimality here implies that the obtained solutions are a subset of the Pareto set.

Note that the optimal distribution of solutions along the front is determined by the choice of reference point, which we fix during the course of the optimization run. If a dynamic reference point would have been used, cyclic behavior could again occur.

3.2 UHV-GOMEA

We use the real-valued single-objective gene-pool optimal mixing evolutionary algorithm (GOMEA) to solve IBMOPs with the UHV as indicator, which we call UHV-GOMEA. We use GOMEA as published in Bouter, Alderliesten et al. (2017), and make only minor adaptations to better align the algorithm with IBMOPs. We discuss the outline of GOMEA here, and refer the reader to Bouter, Alderliesten et al. (2017) for a detailed description with pseudo code.

GOMEA is a model-based evolutionary algorithm that maintains a population of N solutions. For the IBMOP these are the g -solutions $\phi^j \in \mathbb{R}^{p \cdot n}$, $j = 1, \dots, N$. The variation operator in GOMEA is called gene-pool optimal mixing (GOM). GOM was designed specifically to perform variation by adapting only a few decision variables at a time, and thereby exploiting that objective functions can often be quickly updated if only a few

variables change. This notion of *partial evaluations* typically requires that some problem knowledge is known, resulting in a *grey-box* scenario. IBMOPs are by definition grey-box, in the sense that it is known how to update the indicator value when the decision variables corresponding to only one (or a few) MO-solution change, without having to reevaluate all other MO-solutions, or recompute the indicator value from scratch. This means that grey-box properties of the IBMOP can be exploited while still considering f as a black-box.

GOMEA is equipped with a *linkage model* that specifies which subsets of decision variables must be adapted simultaneously. A linkage model L is a subset of the power set $\wp(I)$ of all decision variables, which is $I = \{1, \dots, p \cdot n\}$ in case of an IBMOP with n being the number of decision variables of the multiobjective problem, and p the solution set size. A linkage model is therefore also known as a *family of subsets* (FOS). In GOM, variation is performed by iteratively updating only the decision variables specified by a linkage subset $l \in L$, with $l \subseteq I$. We denote the subset of decision variables in a g -solution ϕ specified by l with $\phi_{(l)} \in \mathbb{R}^{|l|}$. An $|l|$ -dimensional Gaussian distribution is then estimated from the τN best g -solutions (with $\tau = 0.35$ (Bouter, Luong, Witteveen et al., 2017)). From this distribution, to update each g -solution in the population, new values for the decision variables specified by l are sampled. In GOM, only improvements to g -solutions are accepted. Else, the proposed update is discarded.

Originally in GOM, the sample distributions for all linkage subsets $l \in L$ were estimated at the beginning of each generation (Bouter, Alderliesten et al., 2017). Instead, here, we estimate the sample distribution for each linkage subset right before sampling new values for the variables in that subset. By doing so, the sample model is learned based on the most recent set of g -solutions. We found this to improve the rate of convergence, at increased computationally complexity of only $\mathcal{O}((|L| - 1) \cdot N \log N)$ per generation, as selection now needs to be performed for each linkage subset. A Cholesky decomposition (Lay, 1993; Higham, 2008) of the covariance matrix is then performed which is required for the sampling. This decomposition requires a positive definite covariance matrix. Due to numerical errors or small population size, it might be that the decomposition fails (Maree, 2012). If this happens once for a linkage subset, we from then on perform a regularization using the Ledoit–Wolf shrinkage estimator (LWSE) (Ledoit and Wolf, 2012). The LWSE estimates a covariance matrix based on a convex combination of the maximum likelihood estimator (MLE) and a prior matrix, for which we use the diagonal variance matrix (i.e., the diagonal of the MLE). Additionally, when $|l| > \tau N - 1$, we always use a diagonal variance matrix without attempting a decomposition.

3.2.1 IBMOP Linkage Models

Any linkage model can be used in GOMEA, but we discuss three models that we employ in UHV-GOMEA in this work. Let $I_{(i)} = \{\phi_{(i-1)n+1}, \dots, \phi_{in}\}$ for $i = 1, \dots, p$ be the decision variables corresponding to the MO-solution \mathbf{x}_i of g -solution $\phi = [\mathbf{x}_1 \cdots \mathbf{x}_p]$. The three linkage models we consider are then given by,

$$\begin{aligned} Lm &= \{I_{(1)}, I_{(2)}, \dots, I_{(p)}\}, \\ Lf &= \{I_{(1)} \cup I_{(2)} \cup \dots \cup I_{(p)}\} = I, \\ Lt &= \text{UPGMA}(Lm). \end{aligned} \quad (3)$$

In the *marginal linkage model* Lm , only decisions variables corresponding to the same MO-solution are considered to be dependent, that is, $\phi_{(l_i)} = \mathbf{x}_i$ for $l_i \in Lm$. In the *full*

linkage model Lf , the linkage model contains only a single subset, in which all decision variables are considered to be dependent, that is, $\phi_{(l)} = \phi_{(l)} = \phi$. Finally, the *linkage tree model* Lt contains multiple levels of linkage, constructed by hierarchically clustering the subsets of the marginal model Lm using UPGMA (Gronau and Moran, 2007), where the two *nearest* subsets are merged iteratively and added to Lt . Since each subset in $l_i \in Lm$ corresponds to the parameters of a single MO-solution \mathbf{x}_i , we can compute the mean objective values of each l_i , $\mathbf{m}_i = \frac{1}{N} \sum_{j=1}^N \mathbf{f}(\phi_{(l_i)}^j) = \frac{1}{N} \sum_{j=1}^N \mathbf{f}(\mathbf{x}_i^j)$. We use the distance between these means as the distance measure between linkage subsets. Since these change over time, the linkage tree is reconstructed every generation. Linkage tree construction results in $|Lt| = 2p - 1$ subsets, and Lt always contains all elements of Lm as well as Lf , but the other subsets depend on the merge order defined by this distance measure. Additionally, in GOM, all linkage subsets for which $|l| > \tau N - 1$ are skipped, as for these relatively large linkage subsets, dependencies cannot be estimated, as a full-rank covariance matrix cannot be estimated.

3.2.2 Permuting MO-Solutions

Since permuting solutions in \mathcal{S} does not influence its corresponding hypervolume, we aim to reorder the MO-solutions in $\mathcal{S}(\phi)$, and therefore g -solution ϕ , at the beginning of a generation, such that all $\phi_{(l_i)}^j = \mathbf{x}_i^j$ for each g -solution $j = 1, \dots, N$ belong to a similar part of the approximation front. Therefore, we again compute the objective-space means \mathbf{m}_i for $i = 1, \dots, p$. For each g -solution ϕ^j , the MO-solutions are reordered in a greedy fashion by iteratively finding the (next-)nearest MO-solution-mean pair.

3.2.3 Elitist Archive

An elitist archive $\mathcal{E} \subset \mathcal{X}$ is maintained that contains all nondominated MO-solutions $\mathbf{x} \in \mathcal{X}$ that were evaluated during optimization. To keep the archive size tractable, adaptive objective space discretization was used as presented in Luong and Bosman (2012). If, with the archiving scheme, the archive exceeds the target size of $N_{\mathcal{E}} = 1000$ MO-solutions, the objective space is discretized into boxes, and only one MO-solution per box is maintained. Newly obtained nondominated MO-solutions are then only added to the archive when they end up in an empty objective-space box, or when they dominate the MO-solution in that box. This archiving scheme is also used in MO-GOMEA (Bouter, Luong, Alderliesten et al., 2017), and the target size of the elitist archive \mathcal{E} is set to $N_{\mathcal{E}} = 1000$ MO-solutions. The elitist archive is not an essential part of UHV-GOMEA, but merely added to allow for a comparison with archive-based MOEAs.

3.3 Generational Computational Complexity

A main limitation of the use of the hypervolume measure that is often mentioned, is its computational complexity. To compute the hypervolume corresponding to a solution set containing p nondominated solutions with $m \geq 3$ objectives, different algorithms have been proposed (Fonseca et al., 2006; Beume et al., 2009), yielding a computational complexity of $\mathcal{O}(p^{m-2} \log p)$, which becomes the dominating term when m is large. For $m = 2$, however, the hypervolume can be computed by simply sorting the nondominated solutions, resulting in a computational complexity of only $\mathcal{O}(p \log p)$. Additionally, using the sorted approximation set, computing the uncrowded distances can be performed in $\mathcal{O}(mp)$ time. Note also that approximation set size p is typically chosen to be rather small (e.g., $p = 9$).

In contrast, “Fast Nondominated Sorting” in NSGA-II or MO-GOMEA with a population of N MO-solutions is commonly performed in $\mathcal{O}(mN^2)$ time (Deb, Pratap et al.,

2002), however, there, N is typically larger (e.g., $N = 100$). Also, maintaining an elitist archive of size $N_{\mathcal{E}} \leq 1000$ has a computation complexity of $\mathcal{O}(mN_{\mathcal{E}})$, and is essential for the performance of MO-GOMEA (Bouter, Luong, Alderliesten et al., 2017). So, while the hypervolume indicator might not be directly suitable for optimization problems with many objectives, it is well-suited for bi-objective optimization problems. We did, however, not yet explore potential speedups by considering that the hypervolume measure does not have to be computed from scratch every generation, but can be updated efficiently when only one or a few solutions change, which is particularly the case in our problem formulation.

4 Other GOMEA Variants

4.1 Sofomore-GOMEA

While UHV-GOMEA optimizes the UHV by manipulating sets of MO-solutions in a population-based approach, the Sofomore framework (Touré et al., 2019) was introduced to iteratively optimize the UHVI one MO-solution at a time. Specifically, Sofomore performs a search around a single solution set \mathcal{S}_p of fixed size p . For each MO-solution $\mathbf{x}_i \in \mathcal{S}_p$, a single-objective optimizer is initialized that solves the single-objective dynamic optimization problem $h_i : \mathcal{X} \rightarrow \mathbb{R}$ given by $h_i(\mathbf{x}_i | \mathcal{S}_p \setminus \{\mathbf{x}_i\}) = \text{UHVI}(\mathbf{x}_i, \mathcal{S}_p \setminus \{\mathbf{x}_i\})$. Note that the decision space of h_i and the MO problem given by \mathbf{f} are the same, and an h_i -solution is thus also an MO-solution, but with different fitness value assignment. Steps of individual optimizers are then interleaved. Intuitively, each optimizer in turn aims to replace \mathbf{x}_i by the solution with maximal hypervolume contribution with respect to \mathcal{S}_p , while keeping the other MO-solutions in \mathcal{S}_p fixed.

To solve the p dynamic single-objective optimizations problems given by h_i , we again use GOMEA to get a fair comparison basis for comparing different optimization approaches. Our implementation largely agrees with the combination of Sofomore and CMA-ES (called COMO-CMA-ES), as presented in Touré et al. (2019), but a few minor changes were made that we discuss here. We refer the reader to Touré et al. (2019) for further details on Sofomore and COMO-CMA-ES.

We assumed that the MO problem is black-box, and we therefore use a full linkage model for Sofomore-GOMEA, as h_i cannot be partially evaluated. We maintain the default uniform initialization of GOMEA. In this way, Sofomore-GOMEA and HV-GOMEA are initialized from exactly the same set of MO-solutions (when the same random seed is used). Additionally, and most important, the Sofomore approach results in a set of *dynamic* optimization problems. As GOMEA accepts only improvements to solutions, it is beneficial to recompute all fitness values (i.e., h_i) at the beginning of each generation, as they might have changed since the last generation. Note that this requires $h_i(\mathbf{x}_i | \mathcal{S}_p \setminus \{\mathbf{x}_i\}) = \text{UHVI}(\mathbf{x}_i, \mathcal{S}_p \setminus \{\mathbf{x}_i\})$ to be updated, but there is no need to re-compute $\mathbf{f}(\mathbf{x}_i)$ itself. In this work, we evaluate performance based on the number of MO function evaluations (MO-fevals), which is thus unaffected by this reevaluation, although overall computation time will increase. This reevaluation is not required for UHV-GOMEA, as it does not solve a dynamic problem. Finally, we use the same archive for Sofomore-GOMEA as described in Section 3.2.3.

4.2 MO-GOMEA

MO-GOMEA is a domination-based MOEA (Bouter, Luong, Alderliesten et al., 2017). MO-GOMEA optimizes a population of MO-solutions that is aimed to approximate the Pareto front by balancing diversity and proximity. Besides the main population,

MO-GOMEA maintains an elitist archive \mathcal{E} , as described in Section 3.2.3. In MO-GOMEA, MO-solutions are copied back from the archive into the population. It can therefore be roughly said that the aim of MO-GOMEA is to obtain an elitist archive that approximates the Pareto front as good as possible. We discuss the main characteristics of MO-GOMEA here, and refer the reader to Bouter, Luong, Alderliesten et al. (2017) for a full description of the algorithm.

Again, as we assumed the MO problem to be black-box, we use a full linkage model for MO-GOMEA. From a population of N_{mo} MO-solutions, truncation selection is performed based on domination rank, resulting in a selection of size τN . This selection is clustered into K_{mo} clusters, each of size $2\tau N/K_{mo}$. Each cluster models a part of the approximation front, and for this an objective-space based clustering method is used that guarantees overlapping clusters of equal size. For each cluster, a Gaussian distribution is estimated to sample new MO-solutions from. Similar to the single-objective GOMEA, MO-GOMEA only accepts improvements, meaning that offspring need to either dominate the parent, or be accepted into the elitist archive.

To align MO-GOMEA with the other algorithms, we set $N_{mo} = p \cdot N$ and $K_{mo} = 2p$ such that the overall number of MO-solutions in the populations is the same, and all sample distributions are estimated from the same number of MO-solutions. MO-GOMEA estimates its sample distributions in a similar fashion as the single-objective GOMEA which was used in UHV-GOMEA and Sofomore-GOMEA. This makes a comparison between these three approaches most fair. Finally, to be able to compare the limited-size $|S_p| = p$ of UHV-GOMEA and Sofomore-GOMEA, we perform greedy hypervolume subset selection (gHSS) (Guerreiro et al., 2016) to select p solutions from the elitist archive \mathcal{E} .

4.3 A Hybrid Method

MO-GOMEA is expected to perform better initially, but to stagnate in terms of proximity to the Pareto set when the majority of MO-solutions in the population is nondominated. We construct a simple hybrid approach where we initially run MO-GOMEA, which we terminate when it stagnates, that is, when 90% of the MO-solutions in the population are nondominated, or when the elitist archive target size is hit. We then switch to UHV-GOMEA-Lm starting from the elitist archive \mathcal{E} that MO-GOMEA obtained so far. \mathcal{E} is clustered into p clusters of equal size $2|\mathcal{E}|/p$ using the same clustering method that is used in MO-GOMEA. For this, the cluster means are initialized with gHSS, and distances are measured in decision space. If the cluster size is smaller than the desired population size, that is, $2|\mathcal{E}|/p < N$, the remainder of the MO-solutions is sampled uniformly.

5 Experiments

5.1 Experiment 1: Convergence Analysis

In Section 2, we introduced the bi-sphere problem, which we will use to demonstrate convergence properties of various MO optimization algorithms. The bi-sphere problem is a separable problem, defined on the entire real space $\mathcal{X} = \mathbb{R}^n$. We consider an instance of this problem with $n = 10$ MO-decision variables and aim for $p = 9$ MO-solutions, resulting in an IBMOP with 90 decision variables. All algorithms are initialized on the domain $[-100, -50]^n$, away from the Pareto set, which is a straight line between the origin and \mathbf{e}_1 .

With this experiment, we aim to demonstrate convergence toward the Pareto set. As performance metric, we make use of the generational distance (GD) (Zitzler et al., 2003),

$$GD(\mathcal{A}_p, \mathcal{A}^*) = \frac{1}{|\mathcal{A}_p|} \sum_{\mathbf{x} \in \mathcal{A}_p} \min_{\mathbf{y} \in \mathcal{A}^*} \|\mathbf{f}(\mathbf{x}) - \mathbf{f}(\mathbf{y})\|,$$

with $\mathcal{A}_p = A(\mathcal{S}_p)$ and \mathcal{A}^* the Pareto set. Since a parametric expression of the Pareto set \mathcal{A}^* is available for the bi-sphere problem, the GD can be computed analytically using,

$$\mathbf{y}_{\text{bi-sphere}}^*(\mathbf{x}) := \arg \min_{\mathbf{y} \in \mathcal{A}^*} \|\mathbf{f}_{\text{bi-sphere}}(\mathbf{x}) - \mathbf{f}_{\text{bi-sphere}}(\mathbf{y})\| = [\max(\min(x_1, 1), 0), 0, \dots, 0].$$

Note that the analytically computed GD is Pareto compliant for solution sets of fixed size containing only nondominated solutions, and any approximation set containing only Pareto optimal solutions has $GD = 0$. The analytical GD can thus be interpreted as a measure of proximity.

In addition, we measure the distance to the optimal hypervolume $\Delta HV_p = HV(\mathcal{A}_p^*) - HV(\mathcal{A}_p)$, with \mathcal{A}_p^* the approximation set of size p with optimal hypervolume (Auger et al., 2009a). We determined $HV(\mathcal{A}_p^*)$ empirically by solving lower-dimensional problem instances with a large computational budget by various MOEAs. Finally, we keep track of the approximation set size $|\mathcal{A}_p|$, as a verification that p nondominated solutions were obtained.

We include UHV-GOMEA with the three linkage models, Sofomore-GOMEA, and MO-GOMEA. Due to its separability, the bi-sphere problem can be solved with a diagonal covariance matrix, and we therefore use a small population size $N = 31$ for these GOMEA-based algorithms. Additionally, we incorporate various well-known algorithms to demonstrate their convergence to the Pareto set. NSGA-II (Deb, Pratap et al., 2002) and SMS-EMOA (Beume et al., 2007) are included, as implemented in the `evualgos` Python package (Wessing, 2020). For SMS-EMOA, the population size was set to $p = 9$, and NSGA-II makes use of simulated binary crossover and uses a population size of 31 solutions. From `jMetalpy` (Benítez-Hidalgo et al., 2019), MOEA/D (Zhang and Li, 2007), and NSGA-III (Deb and Jain, 2013) were included, ran with recommended population sizes of respectively 300 and 91 solutions. Performing experiments with the default algorithm settings seems most fair, as this is likely how these algorithms will be applied in practice. Only after observing results, one often considers adapting hyperparameters, such as the population size, to improve performance. For all methods that make use of a population of more than p solutions, we use `gHSS` to select an approximation set \mathcal{A}_p of at most p solutions, as described in Section 4.2.

In all experiments, we measure performance in terms of the number of MO-fevals. All experiments in this work are repeated 30 times, and mean results are shown, unless mentioned otherwise. The hypervolume reference point is set to $r = (11, 11)$ for all problems, which is far away from the Pareto front, thereby aiming that the endpoints of the Pareto front are in \mathcal{A}_p^* . However, even by setting the reference point this far, this is not always achieved for all problems.

Results

Figure 3 shows that UHV-GOMEA and Sofomore-GOMEA exhibit linear convergence in terms of both $\log(\Delta HV_p)$ and $\log(GD)$, albeit at different rates. Sofomore-GOMEA performs best, closely followed by UHV-GOMEA-Lm. Due to the small population size, large linkage elements are filtered from the linkage tree in UHV-GOMEA-Lt, which performs the same as UHV-GOMEA-Lm (up to randomness). UHV-GOMEA-Lf can still solve the problem, but it is inefficient. With full linkage, all MO-solutions are updated




Figure 3: Convergence on the bi-sphere problem (with $n = 10$, $p = 9$, and $N = 31$). Lines show mean scores over 30 independent runs, and shaded areas show min/max scores.

and evaluated simultaneously (albeit independently from the diagonal covariance matrix, due to the small population size), and only then the corresponding hypervolume is computed. With marginal linkage however, MO-solutions are updated and evaluated one by one, and after each newly evaluated MO-solution, it is checked if this improved the corresponding hypervolume, which is beneficial here.

The rate of convergence for the hypervolume-based algorithms is constant as soon as $|A_9| = 9$ for UHV-GOMEA and Sofomore-GOMEA, which shows that linear convergence is due to the hypervolume optimization and not due to the uncrowded distance.

In this experimental setup, all initial solutions are dominated by the reference point. The hypervolume-based algorithms are then fully driven by the uncrowded distance aimed to improve up the reference point. The uncrowded distance within UHV-GOMEA and Sofomore-GOMEA is effective in obtaining a set of nondominated solutions, but is less efficient than all other methods. SMS-EMOA is initially the best algorithm, which is mainly because it maintains the smallest population, and updates solutions one by one. This works especially well for this unimodal problem, for which there is no benefit in maintaining a larger population in order to escape local optima. While the hypervolume is used in the search process of SMS-EMOA, it does not obtain the solution set with optimal hypervolume. Additionally, it is initially very good in terms of proximity to the Pareto set (GD), but does stagnate at some point, in contrast to UHV-GOMEA and Sofomore-GOMEA.

In terms of the time required to obtain $p = 9$ nondominated solutions, a clear relation with the population size can be observed, and the smaller the population, the faster p nondominated solutions are obtained. The domination-based methods MO-GOMEA and both versions of NSGA, as well as MOEA/D, perform very similar on this problem. These methods are especially good at quickly obtaining a good front, but show stagnation in terms of ΔHV_p , as was expected. Their obtained populations contain a good distribution of solutions along the front, including solutions close to the endpoints of the Pareto set, but the approximation fronts have a slightly different distribution of solutions due to the gHSS, which partially explains the stagnation in ΔHV_p . More importantly is that the analytic GD plot shows that these methods do not converge to the Pareto set. The distances to the Pareto front are in the order of 10^{-3} , making differences indistinguishable on a Pareto front plot of reasonable size. However, even though these distances to the Pareto front are small, this experiment does demonstrate that running NSGA or MOEA/D much longer after their entire population is nondominated does not

yield structural improvements, which is a fundamental limitation that can be overcome with a hypervolume-based approach.

In the remainder of this article, we consider only the GOMEA-based methods. This allows us to demonstrate the advantages and disadvantages of domination-based and hypervolume-based optimization, while reducing undesirable behavioral differences due to other factors as much as possible.

5.2 Experiment 2: Modeling Dependencies

The bi-sphere problem could be solved efficiently with a small population size N and without dependency modeling. We now perform a number of experiments to investigate if there are scenarios where modeling the dependencies between MO-solutions is beneficial. For this, we use two additional benchmark problems, constructed from two well-known single-objective functions, $f_{\text{elli}}(\mathbf{x}) = \sum_{i=1}^n 10^{\frac{i-1}{n-1}} x_i$, and $f_{\text{Rosenbrock}}(\mathbf{x}) = \sum_{i=1}^{n-1} (100(x_{i+1} - x_i^2)^2 + (1 - x_i)^2)$. These are, similar to the bi-sphere problem, defined for a scalable number of decision variables n and on $\mathcal{X} = \mathbb{R}^n$. From these, we construct two bi-objective optimization problems,

$$\begin{aligned} \mathbf{f}_{\text{sphere-rotatedElli}}(\mathbf{x}) &= [f_{\text{sphere}}(\mathbf{x}); f_{\text{elli}}(R\mathbf{x} - \mathbf{e}_1)], \\ \mathbf{f}_{\text{sphere-Rosenbrock}}(\mathbf{x}) &= \left[\frac{1}{n} f_{\text{sphere}}(\mathbf{x}); \frac{1}{n-1} f_{\text{Rosenbrock}}(\mathbf{x}) \right], \end{aligned} \quad (4)$$

where R is a rotation matrix that defines a rotation around the origin of $\pi/4$ radians in all principal directions. The sphere-rotatedElli problem has the same Pareto front as the bi-sphere problem (but a different, non-linear, Pareto set), and has one nonseparable ill-conditioned objective. The sphere-Rosenbrock problem (also known as BD2s (Bosman et al., 2013)) has pair-wise dependencies and a nonlinear Pareto set. Especially the combination of an easy and a difficult objective can make it difficult to obtain an evenly spread approximation front. All three problems are scaled such that their Pareto fronts have endpoints at $(1, 0)$ and $(0, 1)$. The sphere-rotatedElli problem is initialized on the domain $[-100, -50]^n$, the sphere-Rosenbrock function on $[-5, 5]^n$.

5.2.1 Results: Linkage Tree versus Random Linkage

We compared the linkage tree model, where linkage subsets are formed between neighboring solutions (see Section 3.2.1), with a model in which solutions are merged randomly. Results were compared for 30 runs on sphere-Rosenbrock, with a population size of $N = 200$, $n = 10$ and $p = 9$. The mean number of function evaluations that UHV-GOMEA-Lt required to obtain $\Delta\text{HV}_p < 10^{-10}$ was 1.28×10^6 (range: $1.02\text{--}1.66 \times 10^6$). When we use the same UPGMA method of constructing the linkage tree, but with a random distance between linkage models, the required number of function evaluation increases significantly to 1.59×10^6 (range: $1.35\text{--}1.95 \times 10^6$), as tested with a Wilcoxon rank-sum test with $\alpha = 0.05$. This shows that the intuitively chosen construction approach for the linkage tree has added value, although better linkage models might of course still exist.

5.2.2 Results: Population Size N

The effect of the population size parameter N on the required number of MO-fevals to reach a target accuracy of $\Delta\text{HV}_p < 10^{-5}$, and the corresponding success rate (SR) is shown in Figure 4. The computational budget was set to 10^7 MO-fevals. Performance is rather predictable for $N \geq 200$, but for small population sizes, differences occur.

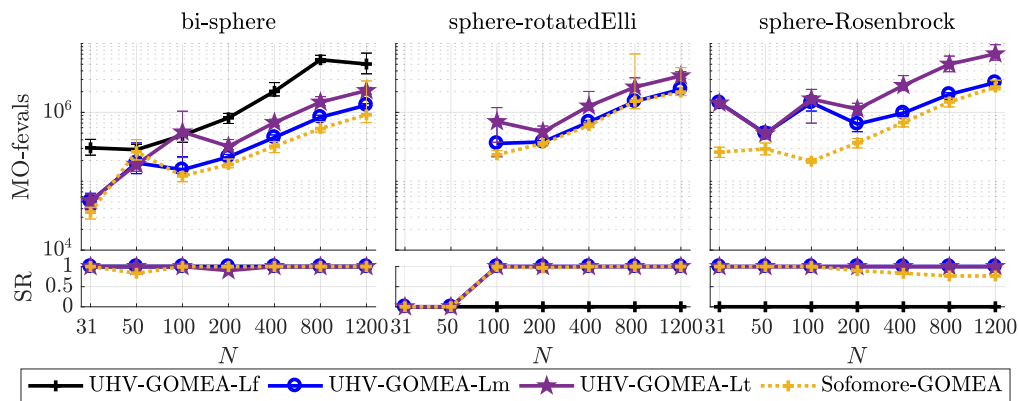


Figure 4: Effect of the population size parameter N on the hypervolume-based algorithms for various problems with $n = 10$ and $p = 9$. Success rate (SR) measures the fraction of runs in which the target accuracy ($\Delta HV_p < 10^{-5}$) was reached, out of 30 runs in total. The top row shows the number of MO-fevals required to reach this accuracy, for all successful runs.

UHV-GOMEA with Lt and Lm perform similar initially, as larger linkage subsets are filtered out. When increasing the population size, UHV-GOMEA-Lt performs slightly worse, although the overhead seems to be a constant factor. This suggests that it is neither beneficial nor harmful to model dependencies in this scenario. UHV-GOMEA-Lf is not able to solve the sphere-rotatedElli, and is significantly slower for the bi-sphere problem. We therefore omit it from further experiments.

Sofomore-GOMEA performs similar to UHV-GOMEA-Lm for larger population sizes, but for smaller population sizes, it clearly outperforms all UHV-GOMEA variants on sphere-Rosenbrock. A relatively large population size is required to estimate a full covariance matrix, which is required to solve sphere-rotatedElli. For smaller population sizes, a diagonal variance matrix or regularized covariance matrix is estimated, which affects performance, making it hard to predict whether it is beneficial to increase or decrease the population size. Interestingly, for large N , the success rate of Sofomore-GOMEA deteriorates on sphere-Rosenbrock, as it converges to a locally optimal distribution of solutions along the front, which can only be escaped if multiple solutions were to be updated simultaneously. This only happens for sphere-Rosenbrock due to the shape of its Pareto front, but also since the population is initialized close to the Pareto set. Since f_1 is easier to solve than f_2 , solutions initially quickly converge toward one end the front, and the larger the population size, the faster this happens, resulting in little to no opportunity for each of the optimizers to adapt for the dynamic nature of its objective function.

5.2.3 Results: Solution Set Size p

When the size p of the solution set \mathcal{S}_p is set larger, resulting MO-solutions will be closer to each other on the front. It can thus be expected that for larger p , dependency modeling becomes more relevant. We let $p = 2^j + 1$ with $j \in \mathbb{N}$, and inspect how many runs obtain a high accuracy of $\Delta HV_p < 10^{-10}$. For this, we let each algorithm run with a large budget of 10^8 MO-fevals, or until it converged, that is, when the standard deviation

Table 1: Success rate (SR) to obtain $\Delta HV_p < 10^{-10}$, together with the MO-fevals per p of successful runs (\pm standard deviation) for the sphere-rotatedElli benchmark problem ($n = 3$), for $N = 50$ and $N = 100$. Bold scores are best obtained scores or those not statistically different from it.

	p	UHV-GOMEA-Lm		UHV-GOMEA-Lt		Sofomore-GOMEA	
		SR	MO-fevals/ p	SR	MO-fevals/ p	SR	MO-fevals/ p
$N = 50$	3	1.00	3.3e+03±1.8e+02	1.00	5.2e+03±7.2e+02	1.00	2.9e+03 ±3.2e+02
	5	1.00	4.3e+03±2.1e+02	1.00	9.1e+03±1.7e+03	1.00	3.6e+03 ±4.1e+02
	9	1.00	7.2e+03 ±5.8e+02	1.00	1.6e+04±2.1e+03	1.00	7.5e+03±3.0e+02
	17	0.73	2.4e+04 ±7.5e+03	0.90	8.2e+04±1.9e+04	1.00	2.2e+04 ±7.0e+02
	33	0.00	-	0.30	3.2e+05 ±1.6e+05	0.00	-
	65	0.00	-	0.00	-	0.00	-
$N = 100$	3	1.00	6.5e+03±2.4e+02	1.00	8.5e+03±5.9e+02	1.00	5.3e+03 ±4.8e+02
	5	1.00	8.3e+03±2.2e+02	1.00	1.3e+04±9.9e+02	1.00	6.1e+03 ±4.2e+02
	9	1.00	1.3e+04 ±6.4e+02	1.00	2.4e+04±2.8e+03	1.00	1.4e+04±8.5e+02
	17	0.63	3.5e+04 ±7.3e+03	1.00	4.0e+04±3.5e+03	1.00	4.1e+04±1.7e+02
	33	0.00	-	0.97	1.5e+05 ±3.9e+04	0.23	1.4e+05 ±4.0e+03
	65	0.00	-	0.97	8.3e+05 ±2.4e+05	0.00	-

of the objective values in the population is less than 10^{-20} . At this point, the machine accuracy becomes an issue, and no further improvements can be obtained. We use the sphere-rotatedElli problem with $n = 3$, as it has dependencies, with the same Pareto front as the bi-sphere problem, so that we can find the empirical $\Delta HV_p(A_p^*)$ by solving this simpler problem with accuracy of 10^{-15} . Success rates and number of MO-fevals to reach the target accuracy are shown in Table 1. Differences were tested for statistical significance using the Wilcoxon rank-sum test at $\alpha = 0.05$.

For $N = 50$ and $p \leq 9$, all algorithms obtain the target accuracy in all runs, but for $p \geq 17$, performance deteriorates for UHV-GOMEA-Lm and Sofomore-GOMEA, while UHV-GOMEA-Lt can still solve the problem. Doubling the population size roughly doubles the required number of MO-fevals for all algorithms for $p \leq 9$. For UHV-GOMEA-Lt, increasing the population size allows it to solve problems with large p , clearly outperforming the other algorithms.

The number of function evaluations GOMEA requires to solve grey-box problems with independent decision variables scales logarithmically in the problem dimensionality (Bouter, Alderliesten et al., 2017), which is for the IBMOP determined by p and n . Here, we see that the required number of MO-fevals grows super-linear in p . This indicates that the strength of the dependencies indeed increases with p , up to the point that UHV-GOMEA-Lm and Sofomore-GOMEA can no longer solve the problem, making dependency modeling essential.

5.3 Experiment 3: Elitist Archive

In Experiment 1, the limited-size S_p obtained by the different algorithms was the basis of comparison together with a notion of proximity of each MO-solution to the Pareto set, which was favorable for the hypervolume-based algorithms. We now compare the elitist archives \mathcal{E} . All algorithms have the same archiving strategy, but only MO-GOMEA uses




Figure 5: Inverse generational distance (IGD) and number of MO-fevals to reach $IGD < 10^{-3}$ in case of success.

the archive as part of the optimization process. As performance indicator, we use the inverse generational distance (IGD) (Bosman and Thierens, 2002),

$$IGD(\mathcal{E}, \mathcal{A}^*) = \frac{1}{|\mathcal{A}^*|} \sum_{y \in \mathcal{A}^*} \min_{x \in \mathcal{A}} \|f(x) - f(y)\|,$$

where we use a finite approximation of the Pareto set $\mathcal{A}^* \approx \mathcal{A}_{5000}^*$ of 5000 MO-solutions. The IGD measures both proximity and diversity, in contrast to the GD, which only measures proximity. We run each algorithm until a target accuracy of $IGD < 10^{-3}$ is obtained, or when it converged before.

Results

Results are shown in Figure 5. Independent of the choice of p , MO-GOMEA was able to obtain the target IGD. Performance seems rather robust against the choice of p , which, in our experiment setup controls its population size $N_{mo} = p \cdot N$ and number of clusters $K_{mo} = 2p$. Furthermore, we see that MO-GOMEA outperforms the other algorithms in all cases by obtaining a better IGD with fewer MO-fevals, except for some runs of the sphere-Rosenbrock problem with $p = 17$ and $p = 33$. In terms of p , $\log(IGD)$ decreases linearly for the hypervolume-based algorithms, and for $p \geq 33$, all algorithms obtain the target IGD, for all three problems.

5.4 Experiment 4: Multimodal MO Problems

As the Sofomore framework performs a form of local search around a single approximation set, it is as expected that it performs well on unimodal functions. Real-world problems are often not that well-behaved. Therefore, we now include two multimodal problems from the ZDT problem set (Deb, Thiele et al., 2002). ZDT3 is characterized by a discontinuous Pareto front. ZDT6 has a concave front, and six local optima in f_1 , resulting in a Pareto set consisting of six subsets that partially overlap. The decision space

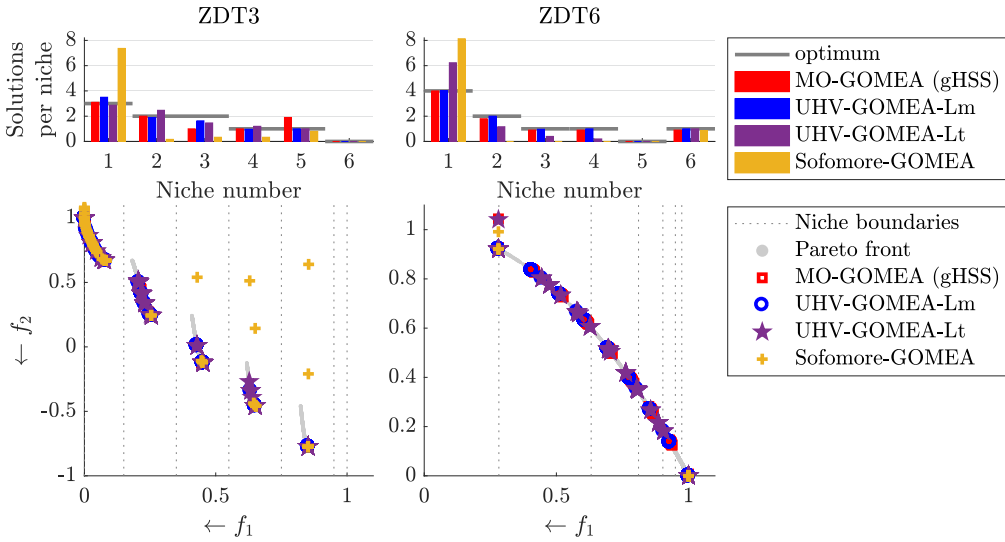


Figure 6: Distribution of solutions over different niches in the final solution sets \mathcal{S}_p for the multimodal functions ZDT3 and ZDT6 ($n = 10$, $p = 9$, $N = 200$). Niche numbers correspond to the objective space regions separated by the dotted lines.

of the ZDT problems is bounded to $[0, 1]^n$. We here use re-sampling as repair mechanism. We furthermore set $p = 9$ and therefore use $N = 200$, as this was found to be a good choice for all algorithms based on the results in Section 5.2.2.

Results

Figure 6 shows the distribution of MO-solutions per niche. UHV-GOMEA-Lm obtains the optimal hypervolume for all runs on ZDT6 and for 60% of the runs on ZDT3. MO-GOMEA obtains an elitist archive with MO-solutions in all niches, but gHSS does not result in the optimal distribution over niches. Sofomore-GOMEA obtains many more MO-solutions in the leftmost niche (Niche 1). It quickly converges to the extremes of the Pareto front in the beginning of the optimization, but is not able to move MO-solutions along the front. For ZDT3, a good spread of MO-solutions within the leftmost niche is obtained, but it is not able to move MO-solutions out of that local optimum to the other subsets of the Pareto front in most runs. When it did not have MO-solutions in one of the subsets of the front, it sometimes ended up with non-optimal MO-solutions within the niches that were obtained in an attempt to fill the resulting gaps in the front. For ZDT6, Sofomore-GOMEA shows the same behavior, but it is not able to move MO-solutions along the front due to the concavity of the front, and only obtains MO-solutions at the extremes of the Pareto front. UHV-GOMEA-Lm obtains $\Delta HV_p < 10^{-5}$ in 60% of the runs for ZDT3 and 100% for ZDT6, while UHV-GOMEA-Lt obtains this accuracy 26% of the runs for ZDT3, and 0% for ZDT6. The other methods were not able to reach this accuracy in any of the runs.

5.5 Experiment 5: A Hybrid Approach

As seen in Figure 3, MO-GOMEA performs best initially, but stagnates in terms of proximity (e.g., GD) when the majority of MO-solutions in the population is nondominated. We construct a simple hybrid approach where we initially run MO-GOMEA, which

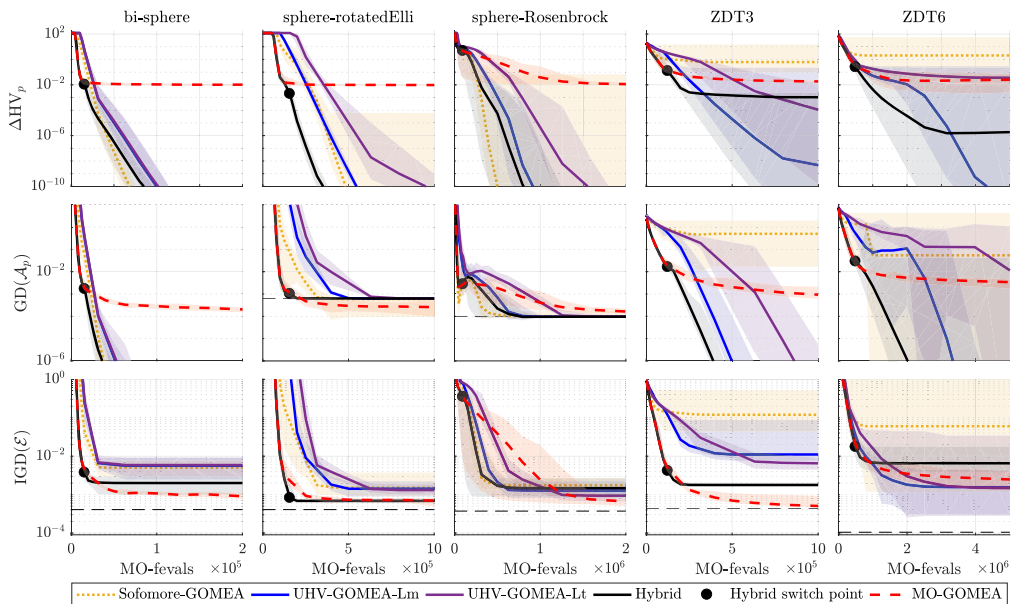


Figure 7: Convergence plots for the algorithms discussed in this work. All problems are $n = 10$, $p = 9$, and run with $N = 200$, except the bi-sphere problem, which is run with $N = 31$. Lines show mean scores; shaded area are min/max scores. Hybrid switch point shows the mean MO-fevals after which MO-GOMEA was terminated and UHV-GOMEA-Lm initialized in the hybrid approach. Dashed-black lines shows the maximally achievable GD scores for problems where a finite approximation of the Pareto front is used, and the maximally achievable IGD with a limited archive size $|\mathcal{E}| \leq 1000$ for the IGD.

we terminate when it is expected to stagnate, that is, when 90% of the MO-solutions in the population are nondominated, or when the elitist archive target size is hit. We then switch to UHV-GOMEA-Lm starting from the elitist archive \mathcal{E} that MO-GOMEA obtained so far. \mathcal{E} is clustered into p clusters of equal size $2|\mathcal{E}|/p$ using the same clustering method that is used in MO-GOMEA. For this, the cluster means are initialized with gHSS, and distances are measured in decision space. If the cluster size is smaller than the desired population size, that is, $2|\mathcal{E}|/p < N$, the remainder of the MO-solutions is sampled uniformly.

Results

Convergence results for the discussed algorithms are shown in Figure 7. As intended, the hybrid approach terminates MO-GOMEA when stagnation occurs and UHV-GOMEA-Lm takes over from there (indicated by the black dot). The hybrid converges to the Pareto set in all cases, as shown by $GD(\mathcal{A}_p)$, and outperformed UHV-GOMEA-Lm, showing that a domination-based initialization is preferable over the uncrowded distance approach. In terms of the archive IGD, the hybrid approach also outperforms UHV-GOMEA-Lm for all problems but ZDT6, where none of the algorithms obtain the maximally achievable IGD. As was previously seen in Figure 5, when p is small, UHV-GOMEA obtains a lower IGD for the elitist archive \mathcal{E} than MO-GOMEA. This explains

why the hybrid method, that initially runs MO-GOMEA, but later switches to UHV-GOMEA, obtains a lower archive IGd than MO-GOMEA.

5.6 Experiment 6: WFG Benchmark

We benchmark the discussed methods on the commonly used WFG Benchmark (Huband et al., 2005). This test suite consists of 9 benchmark functions with different properties. We consider the instances with $m = 2$ objectives, $k_{\text{WFG}} = 4$ position variables, and $l_{\text{WFG}} = 20$ distances variables, resulting in a total of $n = 24$ decision variables. The hypervolume reference point is set to $r = (11, 11)$. Of these problems, WFG1 is separable, but has a flat region in the decision space, which could cause stagnation. WFG2, WFG4, and WFG9 have one or more multimodal objectives. Problems WFG4–9 all have a concave front, WFG1 has a convex front, WFG2 has a disconnected convex front, and WFG3 has a linear front. We solve these benchmark problems with $p = 9$ and a computational budget of 10^7 MO-fevals. We aim in this work to study convergence to the Pareto set, for which the stagnation of domination-based-selection-only methods occurs only later in the optimization process. This is where our method reaches levels of accuracy that are not (often) obtainable with other methods, and therefore, we chose a larger budget so this can be observed. Based on previous results, a population size of $N = 200$ was used for all algorithms. We include two versions of the hybrid method in this experiment. Both first run MO-GOMEA, and Hybrid-Lm then switches to UHV-GOMEA-Lm, while Hybrid-Lt switches to UHV-GOMEA-Lt. All experiments are repeated 30 times. Differences are tested for statistical significance (up to 4 decimals) by a Wilcoxon rank-sum test with $\alpha = 0.05$, pairwise to the best. Ranks (in brackets) are computed based on the mean hypervolume values. All statistics are computed per table.

Results

Results on the WFG benchmark are shown in Table 2. Problem WFG1, which has a plateau in its fitness landscape, is consistently solved better with MO-GOMEA. MO-GOMEA also outperforms the other methods on WFG2, which is multimodal, and has a discontinuous front, which prevents the hypervolume-based methods from obtaining the optimal distribution of solutions, as was shown before for the ZDT3 and ZDT6 problems. For the problems with a concave front, WFG4–9, the optimal hypervolume value is $HV_9^* = 114.40 \dots$, which was obtained by many of the algorithms. Especially WFG3 which has optimal hypervolume $HV_9^* = 116.50 \dots$, and WFG6 and WFG7 seem to be relatively easy, with none of the methods performing particularly worse than the others. Differences occur often at more than two decimals of accuracy, and are therefore not visible in the table.

This experiment setup is unfavorable for MO-GOMEA, as it is not aimed to obtain the optimal distribution of exactly $p = 9$ solutions along the front. Despite this bias, it is still very competitive, obtaining the best scores in three of the problems. Especially the difference between UHV-GOMEA-Lm and UHV-GOMEA-Lt is noteworthy, demonstrating that taking linkage into account greatly improves performance on these more difficult problems. Overall, Hybrid-Lt performs best, although differences between different algorithms are small.

6 Discussion

We formulated the uncrowded hypervolume (UHV) measure, which was used to achieve population-based hypervolume-driven MO optimization using a single-objective problem formulation. We compared this problem formulation with the

Table 2: Results on the WFG benchmark with $p = 9$ MO-solutions, resulting in an $n = 216$ dimensional optimization problem. A computational budget of 10^7 MO-fevals was used. Hypervolume values HV_p are shown (mean, \pm standard deviation (rank)), computed with reference point $r = [11, 11]$. gHSS was used for MO-GOMEA. Bold numbers are best scores per problems, or those not statistically different from it. Bottom row shows the mean rank, together with the overall rank in brackets.

#	Sofomore-GOMEA	UHV-GOMEA-Lm	UHV-GOMEA-Lt	MO-GOMEA	Hybrid-Lm	Hybrid-Lt
1	96.94 \pm 1.74 (5)	93.98 \pm 0.88 (6)	99.22 \pm 2.43 (2)	103.52 \pm 1.72 (1)	98.00 \pm 0.88 (3)	98.00 \pm 0.88 (4)
2	110.16 \pm 0.03 (5)	110.14 \pm 0.01 (6)	110.21 \pm 0.40 (2)	112.28 \pm 3.32 (1)	110.18 \pm 0.00 (4)	110.18 \pm 0.00 (3)
3	116.50 \pm 0.00 (3)	116.50 \pm 0.01 (5)	116.50 \pm 0.00 (1)	116.34 \pm 0.04 (6)	116.50 \pm 0.00 (4)	116.50 \pm 0.00 (2)
4	112.81 \pm 0.60 (5)	112.66 \pm 0.64 (6)	113.23 \pm 0.41 (2)	113.19 \pm 0.63 (3)	112.97 \pm 0.56 (4)	113.50 \pm 0.55 (1)
5	112.27 \pm 0.13 (1)	112.08 \pm 0.38 (6)	112.22 \pm 0.00 (2)	112.17 \pm 0.04 (4)	112.14 \pm 0.28 (5)	112.22 \pm 0.00 (3)
6	114.40 \pm 0.00 (1)	114.39 \pm 0.02 (4)	114.39 \pm 0.02 (3)	114.17 \pm 0.08 (6)	114.29 \pm 0.38 (5)	114.39 \pm 0.01 (2)
7	114.40 \pm 0.00 (3)	114.40 \pm 0.00 (2)	114.40 \pm 0.00 (5)	114.26 \pm 0.04 (6)	114.40 \pm 0.00 (4)	114.40 \pm 0.00 (1)
8	111.74 \pm 0.20 (3)	111.47 \pm 0.29 (5)	111.51 \pm 0.23 (4)	111.11 \pm 0.15 (6)	111.80 \pm 0.04 (2)	111.83 \pm 0.01 (1)
9	111.44 \pm 0.20 (6)	111.49 \pm 0.05 (4)	111.48 \pm 0.04 (5)	112.09 \pm 0.57 (1)	111.57 \pm 0.25 (2)	111.57 \pm 0.25 (3)
	3.56 (3)	4.89 (6)	2.89 (2)	3.78 (5)	3.67 (4)	2.22 (1)

dynamic interleaved Sofomore framework, which is also hypervolume-based, and the MO problem formulation based on Pareto-dominance that is typically used in MOEAs. These three problem formulations were all solved with versions of the gene-pool optimal mixing evolutionary algorithm (GOMEA) (Bouter, Alderliesten et al., 2017; Bouter, Luong, Alderliesten et al., 2017).

We showed that the hypervolume-based methods do not exhibit the stagnation (e.g., in GD) that occurs with domination-based MOEAs, and thereby confirm the results obtained in, for example, Touré et al. (2019). This clearly shows the superiority of hypervolume-based methods when a small number of high-quality solutions is required. However, domination-based MOEAs initially outperform the hypervolume-based methods, especially when the initial population is far away from the Pareto set. A simple hybrid approach, in which a reasonably good approximation set that is obtained with a domination-based MOEAs is used as the initial population of a hypervolume-based algorithm, showed to improve performance compared with the use of both approaches separately. When a large approximation set is required, the difficulty of the hypervolume-based problems increases, and dependency modeling becomes beneficial or even essential for solving them. Additionally, in our experiments, the domination-based MOEA almost always achieves a better elitist archive containing a large number of nondominated solutions (e.g., $|\mathcal{E}| = 1000$, in this work).

MO-GOMEA naturally has diversity-enhancing mechanisms, which might be beneficial for successfully optimizing multimodal problems. The single-objective GOMEA, used here to optimize the hypervolume-based problem formulation, was not particularly developed for this objective function. A better understanding of multiobjective fitness landscapes, of which a first attempt was made in Kerschke and Grimme (2017), might be helpful to adapt GOMEA, or any other single-objective optimizer, for this specific optimization task.

A limitation of the hypervolume-based approach is that a reference point is required, for which a suitable choice could be unknown in a black-box setting. Additionally, the computational complexity of the hypervolume increases when the number of objectives increases. This makes the UHV expensive for MO optimization problems with $m \geq 3$, although the population size in this application is small, and approximation methods could be used (Bader and Zitzler, 2011; Fieldsend, 2019). On the other hand, our IBMOP formulation allows MO problems to be solved with a single-objective optimizer, which provides opportunities to explore techniques such as MO multimodal optimization (Maree et al., 2019; Tanabe and Ishibuchi, 2018), that are well explored for single-objective optimization, but are still upcoming in MO optimization.

7 Conclusion

We introduced a single-objective problem formulation for multiobjective optimization based on the uncrowded hypervolume (UHV). We showed that problems formulated as such can be efficiently solved with GOMEA by exploiting grey-box properties of this problem formulation. We compared the resulting approach with a version of GOMEA that is based on a classical domination-based selection (MO-GOMEA) and a version that is based on hypervolume optimization (Sofomore-GOMEA). We showed that hypervolume-based optimization can overcome the stagnation from which domination-based methods suffer after a while, and that these methods show convergence to the optimal hypervolume, and thereby to a subset of the Pareto set. However, when the multiobjective problem at hand has difficult landscape features such as

multimodality or deceptiveness, the domination-based MO-GOMEA outperformed the hypervolume-based methods.

When the desired approximation set size is small, hypervolume-based methods are generally preferable. When the desired approximation set size is large, domination-based methods obtain a better approximation faster. Additionally, in the latter case, the resulting single-objective optimization problem becomes difficult, and dependency modeling becomes essential to still be able to solve the hypervolume-problem up to high accuracy. Hybrid methods, such as the one proposed in this article, stand the best chance at achieving the overall best performance and being most generally applicable, which also provides a promising area of future research.

Acknowledgments

This work is part of the research programme IPPSI-TA with project number 628.006.003, which is financed by the Dutch Research Council (NWO) and Elekta. We acknowledge financial support of the Nijbakker-Morra Foundation for a high-performance computing system.

References

- Auger, A., Bader, J., Brockhoff, D., and Zitzler, E. (2009a). Theory of the hypervolume indicator. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pp. 87–102.
- Auger, A., Bader, J., Brockhoff, D., and Zitzler, E. (2009b). Theory of the hypervolume indicator: Optimal μ -distributions and the choice of the reference point. In *Proceedings of the Tenth ACM SIGEVO Workshop on Foundations of Genetic Algorithms*, pp. 87–102.
- Bader, J., and Zitzler, E. (2011). Hype: An algorithm for fast hypervolume-based many-objective optimization. *Evolutionary Computation*, 19(1):45–76. 10.1162/EVCO_a_00009, PubMed: 20649424
- Benítez-Hidalgo, A., Nebro, A. J., García-Nieto, J., Oregi, I., and Ser, J. D. (2019). jMetalPy: A Python framework for multi-objective optimization with metaheuristics. *Swarm and Evolutionary Computation*, 100598.
- Berghammer, R., Friedrich, T., and Neumann, F. (2012). Convergence of set-based multi-objective optimization, indicators and deteriorative cycles. *Theoretical Computer Science*, 456:2–17. 10.1016/j.tcs.2012.05.036
- Beume, N., Fonseca, C., López-Ibáñez, M., Paquete, L., and Vahrenhold, J. (2009). On the complexity of computing the hypervolume indicator. *IEEE Transactions on Evolutionary Computation*, 13(5):1075–1082. 10.1109/TEVC.2009.2015575
- Beume, N., Naujoks, B., and Emmerich, M. (2007). SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research*, 181(3):1653–1669. 10.1016/j.ejor.2006.08.008
- Bosman, P. A. N., Grahl, J., and Thierens, D. (2013). Benchmarking parameter-free AMaLGaM on functions with and without noise. *Evolutionary Computation*, 21(3):445–469. 10.1162/EVCO_a_00094, PubMed: 23030365
- Bosman, P. A. N., and Thierens, D. (2002). Multi-objective optimization with diversity preserving mixture-based iterated density estimation evolutionary algorithms. *International Journal of Approximate Reasoning*, 31(3):259–289. 10.1016/S0888-613X(02)00090-7
- Bosman, P. A. N., and Thierens, D. (2003). The balance between proximity and diversity in multi-objective evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 7(2):174–188. 10.1109/TEVC.2003.810761

- Bouter, A., Alderliesten, T., Witteveen, C., and Bosman, P. A. N. (2017). Exploiting linkage information in real-valued optimization with the real-valued gene-pool optimal mixing evolutionary algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pp. 705–712.
- Bouter, A., Luong, N. H., Alderliesten, T., Witteveen, C., and Bosman, P. A. N. (2017). The multi-objective real-valued gene-pool optimal mixing evolutionary algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pp. 537–544.
- Bouter, A., Luong, N. H., Witteveen, C., Alderliesten, T., and Bosman, P. A. N. (2017). The multi-objective real-valued gene-pool optimal mixing evolutionary algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pp. 537–544.
- Deb, K. (2001). *Multi-objective optimization*. Chichester, UK: Wiley.
- Deb, K., and Jain, H. (2013). An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, Part I: Solving problems with box constraints. *IEEE Transactions on Evolutionary Computation*, 18(4):577–601. 10.1109/TEVC.2013.2281535
- Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197. 10.1109/4235.996017
- Deb, K., Thiele, L., and Zitzler, E. (2002). Scalable multi-objective optimization test problems. In *Proceedings of the IEEE Congress on Evolutionary Computation*, Vol. 1, pp. 825–830.
- Emmerich, M., Beume, N., and Naujoks, B. (2005). An EMO algorithm using the hypervolume measure as selection criterion. In *Conference on Evolutionary Multi-Criterion Optimization*, pp. 62–76. Lecture Notes in Computer Science, Vol. 3410. 10.1007/978-3-540-31880-4_5
- Emmerich, M., Deutz, A., and Beume, N. (2007). Gradient-based/evolutionary relay hybrid for computing Pareto front approximations maximizing the S-metric. In *International Workshop on Hybrid Metaheuristics*, pp. 140–156. Lecture Notes in Computer Science, Vol. 4771. 10.1007/978-3-540-75514-2_11
- Fieldsend, J. E. (2019). Efficient real-time hypervolume estimation with monotonically reducing error. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pp. 532–540.
- Fleischer, M. (2003). The measure of Pareto optima. Applications to multi-objective metaheuristics. In *Conference on Evolutionary Multi-Criterion Optimization*, Vol. 2632, pp. 519–533. 10.1007/3-540-36970-8_37
- Fonseca, C., Paquete, L., and López-Ibáñez, M. (2006). An improved dimension-sweep algorithm for the hypervolume indicator. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pp. 1157–1163.
- Gronau, I., and Moran, S. (2007). Optimal implementations of UPGMA and other common clustering algorithms. *Information Processing Letters*, 104(6):205–210. 10.1016/j.ipl.2007.07.002
- Guerreiro, A., Fonseca, C., and Paquete, L. (2016). Greedy hypervolume subset selection in low dimensions. *Evolutionary Computation*, 24(3):521–544. 10.1162/EVCO_a_00188, PubMed: 27303786
- Hansen, M., and Jaszkievicz, A. (1998). *Evaluating the quality of approximations of the non-dominated set*. Technical Report IMM-REP-1998-7.
- Hansen, N., and Ostermeier, A. (2001). Completely derandomized self-adaptation in evolution strategies. *IEEE Computational Intelligence Magazine*, 9(2):159–195.
- Higham, N. (2008). Cholesky factorization. *Wiley Interdisciplinary Reviews: Computational Statistics*, 1(2):251–254. 10.1002/wics.18

- Huband, S., Barone, L., While, L., and Hingston, P. (2005). A scalable multi-objective test problem toolkit. In *International Conference on Evolutionary Multi-Criterion Optimization*, pp. 280–294.
- Igel, C., Hansen, N., and Roth, S. (2007). Covariance matrix adaptation for multi-objective optimization. *Evolutionary Computation*, 15(1):1–28. 10.1162/evco.2007.15.1.1, PubMed: 17388777
- Ishibuchi, H., Masuda, H., Tanigaki, Y., and Nojima, Y. (2015). Modified distance calculation in generational distance and inverted generational distance. In *International Conference on Evolutionary Multi-Criterion Optimization*, pp. 110–125.
- Kerschke, P., and Grimme, C. (2017). An expedition to multimodal multi-objective optimization landscapes. In *9th International Conference on Evolutionary Multi-Criterion Optimization*, Vol. 10173, pp. 329–343. 10.1007/978-3-319-54157-0_23
- Knowles, J. (2002). *Local-search and hybrid evolutionary algorithms for Pareto optimization*. Technical Report, PhD thesis, University of Reading, U.K.
- Knowles, J., Thiele, L., and Zitzler, E. (2006). *A tutorial on the performance assessment of stochastic multiobjective optimization*. Technical Report. Computer Engineering and Networks Laboratory (TIK), ETH Zurich, TIK Report 214.
- Lay, D. (1993). *Linear algebra and its applications*. New York: Addison Wesley.
- Ledoit, O., and Wolf, M. (2012). Nonlinear shrinkage estimation of large-dimensional covariance matrices. *The Annals of Statistics*, 40(2):1024–1060. 10.1214/12-AOS989
- Luong, N. H., and Bosman, P. A. N. (2012). Elitist archiving for multi-objective evolutionary algorithms: To adapt or not to adapt. In *Parallel Problem Solving from Nature*, pp. 72–81.
- Maree, S. C. (2012). *Correcting non-positive definite correlation matrices*. Technical Report. TU Delft.
- Maree, S. C., Alderliesten, T., and Bosman, P. A. N. (2019). Real-valued evolutionary multi-modal multi-objective optimization by hill-valley clustering. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pp. 568–576.
- Mostaghim, S., Branke, J., and Schmeck, H. (2007). Multi-objective particle swarm optimization on computer grids. In *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation (GECCO)*, pp. 869–875.
- Nicolini, M. (2005). A two-level evolutionary approach to multi-criterion optimization of water supply systems. In *Conference on Evolutionary Multi-Criterion Optimization*, pp. 736–751. Lecture Notes in Computer Science, Vol. 3410. 10.1007/978-3-540-31880-4_51
- Tanabe, R., and Ishibuchi, H. (2018). A decomposition-based evolutionary algorithm for multimodal multi-objective optimization. In *Parallel Problem Solving from Nature*, Vol. 5199, pp. 249–261. 10.1007/978-3-319-99253-2_20
- Touré, C., Hansen, N., Auger, A., and Brockhoff, D. (2019). Uncrowded hypervolume improvement: COMO-CMA-ES and the Sofomore framework. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pp. 638–646.
- Wessing, S. (2020). Evoalgos 1.1 Python package.
- Zhang, Q., and Li, H. (2007). MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6):712–731. 10.1109/TEVC.2007.892759
- Zitzler, E., Brockhoff, D., and Thiele, L. (2007). The hypervolume indicator revisited: On the design of Pareto-compliant indicators via weighted integration. In *Conference on Evolutionary Multi-Criterion Optimization*, pp. 862–876. Lecture Notes in Computer Science, Vol. 4403. 10.1007/978-3-540-70928-2_64

- Zitzler, E., and Künzli, S. (2004). Indicator-based selection in multiobjective search. In *International Conference on Parallel Problem Solving from Nature*, pp. 832–842.
- Zitzler, E., Laumanns, M., and Thiele, L. (2001). SPEA2: Improving the strength Pareto evolutionary algorithm for multiobjective optimization. In *Evolutionary Methods for Design, Optimisation and Control with Application to Industrial Problems*, pp. 95–100.
- Zitzler, E., and Thiele, L. (1999). Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach. *IEEE Transactions on Evolutionary Computation*, 3(4):257–271. 10.1109/4235.797969
- Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C., and Da Fonseca, V. (2003). Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on Evolutionary Computation*, 7(2):117–132. 10.1109/TEVC.2003.810758