

Defining and Optimizing Indicator-based Diversity Measures in Multiobjective Search

Tamara Ulrich, Johannes Bader, and Lothar Thiele

Computer Engineering and Networks Laboratory, ETH Zurich
8092 Zurich, Switzerland

`firstname.lastname@tik.ee.ethz.ch`

Abstract. In this paper, we elaborate how decision space diversity can be integrated into indicator-based multiobjective search. We introduce DIOP, the diversity integrating multiobjective optimizer, which concurrently optimizes two set-based diversity measures, one in decision space and the other in objective space. We introduce a possibility to improve the diversity of a solution set, where the minimum proximity of these solutions to the Pareto-front is user-defined. Experiments show that DIOP is able to optimize both diversity measures and that the decision space diversity can indeed be improved if the required maximum distance of the solutions to the front is relaxed.

1 Motivation

The task of evolutionary multiobjective optimization (EMO) includes to find a set of Pareto-optimal solutions which is as diverse as possible to offer the decision maker a good selection of solutions. Traditionally, diversity relates to objective values. Only recently, multiobjective algorithms also aim at finding solutions that are diverse in the decision space. Maintaining multiple solutions which cover different parts of the decision space, e.g. different designs, offers many advantages: first, it enables the decision maker to choose among different designs with the same or at least equally preferable objective values; second, it helps the decision maker to gather information about the problem structure; and third, it can speed up search—for instance by improving exploration and preventing premature convergence.

Many algorithms have been proposed to promote diversity of solutions also in the decision space. However, the exact optimization goal is often far from clear. The Omni-Optimizer [4] for example is based on a crowding distance, which prefers solutions with large distance to the remaining solutions and alternates between the distance in the objective space and in the decision space. In this setting, the optimal set of solutions is not well-defined, nor is it easily possible to specify the desired tradeoff between diversity in the objective space and diversity in the decision space.

We here make the following assumptions about the preference of a decision maker:

1. The decision maker is interested in a set of solutions.
2. Each solution in this so-called *target population* should be close to optimal, i.e., not “far” from the Pareto-front in objective space.

3. The target population should cover large parts of the Pareto-front or regions nearby and should therefore offer objective space diversity.
4. The target population should cover large parts of the decision space, i.e. offering decision space diversity.

Diversity is inherently a property of sets of solutions rather than single solutions as individual solutions can only be diverse with respect to others. Therefore, optimizing diversity is closely linked to the set-based view on multiobjective optimization as proposed in SPAM [21] for example. The advantages of formalizing the optimization goal by a set-based preference relation are twofold: A preference relation defines which of two sets is preferred, and therefore, the optimization goal of the multiobjective search is clearly defined. In addition, the convergence of algorithms using this preference relation can be proven under certain conditions.

This study makes the following contributions to optimization considering diversity:

- A new diversity measure of sets in decision space is proposed which has not been used in evolutionary multiobjective optimization before. This measure imposes less strict requirements on the decision space properties than other commonly used diversity measures. An efficient procedure is presented to use this set diversity as a selection criterion for solutions during the optimization process.
- We introduce the possibility of predefining a maximal distance to the Pareto-front, that must not be exceeded by any solution. This mechanism enables a decision maker to explore the tradeoff between diversity in decision space and solution optimality.
- We provide experimental results which compare the proposed method to the well-established Omni-Optimizer [4] and which show the influence of the different parameters involved.

2 Background and Notation

Consider a multiobjective optimization problem with a decision space X and an objective space $Z \subseteq \mathbb{R}^n = \{f(a) \mid a \in X\}$, where $f : X \rightarrow Z$ denotes a mapping from the decision space to the objective space with n objective functions $f = \{f_1, \dots, f_n\}$ which are to be minimized. An element $a \in X$ of the decision space is also named a solution.

The underlying preference relation is weak Pareto dominance, where a solution $a \in X$ weakly dominates another solution $b \in X$, denoted $a \preceq b$, if and only if solution a is better or equal than b in all objectives, i.e., $a \preceq b$ iff $f(a) \leq f(b)$ or equivalently, iff $f_i(a) \leq f_i(b), \forall i \in \{1, \dots, n\}$. Furthermore, we will use the notion of weak ε -Pareto-dominance defined as $a \preceq_\varepsilon b$ iff $f(a) - \varepsilon \leq f(b)$. In other words, suppose that we improve solution a in any objective by ε . Then $a \preceq_\varepsilon b$ iff the improved solution weakly dominates solution b .

Let $X^* \subseteq X$ denote the Pareto-optimal set, $X^* = \{x \mid \nexists a \in X : a \preceq x \wedge x \not\preceq a\}$, let $T \subset X$ denote a target population of solutions, and let $q_{X^*} : X \rightarrow \mathbb{R}^{\geq 0}$ measure for each solution $x \in X$ the distance $q_{X^*}(x)$ to the Pareto-optimal set X^* . Let $D_o(T) : 2^X \rightarrow \mathbb{R}^{\geq 0}$ and $D_d(T) : 2^X \rightarrow \mathbb{R}^{\geq 0}$ measure the diversity of a set of solutions $T \subseteq X$ in the objective space ($D_o(T)$) and in the decision space ($D_d(T)$), respectively. Given

this notation, the four optimization assumptions provided in Section 1 can be formalized as follows:

1. We are interested in a target population of solutions $T \subseteq X$, $|T| = \mu$, where μ denotes its size.
2. Optimality: $\forall t \in T : q_{X^*}(t) \leq \varepsilon$, where ε is given bound on the optimality of solutions in T .
3. Diversity in objective space: Determine T such that $D_o(T)$ is maximal among all possible target populations.
4. Diversity in decision space: Determine T such that $D_d(T)$ is maximal.

As a consequence, we are dealing with a *bi-objective optimization problem on sets of solutions*. Given this setting, different problems arise:

- In order to determine $q_{X^*}(T)$ one needs the knowledge of the Pareto-optimal set of solutions X^* , which in general is not known.
- The problem of optimizing diversity in objective and decision space is a bi-objective problem on the set of all possible populations. It is not clear which tradeoff the decision maker is interested in and how to express these tradeoffs in an optimization method.
- There are many choices for the distance and diversity measures q_{X^*} , D_o and D_d . Guidelines are necessary to choose appropriate measures (see the following Sec. 3).

3 Measuring Diversity—Approaches in Biology and in EAs

Typically, measures for the diversity of a set are based on the definition of a pairwise distance between any two elements. Therefore, we assume that we are given a distance measure $d : X^2 \rightarrow \mathbb{R}^{\geq 0}$ on the decision space. Here, we are often confronted with many different classes of decision spaces, such as vectors, graphs, trees or even programs. In order to be applicable to a large class of optimization domains, we would like to place as few restrictions on the structure of the decision space as possible, i.e. we do not require that X is an Euclidean space or that the triangle inequality is satisfied. Instead, we just assume X to be a semimetric space, i.e., $\forall a, b \in X : d(a, b) \geq 0$ (*non-negativity*), $d(a, b) = d(b, a)$ (*symmetry*), $d(a, a) = 0$ (*identity of indiscernibles*). Given such a distance measure, we now would like to define a set diversity measure $D : 2^X \rightarrow \mathbb{R}^{\geq 0}$ which assigns to each subset of the decision space a real value, i.e. its diversity.

There are many possible interpretations and concepts of set diversity, i.e. how a given number of solutions should be distributed such that they achieve an optimal set diversity. In order to get a first insight, let us consider a very simple example. Figure 1 shows the optimized distribution of 100 points in a two dimensional Euclidean space $X = [0, 1]^2$ for two diversity measures, namely the sum of all pairwise distances and the Solow-Polasky [12] measure. While the Solow-Polasky measure gives a grid-like structure, the sum of pairwise distance measure distributes all 100 solutions into the four corners. As a result, it appears that we need to define a set of formal requirements for a useful diversity measure.




Fig. 1. Best distributions found by the hill-climber, for the sum of pairwise distances diversity measure (left) and the Solow-Polasky measure (right).

Measuring diversity of sets is much-discussed in biology, more specifically in the field of biodiversity. Just as for the decision maker’s preference, no generally agreed-on definition exists neither in biology nor in the field of evolutionary algorithms. In the following, we discuss the most prominent classes of existing biodiversity measures with respect to their applicability to EAs. In particular we consider the following three requirements to a diversity measure D , first proposed by Solow and Polasky [12]:

- P1: Monotonicity in Varieties** The diversity of a set of solutions A should increase when adding an individual b not yet in A , i.e., $D(A \cup b) > D(A)$ if $\min_{a \in A} d(a, b) > 0$. This fundamental property assures that increased species richness is reflected by the diversity measure [6].
- P2: Twinning** Diversity should stay constant when adding an individual c already in A , i.e., $D(A \cup c) = D(A)$. Intuitively, if diversity is understood as the coverage of a space by a set of solutions [17], adding duplicates should not increase the coverage and the chosen diversity measure should reflect that property.
- P3: Monotonicity in Distance** The diversity of set A should not decrease if all pairs of solutions are at least as dissimilar (measured by d) as before $D(A') \geq D(A)$, iff $d(a'_i, a'_j) \geq d(a_i, a_j), \forall a_i, a_j \in A, a'_i, a'_j \in A'$. So the more dissimilar solutions are, the better.

One straightforward way of measuring diversity is based on the *relative abundance* of each solution present in set A , e.g. [5]. But the degree of dissimilarity between individuals has no influence and the twinning property is not fulfilled. The second group of diversity measures is based on *taxonomy*, e.g. [19], but unfortunately building the taxonomic tree has a runtime which is exponential in the number of individuals. A very simple way of *aggregating the dissimilarity* information into a diversity measure is to sum up the values, $D(A) = \sum_{a \in A} \sum_{b \in B} d(a, b)$ [6]. Shir et al. for instance used this measure in their EA [11], while the Omni-Optimizer considers the distance d to the closest neighbors of a solution. However, these measures do not meet the twinning requirement and they promote having only two solutions with large distance duplicated multiple times. A completely new approach has been presented by Solow and Polasky [12]. Their measure is based on an *utilitarian view on individuals*, where the function $u : X \rightarrow \mathbb{R}^{\geq 0}$ defines the utility of any subset of solutions. This view of utility is equivalent to the method proposed in a previous study of the authors [17], where instead of utility the area covered by individuals has been considered. All three above requirements are fulfilled.

In the *evolutionary algorithm literature*, decision space diversity has often been used to prevent premature convergence. Examples of measures can be found in [16], [10], [13], [18, 17], [4], [14], [20], [7] or [8, 4]. Most of these measures either require a

Table 1. Comparison of different diversity metrics with respect to the three properties: monotonicity in varieties (P1), twinning (P2), and monotonicity in distance (P3)

class	method	P1	P2	P3
relative abundance	Simpson, Shannon, Berger-Parker	no	no	yes
taxonomy	clustering Weitzman	no yes	yes yes	no no
functions of distance	sum crowding distance	yes no	no no	yes yes
utilitarian	Solow-Polasky	yes	yes	yes

specific structure of the decision space, they do not define a measure on sets, they make assumptions about the Pareto-front or the problem landscape or they do not satisfy the required properties.

Table 1 summarizes the different diversity measures in context of the three requirements $P1$, $P2$ and $P3$. As can be seen, only the measure by Solow-Polasky satisfies all three requirements, so we will apply this measure in the experimental study (Sec. 5). However, the algorithmic framework presented in this paper is also compatible with other measures.

4 Optimizing Diversity—A Novel Set-Based Algorithm

Now that we have presented some possibilities to measure diversity, be it to determine $D_o(A)$ or $D_d(A)$, the decision maker’s preference 3 and 4 stated in Sec. 1 can be formally expressed. Optimizing those indicator-based set preferences can be accomplished within the SPAM framework [22]. There remain, however, a number of issues to be resolved which we are going to tackle with DIOP (Diversity Integrating Optimizer).

As the *Pareto-optimal set X^* in general is unknown*, we propose using a helper set, called the *archive A* , which approximates X^* . We therefore have two concurrent EAs, one which optimizes the target population and one which optimizes the archive population. This offers the advantage that the quality constraint (decision maker preference 2, Sec. 1) continuously tightens as the archive population improves. In order to benefit from one another, the two sets can exchange solutions, therefore improving the diversity in the archive and producing more solutions that satisfy the quality constraint in the target. This is useful as experiments have indicated that considering diverse solutions might speed up search for some problems [17].

Having an approximation A of the Pareto-optimal set X^* , a *distance metric q_A* has to be defined. We here propose to use \preceq_ε to define the distance as the smallest ε to reach ε -dominance of any solution in A , i.e.,

$$q_A(x) := \min\{\varepsilon \mid \exists y \in A : x \preceq_\varepsilon y\} . \quad (1)$$

As the decision maker is only interested in solutions not exceeding a predefined distance ε to the Pareto-front, the diversity measures of an arbitrary set P is only calculated for those solutions $P^\varepsilon \subseteq P$ not exceeding the distance ε from the front approximation A , $P^\varepsilon = \{p \in P \mid q_A(p) \leq \varepsilon\}$.

Algorithm 1 DIOP algorithm. Takes a parameter ε , an archive size μ^a , a target size μ^t , and a decision space X . Returns the optimized target set.

```

function DIOP( $\varepsilon, \mu^a, \mu^t$ )
   $A = \{x_1, \dots, x_{\mu^a}\}, x_i \in X$  /* randomly initialize archive */
   $T = \{x_1, \dots, x_{\mu^t}\}, x_i \in X$  /* randomly initialize target */
  while stopping criterion not met do
     $A' = \text{variate}(A \cup T)$  /* generate archive offspring */
     $A'' = \text{archiveSelect}(A \cup A' \cup T, \mu^a)$  /* select  $\mu^a$  new individuals */
    /* Only use new archive if its  $D_o$  value is better */
    if  $D_o(A'') > D_o(A)$  then
       $A = A''$ 
    end if
     $T' = \text{variate}(A \cup T)$  /* generate target offspring */
     $T'' = \text{targetSelect}(A, T \cup T' \cup A, \mu^t, \varepsilon)$  /* select  $\mu^t$  new individuals */
    /* Only use new target if its  $D_d$  value is better */
    if  $G(T'') > G(T)$  then
       $T = T''$ 
    end if
  end while
  Return  $T$ 
end function

```

In contrast to the framework of SPAM [22], DIOP needs to *maximize two indicators* D_o and D_d instead of one. Therefore, two sets can no longer be unambiguously compared in general, as we are dealing with a biobjective problem. Many other studies have implicitly tackled this tradeoff, however, to the best of the authors' knowledge, none of these approaches explicitly set the tradeoff, but use subpopulations [9, 16], adapt mutation [18], use the diversity to the best single objective solution [10], use the contribution to the set diversity [13], use nondominated sorting [14], use a sequence of indicators [22], alternate between decision space and objective space diversity [4], use an unweighted sum of both measures [11], use diversity as an additional objective [15], adapt the variation process [20] or integrate diversity into the hypervolume indicator [17].

In this study, we propose to consider a weighted sum of the two diversity indicators:

$$G(T) := w_o \cdot D_o(T) + w_d \cdot D_d(T), |T| = \mu \text{ with } q_A(t) \leq \varepsilon \forall t \in T, w_o + w_d = 1 \quad (2)$$

This enables a flexible tradeoff between the two diversity indicator values D_o and D_d by using different weights.

The DIOP algorithm simultaneously evolves two population, namely the archive A which approximates X^* , and the target population T which maximizes $G(T)$. Offspring is always generated from the union of both sets, whereas the selection procedure uses different indicators for the archive and target. In each generation, a selected subset is only accepted if the corresponding indicator value is larger than the one of the parent population. The pseudocode of the proposed algorithm is shown in Algorithm 1.

The function $A'' = \text{archiveSelect}(A, \mu^a)$, selects μ^a solutions A'' from a set A . The selection goal is to maximize $D_o(A'')$. The function $P' = \text{variate}(P, m)$ generates m offspring P' from a given set P . The method $T' = \text{targetSelect}(A, T, \mu^t, \varepsilon)$ selects μ^t solutions T' from set T . The goal here is to maximize $G(\{t \in T : q_A(t) \leq \varepsilon\})$.

5 Experimental Results

In this section, two main questions are investigated: first, how do the parameters of DIOP, i.e. ε and w_o , influence the obtained target population in terms of the two diversity measures D_d and D_o ? Second, we compare DIOP on two test problems to the Omni-Optimizer [4] to assess its performance.

Experimental Setup: The method $variate(P, m)$ selects $m/2$ random pairs of solutions from P to generate the offspring population. These pairs are then recombined by the SBX crossover operator [2] and mutated by adding a new normally distributed value with standard deviation $1/\eta_m$. Solow-Polasky with $\eta_{SP} = 10$ is used to measure the decision space diversity D_d . To determine the objective space diversity D_o , the hypervolume indicator is used with iterative greedy environmental selection as described in [22]. To perform the target selection $targetSelect(T, n)$ according to $G(T)$ (Eq. 2), the following wide-spread greedy strategy is used: Starting with an empty set $T' = \{\}$, iteratively the solution $t_i \in T$ is added to T' which leads to the largest indicator increase $\Delta_{t_i}G(T') := w_o(D_o(T' \cup t_i) - D_o(T')) + w_d(D_d(T' \cup t_i) - D_d(T'))$. Since determining the diversity measure of Solow-Polasky is costly (involving matrix inverses [12]), we use the following approximation: $D_d(T' \cup t_i) - D_d(T') \approx \min_{a \in T' \setminus t_i} d(t_i, a)$, i.e., we take the utility lost with respect to the closest individual as the overall utility loss.

Influence of ε and w_o : To assess the influence of the parameters ε and w_o , DIOP is run on DTLZ2 [3] with 3 objectives and $d = 7$ decision variables. DTLZ2 was chosen as it is a well-known problem, its results are easy to interpret as the connection between decision space values and objective space values is known, and the true Pareto-front is known. Note though that DIOP can also be run on real-world problems with more complex decision spaces that are not metric. The variation parameters are set according to [3] with a crossover probability of 1 with $\eta_c = 15$ and a variable exchange probability of 0.5, as well as a mutation probability of $1/d$ with $\eta_m = 20$. We chose the archive and target size to be 50 and run the algorithm for 1000 generations. The parameter ε takes the values $\{0, 0.0865, 1\}$, the weights $w_o = \{0, 0.7692, 0.9091, 0.9677, 1\}$ are logarithmically spaced with $w_d = 1 - w_o$. The results are shown in Fig. 2 on the left hand side.

It can be seen that with an increasing ε and an increasing w_d value, the achievable diversity increases, while the hypervolume decreases. This illustrates how the tradeoff between hypervolume and decision space diversity can be set by the user. Figure 3 shows the non-dominated solutions for one run with $\varepsilon = 0$ and $w_o = \{0, 0.7699, 1\}$. The higher w_o is, the more solutions lie on the Pareto-front (50/8/1 out of 50 solutions lie on the front for $w_o = 1/0.7699/0$, respectively). The dominated solutions do not contribute to the hypervolume and are distributed within the quality constraint set by ε and A in such a way that they optimize diversity. This indicates how the tradeoff is set in practice: A subset of the final target population is distributed on the Pareto-front and optimizes the hypervolume, whereas the remaining solutions optimize decision space diversity. The number of solutions that optimize the hypervolume increases with w_o .

Comparison to the Omni-Optimizer While the Omni-optimizer uses the same SBX crossover operator as DIOP, it uses an adaption of polynomial mutation with $\eta_m = 20$ [4] instead of the Gaussian mutation employed by DIOP.




Fig. 2. Left: Influence of the ϵ values on the diversity according to Solow-Polasky (x-axis) and the hypervolume (y-axis) for $\epsilon \in \{0, 0.86, 1\}$. For each ϵ , the weight w_o is decreased from 1 (leftmost data points) to 0 (rightmost data points). Right: Decision space diversity (y-axis) of all solutions within a certain distance (x-axis) of the true Pareto-front.

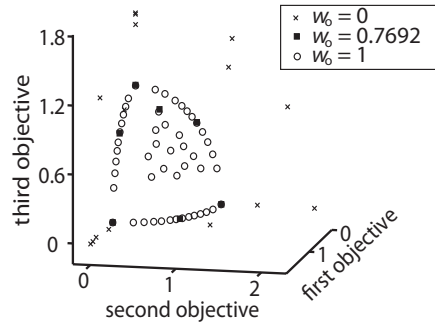


Fig. 3. Non-dominated solutions of one DTLZ2 run for three different weights.

As the first test problem we use the Omni-Test as described in [11] with 5 decision variables. The Omni-Test was chosen because it allows for an additional intuitive problem-specific diversity measure, which exploits the fact that the Pareto-optimal solutions are distributed over a total of 3^d clusters, where d is the number of decision variables. Therefore, the additional diversity measure can be defined as the number of clusters found by the algorithm. For optimization, we use the parameters from [11] with a population size of 50, 1000 generations, and $\epsilon = 0$. As we use an archive of size 50 in addition to the target of size 50, we require more fitness evaluations than the Omni-Optimizer. In order to compensate for that, the Omni-Optimizer is run for twice as many generations, i.e. 2000. For the variation operators, we use the parameters from [4] with a crossover probability of 0.9 with $\eta_c = 1$, where the variables are exchanged with a probability of 0.5, and a mutation probability of $1/n$ with $\eta_m = 1$. Each algorithm was run 15 times with different random seeds. To test the two algorithms for statistically significant differences, the Kruskal-Wallis with post-hoc Conover-Inman procedure [1] is applied with a significance level of 5%. The results are given in Table 2. It can be seen that DIOP achieves significantly better hypervolume values than the Omni-optimizer for all weight combinations, even if only decision space diversity is optimized. This is due to the fact that the solutions, while optimizing decision space diversity, must not be

	Hypervolume	Diversity (Pairs)	Diversity (Solow)	Found Clusters
DIOP $w_o = 0.00$	$30.04 \pm 0.10^+$	$0.63 \pm 0.05^-$	$46.7 \pm 2.0^+$	$33.2 \pm 5.2^+$
DIOP $w_o = 0.77$	$30.21 \pm 0.03^+$	$0.64 \pm 0.05^-$	$47.9 \pm 2.3^+$	$37.3 \pm 5.2^+$
DIOP $w_o = 0.91$	$30.25 \pm 0.03^+$	$0.66 \pm 0.06^-$	$48.7 \pm 0.9^+$	$39.9 \pm 3.7^+$
DIOP $w_o = 0.97$	$30.31 \pm 0.02^+$	$0.65 \pm 0.06^-$	$48.5 \pm 1.0^+$	$39.5 \pm 4.0^+$
DIOP $w_o = 1.00$	$30.42 \pm 0.00^+$	$0.43 \pm 0.11^-$	$16.1 \pm 2.5^-$	$9.8 \pm 2.6^-$
Omni	29.94 ± 0.05	0.70 ± 0.04	34.7 ± 1.1	23.8 ± 1.4

Table 2. Omni-Test problem: Four measures, all to be maximized. Statistically significantly better/worse results of DIOP compared to the Omni-optimizer are marked with a $^+/-$.

dominated by any archive solutions ($\varepsilon = 0$). Even though DIOP finds twice as many clusters as the Omni-Optimizer (except for $w_o = 1$, i.e. when the decision space diversity is not optimized at all), its pairwise distance measure is significantly worse than the Omni-Optimizers. This indicates that the pairwise distance measure does not accurately reflect the number of found clusters. DIOP’s Solow-Polasky values, on the other hand, are significantly better than the Omni-optimizer’s, as expected.

As a second test problem, we selected DTLZ2 with 3 objectives and 7 decision variables. In this test problem, the last 5 decision variables of a Pareto-optimal solution are equal to 0.5, whereas the first two variables define its location on the front. Solutions with values that differ from 0.5 in the last 5 variables are not Pareto-optimal. The population sizes and generation numbers are the same as for the Omni-Test problem. DIOP was run for $\varepsilon = \{0, 0.0856, 1\}$, with the weights set to $w_o = 0.9677$, $w_d = 0.0333$. The algorithms were again run 15 times with different seeds. The results are shown in Figure 2 on the right hand side. At each point in the figure, all solutions that are within the distance given on the x-axis from the true Pareto-front are used to calculate the decision space diversity D_d , which gives the corresponding y-axis value. The results show that the Omni-Optimizer has problems approximating the Pareto-front. Its diversity remains close to zero until about a distance of 0.5 from the front, which is due to the fact that it finds only few solutions that are closer than 0.5 to the true front. The DIOP population reaches its maximum diversity at a distance of around 0.2 from the front, which is an effect of the fact that not the true front but an approximation thereof is used during the optimization. For $\varepsilon = 0$ and $\varepsilon = 0.0856$, DIOPs population has a better diversity than the Omni-optimizer no matter what distance from the front is considered. For $\varepsilon = 1.0$, the solutions seem to be located in a distance interval between 0.5 and 1.2 from the front, which indicates that solutions further away from the front are more diverse than those close to the front. This matches the DTLZ2 problem; the further the last 5 decision variables are from their optimal value of 0.5, the better the diversity and the larger the distance to the front gets.

6 Conclusions

In this paper we investigate how decision space diversity can be integrated into indicator-based search. Experiments show that the algorithm can generate various tradeoffs between objective and decision space diversity, adjustable by the user. Furthermore, it is shown that the algorithm performs well when compared to the well-known Omni-Optimizer. In the future, DIOP should be tested on more complex, non-Euclidean prob-

lems. Also, it could be compared to other state-of-the-art multiobjective optimizers that do not optimize diversity, in order to quantify the increase in diversity that can be gained from using DIOP.

References

1. W. J. Conover. *Practical Nonparametric Statistics*. John Wiley, 3rd edition, 1999.
2. K. Deb and S. Agrawal. A niched-penalty approach for constraint handling in genetic algorithms. In *ICANNGA*, 1999.
3. K. Deb, L. Thiele, M. Laumanns, and E. Zitzler. Scalable Test Problems for Evolutionary Multi-Objective Optimization. TIK Report 112, ETH Zurich, 2001.
4. K. Deb and S. Tiwari. Omni-optimizer: A generic evolutionary algorithm for single and multi-objective optimization. *EJOR*, 185(3):1062–1087, 2008.
5. K. J. Gaston and J. I. Spicer. *Biodiversity: An Introduction*. Wiley-Blackwell, 2 edition, 2004.
6. J. Izsák and L. Papp. A link between ecological diversity indices and measures of biodiversity. *Ecological Modelling*, 130(1-3):151 – 156, 2000.
7. X. Li, J. Zheng, and J. Xue. A Diversity Metric for Multi-objective Evolutionary Algorithms. In *ICNC*, volume 3612 of *LNCS*, pages 68–73, 2005.
8. G. Rudolph, B. Naujoks, and M. Preuss. Capabilities of emoa to detect and preserve equivalent pareto subsets. In *EMO*, 2007.
9. J. Sarma and K. A. D. Jong. An analysis of the effects of neighborhood size and shape on local selection algorithms. In *PPSN IV*, pages 236–244, 1996.
10. H. Shimodaira. A diversity-control-oriented genetic algorithm (dcga): Performance in function optimization. In *Genetic and Evolutionary Computation Conference*, page 366, 2000.
11. O. M. Shir, M. Preuss, B. Naujoks, and M. Emmerich. Enhancing decision space diversity in evolutionary multiobjective algorithms. In *EMO*, pages 95–109, 2009.
12. A. R. Solow and S. Polasky. Measuring biological diversity. *Environmental and Ecological Statistics*, 1(2):95–103, 1994.
13. G. Squillero and A. P. Tonda. A novel methodology for diversity preservation in evolutionary algorithms. In *GECCO*, pages 2223–2226. ACM, 2008.
14. N. Srinivas and K. Deb. Multiobjective Optimization Using Nondominated Sorting in Genetic Algorithms. *Evolutionary Computation*, 2(3):221–248, 1994.
15. A. Toffolo and E. Benini. Genetic diversity as an objective in multi-objective evolutionary algorithms. *Evolutionary Computation*, 11(2):151–167, 2003.
16. S. Tsutsui, Y. Fujimoto, and A. Ghosh. Forking genetic algorithms: Gas with search space division schemes. *Evolutionary Computation*, 5(1):61–80, 1997.
17. T. Ulrich, J. Bader, and E. Zitzler. Integrating Decision Space Diversity into Hypervolume-based Multiobjective Search. In *Genetic and Evolutionary Computation Conference*, 2010. to appear.
18. R. K. Ursem. Diversity-guided evolutionary algorithms. In *Congress on Evolutionary Computation*, pages 1633–1640, 2002.
19. M. Weitzman. On diversity. *The Quarterly Journal of Economics*, 107(2):363–405, 1992.
20. A. Zhou, Q. Zhang, and Y. Jin. Approximating the set of pareto optimal solutions in both the decision and objective spaces by an estimation of distribution algorithm. *IEEE Transactions on Evolutionary Computation*, 2009. Accepted.
21. E. Zitzler, L. Thiele, and J. Bader. On Set-Based Multiobjective Optimization (Revised Version). TIK Report 300, ETH Zurich, 2008.
22. E. Zitzler, L. Thiele, and J. Bader. On Set-Based Multiobjective Optimization. *IEEE Transactions on Evolutionary Computation*, 2009. to appear.