

Study of Preference Relations in Many-Objective Optimization

Antonio López Jaimes
CINVESTAV-IPN
Computer Science Department
Av. IPN No. 2508
México, D.F. 07360
tonio.jaimes@gmail.com

Carlos A. Coello Coello^{*}
CINVESTAV-IPN
Computer Science Department
Av. IPN No. 2508
México, D.F. 07360
ccoello@cs.cinvestav.mx

ABSTRACT

This paper presents a quantitative analysis of different preference relations proposed to deal with problems with a high number of objectives. Since the relations stress different subsets of the Pareto front, we based the comparison on the Tchebycheff distance of the approximation set to the “knee” of the Pareto front. Additionally, the convergence induced by the preference relations is studied by analyzing the generational distance observed at each generation of the search. The results show that some preference relations contribute to converge quickly to the Pareto front, but they promote the generation of solutions far from the knee region. Moreover, even if a preference relation generates solutions near the knee, there exists a trade-off between convergence and the extension of the Pareto front covered.

Categories and Subject Descriptors

I.2.8 [Computing Methodologies]: Artificial Intelligence—*Problem Solving, Control Methods and Search*

General Terms

Algorithms, Performance, Experimentation

Keywords

Multiobjective optimization, many-objective optimization, preference relations

1. INTRODUCTION

Multiobjective evolutionary algorithms (MOEAs) rely on preference relations to identify high-potential regions of the search space in order to converge to the optimal set. A preference relation is the mechanism to decide if a solution \mathbf{x} is

^{*}The second author is also affiliated to the UMI-LAFMIA 3175 CNRS.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '09, July 8–12, 2009, Montréal Québec, Canada.
Copyright 2009 ACM 978-1-60558-325-9/09/07 ...\$5.00.

preferable over \mathbf{y} in the search space. In single-objective optimization, the determination of the optimum among a set of given solutions is clear. However, in the absence of preference information, in multiobjective optimization there does not exist a unique or straightforward preference relation to determine if a solution is better than other. The preference relation most commonly adopted is the one called *Pareto dominance relation* [13], which leads to trade-offs among the objectives. Thus, by using this relation, it is not possible to obtain a single optimal solution, but instead, we produce a set of them. This set is called the *Pareto optimal set* (P_{opt}) and its image in objective space is known as the *Pareto front* (PF_{opt}). Formally, we say that a solution \mathbf{x} Pareto-dominates solution \mathbf{y} , denoted by $\mathbf{x} \prec_{\text{pareto}} \mathbf{y}$, if and only if:

$$\forall i : f_i(\mathbf{x}) \leq f_i(\mathbf{y}) \text{ and } \exists i : f_i(\mathbf{x}) < f_i(\mathbf{y})$$

Nowadays, MOEAs have shown a remarkable performance in many real-life problems with 2 or 3 objectives. However, recent experimental [10, 18, 14] and analytical [11] studies have shown that MOEAs based on the Pareto dominance relation scale poorly in multiobjective optimization problems (MOPs) with a high number of objectives (4 or more).

One of the reasons for this limitation is that the proportion of nondominated solutions (i.e., incomparable solutions regarding Pareto dominance) in a population increases rapidly with the number of objectives (see e.g., [9]). As a result, in MOPs with a large number of objectives (known as many-objective optimization problems) the Pareto dominance relation is incapable of providing the necessary information to select the correct solutions in order to steer the search towards the Pareto optimal set. Although this limitation seems to affect only to Pareto-based MOEAs, many-objective problems pose some other difficulties common to any other multiobjective optimizer. For instance, the exponential growth of the number of points required to represent accurately a Pareto front with respect to the number of objectives, and the difficulty to visualize the Pareto front in more than 3 dimensions.

In the current literature, we can identify two approaches commonly adopted to cope with many-objectives problems, namely: *i*) adopt or propose a preference relation that induces a finer grain order on the solutions than that induced by the Pareto dominance relation [7, 9, 16, 15], and *ii*) reduce the number of objectives of the problem during the search process [3] or, a posteriori, during the decision making process [6, 2, 12].

Current works in many-objective optimization have analyzed some preference relations qualitatively by using distribution of solutions' ranks (i.e., the number of possible ranks and the number of solutions in each rank), or quantitatively by adopting quality indicators commonly used in the field of evolutionary multiobjective optimization (i.e., hypervolume, coverage, or generational distance). However, we believe that using directly standard quality indicators is not appropriate to compare preference relations since their optimal sets are, roughly speaking, different subsets of the Pareto optimal set. In other words, the preference relations prefer different regions of the Pareto optimal front.

In the absence of particular decision maker's preferences, the generally accepted assumption is that the most interesting solution is the "knee" of the Pareto front, i.e., the region of maximum bulge on the Pareto curve [5]. For this reason, a Pareto front approximation (PF_{approx}) generated using a preference relation should be preferred over another approximation if it generates more solutions around the knee than the other approximation. Therefore, in this paper we present a comparative study that analyzes the performance of some preference relations based on the distance of their approximation sets to the knee of the Pareto front. The goal of this study is to reveal the advantages and disadvantages of the preference relations incorporated into a MOEA.

The remainder of this paper has the following structure. The preference relations included in this study are described in Section 2. In Section 3 we describe the methodology to compare the preference relations, and the results obtained in the comparison are also presented and discussed. Finally, in Section 4 we draw some conclusions about the preference relations studied, as well as some paths for future research.

2. PREFERENCE RELATIONS STUDIED

2.1 Average and Maximum Ranking Methods

Although without a specific interest in many-objective problems, Bentley and Wakefield [1] proposed the *average ranking* (AR) and the *maximum ranking* (MR) preference relations. The AR relation computes for each solution a different rank considering each objective independently. The final rank of a solution is obtained by summing all their ranks on each objective. Let $R_{avg}(\mathbf{x}) = \sum_{1 \leq i \leq k} \text{rank}_{f_i}(\mathbf{x})$ be the average rank assigned to solution \mathbf{x} , where k is the number of objectives. Then, a solution \mathbf{x} dominates solution \mathbf{y} with respect to the average relation, denoted by $\mathbf{x} \prec_{avg} \mathbf{y}$, if and only if $R_{avg}(\mathbf{x}) < R_{avg}(\mathbf{y})$.

Table 1 illustrates the AR method with a small example considering 3-objective solutions.

Table 1: An example of the Average (AR) and Maximum (MR) preference relations.

(f_1, f_2, f_3)	rank 1	rank 2	rank 3	AR	MR
(4, 3, 5)	3	3	4	10	3
(1, 4, 7)	1	4	6	11	1
(6, 2, 2)	4	2	1	7	1
(7, 7, 6)	5	5	5	15	5
(8, 1, 3)	6	1	2	9	1
(3, 8, 4)	2	6	3	11	2

In turn, the MR relation takes the best rank as the global rank for each solution (see Table 1). Clearly this method fa-

vours extreme solutions, i.e., solutions with high performance in some of the objectives, although with poor overall performance. Let $R_{max}(\mathbf{x}) = \min_{1 \leq i \leq k} \{\text{rank}_{f_i}(\mathbf{x})\}$ be the best rank assigned to solution \mathbf{x} . Then, a solution \mathbf{x} dominates solution \mathbf{y} with respect to the MR, denoted by $\mathbf{x} \prec_{max} \mathbf{y}$, if and only if $R_{max}(\mathbf{x}) < R_{max}(\mathbf{y})$.

In the same way as in the study by Corne and Knowles [4], we apply the AR and MR methods on the PF_{approx} set of the current population only. The final rank of the dominated solutions will be the sum of its Pareto rank and the worst rank assigned to the nondominated solutions according to AR or MR, respectively.

2.2 Favour Ranking

In the *favour relation* (FR), proposed by Drechsler et al. [8], a solution \mathbf{x} dominates solution \mathbf{y} with respect to the favour relation, denoted by $\mathbf{x} \prec_{favour} \mathbf{y}$, if and only if:

$$\begin{aligned} &|\{i : f_i(\mathbf{x}) < f_i(\mathbf{y}), 1 \leq i \leq k\}| > \\ &|\{j : f_j(\mathbf{x}) > f_j(\mathbf{y}), 1 \leq j \leq k\}| \end{aligned}$$

This means that \mathbf{x} is *favoured* to \mathbf{y} iff it outperforms \mathbf{y} in more objectives than those in which \mathbf{y} outperforms \mathbf{x} . For example, given $f(\mathbf{x}^{(1)}) = (5, 3, 1)$ and $f(\mathbf{x}^{(2)}) = (1, 1, 2)$, then we have that $\mathbf{x}^{(2)} \prec_{favour} \mathbf{x}^{(1)}$.

2.3 Preference Order Ranking

The *preference order relation* (POR), developed by di Pierro [7], is based on the concept of *efficiency of order*, which states that:

A solution \mathbf{x}^ is considered efficient of order q if it is not dominated by any other solution considering all the $\binom{k}{q}$ objective subsets of size q , where k is the number of objectives.*

The process for assigning ranks to a set of solutions is based on nondominated sorting. First, the solutions in the first nondominated front are ranked according to a strategy based on the preference order relation (see [7] for details); let w be the worst given rank in this process. Next, the solutions in subsequent nondominated fronts receive a rank equal to $w + s$, where s is the number of the nondominated front.

2.4 Expansion relation

Sato et al. [15] proposed a preference relation to control the dominance area of solutions. This method can control the degree of expansion or contraction of the dominance area adopting a user-defined vector $\mathbf{S} = [S_1, \dots, S_k]$. To do so, the value for each objective function is modified using the values S_i in the following manner:

$$f'_i(\mathbf{x}) = \frac{r \cdot \sin(\omega_i + S_i \cdot \pi)}{\sin(S_i \cdot \pi)} \quad \forall i = 1, 2, \dots, k$$

where r is the norm of $f(\mathbf{x})$, $f_i(\mathbf{x})$ is the evaluation of the i -th objective, and ω_i is the angle between $f(\mathbf{x})$ and $f_i(\mathbf{x})$, i.e., $\omega_i = \cos^{-1}(f_i(\mathbf{x})/r)$.

The possible values for S_i range from 0.25 to 0.75. If the user adopts values $S_i < 0.5$, the dominance area is expanded and produces a more fine grained ranking of solutions and would strengthen selection. On the other hand, if the user sets $S_i > 0.5$, the dominance area is contracted and produces a coarser ranking of solutions, weakening the selection procedure. By setting $S_i = 0.5$ the usual Pareto dominance relation is obtained. Since we are interested in

producing a fine grained order of the solutions we will use only values between $0.25 \leq S_i < 0.5$. We will refer to this configuration as the *expansion relation* (ER). Thus, we can say that solution \mathbf{x} dominates solution \mathbf{y} with respect to the expansion relation, denoted by $\mathbf{x} \prec_{\text{expansion}} \mathbf{y}$, if and only if $\forall i : f'_i(\mathbf{x}) \leq f'_i(\mathbf{y})$ and $\exists i : f'_i(\mathbf{x}) < f'_i(\mathbf{y})$.

3. QUANTITATIVE ANALYSIS OF THE PREFERENCE RELATIONS

3.1 Quality Indicators and Methods Used

As noticed earlier, the optimal solution set of each preference relation is, roughly speaking, a subset of PF_{opt} . As a consequence, although a preference relation is only applied on the current PF_{approx} and the archive is maintained using Pareto dominance, the preferred solutions by the preference relation in the primary population belong to a portion of PF_{opt} . Thus, in spite of the fact that the final PF_{approx} set may contain solutions over all the Pareto optimal front, the solutions included in the optimal solution set of the given preference relation are constantly exploited, and the remainder of the solutions may be suboptimal. We can see this situation by comparing the two PF_{approx} sets presented in Figure 1. These sets were obtained using the relations ER and AR. Clearly, these preference relations promote solutions in different regions of the Pareto front.




Figure 1: Pareto front approximations obtained by the *expansion* and the *average ranking* relations.

This poses a challenge to compare Pareto front approximations achieved using different preference relations. For instance, let us suppose we want to compare two preference relations, one that finds solutions in a small region in the middle of the Pareto front, and another that finds solutions in a larger region, but in extreme regions of the Pareto front. If we use, for example, the hypervolume indicator, the preference relation with the larger region will have an inherent advantage over the other relation.

It is commonly accepted that decision makers often select a solution located in the middle of the Pareto front [5], i.e., the knee of the Pareto front. Therefore, we believe that one natural criterion to evaluate preference relations is measuring the distance between the knee and the points in the PF_{approx} set generated using the preference relation. There exist different characterizations of the knee in the literature. Nonetheless, in this paper we will consider that the knee of the Pareto front is the point with the minimum Tchebycheff distance to the ideal point, \mathbf{z}^* , or an approximation of it.

The weighted Tchebycheff distance to \mathbf{z}^* is defined in the following way:

$$d(\mathbf{z}, \mathbf{z}^*, \lambda) = \max_{1 \leq j \leq k} \{\lambda_j |z_j^* - z_j|\},$$

where k is the number of objectives. Defined this way, the knee of the Pareto front is the point in the feasible objective space, Λ , which corresponds to $\min_{\mathbf{z} \in \Lambda} d(\mathbf{z}, \mathbf{z}^*, \lambda)$. Here, we assume that $\lambda_i = \frac{1}{R_i}$, where R_i is the range of the i -th objective in PF_{opt} . Figure 1 shows the knee for the problem DTLZ2, which is located at the point $(\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}})$.

This way, a preference relation is better than other relation if its PF_{approx} set contains more solutions around the knee than the PF_{approx} set achieved by the other relation. In order to evaluate this situation we will plot the distribution of the Tchebycheff distance from the ideal point to the points in a given PF_{approx} . The desired distribution is one with a peak near zero and that decays slowly towards the right since it will indicate that most of the solutions are situated near the knee.

Similar to the approach followed by Corne and Knowles [4], we employed a simple MOEA to evaluate the preference relations included in this study. This way, we try to minimize the effect of some specialized techniques in such a way that the performance of the MOEA can be mainly attributed to the preference relation. Accordingly, the MOEA uses binary encoding, two point crossover, and uniform mutation. To select the parents we used a binary tournament based on the ranks assigned by the given preference relation. The MOEA is equipped with an archive that is truncated by removing a solution selected at random to introduce a new Pareto nondominated solution when the archive is full.

Additionally, for each preference relation we will plot the online generational distance achieved by the current nondominated set generated by the MOEA using the given relation. This way, we can figure out how fast the MOEA converges towards the Pareto front regardless of the spread of the solutions. This information can be useful if we want, for instance, to use a relaxed preference relation to quickly reach the Pareto front and, afterwards, to employ other preference relation to cover a broader extension of the Pareto front. The *generational distance* (GD) [17] is defined by $GD = \left(\sqrt{\sum_{i=1}^n d_i^2} \right) / n$, where n is the size of PF_{approx} and d_i is the Euclidian distance between each vector in PF_{approx} and the nearest member of PF_{opt} . Finally, to assess the distribution of the nondominated set obtained by the MOEA we use the *inverted generational distance* (IGD), which is obtained by interchanging the roles of PF_{opt} and PF_{approx} in the GD's definition.

3.2 Experimental Settings

We adopt the problems DTLZ2 and DTLZ7 to evaluate the performance of the preference relation. The problem DTLZ7 was selected to test the ability of the preference relations to converge towards the knee on problems with a non-convex and disconnected Pareto front¹. We used 3, 5, 8, 10 and 15 objectives in each problem. Regarding the MOEA, in all the simulations we employed a crossover probability of 0.9 and a mutation probability of $1/\ell$, where ℓ is the

¹Objectives in DTLZ2 have the same range, however for DTLZ7 we needed to normalize its objectives using the minimum and maximum values of PF_{opt} .

length of the binary string needed to encode solutions with 5 digits of precision. For each preference relation the MOEA was run 30 times. In each run we used a population of 200 individuals during 300 generations, and an archive of size 300. The reported values of GD and IGD correspond to the average of the 30 runs, whereas the distributions of the Tchebycheff distance were calculated using the union of the PF_{approx} sets achieved by each preference relation.

Since the expansion relation requires a user-defined parameter, S , in the next section we present a preliminary analysis to investigate the influence of that parameter on the performance of ER.

3.3 Analysis of the Expansion Preference Relation

In order to investigate the influence of the parameter S on the expansion relation we solve DTLZ2 using three different values of S , namely $S = 0.3$, $S = 0.35$, $S = 0.4$. Among these values, $S = 0.3$ is the one that expands more the dominance area, and, therefore, it may allow the MOEA to reach faster the Pareto front. In each run of these experiments we used a population of 200 individuals during 200 generations, and an archive size of 200.

For the sake of clarity, the online GD values presented in Figures 2-4 are plotted on a semi-logarithmic scale. From these plots it is clear that with $S = 0.3$ it is obtained the fastest convergence to the Pareto front, while with $S = 0.4$ it is obtained the slowest one. Nevertheless, we have to note that the smaller the S value, the smaller the Pareto region covered by the MOEA using the expansion relation.

In order to study the distribution of the solutions around the knee region, we use the distribution of the Tchebycheff distances with respect to the origin. Figure 5 shows that for 5 objectives, the larger peak of the distributions for $S = 0.3$ and $S = 0.35$ is around 1. That is, most of the solutions are clustered in extreme regions of the Pareto front. However, with $S = 0.35$ there is a considerable number of solutions near the knee of the front. In fact, this is the value that achieves the largest number of solutions around the knee. With respect to 10 and 15 objectives (see Figure 6), for the three values of S , the solutions are concentrated at similar distances from the knee. Since from the three values considered, $S = 0.35$ represents the best trade-off between convergence and distribution around the knee, we used this value for the rest of the experiments.




Figure 2: Online GD achieved with the expansion relation using different values for S in DTLZ2 with 5 objectives.




Figure 3: Online GD achieved with the expansion relation using different values for S in DTLZ2 with 10 objectives.




Figure 4: Online GD achieved with the expansion relation using different values for S in DTLZ2 with 15 objectives.




Figure 5: Distribution of the Tchebycheff distance obtained with the expansion relation and using different values of S in DTLZ2 with 5 objectives.

3.4 Analysis of All the Preference Relations

Like in the previous analysis, in the experiments of this section we used the online generational distance and the distribution of the Tchebycheff distances. However, we also used the inverted generational distance to measure both the spread and convergence to the Pareto front.




Figure 6: Distribution of the Tchebycheff distance obtained with the expansion relation and using different values of S in DTLZ2 with 15 objectives.

The results of the online generational distance on problem DTLZ2 are presented in Figures 7-10. From these plots we can clearly see that, on the one hand, the MOEA with the Pareto relation achieves the worst convergence to the Pareto front during all the search (except for the favour relation). On the other hand, by employing the expansion relation, the MOEA converges very fast during the first 50 and 150 generations for 10 and 15, and 5 objectives, respectively. Then, AR, POR, and FR achieve a closer approximation to the Pareto front.

The second best convergence is obtained using AR, which at the end of the search achieves a better convergence than the expansion relation. The convergence obtained by the favour relation presents an interesting behavior. In the first half of the optimization it achieves a poor convergence. However in the second half, it improves dramatically the convergence and towards the end of the search it produces the best convergence. This behavior is explained by analyzing the IGD results presented in Figure 11 and the distribution of the Tchebycheff distance shown in Figures 12-15. That is to say, Tchebycheff distances show that most of the solutions generated by the MOEA using the favour relation are located far from the knee region (solutions with a Tchebycheff distance of 1). In addition, the IGD value for every number of objectives is very poor on problem DTLZ2 (see Fig. 11). These two facts suggest that the solutions promoted by the favour relation are concentrated in a small region of the Pareto front and, consequently, after a certain number of generations the solutions overexploit that region achieving very small values on the GD indicator but large values on IGD. This behavior, which is worsened with the number of objectives, is also presented using AR and POR.

On the other hand, MR, ER and the Pareto dominance relations present a wider distribution of the Tchebycheff distances for any number of objectives. It is noticeable, though, that the MOEA with the expansion relation finds the closest solutions to the front's knee for any number of objectives. This can be checked by observing that the left tail of its distribution is closer to zero.

With respect to the problem DTLZ2 we can conclude that the best preference relation is the expansion relation since it helps the MOEA to converge quickly to the Pareto front and to maintain more solutions near to the knee of the Pareto

front. Although AR and POR provide good convergence they promote solutions away from the knee.




Figure 7: Online GD achieved by the preference relations in DTLZ2 with 3 objectives.




Figure 8: Online GD achieved by the preference relations in DTLZ2 with 5 objectives.

In the problem DTLZ7 we only analyzed the distribution of the Tchebycheff distances and the results of the IGD indicator. With 3 objectives, the distribution of the Tchebycheff distances achieved by the expansion relation (Fig. 16) shows that most of the solutions are located around a distance of 0.8, which is far from the knee of the front. The other preference relations present similar distributions where the solutions are concentrated around 0.75.

Nevertheless, for more than 3 objectives (Figs. 17-19), most solutions achieved by the expansion relation are near the knee of the Pareto front and they are the closest solutions to the knee. On the other hand, the MOEA with the maximum ranking and Pareto dominance relations generates solutions far away from the true Pareto front (the extreme solutions remain at a distance of 1 since the Tchebycheff distance is normalized). As it can be seen, with these two relations the convergence is worsened as the number of objectives is increased. The results of the IGD indicator shown in Figure 20 confirm this observation since the maximum ranking and Pareto dominance relations obtain the worst values in this indicator. Since most of the solutions obtained by ER are clustered around the knee region, it ob-




Figure 9: Online generational distance achieved by the preference relations in the problem DTLZ2 with 10 objectives.




Figure 10: Online generational distance achieved by the preference relations in the problem DTLZ2 with 15 objectives.




Figure 11: IGD achieved using the preference relations on DTLZ2.




Figure 12: Distribution of the Tchebycheff distance over the solutions generated using each preference relation in the problem DTLZ2 with 3 objectives.




Figure 13: Distribution of the Tchebycheff distance over the solutions generated using each preference relation in the problem DTLZ2 with 5 objectives.




Figure 14: Distribution of the Tchebycheff distance over the solutions generated using each preference relation in the problem DTLZ2 with 10 objectives.




Figure 15: Distribution of the Tchebycheff distance over the solutions generated using each preference relation in the problem DTLZ2 with 15 objectives.




Figure 16: Distribution of the Tchebycheff distance over the solutions generated using each preference relation in the problem DTLZ7 with 3 objectives.




Figure 17: Distribution of the Tchebycheff distance over the solutions generated using each preference relation in the problem DTLZ7 with 5 objectives.

tained poor values in IGD. In turn, AR and POR obtained the best values in IGD.

In the distribution of the Tchebycheff distance from these two relations we can see that there are three peaks in their distributions (specially in Figs. 18 and 19), one close the knee, another in the middle of the distribution, and a third one on the right of the distribution. This suggests that AR and POR yielded a diverse approximation set concentrated in three regions of the Pareto front, and hence their good performance with respect to IGD.

4. CONCLUSIONS AND FUTURE WORK

This paper presented a quantitative analysis of different preference relations proposed to solve many-objective problems. The analysis was mainly based on the distribution of the Tchebycheff distances between the ideal solution and each solution generated using a given preference relation. This distribution reveals the closeness of the solutions achieved to the knee of the Pareto front.

The study revealed that, in spite of the fact that some preference relations contribute to converge faster to the Pareto front than the Pareto dominance relation, they also stress the generation of solutions far from the knee region.




Figure 18: Distribution of the Tchebycheff distance over the solutions generated using each preference relation in the problem DTLZ7 with 10 objectives.




Figure 19: Distribution of the Tchebycheff distance over the solutions generated using each preference relation in the problem DTLZ7 with 15 objectives.




Figure 20: IGD achieved using the preference relations on DTLZ7.

This behavior was observed, for example, in the average ranking and preference order relations in problem DTLZ2. This fact limits the applicability of these relations since, in the general case, it is commonly assumed that the decision maker prefers solutions on the knee region. The expansion relation, on the other hand, presented a remarkable performance. In both problems this relation produced a fast convergence to the Pareto front and, in both problems, it achieved solutions very close to the knee region. The second best preference relation was the average ranking relation followed by the preference order and the favour relations. In terms of convergence this result agrees with the conclusions obtained in [4].

Although, the expansion relation helped to produce solutions near the knee of the Pareto front, in problem DTLZ7, the solutions were concentrated in a small region around the knee. This introduces a trade-off between convergence and the size of the region covered. The parameter of the expansion preference relation opens interesting applications to the relation in MOEAs. For instance, it can be incremented gradually during the search in order to approach quickly the Pareto front during the first half of the search and then cover the rest of the Pareto front in the second half of the search.

We believe that this preliminary study introduced an interesting methodology to compare preference relations that find subsets of the Pareto front. Nonetheless, as part of our future work, we plan to incorporate more test problems to the study, specially real-life many-objective problems. Furthermore, we want to investigate the k -optimality relation [9], which is another preference relation with a user-defined parameter to control the relaxation of the relation.

5. REFERENCES

- [1] P. J. Bentley and J. P. Wakefield. Finding Acceptable Solutions in the Pareto-Optimal Range using Multiobjective Genetic Algorithms. In P. K. Chawdhry, R. Roy, and R. K. Pant, editors, *Soft Computing in Engineering Design and Manufacturing*, pages 231–240, London, June 1997. Springer Verlag.
- [2] D. Brockhoff and E. Zitzler. Are All Objectives Necessary? On Dimensionality Reduction in Evolutionary Multiobjective Optimization. In T. P. Runarsson, H.-G. Beyer, E. Burke, J. J. Merelo-Guervós, L. D. Whitley, and X. Yao, editors, *Parallel Problem Solving from Nature - PPSN IX, 9th International Conference*, pages 533–542. Springer. Lecture Notes in Computer Science Vol. 4193, Reykjavik, Iceland, September 2006.
- [3] D. Brockhoff and E. Zitzler. Improving Hypervolume-based Multiobjective Evolutionary Algorithms by Using Objective Reduction Methods. In *2007 IEEE Congress on Evolutionary Computation (CEC'2007)*, pages 2086–2093, Singapore, September 2007. IEEE Press.
- [4] D. Corne and J. Knowles. Techniques for Highly Multiobjective Optimisation: Some Nondominated Points are Better than Others. In D. Thierens, editor, *2007 Genetic and Evolutionary Computation Conference (GECCO'2007)*, volume 1, pages 773–780, London, UK, July 2007. ACM Press.
- [5] I. Das. On characterizing the “knee” of the Pareto curve based on Normal-Boundary Intersection. *Structural Optimization*, 18(2–3):107–115, October 1999.
- [6] K. Deb and D. Saxena. On finding pareto-optimal solutions through dimensionality reduction for certain large-dimensional multi-objective optimization problems. *Kangal report*, 2005.
- [7] F. di Pierro, S.-T. Khu, and D. A. Savić. An Investigation on Preference Order Ranking Scheme for Multiobjective Evolutionary Optimization. *IEEE Transactions on Evolutionary Computation*, 11(1):17–45, February 2007.
- [8] N. Drechsler, R. Drechsler, and B. Becker. Multi-Objected Optimization in Evolutionary Algorithms Using Satisfiability Classes. In B. Reusch, editor, *International Conference on Computational Intelligence, Theory and Applications, 6th Fuzzy Days*, pages 108–117, Dortmund, Germany, 1999. Springer-Verlag. Lecture Notes in Computer Science Vol. 1625.
- [9] M. Farina and P. Amato. On the Optimal Solution Definition for Many-criteria Optimization Problems. In *Proceedings of the NAFIPS-FLINT International Conference'2002*, pages 233–238, Piscataway, New Jersey, June 2002. IEEE Service Center.
- [10] E. J. Hughes. Evolutionary Many-Objective Optimisation: Many Once or One Many? In *2005 IEEE Congress on Evolutionary Computation (CEC'2005)*, volume 1, pages 222–227, Edinburgh, Scotland, September 2005. IEEE Service Center.
- [11] J. Knowles and D. Corne. Quantifying the Effects of Objective Space Dimension in Evolutionary Multiobjective Optimization. In S. Obayashi, K. Deb, C. Poloni, T. Hiroyasu, and T. Murata, editors, *Evolutionary Multi-Criterion Optimization, 4th International Conference, EMO 2007*, pages 757–771, Matshushima, Japan, March 2007. Springer. Lecture Notes in Computer Science Vol. 4403.
- [12] A. López Jaimes, C. A. Coello Coello, and D. Chakraborty. Objective Reduction Using a Feature Selection Technique. In *2008 Genetic and Evolutionary Computation Conference (GECCO'2008)*, pages 674–680, Atlanta, USA, July 2008. ACM Press. ISBN 978-1-60558-131-6.
- [13] V. Pareto. *Cours D'Economie Politique*. F. Rouge, 1896.
- [14] K. Praditwong and X. Yao. How Well Do Multi-Objective Evolutionary Algorithms Scale to Large Problems. In *2007 IEEE Congress on Evolutionary Computation (CEC'2007)*, pages 3959–3966, Singapore, September 2007. IEEE Press.
- [15] H. Sato, H. E. Aguirre, and K. Tanaka. Controlling Dominance Area of Solutions and Its Impact on the Performance of MOEAs. In S. Obayashi, K. Deb, C. Poloni, T. Hiroyasu, and T. Murata, editors, *Evolutionary Multi-Criterion Optimization, 4th International Conference, EMO 2007*, pages 5–20, Matshushima, Japan, March 2007. Springer. Lecture Notes in Computer Science Vol. 4403.
- [16] A. Sülflow, N. Drechsler, and R. Drechsler. Robust Multi-objective Optimization in High Dimensional Spaces. In S. Obayashi, K. Deb, C. Poloni, T. Hiroyasu, and T. Murata, editors, *Evolutionary Multi-Criterion Optimization, 4th International Conference, EMO 2007*, pages 715–726, Matshushima, Japan, March 2007. Springer. Lecture Notes in Computer Science Vol. 4403.
- [17] D. A. Van Veldhuizen. *Multiobjective Evolutionary Algorithms: Classifications, Analyses, and New Innovations*. PhD thesis, Department of Electrical and Computer Engineering. Graduate School of Engineering. Air Force Institute of Technology, Wright-Patterson AFB, Ohio, May 1999.
- [18] T. Wagner, N. Beume, and B. Naujoks. Pareto-, Aggregation-, and Indicator-Based Methods in Many-Objective Optimization. In S. Obayashi, K. Deb, C. Poloni, T. Hiroyasu, and T. Murata, editors, *Evolutionary Multi-Criterion Optimization, 4th International Conference, EMO 2007*, pages 742–756, Matshushima, Japan, March 2007. Springer. Lecture Notes in Computer Science Vol. 4403.