

# Connectedness, Regularity and the Success of Local Search in Evolutionary Multi-objective Optimization

**Yaochu Jin**

Honda Research Institute Europe  
Carl-Legien Strasse 30,  
63073 Offenbach/M, Germany  
yaochu.jin@honda-ri.de

**Bernhard Sendhoff**

Honda Research Institute Europe  
Carl-Legien Strasse 30,  
63073 Offenbach/M, Germany  
bernhard.sendhoff@honda-ri.de

**Abstract-** Local search techniques have proved to be very efficient in evolutionary multi-objective optimization(MOO). However, the reasons behind the success of local search in MOO have not yet been well discussed. This paper attempts to investigate empirically the main factors that may have contributed significantly to the success of local search in MOO. It is found that for many widely used test problems, the Pareto optimal solutions are connected both in objective space and parameter space. Besides, the Pareto-optimal solutions often distribute so regularly in parameter space that they can be defined by piecewise linear functions. By constructing an approximate model using the solutions produced by an optimizer, the quality of the non-dominated solution set can be further improved.

The evolutionary dynamic weighted aggregation (EDWA) method has been adopted as a local search technique in finding Pareto-optimal solutions. Its effectiveness for MOO is demonstrated on a number of two or three objective optimization problems.

## 1 Introduction

Local search techniques for evolutionary multi-objective optimization have received increasing attention in the recent years. Generally, multi-objective local search can be divided into three categories. First, multi-objective search is carried out by heuristic local search methods, such as simulated annealing [4, 25], Tabu search [9], particle swarm optimization [21] or other techniques [10]. Second, local search is combined with global search either during the optimization [11, 17, 12, 18] or at the end of optimization [16, 26]. Third, one or a part of the Pareto-optimal solutions are found first and then the further search is continued starting from these solutions [22, 20, 23].

The dynamic weighted aggregation (DWA) method [13, 14] has shown to be effective for multi-objective optimization using evolution strategies [15] and particle swarm optimization [21]. Evolution strategies (ES) are not strictly local, however, as the evolution proceeds, the step-sizes become smaller and smaller and the search becomes more

local. Therefore, DWA in combination with an evolution strategy can also be regarded as a local search method for MOO.

The reasons for the success of local search techniques have not been fully discussed. One of the first research in investigating the basis of the success of local search in multi-objective combinatorial optimization (MOCO) is reported in [2]. It is pointed out that the success of local search in MOCO can be attributed to the phenomenon known as *global convexity*. By global convexity, it is understood that in MOCO, the Pareto optimal solutions are concentrated in a small fraction of parameter space and those are in the neighborhood in objective space are also in the neighborhood in parameter space. This is also termed as *connectedness* in MOCO [8]. The connectedness is of particular interest because if this is true, a local search technique can find Pareto optimal solutions without exploring the dominated regions when one Pareto optimal solution has been found. However, as it is pointed out in [8], connectedness does not hold for all MOCO problems.

Inspired from the above discussions, it has been speculated that the success of the DWA method can also partly be attributed to the connectedness of the Pareto optimal solutions for many continuous optimization problems [14]. Another important factor that makes the DWA successful is that evolution strategies are able to conduct local search at the late stage of the optimization when the population has reached the Pareto front. On the contrary, canonical genetic algorithms are not able to conduct local search in multi-objective optimization even when they are converged to the Pareto front. To enable genetic algorithms to carry out local search, crossover should be limited to parents in neighborhood only.

This paper aims to show empirically that connectedness holds also for many continuous optimization problems. Besides, it demonstrates that the distribution of the Pareto optimal solutions in parameter space exhibits high simplicity and regularity. In some cases, the Pareto optimal solutions can be fully defined by a number of piecewise linear functions. Once a model is constructed from the solutions generated by an optimizer, we are able to refine the quality of

solutions significantly.

In the following Section, the local search method used in the empirical studies in this paper, i.e., the DWA using an evolution strategy, will be briefly reviewed. Illustrative simulation results are provided in Section 3 with observations of typical behavior of individuals on convex and concave Pareto fronts. In Section 4, the connectedness of Pareto optimal solutions is discussed by showing empirical results from some test functions. The regularity in distribution of Pareto optimal solutions is addressed in Section 5, where piecewise linear models are constructed to approximate Pareto optimal solutions in parameter space. Section 6 concludes the paper with discussions of some open issues.

## 2 Dynamic Weighted Aggregation

The central idea behind the dynamic weighted aggregation method is to use variable weights during the evolutionary optimization to obtain multiple non-dominated solutions in one run. A simple approach to realizing this is to change weights gradually from generation to generation during optimization, which is termed as dynamic weighted aggregation (DWA) in [14]. Although there are no specific requirements on how the weights should be changed, the following two factors should be taken into account:

- Each weight should be changed from 0 to 1 or from 1 to 0 once to ensure that the whole Pareto front can be searched.
- The speed of the weight change should be sufficiently slow so that the evolutionary algorithm is able to converge to the Pareto front. One practice is to fix the weight first and let the population converge to one point of the Pareto front. However, if evolution strategies are used, the step-sizes should be prevented from converging to zero so that the population can move fast enough to follow the change of the weights.

A simple mode for weight change for two-objective optimization problems is as follows :

$$\begin{aligned} w_1(t) &= |\sin(2\pi t/F)| \\ w_2(t) &= 1.0 - w_1(t), \end{aligned} \quad (1)$$

where  $t$  is the generation index,  $|\cdot|$  takes the absolute value,  $F$  is the weight change frequency, usually between 50 and 200. For example, if  $F = 200$ , the weights will change twice between 0 and 1 within 100 generations.

The dynamic weighted aggregation method can be extended to three-objective problems in a straightforward way. An example of the weight change is as follows:

$$\text{Let } t_1 = 0, t_2 = 0, \text{ and } t = t_1 + t_2$$

```

for  $t_1 = 0$  to  $F/2$ 
   $w_1(t) = |\sin(2\pi t_1/F)|$ ;
  for  $t_2 = 0$  to  $F/2$ 
     $w_2(t) = (1.0 - w_1(t)) |\sin(2\pi t_2/F)|$ ;
     $w_3(t) = 1.0 - w_1(t) - w_2(t)$ ;
  end;
end.

```

A special case of the DWA method is to switch the weights between 0 and 1, which has been termed the bang-bang weighted aggregation (BWA) in [14]. Empirical studies have shown that the BWA method exhibits very good performance if the Pareto front is concave.

In DWA, the population is not able to keep all found non-dominated solutions. Therefore, it is necessary to store the non-dominated solutions found so far. However, the archive has no influence on the search behavior.

## 3 Simulations and Observations

### 3.1 Evolution Strategy

A standard ES [24] with comma selection is adopted in the simulations. The parent and offspring population sizes are set to 15 and 100, respectively. The initial step-sizes are set to 0.1 and 400 generations are run for each test function. The weights are changed periodically according to equation (1), where  $F$  is set to 200. No recombination operation has been implemented. During optimization, the step-sizes are checked in each generation and are re-initialized if they are smaller than a prescribed value. This is very important because if the step-sizes converge to 0 or a very small value, the population will lose its ability to move.

### 3.2 Test Functions

To show the effectiveness of the method, two bi-objective and two three-objective test functions have been used in this paper.  $F_1$  is the **Schaffer function** with a connected convex Pareto front.  $F_2$  is the **Fonseca (2) function** with a connected and concave Pareto front.  $F_3$  is a three-objective **Viennet function** of which the Pareto solutions form a convex surface and  $F_4$  is the three-objective **Viennet (3) function**, whose Pareto front consists of disconnected pieces of curves. Refer to [3] for the detailed descriptions of the test functions.

### 3.3 Simulation Results and Observations

The found non-dominated solutions for the four test functions are provided in Figures 1, 2, 3 and 4, respectively.

The following two important observations can be made from the optimization processes using the EDWA method.




Figure 1: Achieved non-dominated solutions for  $F_1$ .




Figure 2: Achieved non-dominated solutions for  $F_2$ .




Figure 3: Achieved non-dominated solutions from  $F_3$ .




Figure 4: Achieved non-dominated solutions from  $F_4$ .

- If the Pareto front is convex, the population will first converge to a certain point on the Pareto front. To which point the population converges depends on the problem difficulty (convergence speed) and on the speed of weight change. Once the population has reached the Pareto front, it will keep moving on it as the weights change. The speed of movement can be controlled by the frequency of the weights change. Fig.5 shows a typical trace of the parent population during optimization of the test function  $F_1$  in the first 100 generations represented by the mean of the population. Using different strategy parameters, only minor changes will occur in the starting phase of the trace.




Figure 5: Movements of the population with time (generation) on a convex Pareto front. The arrows denote the direction of the movement and the numbers the generation.

- If the Pareto front is concave, the population will converge to one of the two ends of the concave Pareto front. Thereafter, the population will remain on this Pareto optimal point until the weights are changed to a threshold value, which is determined by the characteristics of the Pareto front. Then, the population will move along or close to the Pareto front very quickly to the other end of the Pareto front. To illustrate this, the first 100 moves of the parent population during optimization of test function  $F_2$  are presented in Fig. 6. From the figure, it can be seen that the population converges to the Pareto point  $(f_1, f_2)=(0.98,0)$  in a few generation. As the weights change, the population remains on the same point until generation 29. Then, the population moves along the Pareto front quickly to the other end of the Pareto front, where  $(f_1, f_2)=(0,0.98)$ . Again, the population remains on this point until generation 91.

The effectiveness of the DWA method, as well as other local search techniques, may probably be attributed to the connectedness and regularity of Pareto optimal solutions in many MOO problems, which will be discussed in more details in the following sections.




Figure 6: Movements of the population with time (generation) on a concave Pareto front. The arrows denote the direction of the movement and the numbers the generation.

#### 4 Connectedness and Search Efficiency

Consider the multi-objective optimization problem:

$$\min_{x \in S} (f_1(x), \dots, f_m(x)), \quad (2)$$

where  $S \subseteq R^n$  is non-empty, compact, and convex, the functions  $f_j : S \rightarrow R, j = 1, \dots, m$  are continuous, where  $n$  and  $m$  are the dimension of parameter space and number of objectives, respectively. Then the Pareto-optimal solutions are connected if all objective functions are convex [19]. Fig. 7 shows a connected and non-connected Pareto fronts.




Figure 7: Connectedness of Pareto optimal solutions. (a) The Pareto front consists of two disconnected sections,  $AB$  and  $CD$ . (b) The Pareto front consists of one connected section  $AB$ .

The convexity condition for the connectedness can be relaxed. For example, it has been shown that the connectedness holds if the objective functions are semi-convex [5]. Refer to [5, 7] and the references therein for detailed discussions on connectedness.

If connectedness holds, a local search algorithm will be able to get all Pareto optimal solutions very efficiently starting from a single Pareto optimal point. However, the above definition of connectedness is very strict. Fortunately, a local search algorithm can benefit much even if the strict connectedness does not hold. For example, a local search al-

gorithm is able to find all solutions efficiently that are connected, although the whole Pareto front is disconnected. In case a Pareto front is composed of a finite number of disconnected sections, we say that the Pareto optimal solutions are *loosely connected*.

It is also found in evolutionary DWA based search, the population is able to move along the boundary of the feasible region, such as section  $BC$  in Fig. 7(a), which connects different sections of the Pareto front. This is very interesting because it means that the optimizer is able to find a pathway that bridges the gap between the disconnected pieces of the Pareto front. To show this phenomenon, the DWA algorithm is applied to the third test function used in [28], whose Pareto optimal front consists of five disconnected sections. It can be seen from Fig. 8 that the population is moving along the boundaries that connect the Pareto optimal sections, although solutions on these boundaries are not Pareto optimal themselves.




Figure 8: The population is moving between the disconnected Pareto sections.

#### 5 Regularity and Approximation of the Pareto Optimal Solutions

##### 5.1 Conjectures on Regularity

One of the important aspects of global convexity is that Pareto-optimal solutions are concentrated in a very small region of parameter space. In addition, Pareto optimal solutions are loosely connected for a large number of MOO problems, which can be taken advantage of by local search methods.

In this section, we will look into regularity of the distribution of Pareto optimal solutions in parameter space. Before going into details, we start with the following conjectures concerning the properties of the distribution of Pareto-optimal solutions. These conjectures can be empirically verified on many popular multi-objective test problems.

- If the Pareto front is a curve in objective space, the Pareto optimal solutions can also be defined by a curve in parameter space. If the Pareto front is a sur-

face in objective space, the Pareto optimal solutions are also a surface.

- The order of the function that defines the Pareto optimal set in parameter space is equal to or lower than that of the function describing the Pareto front in objective space. Thus, if the Pareto front consists of lower order curves, the solutions in parameter space can most probably be described by piecewise linear functions.

Note that it can be shown that a Pareto optimal surface is at most an  $m - 1$  dimensional surface, where  $m$  is the number of objectives. Therefore, if  $m$  is less than or equal to 3, the Pareto front is at most a two-dimensional surface.

If the conjectures about the regularity hold for an unknown optimization task, it will be very helpful in improving the quality of the solutions obtained from an optimizer. Besides, it will be theoretically possible to recover more Pareto optimal solutions from a limited number of Pareto solutions achieved by an optimizer. Since the conjectures suggest that the complexity of the function that defines the Pareto optimal solutions in the parameter space (we will call it *definition function* hereafter) is simpler than that defines the Pareto front in objective space, it will be more effective to approximate the definition function than to approximate the Pareto front as done in [27].

In the following, approximate models for the definition function will be constructed for the test problems on the basis of the previous conjectures. It will be shown that in all the cases, the accuracy of the solutions will be improved significantly. Furthermore, missing solutions could be recovered through such analysis.

### 5.2 Test Function $F_1$

Using the solutions obtained in Section 3, it is straightforward to get the following linear model, which is illustrated in Fig. 9 (a). The Pareto front directly generated from the approximate model as well as the solutions achieved by the optimizer is plotted in Fig. 9 (b). Obviously, the accuracy of the solutions has been improved through the linear approximation.

$$x_2 = 0.0296 + 0.981x_1. \quad (3)$$

### 5.3 Test Function $F_2$

Similarly, the following linear model has been obtained for  $F_2$  based on the points achieved by the DWA:

$$x_2 = -0.0397 + 0.9035x_1, -0.70 \leq x_1 \leq 0.70. \quad (4)$$

The solutions achieved by the optimizer, its approximate model and the Pareto front generated from the model are shown in Fig. 10.




Figure 9: Test function  $F_1$ : (a) Distribution of the solutions in parameter space and the estimated definition function; (b) Pareto front generated from the estimated definition function (solid line) and the achieved solutions (dots).




Figure 10: Test function  $F_2$ : (a) The achieved solutions in parameter space and its estimated model; (b) Pareto front generated from the estimated definition function (solid line) and the achieved solutions (dots).

### 5.4 Test Function $F_3$

The regularity of the distribution of the solutions in parameter space can easily be observed. It is found that the boundary of Pareto optimal solution region can be defined by the following three lines:

$$x_2 = x_1 - 1.0, \quad 0 \leq x_1 \leq 1, \quad (5)$$

$$x_2 = -x_1 + 1.0, \quad 0 \leq x_1 \leq 1, \quad (6)$$

$$x_1 = 0, \quad -1 \leq x_2 \leq 1. \quad (7)$$

Using these linear curves, it is straightforward to get the boundary of the Pareto surface, as shown in Fig. 11, where the approximated solutions are also provided.

In this way, it is easy to get the solutions that are not Pareto-optimal. Meanwhile, it is interesting to find out that the boundary of the Pareto optimal solutions in parameter space corresponds nicely to that in objective space.

### 5.5 Test Function $F_4$

From the distribution of the obtained solutions in parameter space, refer to Fig. 12, the definition function of this test problem consists of more than one linear sections plus a




Figure 11: Boundary of the true Pareto surface. (a)  $f_1-f_2$ . (b)  $f_2-f_3$  and (c)  $f_3-f_1$ .

separate point, while the Pareto front appears to be a continuous curve. From the obtained solutions, it is observed that the Pareto optimal solutions is composed of four sections. The following approximate model can be used to describe the solutions:

$$S1: x_2 = 0.001 - 0.62x_1, \quad -0.95 \leq x_1 \leq 0, \quad (8)$$

$$S2: x_2 = 1.94 + 1.47x_1, \quad -1.2 \leq x_1 \leq -0.95, \quad (9)$$

$$S3: x_2 = 1.94 + 1.47x_1, \quad -3 \leq x_1 \leq -1.95, \quad (10)$$

$$S4: x_2 = 30.8 + 11.19x_1, \quad -3 \leq x_1 \leq -2.89. \quad (11)$$

In the obtained solutions, section 4 consists of only a few solutions. However, it is noticed in model construction that the range of this section can be extended and the solutions are still Pareto optimal. This results in finding an additional section of Pareto fronts that has often been missing in the existing literature, refer to Fig. 20. The obtained solutions together with their approximation are shown in Fig. 12.




Figure 12: Distribution of the solutions and the approximation in parameter space.

The Pareto front generated by line section 1 is shown in Fig. 13 and the two sections of the Pareto front generated by line sections 2 and 3 are illustrated in Figures 14, 15 and 16 respectively.

When we take a closer look at the solutions reconstructed from the approximate definition functions, we notice that the Pareto front has richer features than what has been obtained in existing work [6, 3]. The Pareto solutions gen-




Figure 13: Pareto front generated by line section 1.




Figure 14: Pareto front generated by section 2 and 3:  $f_1-f_2$ .

erated by line section 2 are so close to those generated by line section 3 that it is difficult to distinguish them without zooming in. The solutions generated by line section 4 are illustrated in Figures 17, 18 and 19.

Finally, the complete Pareto front generated from the approximate definition function in parameter space is presented in Fig. 20. Compared to the results shown in Fig. 4(b), it can be seen that the Pareto-optimal solutions generated by line section 4 have been completely missing in the Pareto front obtained using the optimization algorithm.

## 6 Discussions and Conclusions

The main target of the paper is to suggest that connectedness and regularity can be the most important reasons behind the success of local search algorithms in MOO. It is shown that local search algorithms, such as the DWA method, can effectively take advantage of the connectedness and consequently its search efficiency is high. It is also shown that the distribution of the Pareto optimal set exhibits surprising regularity and simplicity in parameter space, which is very interesting and could be very helpful. By taking advantage of such regularities, it is possible to build simple models from the obtained Pareto-optimal solutions. Such an approximate model can be of great significance in the following aspects.




Figure 15: Pareto front generated by section 2 and 3:  $f_2$ - $f_3$ .




Figure 16: Pareto front generated by section 2 and 3:  $f_3$ - $f_1$ .

- It allows to improve the accuracy and distribution of Pareto optimal solutions obtained by an optimizer.
- It could alleviate difficulties in multi-objective optimization. If the whole Pareto front can be generated from a few Pareto optimal solutions, requirements such as distribution and spread on the optimizer can be alleviated.

It should be pointed out that strict connectedness holds only for some particular class of MOO problems. Whether connectedness in a loose sense holds for all MOO problems is still unclear. Besides, the regularity and simplicity observed in the test functions needs to be verified for a wider range of MOO problems. It is thus important to check if most real-world applications also show connectedness and regularity. If this is not the case, it would be a critical issue in constructing test functions to avoid this kind of regularity.

Another issue that deserves further research is to investigate whether self-adaptation mechanism of the ES works properly for multi-objective optimization. When ES is combined with DWA, it can be seen as to track a dynamic minimum moving during the optimization, in which case the self-adaptation behavior is not well understood yet [1]. Ideally, in tracking a convex Pareto front, the step-sizes should enable the population to follow the movement of the minimum. If the Pareto front is concave, the step-sizes should be prevented from converging to 0. Meanwhile, they should




Figure 17: Pareto front generated by section 4:  $f_1$ - $f_2$ .




Figure 18: Pareto front generated by section 4:  $f_2$ - $f_3$ .

not be too large so that the population can keep moving along the Pareto front.

**Acknowledgment** The authors are grateful to Prof. E. Körner for his support.

## Bibliography

- [1] D.V. Arnold and H.-G. Beyer. Random dynamic optimum tracking with evolution strategies. In *Parallel Problem Solving from Nature*, pages 3–12, Berlin, 2002. Springer.
- [2] P.C. Borges and M.P. Hansen. A basis for future successes in multi-objective combinatorial optimization. Technical Report IMM-REP-1998-8, Department of Mathematical Modeling, Technical University of Denmark, 1998.
- [3] C. Coello Coello, D. Van Veldhuizen, and G. Lamont. *Evolutionary Algorithms for Solving Multi-Objective Problems*, chapter 3, page 114. Kluwer, New York, 2002.
- [4] P. Czyzak and A. Jaszkiewicz. Pareto simulated annealing - A meta-heuristic techniques for multiple-objective combinatorial optimization. *Journal of Multi-Criteria Decision Analysis*, 7:34–47, 1998.
- [5] A. Daniilidis, N. Hadjisavvas, and S. Schable. Connectedness of the efficient set for three-objective quasiconcave maximization problems. *Journal of Optimization Theory and Applications*, 93(3):517–524, 1997.
- [6] K. Deb. *Multi-objective Optimization Using Evolutionary Algorithms*, chapter 8, page 331. Wiley, Chichester, 2001.
- [7] M. Ehrgott. *Multi-criteria Optimization*. Springer, Berlin, 2000.




Figure 19: Pareto front generated by section 4:  $f_3$ - $f_1$ .




Figure 20: The Pareto front reconstructed from the approximate models of the definition function.

- [8] M. Ehrgott and K. Klamroth. Connectedness of efficient solutions in multiple criteria combinatorial optimization. *European Journal of Operations Research*, 97:159–166, 1997.
- [9] X. Grandbleux, N. Mezdaoui, and A. Freville. A tabu search procedure to solve multi-objective combinatorial optimization problems. In *Advances in Multiple Objective and Goal Planning*, pages 291–300. Springer, 1997.
- [10] M.P. Hansen. Use of substitute scalarizing functions to guide a local search base heuristic: The case of moTSP. *Journal of Heuristics*, 6(3):419–431, 2000.
- [11] H. Ishibuchi and T. Murada. A multi-objective genetic local search algorithm and its application to flow shop scheduling. *IEEE Transactions on Systems, Man, and Cybernetics - Part C: Applications and Reviews*, 28(3):392–403, 1998.
- [12] A. Jaszkiewicz. Genetic local search for multiple objective combinatorial optimization. *European Journal of Operational Research*, 137(1):50–71, 2002.
- [13] Y. Jin, T. Okabe, and B. Sendhoff. Adapting weighted aggregation for multi-objective evolution strategies. In *Proceedings of The 1st Int. Conf. on Evolutionary Multi-Criterion Optimization*, pages 96–110, Berlin, 2001. Springer.
- [14] Y. Jin, M. Olhofer, and B. Sendhoff. Evolutionary dynamic weighted aggregation for multi-objective optimization: Why does it work and how? In *Proceedings of Genetic and Evolutionary Computation Conference*, pages 1042–1049, San Francisco, CA, 2001.
- [15] Y. Jin and B. Sendhoff. Trade-off between performance and robustness: An evolutionary multi-objective approach. In *Proceedings of the 2nd Int. Conf. on Evolutionary Multi-Criterion Optimization*, pages 237–251, Berlin, 2003. Springer.
- [16] K. and T. Goel. A hybrid multi-objective evolutionary approach to engineering shape design. In *Proceedings of the 1st Int. Conf. on Multi-Criterion Optimization*, pages 385–399, Berlin, 2001. Springer.
- [17] J.D. Knowles. M-PAES: A memetic algorithm for multi-objective optimization. In *Proceedings of the 2000 Congress on Evolutionary Computation*, volume 1, pages 325–332, Piscataway, 2000. IEEE Press.
- [18] T. Murata, H. Nozawa, Y. Tsujimura, and Y. Gen. Effects of local search on the performance of cellular multi-objective genetic algorithms for designing fuzzy rule-based classification systems. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 663–668, San Francisco, 2002. Morgan Kaufmann.
- [19] P.H. Naccache. Connectedness of the set of non-dominated outcomes in multi-criteria optimization. *Journal of Optimization Theory and Applications*, 25:459–467, 1978.
- [20] L. Paquete. A two-phase local search for the bi-objective traveling salesman problem. In *Proceedings of the Second Int. Conf. on Evolutionary Multi-Criterion Optimization*, pages 479–493. Springer, 2003.
- [21] K.E. Parsopoulos and M.N. Vrahatis. Particle swarm optimization method in multi-objective problems. In *Proceedings of the 2002 ACM Symposium on Applied Computing*, pages 603–607, 2003.
- [22] S. Ranjithan, K. Chetan, and H. Dakshina. Constraint method based evolutionary algorithm for multi-objective optimization. In *Proceedings of 1st Int. Conf. on Evolutionary Multi-Criterion Optimization*, pages 299–313, Berlin, 2001. Springer.
- [23] O. Schütze, S. Mostaghim, M. Dellnitz, and J. Teich. Covering Pareto sets by multi-level evolutionary subdivision techniques. In *Proceedings of the Second Int. Conf. on Evolutionary Multi-Criterion Optimization*, pages 118–132, Berlin, 2003. Springer.
- [24] H.-P. Schwefel. *Evolution and Optimum Seeking*. John Wiley & Sons, Inc., 1994.
- [25] P. Serafini. Simulated annealing for multi-objective optimization problems. In *Proceedings of the Tenth Int. Conf. on Multiple Criteria Decision Making*, volume 1, pages 87–96, 1992.
- [26] E. Talbi, M. Rahoual, M.H. Mabed, and C. Dhaenens. A hybrid evolutionary approach for multi-criteria optimization problems: Application to the flow shop. In *Proceedings of the 1st Int. Conf. on Multi-Criterion Optimization*, Berlin, 2001. Springer.
- [27] B. Wilson, D. Cappelleri, T. Simpson, , and M. Frecker. Efficient Pareto frontier exploration using surrogate approximations. *Optimization and Engineering*, 2:31–50, 2001.
- [28] E. Zitzler, K. Deb, and L. Thiele. Comparison of multi-objective evolution algorithms: empirical results. *Evolutionary Computation*, 8(2):173–195, 2000.