Twenty Years of Evolutionary Multi-Objective Optimization: A Historical View of the Field

Carlos A. Coello Coello
CINVESTAV-IPN
Evolutionary Computation Group
Departamento de Ingeniería Eléctrica
Sección de Computación
Av. Instituto Politécnico Nacional No. 2508
Col. San Pedro Zacatenco
México, D. F. 07300
ccoello@cs.cinvestav.mx

November 11, 2005

Abstract

This article provides a general overview of the field now known as "evolutionary multi-objective optimization", which refers to the use of evolutionary algorithms to solve problems with two or more (often conflicting) objective functions. Using as a framework the history of this discipline, we discuss some of the most representative algorithms that have been developed so far, as well as some of their applications. Also, we discuss some of the methodological issues related to the use of multi-objective evolutionary algorithms, as well as some of the current and future research trends in the area.

1 Introduction

Optimization using metaheuristics has become a very popular research topic in the last few years [14]. Optimization refers to finding the best possible solution to a problem given a set of limitations (or constraints). When dealing with a single objective to be optimized (e.g., the cost of a design), we aim to find the best possible solution available (called "global optimum"), or at least a good approximation of it.

However, when devising optimization models for a problem, it is frequently the case that there is not one but several objectives that we would like to optimize. In fact, it is normally the case that these objectives are in conflict with each other. For example, when designing a bridge, we would like to minimize its cost while maximizing its safety. These problems with two or more objective functions are called "multi-objective" and require different mathematical and algorithmic tools than those adopted to solve single-objective optimization problems. In fact, even the notion of "optimality" changes when dealing with multi-objective optimization problems. The general nonlinear single-objective optimization problem is as much an open problem as the general nonlinear multi-objective optimization problem. Therefore, the use of metaheuristics is a valid choice which, in fact, has rapidly gained acceptance among researchers from a wide variety of disciplines. From the several metaheuristics currently available, evolutionary algorithms (which are based on the emulation of the mechanism of natural selection) are among the most popular [35, 28]. Evolutionary algorithms have been popular in single-objective optimization and, more recently, have also become common in multi-objective optimization, in which they present several advantages with respect to other techniques, as we will see later on. In this article, we will provide an overview of the field now called "evolutionary multi-objective optimization", which refers to the use of evolutionary algorithms to solve multi-objective optimization problems. The overview will not be comprehensive nor will discuss in detail the many approaches currently available (more technical surveys with that sort of information already exist [5, 85, 88]). Instead, it will be a historical tour that will try to illustrate the rapid development of this field.

2 Basic Concepts

In an attempt for avoiding the use of cumbersome mathematical terms, we will provide only informal definitions of the main concepts required to understand the rest of this article. It is assumed, however, that the reader is familiar with the generalities of evolutionary algorithms (if more background on this topic is required, there are several good references that the reader is invited to consult, such as [35, 28, 25]).

First, we need to define a multi-objective optimization problem (MOP). A MOP is a problem which has two or more objectives that we need to optimize simultaneously. It is important to mention that there might be constraints imposed on the objectives. It is also important to emphasize that it is normally the case that the objectives of the MOP are in conflict with each other. If this is not the case, then a single solution exists for the MOP, and what we will discuss in the remainder of this article no longer applies, because the objectives can be optimized one by one, in sequential order to find this single solution.

Most MOPs, however, do not lend themselves to a single solution and have, instead, a set of solutions. Such solutions are really "trade-offs" or good compromises among the objectives. In order to generate these trade-off solutions, an old notion of optimality is normally adopted. This notion of optimality was originally introduced by Francis Ysidro Edgeworth in 1881 [23] and later generalized by Vilfredo Pareto in 1896 [69]. It is called *Edgeworth-Pareto optimum* or, simply, *Pareto optimum*. In words, this definition says that a solution to a MOP is Pareto optimal if there exists no other feasible solution which would decrease some criterion without causing a simultaneous increase in at least one other criterion. It should not be difficult to realize that the use of this concept almost always gives not a single solution but a set of them, which is called the *Pareto optimal set*. The vectors of the decision variables corresponding to the solutions included in the Pareto optimal set are called *nondominated*. The plot of the objective functions whose nondominated vectors are in the Pareto optimal set is called the *Pareto front*.

3 Why Evolutionary Algorithms?

The Operations Research community has developed approaches to solve MOPs since the 1950s. Currently, a wide variety of mathematical programming techniques to solve MOPs are available in the specialized literature (see for example [59, 24]). However, mathematical programming techniques have certain limitations when tackling MOPs. For example, many of them are susceptible to the shape of the Pareto front and may not work when the Pareto front is concave or disconnected. Others require differentiability of the objective functions and the constraints. Also, most of them only generate a single solution from each run. Thus, several runs (using different starting points) are required in order to generate several elements of the Pareto optimal set [59]. In contrast, evolutionary algorithms deal simultaneously with a set of possible solutions (the so-called population) which allows us to find several members of the Pareto optimal set in a single run of the algorithm. Additionally, evolutionary algorithms are less susceptible to the shape or continuity of the Pareto front (e.g., they can easily deal with discontinuous and concave Pareto fronts).

4 The Origins of the Field

The first hint regarding the possibility of using evolutionary algorithms to solve a MOP appears in a PhD thesis from 1967 [74] in which, however, no actual multi-objective evolutionary algorithm (MOEA) was developed (the multi-objective problem was restated as a single-objective problem and solved with a genetic algorithm). Although there is a rarely mentioned attempt to use a genetic algorithm to solve a multi-objective optimization problem from 1983 (see [43]), David Schaffer is normally considered to be the first to have designed a MOEA during the mid-1980s [76, 77]. Schaffer's approach, called Vector Evaluated Genetic Algorithm (VEGA) consists of a simple genetic algorithm with a modified selection mechanism. At each generation, a number of sub-populations were generated by performing proportional selection according to each objective function in turn. These sub-populations would then be shuffled together to obtain a new population, on which the GA would apply the crossover and mutation operators in the usual way. VEGA had a number of problems, from which the main one had to do with its inability to retain solutions with acceptable performance, perhaps above average, but not outstanding for any of the objective functions. These solutions were perhaps good candidates for becoming nondominated solutions, but could not survive under the selection scheme of this approach.

5 The First Generation: Emphasis on Simplicity

After VEGA, researchers adopted for several years other naive approaches. The most popular were the linear aggregating functions, which consists in adding all the objective functions into a single value which is directly adopted as the fitness of an evolutionary algorithm [26]. Nonlinear aggregating functions were also popular [6]. Lexicographic ordering was another interesting choice. In this case, a single objective (which is considered the most important) is chosen and optimized without considering any of the others. Then, the second objective is optimized but without decreasing the quality of the solution obtained for the first objective. This process is repeated for all the remaining objectives [33].

Despite all these early efforts, the direct incorporation of the concept of Pareto optimality into an evolutionary algorithm was first hinted by David E. Goldberg in his seminal book on genetic algorithms [35]. While criticizing Schaffer's VEGA, Goldberg suggested the use of nondominated ranking and selection to move a population towards the Pareto front in a multi-objective optimization problem. The basic idea is to find the set of solutions in the population that are Pareto nondominated by the rest of the population. These solutions are then assigned the highest rank and eliminated from further contention. Another set of Pareto nondominated solutions is determined from the remaining population and are assigned the next highest rank. This process continues until all the population is suitably ranked. Goldberg also suggested the use of some kind of niching technique to keep the GA from converging to a single point on the front. A niching mechanism such as fitness sharing [36] would allow the evolutionary algorithm to maintain individuals all along the nondominated frontier. Goldberg did not provide an actual implementation of his procedure, but practically all the MOEAs developed after the publication of his book were influenced by his ideas. From the several MOEAs developed from 1989 to 1998, the most representative are the following:

1. **Nondominated Sorting Genetic Algorithm** (NSGA): This algorithm was proposed by Srinivas and Deb [80], and was the first to be published in a specialized journal (*Evolutionary Computation*). The NSGA is based on several layers of classifications of the individuals as suggested by Goldberg [35]. Before selection is performed, the population is ranked on the basis of nondomination: all nondominated individuals are classified into one category (with a dummy fitness value, which is proportional to the population size, to provide an equal reproductive potential for these individuals). To maintain the diversity of the population, these classified individuals are shared with their dummy fitness values. Then this group of classified individuals is ignored and another layer of nondominated individuals is considered. The process continues until all individuals in the population are classified. Since individuals in the first front have the maximum fitness value, they always get more copies than the rest of the population. The algorithm of the NSGA is not very efficient, because Pareto ranking has to be repeated over an over again. Evidently, it is possible to achieve the same goal in a more efficient way.

2. **Niched-Pareto Genetic Algorithm** (NPGA): Proposed in [40]. The NPGA uses a tournament selection scheme based on Pareto dominance. The basic idea of the algorithm is quite clever: Two individuals are randomly chosen and compared against a subset from the entire population (typically, around 10% of the population). If one of them is dominated (by the individuals randomly chosen from the

Figure 1: Masahiro Tanaka.

population) and the other is not, then the nondominated individual wins. All the other situations are considered a tie (i.e., both competitors are either dominated or nondominated). When there is a tie, the result of the tournament is decided through fitness sharing.

3. **Multi-Objective Genetic Algorithm** (MOGA): Proposed in [29]. In MOGA, the rank of a certain individual corresponds to the number of chromosomes in the current population by which it is dominated. All nondominated individuals are assigned the highest possible fitness value (all of them get the same fitness, such that they can be sampled at the same rate), while dominated ones are penalized according to the population density of the corresponding region to which they belong (i.e., fitness sharing is used to verify how crowded is the region surrounding each individual).

During the first generation, few people performed comparative studies among different MOEAs. However, those who compared the three previous MOEAs unanimously agreed on the superiority of MOGA, followed by the NPGA and by the NSGA (in a distant third place) [84]. This period was characterized by the simplicity of the algorithms proposed and by the lack of methodology to validate them. No standard test functions were available and comparisons were normally done visually (no performance measures were available).

There is, however, an important result during this period that is normally disregarded. Masahiro Tanaka (see Figure 1) [82] developed the first scheme to incorporate user's preferences into a MOEA. This is a very important topic, since in real-world applications it is normally the case that the user does not need the entire Pareto optimal set, but only a small portion of it (or perhaps only a single solution). Thus, it is normally desirable that the user can define certain preferences that can narrow the search and that can magnify certain portions of the Pareto front. For many years, however, few researchers in this area paid attention to this issue (see for example [12]).

Another important event during the first generation, was the publication of the first survey of the field. Fonseca and Fleming published such survey in the journal *Evolutionary Computation* in 1995 [30]. Carlos M. Fonseca (see Figure 2) also proposed the first performance measure that did not require the true Pareto front of the problem beforehand (see [31]), and was also the first to suggest a way of modifying the Pareto

Figure 2: Carlos M. Fonseca.

Figure 3: Eckart Zitzler.

dominance relationship in order to handle constraints [32]. The main lesson learnt from the first generation was that a successful MOEA had to combine a good mechanism to select nondominated individuals (perhaps, but not necessarily, based on the concept of Pareto optimality) combined with a good mechanism to maintain diversity (fitness sharing was a choice, but not the only one). The question was: can we design more efficient algorithms while keeping at least the effectiveness achieved by first generation MOEAs?

6 The Second Generation: Emphasis on Efficiency

From the author's perspective, a second generation of MOEAs started when elitism became a standard mechanism. Although there were some early studies that considered the notion of elitism in a MOEA (see for example [42]), most authors credit Eckart Zitzler (see Figure 3) with the formal introduction of this concept in a MOEA, mainly because his *Strength Pareto Evolutionary Algorithm* (SPEA) was published in a specialized journal (the *IEEE Transactions on Evolutionary Computation*), [92] which made it a landmark in the field. Needless to say, after the publication of this paper, most researchers in the field started to incorporate external populations in their MOEAs and the use of this mechanism (or an alternative form of elitism) became a common practice. In fact, the use of elitism is a theoretical requirement in order to guarantee convergence of a MOEA and therefore its importance [75].

In the context of multi-objective optimization, elitism usually (although not necessarily) refers to the use of an external population (also called secondary population) to retain the nondominated individuals found along the evolutionary process. The main motivation for this mechanism is the fact that a solution that is nondominated with respect to its current population is not necessarily nondominated with respect to all the populations that are produced by an evolutionary algorithm. Thus, what we need is a way of guaranteeing that the solutions that we will report to the user are nondominated with respect to every other solution that our algorithm has produced. Therefore, the most intuitive way of doing this is by storing in an external memory (or archive) all the nondominated solutions found. If a solution that wishes to enter the archive is dominated by its contents, then it is not allowed to enter. Conversely, if a solution dominates anyone stored in the file, the dominated solution must be deleted. The use of this external file raises several questions:

- How does the external file interact with the main population? In other words, do we select individuals from the union of the main population and the external file?, or do we select only from the main population, ignoring the contents of the external file?

- What do we do when the external file is full (assuming that the capacity of the external file is bounded)? Since memory capabilities are always limited, this issue deserves special attention.

- Do we impose additional criteria for a nondominated solution to be allowed to enter the file rather than just using Pareto dominance (e.g., use the distribution of solutions as an additional criterion)?

These and some other issues related to external archives (also called “elite” archives) have been studied both from an empirical and from a theoretical perspective (see for example [49, 27]). Besides the use of an external file, elitism can also be introduced through the use of a $(\mu+\lambda)$-selection in which parents compete with their children and those which are nondominated (and possibly comply with some additional criterion such as providing a better distribution of solutions) are selected for the following generation. Many MOEAs have been proposed during the second generation (which we are still living today). However, most researchers will agree that few of these approaches have been adopted as a reference or have been used by others. The author considers that the most representative MOEAs of the second generation are the following:

1. **Strength Pareto Evolutionary Algorithm** (SPEA): This algorithm was introduced in [91, 92]. This approach was conceived as a way of integrating different MOEAs. SPEA uses an archive containing nondominated solutions previously found (the so-called *external nondominated set*). At each generation, nondominated individuals are copied to the external nondominated set. For each individual in this external set, a *strength* value is computed. This strength is similar to the ranking value of MOGA [29], since it is proportional to the number of solutions to which a certain individual dominates. In SPEA, the fitness of each

Figure 4: Joshua D. Knowles.

member of the current population is computed according to the strengths of all external nondominated solutions that dominate it. The fitness assignment process of SPEA considers both closeness to the true Pareto front and even distribution of solutions at the same time. Thus, instead of using niches based on distance, Pareto dominance is used to ensure that the solutions are properly distributed along the Pareto front. Although this approach does not require a niche radius, its effectiveness relies on the size of the external nondominated set. In fact, since the external nondominated set participates in the selection process of SPEA, if its size grows too large, it might reduce the selection pressure, thus slowing down the search. Because of this, the authors decided to adopt a technique that prunes the contents of the external nondominated set so that its size remains below a certain threshold.

2. **Strength Pareto Evolutionary Algorithm 2** (SPEA2): SPEA2 has three main differences with respect to its predecessor [89]: (1) it incorporates a fine-grained fitness assignment strategy which takes into account for each individual the number of individuals that dominate it and the number of individuals by which it is dominated; (2) it uses a nearest neighbor density estimation technique which guides the search more efficiently, and (3) it has an enhanced archive truncation method that guarantees the preservation of boundary solutions.

3. **Pareto Archived Evolution Strategy** (PAES): This algorithm was introduced in [52]. PAES consists of a (1+1) evolution strategy (i.e., a single parent that generates a single offspring) in combination with a historical archive that records the nondominated solutions previously found. This archive is used as a reference set against which each mutated individual is being compared. Such a historical archive is the elitist mechanism adopted in PAES. However, an interesting aspect of this algorithm is the procedure used to maintain diversity which consists of a crowding procedure that divides objective space in a recursive manner. Each solution is placed in a certain grid location based on the values of its objectives (which are used as its "coordinates" or "geographical location"). A map of such grid is maintained, indicating the number of solutions that reside in each grid location. Since the procedure is adaptive, no extra parameters are required (except for the number of divisions of the objective space). This adaptive grid (or vari-

ations of it) has been adopted by several modern MOEAs (e.g., [10]), and is the contribution by which Joshua D. Knowles (see Figure 4) is more well-known, although he has made several other significant contributions to the field (e.g., [49, 50, 51]).

4. **Nondominated Sorting Genetic Algorithm II** (NSGA-II): This approach was introduced in [19, 20] as an improved version of the NSGA [80].[1] In the NSGA-II, for each solution one has to determine how many solutions dominate it and the set of solutions to which it dominates. The NSGA-II estimates the density of solutions surrounding a particular solution in the population by computing the average distance of two points on either side of this point along each of the objectives of the problem. This value is the so-called *crowding distance*. During selection, the NSGA-II uses a crowded-comparison operator which takes into consideration both the nondomination rank of an individual in the population and its crowding distance (i.e., nondominated solutions are preferred over dominated solutions, but between two solutions with the same nondomination rank, the one that resides in the less crowded region is preferred). The NSGA-II does not use an external memory as the other MOEAs previously discussed. Instead, the elitist mechanism of the NSGA-II consists of combining the best parents with the best offspring obtained (i.e., a $(\mu + \lambda)$-selection). Due to its clever mechanisms, the NSGA-II is much more efficient (computationally speaking) than its predecessor, and its performance is so good, that it has become very popular in the last few years, becoming a landmark against which other multi-objective evolutionary algorithms have to be compared.

Many other algorithms have been proposed during the second generation (see for example [10, 16, 15]). Also, fitness sharing is no longer the only alternative to maintain diversity, since several other approaches have been proposed. Besides the adaptive grid from PAES, researchers have adopted clustering techniques [61], crowding [20], entropy [48], and geometrically-based approaches [86], among other mechanisms. Additionally, some researchers have also adopted mating restriction schemes [62]. More recently, the use of relaxed forms of Pareto dominance has been adopted as a mechanism to encourage more exploration and, therefore, to provide more diversity. From these mechanisms, ϵ-dominance has become increasingly popular, not only because of its effectiveness, but also because of its sound theoretical foundation [56]. Also, new surveys on evolutionary multi-objective optimization were published during this period, as well as several monographic books [18, 12, 13, 81, 67].

During the second generation, many other aspects were emphasized. The main one has been, with no doubt, efficiency. Researchers raised concerns about efficiency both at an algorithmic level and at the data structures level [45, 51]. The increasing number of publications during the second generation (from the end of 1998 to date) makes us wonder what will the third generation have to offer.

[1]Note however that the differences between the NSGA-II and the NSGA are so significant that they are considered as two completely different algorithms by several researchers.

Figure 5: Kalyanmoy Deb.

7 Methodological Issues

During the second generation, several researchers proposed a variety of performance measures to allow a quantitative (rather than only qualitative) comparison of results [87, 31, 92]. Zitzler et al. [87] stated that, when assessing performance of a MOEA, one was interested in measuring three things:

1. Maximize the number of elements of the Pareto optimal set found.

2. Minimize the distance of the Pareto front produced by our algorithm with respect to the global Pareto front (assuming we know its location).

3. Maximize the spread of solutions found, so that we can have a distribution of vectors as smooth and uniform as possible.

This, however, raised some issues. First, it was required to know beforehand the exact location of the true Pareto front of a problem in order to use a performance measure. This may not be possible in real-world problems in which the location of the true Pareto front is unknown. The second issue was that it is unlikely that a single performance measure can assess the three things indicated by Zitzler et al. [87]. In other words, assessing the performance of a MOEA is, also, a MOP! Once researchers started to propose performance measures, criticisms arose. Some researchers realized (empirically) that many of the new performance measures were biased. In other words, there were cases in which they provided results that did not correspond to what we could see from the graphical representation of the results (see for example [84]). Ironically, many researchers went back to the graphical comparisons when suspected that something was wrong with the numerical results produced from applying the performance measures available. Although slowly, researchers started to proposed a different type of performance measures that considered not one algorithm at a time, but two [31, 92]. These performance measures were called "binary" (in contrast to those that assess performance of a single algorithm at a time, which were called "unary"). In 2002, the truth was finally in the open: Unary performance measures are not compliant with Pareto dominance and, therefore, are not reliable [90, 93]. Not everything is lost, however, since binary performance measures can overcome this limitation [88, 93].

Concurrently with the research on performance measures, other researchers were designing test functions. The most remarkable work in this regard is due to Kalyanmoy Deb (see Figure 5) who proposed in 1999 a methodology to design MOPs that was widely used during several years [17]. Later on, an alternative set of test functions was proposed, but this time, due to their characteristics, no enumerative process was required to generate their true Pareto front [21, 22]. Considering that these test functions are also scalable, their use has become widespread. So, today, researchers in the field normally validate their MOEAs with problems having three or more objective functions, and 10 or more decision variables.

8 Applications

MOEAs have become increasingly popular in a wide variety of application domains, as reflects a recent book entirely devoted to this topic [9]. In order to provide a rough idea of the sort of applications that are being tackled in the current literature, we will classify the applications in three large groups: engineering, industrial and scientific. Some specific areas within each of these groups are indicated next. We will start with the engineering applications, which are, by far, the most popular in the literature, perhaps due to the fact that engineering problems have well-studied mathematical models. A representative sample of engineering applications is the following:

- Electrical engineering [73]
- Hydraulic engineering [63]
- Structural engineering [37]
- Aeronautical engineering [4]
- Robotics and Control [70, 60]

A representative sample of industrial applications is the following:

- Design and manufacture [47]
- Scheduling [39]
- Management [66]

Finally, we have a variety of scientific applications:

- Chemistry [57]
- Physics [58]
- Medicine [55]
- Computer science [34]

The strong interest for using MOEAs in so many different disciplines reinforces the idea stated at the beginning of this article regarding the multi-objective nature of many real-world problems. However, some application domains have received relatively little attention from researchers are represent areas of opportunity. For example: cellular automata [64], pattern recognition [65], data mining [68], bioinformatics [41], and financial applications [78].

9 Current Research Trends

From the author's perspective, researchers haven't produced another breakthrough that is so significant (as elitism) as to redirect most of the research into a new direction. Thus, the third generation is yet to appear. However, there are several interesting ideas that have certainly influenced a lot of the work being done these days and which deserve closer attention. Some examples are the following:

- The use of relaxed forms of Pareto dominance has become popular as a mechanism to regulate convergence of a MOEA. From these mechanisms, ϵ-dominance is, with no doubt, the most popular [56], but it is not the only mechanism of this type (see for example [54]). ϵ-dominance allows to control the granularity of the approximation of the Pareto front obtained. As a consequence, it is possible to accelerate convergence using this mechanism (if we are satisfied with a very coarse approximation of the Pareto front).

- The transformation of single-objective problems into a multi-objective form that somehow facilitates their solution. For example, some researchers have proposed the handling of the constraints of a problem as objectives [8], and others have proposed the so-called "multi-objectivization" by which a single-objective optimization problem is decomposed into several subcomponents considering a multi-objective approach [44, 53]. This procedure has been found to be helpful in removing local optima from a problem and has attracted a lot of attention in the last few years.

- The use of alternative bio-inspired heuristics for multi-objective optimization. The most remarkable examples are particle swarm optimization [46] and differential evolution [71], whose use has become increasingly popular in multi-objective optimization (see for example [1, 11]). However, other bio-inspired algorithms such as artificial immune systems and ant colony optimization have also been used to solve multi-objective optimization problems [7, 38].

10 Future Research Trends

There are several topics involving challenges that will keep busy to the researchers in this area for the next few years. Some of them are the following:

- Parameter control is certainly a topic that has been only scarcely explored in MOEAs. Is it possible to design a MOEA that self-adapts its parameters such that the user doesn't have to fine-tune them by hand? Some researchers have proposed a few self-adaptation and on-line adaptation procedures for MOEAs (see for example [83, 3]), but recently, not much work seems to be going in this direction.

- What is the minimum number of fitness function evaluations that are actually required to achieve a minimum desirable performance with a MOEA? Recently, some researchers have proposed the use of black-box optimization techniques normally adopted in engineering to perform an incredibly low number of fitness function evaluations while still producing reasonably good solutions (see for example [51]). However, this sort of approach is inherently limited to problems of low dimensionality. So, the question is: are there any other ways of reducing the number of evaluations without sacrificing dimensionality?

- The development of implementations of MOEAs that are independent of the platform and programming language in which they were developed is an important step towards a common platform that can be used to validate new algorithms. In this direction, PISA (A Platform and programming language independent Interface for Search Algorithms) [2] constitutes an important step, and more work is expected in this direction.

- How to deal with problems that have "many" objectives? Some recent studies have shown that traditional Pareto ranking schemes do not behave well in the presence of many objectives (where "many" is normally a number above 3 or 4) [72].

- There are plenty of fundamental questions that remain unanswered. For example: what are the sources of difficulty of a multi-objective optimization problem for a MOEA? What are the dimensionality limitations of current MOEAs? Can we use alternative mechanisms into an evolutionary algorithm to generate nondominated solutions without relying on Pareto ranking (e.g., adopting concepts from game theory [79])?

11 Conclusions

Using as a basis a historical framework, we have attempted to provide a general overview of the work that has been done in the last twenty years in evolutionary multi-objective optimization. Many details were left out due to obvious space limitations, but the most significant achievements to date were at least mentioned. Our discussion included algorithms, methods to maintain diversity, methodological issues and applications. Some of the most representative current research trends were also discussed, and in the last part of the article, the author provided his own insights regarding the future of the field. We are still awaiting for the third generation to arrive. As more papers get published in this area[2], it gets harder every day to produce new contributions that are truly sig-

[2] The author maintains the EMOO repository, which currently holds over 2100 bibliographic references. The EMOO repository is located at: `http://delta.cs.cinvestav.mx/~EMOO`

nificant. So, we wonder what sort of change will make it possible to shift the research trends in an entirely new direction.

Acknowledgements

The author acknowledges support from CONACyT through project 42435-Y.

References

[1] Hussein A. Abbass and Ruhul Sarker. The Pareto Differential Evolution Algorithm. *International Journal on Artificial Intelligence Tools*, 11(4):531–552, 2002.

[2] Stefan Bleuler, Marco Laumanns, Lothar Thiele, and Eckart Zitzler. PISA—A Platform and Programming Language Independent Interface for Search Algorithms. In Carlos M. Fonseca et al., editor, *Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003*, pages 494–508, Faro, Portugal, April 2003. Springer. Lecture Notes in Computer Science. Volume 2632.

[3] Dirk Büche, Sibylle Müller, and Petro Koumoutsakos. Self-Adaptation for Multi-objective Evolutionary Algorithms. In Carlos M. Fonseca et al., editor, *Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003*, pages 267–281, Faro, Portugal, April 2003. Springer. Lecture Notes in Computer Science. Volume 2632.

[4] Kazuhisa Chiba, Shigeru Obayashi, Kazuhiro Nakahashi, and Hiroyuki Morino. High-Fidelity Multidisciplinary Design Optimization of Wing Shape for Regional Jet Aircraft. In Carlos A. Coello Coello et al., editor, *Evolutionary Multi-Criterion Optimization. Third International Conference, EMO 2005*, pages 621–635, Guanajuato, México, March 2005. Springer. Lecture Notes in Computer Science Vol. 3410.

[5] Carlos A. Coello Coello. A Comprehensive Survey of Evolutionary-Based Multiobjective Optimization Techniques. *Knowledge and Information Systems. An International Journal*, 1(3):269–308, August 1999.

[6] Carlos A. Coello Coello and Alan D. Christiansen. Two New GA-based methods for multiobjective optimization. *Civil Engineering Systems*, 15(3):207–243, 1998.

[7] Carlos A. Coello Coello and Nareli Cruz Cortés. Solving Multiobjective Optimization Problems using an Artificial Immune System. *Genetic Programming and Evolvable Machines*, 6(2):163–190, June 2005.

[8] Carlos A. Coello Coello and Arturo Hernández Aguirre. Design of Combinational Logic Circuits through an Evolutionary Multiobjective Optimization Approach. *Artificial Intelligence for Engineering, Design, Analysis and Manufacture*, 16(1):39–53, January 2002.

[9] Carlos A. Coello Coello and Gary B. Lamont, editors. *Applications of Multi-Objective Evolutionary Algorithms*. World Scientific, Singapore, 2004. ISBN 981-256-106-4.

[10] Carlos A. Coello Coello and Gregorio Toscano Pulido. Multiobjective Optimization using a Micro-Genetic Algorithm. In Lee Spector et al., editor, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'2001)*, pages 274–282, San Francisco, California, 2001. Morgan Kaufmann Publishers.

[11] Carlos A. Coello Coello, Gregorio Toscano Pulido, and Maximino Salazar Lechuga. Handling Multiple Objectives With Particle Swarm Optimization. *IEEE Transactions on Evolutionary Computation*, 8(3):256–279, June 2004.

[12] Carlos A. Coello Coello, David A. Van Veldhuizen, and Gary B. Lamont. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Kluwer Academic Publishers, New York, May 2002. ISBN 0-3064-6762-3.

[13] Yann Collette and Patrick Siarry. *Multiobjective Optimization. Principles and Case Studies*. Springer, August 2003.

[14] David Corne, Marco Dorigo, and Fred Glover, editors. *New Ideas in Optimization*. McGraw-Hill, London, UK, 1999.

[15] David W. Corne, Nick R. Jerram, Joshua D. Knowles, and Martin J. Oates. PESA-II: Region-based Selection in Evolutionary Multiobjective Optimization. In Lee Spector et al., editor, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'2001)*, pages 283–290, San Francisco, California, 2001. Morgan Kaufmann Publishers.

[16] David W. Corne, Joshua D. Knowles, and Martin J. Oates. The Pareto Envelope-based Selection Algorithm for Multiobjective Optimization. In Marc Schoenauer et al., editor, *Proceedings of the Parallel Problem Solving from Nature VI Conference*, pages 839–848, Paris, France, 2000. Springer. Lecture Notes in Computer Science No. 1917.

[17] Kalyanmoy Deb. Multi-Objective Genetic Algorithms: Problem Difficulties and Construction of Test Problems. *Evolutionary Computation*, 7(3):205–230, Fall 1999.

[18] Kalyanmoy Deb. *Multi-Objective Optimization using Evolutionary Algorithms*. John Wiley & Sons, Chichester, UK, 2001. ISBN 0-471-87339-X.

[19] Kalyanmoy Deb, Samir Agrawal, Amrit Pratab, and T. Meyarivan. A Fast Elitist Non-Dominated Sorting Genetic Algorithm for Multi-Objective Optimization: NSGA-II. In Marc Schoenauer et al., editor, *Proceedings of the Parallel Problem Solving from Nature VI Conference*, pages 849–858, Paris, France, 2000. Springer. Lecture Notes in Computer Science No. 1917.

[20] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. A Fast and Elitist Multiobjective Genetic Algorithm: NSGA–II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, April 2002.

[21] Kalyanmoy Deb, Lothar Thiele, Marco Laumanns, and Eckart Zitzler. Scalable Multi-Objective Optimization Test Problems. In *Congress on Evolutionary Computation (CEC'2002)*, volume 1, pages 825–830, Piscataway, New Jersey, May 2002. IEEE Service Center.

[22] Kalyanmoy Deb, Lothar Thiele, Marco Laumanns, and Eckart Zitzler. Scalable Test Problems for Evolutionary Multiobjective Optimization. In Ajith Abraham, Lakhmi Jain, and Robert Goldberg, editors, *Evolutionary Multiobjective Optimization. Theoretical Advances and Applications*, pages 105–145. Springer, USA, 2005.

[23] Francis Ysidro Edgeworth. *Mathematical Physics*. P. Keagan, London, England, 1881.

[24] Matthias Ehrgott. *Multicriteria Optimization*. Springer, second edition, 2005.

[25] A.E. Eiben and J.E. Smith. *Introduction to Evolutionary Computing*. Springer, Berlin, 2003. ISBN 3-540-40184-9.

[26] Neil H. Eklund and Mark J. Embrechts. GA-Based Multi-Objective Optimization of Visible Spectra for Lamp Design. In Cihan H. Dagli et al., editor, *Smart Engineering System Design: Neural Networks, Fuzzy Logic, Evolutionary Programming, Data Mining and Complex Systems*, pages 451–456, New York, November 1999. ASME Press.

[27] M. Fleischer. The Measure of Pareto Optima. Applications to Multi-objective Metaheuristics. In Carlos M. Fonseca et al., editor, *Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003*, pages 519–533, Faro, Portugal, April 2003. Springer. Lecture Notes in Computer Science. Volume 2632.

[28] Lawrence J. Fogel. *Artificial Intelligence through Simulated Evolution. Forty Years of Evolutionary Programming*. John Wiley & Sons, Inc., New York, 1999.

[29] Carlos M. Fonseca and Peter J. Fleming. Genetic Algorithms for Multiobjective Optimization: Formulation, Discussion and Generalization. In Stephanie Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 416–423, San Mateo, California, 1993. University of Illinois at Urbana-Champaign, Morgan Kauffman Publishers.

[30] Carlos M. Fonseca and Peter J. Fleming. An Overview of Evolutionary Algorithms in Multiobjective Optimization. *Evolutionary Computation*, 3(1):1–16, Spring 1995.

[31] Carlos M. Fonseca and Peter J. Fleming. On the Performance Assessment and Comparison of Stochastic Multiobjective Optimizers. In Hans-Michael Voigt et al., editor, *Parallel Problem Solving from Nature—PPSN IV*, Lecture Notes in Computer Science, pages 584–593, Berlin, Germany, September 1996. Springer-Verlag.

[32] Carlos M. Fonseca and Peter J. Fleming. Multiobjective Optimization and Multiple Constraint Handling with Evolutionary Algorithms—Part I: A Unified Formulation. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*, 28(1):26–37, 1998.

[33] Louis Gacôgne. Research of Pareto Set by Genetic Algorithm, Application to Multicriteria Optimization of Fuzzy Controller. In *5th European Congress on Intelligent Techniques and Soft Computing EUFIT'97*, pages 837–845, Aachen, Germany, September 1997.

[34] Nicolás García-Pedrajas, César Hervás-Martínez, and Domingo Ortiz-Boyer. Cooperative Coevolution of Artificial Neural Network Ensembles for Pattern Classification. *IEEE Transactions on Evolutionary Computation*, 9(3):271–302, June 2005.

[35] David E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1989.

[36] David E. Goldberg and Jon Richardson. Genetic algorithm with sharing for multimodal function optimization. In John J. Grefenstette, editor, *Genetic Algorithms and Their Applications: Proceedings of the Second International Conference on Genetic Algorithms*, pages 41–49, Hillsdale, New Jersey, 1987. Lawrence Erlbaum.

[37] David Greiner, Gabriel Winter, José M. Emperador, and Blas Galván. Gray Coding in Evolutionary Multicriteria Optimization: Application in Frame Structural Optimum Design. In Carlos A. Coello Coello et al., editor, *Evolutionary Multi-Criterion Optimization. Third International Conference, EMO 2005*, pages 576–591, Guanajuato, México, March 2005. Springer. Lecture Notes in Computer Science Vol. 3410.

[38] Michael Guntsch. *Ant Algorithms in Stochastic and Multi-Criteria Environments*. PhD thesis, Department of Economics and Business Engineering, University of Karlsruhe, Germany, 2004.

[39] T. Hanne and S. Nickel. A multiobjective evolutionary algorithm for scheduling and inspection planning in software development projects. *European Journal of Operational Research*, 167(3):663–678, December 2005.

[40] Jeffrey Horn, Nicholas Nafpliotis, and David E. Goldberg. A Niched Pareto Genetic Algorithm for Multiobjective Optimization. In *Proceedings of the First IEEE Conference on Evolutionary Computation, IEEE World Congress on Computational Intelligence*, volume 1, pages 82–87, Piscataway, New Jersey, June 1994. IEEE Service Center.

[41] R.M. Hubley, E. Zitzler, and J.C. Roach. Evolutionary algorithms for the selection of single nucleotide polymorphisms. *BMC Bioinformatics*, 4(30), July 2003.

[42] Phil Husbands. Distributed Coevolutionary Genetic Algorithms for multi-Criteria and Multi-Constraint Optimisation. In Terence C. Fogarty, editor, *Evolutionary Computing. AIS Workshop. Selected Papers*, Lecture Notes in Computer Science Vol. 865, pages 150–165. Springer Verlag, April 1994.

[43] K. Ito, S. Akagi, and M. Nishikawa. A Multiobjective Optimization Approach to a Design Problem of Heat Insulation for Thermal Distribution Piping Network Systems. *Journal of Mechanisms, Transmissions and Automation in Design (Transactions of the ASME)*, 105:206–213, June 1983.

[44] Mikkel T. Jensen. Guiding Single-Objective Optimization Using Multi-objective Methods. In Günther Raidl et al., editor, *Applications of Evolutionary Computing. Evoworkshops 2003: EvoBIO, EvoCOP, EvoIASP, EvoMUSART, EvoROB, and EvoSTIM*, pages 199–210, Essex, UK, April 2003. Springer. Lecture Notes in Computer Science Vol. 2611.

[45] Mikkel T. Jensen. Reducing the Run-Time Complexity of Multiobjective EAs: The NSGA-II and Other Algorithms. *IEEE Transactions on Evolutionary Computation*, 7(5):503–515, October 2003.

[46] James Kennedy and Russell C. Eberhart. *Swarm Intelligence*. Morgan Kaufmann Publishers, San Francisco, California, 2001.

[47] Timoleon Kipouros, Daniel Jaeggi, Bill Dawes, Geoff Parks, and Mark Savill. Multi-objective Optimisation of Turbomachinery Blades Using Tabu Search. In Carlos A. Coello Coello et al., editor, *Evolutionary Multi-Criterion Optimization. Third International Conference, EMO 2005*, pages 897–910, Guanajuato, México, March 2005. Springer. Lecture Notes in Computer Science Vol. 3410.

[48] Hajime Kita, Yasuyuki Yabumoto, Naoki Mori, and Yoshikazu Nishikawa. Multi-Objective Optimization by Means of the Thermodynamical Genetic Algorithm. In Hans-Michael Voigt et al., editor, *Parallel Problem Solving from Nature—PPSN IV*, Lecture Notes in Computer Science, pages 504–512, Berlin, Germany, September 1996. Springer-Verlag.

[49] Joshua Knowles and David Corne. Properties of an Adaptive Archiving Algorithm for Storing Nondominated Vectors. *IEEE Transactions on Evolutionary Computation*, 7(2):100–116, April 2003.

[50] Joshua Knowles and David Corne. Bounded Pareto Archiving: Theory and Practice. In Xavier Gandibleux, Marc Sevaux, Kenneth Sörensen, and Vincent T'kindt, editors, *Metaheuristics for Multiobjective Optimisation*, pages 39–64, Berlin, 2004. Springer. Lecture Notes in Economics and Mathematical Systems Vol. 535.

[51] Joshua Knowles and Evan J. Hughes. Multiobjective Optimization on a Budget of 250 Evaluations. In Carlos A. Coello Coello et al., editor, *Evolutionary Multi-Criterion Optimization. Third International Conference, EMO 2005*, pages 176–190, Guanajuato, México, March 2005. Springer. Lecture Notes in Computer Science Vol. 3410.

[52] Joshua D. Knowles and David W. Corne. Approximating the Nondominated Front Using the Pareto Archived Evolution Strategy. *Evolutionary Computation*, 8(2):149–172, 2000.

[53] Joshua D. Knowles, Richard A. Watson, and David W. Corne. Reducing Local Optima in Single-Objective Problems by Multi-objectivization. In Eckart Zitzler et al., editor, *First International Conference on Evolutionary Multi-Criterion Optimization*, pages 268–282. Springer-Verlag. Lecture Notes in Computer Science No. 1993, 2001.

[54] Ikeda Kokolo, Kita Hajime, and Kobayashi Shigenobu. Failure of Pareto-based MOEAs: Does Non-dominated Really Mean Near to Optimal? In *Proceedings of the Congress on Evolutionary Computation 2001 (CEC'2001)*, volume 2, pages 957–962, Piscataway, New Jersey, May 2001. IEEE Service Center.

[55] Michael Lahanas. Application of Multiobjective Evolutionary Optimization Algorithms in Medicine. In Carlos A. Coello Coello and Gary B. Lamont, editors, *Applications of Multi-Objective Evolutionary Algorithms*, pages 365–391. World Scientific, Singapore, 2004.

[56] Marco Laumanns, Lothar Thiele, Kalyanmoy Deb, and Eckart Zitzler. Combining Convergence and Diversity in Evolutionary Multi-objective Optimization. *Evolutionary Computation*, 10(3):263–282, Fall 2002.

[57] S. Majumdar, K. Mitra, and G. Sardar. Kinetic Analysis & Optimization for the Catalytic Esterification Step of PPT Polymerization. *Macromolecular Theory and Simulations*, 14:49–59, 2005.

[58] Steven Manos, Leon Poladian, Peter Bentley, and Maryanne Large. Photonic Device Design Using Multiobjective Evolutionary Algorithms. In Carlos A. Coello Coello et al., editor, *Evolutionary Multi-Criterion Optimization. Third International Conference, EMO 2005*, pages 636–650, Guanajuato, México, March 2005. Springer. Lecture Notes in Computer Science Vol. 3410.

[59] Kaisa M. Miettinen. *Nonlinear Multiobjective Optimization*. Kluwer Academic Publishers, Boston, Massachusetts, 1999.

[60] A. Molina-Cristobal, I.A. Griffin, P.J. Fleming, and D.H. Owens. Multiobjective Controller Design: Optimising Controller Structure with Genetic Algorithms. In *Proceedings of the 2005 IFAC World Congress on Automatic Control*, Prague, Czech Republic, July 2005.

[61] Sanaz Mostaghim and Jürgen Teich. The Role of ε-dominance in Multi Objective Particle Swarm Optimization Methods. In *Proceedings of the 2003 Congress on Evolutionary Computation (CEC'2003)*, volume 3, pages 1764–1771, Canberra, Australia, December 2003. IEEE Press.

[62] Tadahiko Murata, Hisao Ishibuchi, and Mitsuo Gen. Specification of Genetic Search Directions in Cellular Multi-objective Genetic Algorithms. In Eckart Zitzler et al., editor, *First International Conference on Evolutionary Multi-Criterion Optimization*, pages 82–95. Springer-Verlag. Lecture Notes in Computer Science No. 1993, 2001.

[63] Matteo Nicolini. A Two-Level Evolutionary Approach to Multi-criterion Optimization of Water Supply Systems. In Carlos A. Coello Coello et al., editor, *Evolutionary Multi-Criterion Optimization. Third International Conference, EMO 2005*, pages 736–751, Guanajuato, México, March 2005. Springer. Lecture Notes in Computer Science Vol. 3410.

[64] Gina M. B. Oliveira, José C. Bortot, and Pedro P.B. de Oliveira. Multiobjective evolutionary search for one-dimensional cellular automata in the density classification task. In R. Standish, M. Bedau, and H. Abbass, editors, *Artificial Life VIII: The 8th International Conference on Artificial Life*, pages 202–206, Cambridge, Massachusetts, 2002. MIT Press.

[65] L.S. Oliveira, R. Sabourin, F. Bortolozzi, and C.Y. Suen. Feature Selection Using Multi-Objective Genetic Algorithms for Handwritten Digit Recognition. In *Proceedings of the 16th International Conference on Pattern Recognition (ICPR'2002)*, volume 1, pages 568–571, Quebec City, Canada, August 2002. IEEE Computer Society Press.

[66] M.S. Osman, M.A. Abo-Sinna, and A.A. Mousa. An effective genetic algorithm approach multiobjective resource allocation problems (MORAPs). *Applied Mathematics and Computation*, 163(2):755–768, April 2005.

[67] Andrzej Osyczka. *Evolutionary Algorithms for Single and Multicriteria Design Optimization*. Physica Verlag, Germany, 2002. ISBN 3-7908-1418-0.

[68] Gisele L. Pappa, Alex A. Freitas, and Celso A.A. Kaestner. Multi-Objective Algorithms for Attribute Selection in Data Mining. In Carlos A. Coello Coello and Gary B. Lamont, editors, *Applications of Multi-Objective Evolutionary Algorithms*, pages 603–626. World Scientific, Singapore, 2004.

[69] Vilfredo Pareto. *Cours D'Economie Politique*, volume I and II. F. Rouge, Lausanne, 1896.

[70] Eduardo José Solteiro Pires, Paulo B. de Moura Oliveira, and José António Tenreiro Machado. Multi-objective Genetic Manipulator Trajectory Planner. In Günther R. Raidl et al., editor, *Applications of Evolutionary Computing. Proceedings of Evoworkshops 2004: EvoBIO, EvoCOMNET, EvoHOT, EvoIASP, EvoMUSART, and EvoSTOC*, pages 219–229, Coimbra, Portugal, April 2004. Springer. Lecture Notes in Computer Science Vol. 3005.

[71] Kenneth V. Price. An Introduction to Differential Evolution. In David Corne, Marco Dorigo, and Fred Glover, editors, *New Ideas in Optimization*, pages 79–108. McGraw-Hill, London, UK, 1999.

[72] Robin Charles Purshouse. *On the Evolutionary Optimisation of Many Objectives*. PhD thesis, Department of Automatic Control and Systems Engineering, The University of Sheffield, Sheffield, UK, September 2003.

[73] Francisco Rivas-Dávalos and Malcolm R. Irving. An Approach Based on the Strength Pareto Evolutionary Algorithm 2 for Power Distribution System Planning. In Carlos A. Coello Coello et al., editor, *Evolutionary Multi-Criterion Optimization. Third International Conference, EMO 2005*, pages 707–720, Guanajuato, México, March 2005. Springer. Lecture Notes in Computer Science Vol. 3410.

[74] R. S. Rosenberg. *Simulation of genetic populations with biochemical properties*. PhD thesis, University of Michigan, Ann Harbor, Michigan, 1967.

[75] Günter Rudolph and Alexandru Agapie. Convergence Properties of Some Multi-Objective Evolutionary Algorithms. In *Proceedings of the 2000 Conference on Evolutionary Computation*, volume 2, pages 1010–1016, Piscataway, New Jersey, July 2000. IEEE Press.

[76] J. David Schaffer. *Multiple Objective Optimization with Vector Evaluated Genetic Algorithms*. PhD thesis, Vanderbilt University, 1984.

[77] J. David Schaffer. Multiple Objective Optimization with Vector Evaluated Genetic Algorithms. In *Genetic Algorithms and their Applications: Proceedings of the First International Conference on Genetic Algorithms*, pages 93–100. Lawrence Erlbaum, 1985.

[78] Frank Schlottmann and Detlef Seese. Financial Applications of Multi-Objective Evolutionary Algorithms: Recent Developments and Future Research Directions. In Carlos A. Coello Coello and Gary B. Lamont, editors, *Applications of Multi-Objective Evolutionary Algorithms*, pages 627–652. World Scientific, Singapore, 2004.

[79] M. Sefrioui and J. Periaux. Nash Genetic Algorithms: examples and applications. In *2000 Congress on Evolutionary Computation*, volume 1, pages 509–516, San Diego, California, July 2000. IEEE Service Center.

[80] N. Srinivas and Kalyanmoy Deb. Multiobjective Optimization Using Nondominated Sorting in Genetic Algorithms. *Evolutionary Computation*, 2(3):221–248, Fall 1994.

[81] K.C. Tan, E.F. Khor, and T.H. Lee. *Multiobjective Evolutionary Algorithms and Applications*. Springer-Verlag, London, 2005. ISBN 1-85233-836-9.

[82] Masahiro Tanaka and Tetsuzo Tanino. Global optimization by the genetic algorithm in a multiobjective decision support system. In *Proceedings of the 10th International Conference on Multiple Criteria Decision Making*, volume 2, pages 261–270, 1992.

[83] Gregorio Toscano Pulido and Carlos A. Coello Coello. The Micro Genetic Algorithm 2: Towards Online Adaptation in Evolutionary Multiobjective Optimization. In Carlos M. Fonseca et al., editor, *Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003*, pages 252–266, Faro, Portugal, April 2003. Springer. Lecture Notes in Computer Science. Volume 2632.

[84] David 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.

[85] David A. Van Veldhuizen and Gary B. Lamont. Multiobjective Evolutionary Algorithms: Analyzing the State-of-the-Art. *Evolutionary Computation*, 8(2):125–147, 2000.

[86] Mario Alberto Villalobos-Arias, Gregorio Toscano Pulido, and Carlos A. Coello Coello. A Proposal to Use Stripes to Maintain Diversity in a Multi-Objective Particle Swarm Optimizer. In *2005 IEEE Swarm Intelligence Symposium (SIS'05)*, pages 22–29, Pasadena, California, USA, June 2005. IEEE Press.

[87] Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. *Evolutionary Computation*, 8(2):173–195, Summer 2000.

[88] Eckart Zitzler, Marco Laumanns, and Stefan Bleuler. A Tutorial on Evolutionary Multiobjective Optimization. In Xavier Gandibleux et al., editor, *Metaheuristics for Multiobjective Optimisation*, pages 3–37, Berlin, 2004. Springer. Lecture Notes in Economics and Mathematical Systems Vol. 535.

[89] Eckart Zitzler, Marco Laumanns, and Lothar Thiele. SPEA2: Improving the Strength Pareto Evolutionary Algorithm. In K. Giannakoglou et al., editor, *EUROGEN 2001. Evolutionary Methods for Design, Optimization and Control with Applications to Industrial Problems*, pages 95–100, Athens, Greece, 2002.

[90] Eckart Zitzler, Marco Laumanns, Lothar Thiele, Carlos M. Fonseca, and Viviane Grunert da Fonseca. Why Quality Assessment of Multiobjective Optimizers Is Difficult. In W.B. Langdon et al., editor, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'2002)*, pages 666–673, San Francisco, California, July 2002. Morgan Kaufmann Publishers.

[91] Eckart Zitzler and Lothar Thiele. Multiobjective Optimization Using Evolutionary Algorithms—A Comparative Study. In A. E. Eiben, editor, *Parallel Problem Solving from Nature V*, pages 292–301, Amsterdam, September 1998. Springer-Verlag.

[92] Eckart Zitzler and Lothar Thiele. Multiobjective Evolutionary Algorithms: A Comparative Case Study and the Strength Pareto Approach. *IEEE Transactions on Evolutionary Computation*, 3(4):257–271, November 1999.

[93] Eckart Zitzler, Lothar Thiele, Marco Laumanns, Carlos M. Fonseca, and Viviane Grunert da Fonseca. Performance Assessment of Multiobjective Optimizers: An Analysis and Review. *IEEE Transactions on Evolutionary Computation*, 7(2):117–132, April 2003.