20 Years of Evolutionary Multi-Objective Optimization: What Has Been Done and What Remains To Be Done

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

June 18, 2006

Abstract

Evolutionary algorithms have been successfully used to solve problems with 2 or more objective functions (called "multi-objective") during the last 20 years.

This field is now called "Evolutionary Multi-Objective Optimization" and has become a very active research area, giving rise to a wide variety of algorithms, techniques to maintain diversity, selection mechanisms, archiving schemes, and applications, among other important contributions. In this paper, we will provide a general overview of this area, emphasizing the main research findings that have shaped the field, as well as its current research trends and its future challenges.

1 Introduction

Problems with two or more objectives (called "multi-objective" or "multi-criteria") are very common in engineering and many other disciplines. The solution of such problems is difficult because their objectives tend to be in conflict with each other, which makes necessary a new notion of optimality.

In the late XIX century, a notion of optimality was developed for these problems, in the context of economics. Later on, such a notion was formally introduced in Operations Research, originating several methods to solve multi-objective problems. Over the years, this research area grew until become practically a separate branch of Operations Research.

Historically, evolutionary algorithms' experts first met multi-objective optimization in the 1960s. However, the first actual algorithmic contribution came until 1985. Since then, this field (now called "evolutionary multi-objective optimization", or EMO) has experienced a significant growth.

This paper presents a general overview of EMO, seen from the perspective of the key contributions that have shaped this field. The remainder of this paper is organized as follows. Section 2 presents some basic concepts required to make this paper self-contained. In Section 3, we describe the origins of multi-objective optimization, and a brief motivation of the use of evolutionary algorithms in this area. Section 4 describes the initial period of EMO, which consists of approximately 13 years, and includes approaches characterized by their simplicity. Section 5 describes the second period, which includes our current days. Finally, Section 6 provides some of the topics that, from the author's perspective, will keep busy to EMO researchers in the next few years. Some conclusions are drawn in Section 7.

2 Basic Concepts

The emphasis of this paper is the solution of multiobjective optimization problems (MOPs) of the form:

$$\text{minimize } [f_1(\vec{x}), f_2(\vec{x}), \ldots, f_k(\vec{x})] \tag{1}$$

subject to the m inequality constraints:

$$g_i(\vec{x}) \leq 0 \quad i = 1, 2, \ldots, m \tag{2}$$

and the p equality constraints:

$$h_i(\vec{x}) = 0 \quad i = 1, 2, \ldots, p \tag{3}$$

where k is the number of objective functions $f_i : \mathbb{R}^n \to \mathbb{R}$. We call $\vec{x} = [x_1, x_2, \ldots, x_n]^T$

the vector of decision variables. We wish to determine from among the set $\mathcal{F}$ of all vectors which satisfy (2) and (3) the particular set of values $x_1^*, x_2^*, \ldots, x_n^*$ which yield the optimum values of all the objective functions.

2.1 Pareto optimality

It is rarely the case that there is a single point that simultaneously optimizes all the objective functions. Therefore, we normally look for "trade-offs", rather than single solutions when dealing with multiobjective optimization problems. The notion of "optimality" is therefore, different. The most commonly adopted notion of optimality is that originally proposed by Francis Ysidro Edgeworth [23] and later generalized by Vilfredo Pareto [58]. Although some authors call this notion *Edgeworth-Pareto optimality* (see for example [68]), we will use the most commonly accepted term: *Pareto optimality*.

We say that a vector of decision variables $\vec{x}^* \in \mathcal{F}$ is *Pareto optimal* if there does not exist another $\vec{x} \in \mathcal{F}$ such that $f_i(\vec{x}) \leq f_i(\vec{x}^*)$ for all $i = 1, \ldots, k$ and $f_j(\vec{x}) < f_j(\vec{x}^*)$ for at least one j.

In words, this definition says that $\vec{x}^*$ is Pareto optimal if there exists no feasible vector of decision variables $\vec{x} \in \mathcal{F}$ which would decrease some criterion without causing a simultaneous increase in at least one other criterion. Unfortunately, this concept almost always gives not a single solution, but rather a set of solutions called the *Pareto optimal set*. The vectors $\vec{x}^*$ correspoding to the solutions included in the Pareto optimal set are called *nondominated*. The image of the Pareto optimal set under the objective functions is called *Pareto front*.

3 The Origins of the Field

The Operations Research (OR) community has developed approaches to solve MOPs since the late 1950s. Currently, a wide variety of mathematical programming techniques designed to solve MOPs are available in the OR literature (see for example [52, 24]). However, mathematical programming techniques have certain limitations when tackling MOPs. For example, many of them are susceptible to the shape and/or continuity 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. Additionally, an initial point is required to perform a run of a mathematical programming technique, and the type of search normally performed is such that the final solution is located relatively close to this initial point. Thus, mathematical programming techniques are normally very susceptible to their initialization. Also, the outcome of each run of a mathematical programming technique is normally a single nondominated solution. So, in order to obtain several elements of the Pareto optimal set, several runs, departing from different initial points, are required [52].

Evolutionary Algorithms (EAs) have been found to be very successful in a variety of (single-objective) optimization problems [37, 29, 66, 25]. The use of EAs for solving MOPs seems like a natural choice if we consider some of their main features. EAs operate on a set of solutions (called "population"), which make us think of the possibility of finding several members of the Pareto optimal set in a single run of an EA. Additionally, EAs are also less susceptible to the shape or continuity of the Pareto front (i.e., they can easily deal with discontinuous and concave Pareto fronts), and do not require any information about the derivatives of the objectives or the constraints.

The first hint regarding the possibility of using evolutionary algorithms to solve a MOP appears in a PhD thesis from 1967 [60] 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 [42]), David Schaffer is normally considered to be the first to have designed a MOEA during the mid-1980s [64, 65]. 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.

4 The Old Days: Naive and Simple Approaches

After VEGA, researchers adopted for several years other naive approaches. The most popular were the **linear aggregating functions**, which consist in adding all the objective functions into a single value which is directly adopted as the fitness of an evolutionary algorithm [22, 26]. This sort of aggregating approaches are, in fact, the oldest mathematical programming methods used for multi-objective optimization, since they can be derived from the Kuhn-Tucker conditions for nondominated solutions [50].

Nonlinear aggregating functions were also popular [77, 62, 39, 4], but were severely criticized despite the fact that they normally do not have the main limitation of linear aggregating techniques (i.e., nonlinear aggregating functions can normally generate non-convex portions of the Pareto front, whereas linear aggregating functions cannot).

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 [34]. Lexicographic ordering is still used today, particularly in applications in which certain objective is known to be more important than the others (see for example [35, 56]).

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 [37]. Goldberg suggested the use of nondominated ranking and selection to move a population toward the Pareto front in a multiobjective optimization problem. This mechanism was called **Pareto ranking**. The basic idea is to find the set of strings in the population that are Pareto nondominated by the rest of the population. These strings are then assigned the highest rank and eliminated from further contention. Another set of Pareto nondominated strings are determined from the remaining population and are assigned the next highest rank. This process continues until 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 [15]. A niching mechanism such as fitness sharing [38] would allow the EA to maintain individuals all along the nondominated frontier. The basic expression adopted in fitness sharing is the following:

$$\phi(d_{ij}) = \begin{cases} 1 - \left(\frac{d_{ij}}{\sigma_{sh}}\right)^{\alpha}, & d_{ij} < \sigma_{share} \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

where $\alpha = 1$, d_{ij} indicates the distance between solutions i and j, and σ_{share} is the niche radius (or sharing threshold). By using this parameter, the fitness of the individual i is modified as:

$$f_{s_i} = \frac{f_i}{\sum_{j=1}^{M} \phi(d_{ij})} \tag{5}$$

where M is the number of individuals that are located in the neighborhood of the i-th individual.

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.

Besides **VEGA**, the most remarkable MOEAs from the early days of evolutionary multi-objective optimization are: the Nondominated Sorting Genetic Algorithm (**NSGA**) [67], the Niched-Pareto Genetic Algorithm (**NPGA**) [40], and the Multi-Objective Genetic Algorithm (**MOGA**) [30].

During this early period, few researchers reported comparative studies among different MOEAs, since the main focus was to introduce new approaches which were normally compared to single-objective EAs. However, those who compared VEGA, NSGA, NPGA and MOGA unanimously agreed on the superiority of MOGA, followed by the NPGA, the NSGA, and VEGA (which every other MOEA could outperform) [10, 72]. These early days were characterized by the design of simple (and even naive) algorithms, the lack of a methodology to validate them, and the lack of a benchmark that other researchers could use as a reference. Comparisons were visual in most cases and most problems tackled were bi-objective.

One of the most remarkable outcomes of these early days was the development of the first scheme to incorporate user's preferences into a MOEA, which is due to Masahiro Tanaka [69]. The incorporation of user's preferences into a MOEA is a topic commonly disregarded in the evolutionary multi-objective optimization literature (even today), but it's a very important issue when dealing with real-world applications [9]. It is normally the case that in real-world problems, the entire Pareto front is not needed, but only a portion of it. So, if we knew the sort of trade-offs that the user requires, it would be possible to magnify the portions of the Pareto front that the user is more interested on.

Another important event during these early days was the publication of the first survey of the field. Fonseca and Fleming published such a survey in the journal *Evolutionary Computation* in 1995 [31]. Less known are two other important contributions from Carlos M. Fonseca that was made in those days: (1) proposed the first performance measure that did not require the true Pareto front of the problem beforehand (he called it "attainment surfaces") [32], and (2) he was the first to suggest a way of modifying the Pareto dominance relationship in order to handle constraints [33].

5 The Second Period: The Growing Pains

Towards the end of the 1990s, things started to change regarding the trends in evolutionary multi-objective optimization. However, the changes were so fast, and some of them are still not fully absorbed by researchers working in this field.

In 1998, Eckart Zitzler proposed a MOEA called Strength Pareto Evolutionary Algorithm (*SPEA*) [82]. An extended version of this work was published in 1999 in the *IEEE Transactions on Evolutionary Computation* [83]. This paper is particularly important, because it contains several elements that gave the first indications of the new period coming. First, SPEA popularized the notion of using elitism in MOEAs. The idea of retaining the nondominated solutions found along the evolutionary process (the notion of elitism in evolutionary multi-objective optimization) wasn't new (see for example [41, 57]). However, it was until the publication of SPEA that the use of elitism started to become common.[1] To retain the nondominated solutions previously found, SPEA uses an archive that is called the *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 [30], since it is proportional to the number of solutions to which a certain individual dominates. In SPEA, the fitness of each 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

[1] In fact, the use of elitism is a theoretical requirement in order to guarantee convergence of a MOEA and therefore its importance [61].

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 clustering technique that prunes the contents of the external nondominated set so that its size remains below a certain threshold [53].

Another important aspect of the paper on SPEA is the introduction of two performance measures to allow a comparison of different MOEAs. The notion of using standard test functions was also indicated in the paper (Zitzler adopted 0/1 knapsack problems). This same idea was later developed by Kalyanmoy Deb and by Zitzler himself (in collaboration with other researchers), who provided different methodologies to construct multi-objective test functions [13, 19, 78, 20].

5.1 Performance Measures

Performance measures are, with no doubt, an important topic in evolutionary multi-objective optimization. Thus, we will discuss next a little bit about their development. In a paper from 2000, Zitzler et al., summarized the three most important aspects that we aimed to assess when measuring the performance of a MOEA [78]:

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.

Note however, that some of these aspects require that we know beforehand the exact location of the true Pareto front. This is certainly not possible in most real-world problems. Another interesting issue was that, given the nature of the three above aspects, it is unlikely that a single performance measure can assess such aspects at the same time. In other words, assessing the performance of a MOEA is also a MOP!

Performance measures were later studied by David A. Van Veldhuizen who empirically identified several of their main weaknesses [72, 74]. Several researchers realized that most performance measures were biased. In other words, some times they provided results that didn't correspond to what we could see from the graphical representation of the results. 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 [32, 83]. These performance measures were called "binary" (in contrast to those that assess performance of a single algorithm at a time, which were called "unary"). The first formal studies of performance assessment measures were published in 2002 [47, 81], and we soon found out what was wrong with some of them: Unary performance measures are not compliant with Pareto dominance and, therefore, are not reliable for assessing performance [81, 84]. Not everything is lost, however, since binary performance measures can overcome this limitation [80, 84].

5.2 Archiving

We have already mentioned the use of an external archive in SPEA, which made popular this form of elitism. The main motivation for adopting a mechanism of this sort 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, the use of such a type of mechanism guarantees that the solutions that we will report to the user are nondominated with respect to every other solution that our algorithm has produced. An archive is the most intuitive way of retaining all the nondominated solutions found along a run of a MOEA. 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. Note however, that the use of this external file raises several questions:

- Is there any interaction between the main population and the external archive (also called "secondary population")?
- Do we impose bounds on the size of the external archive? If so, what do we do when the archive is full?

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 [45, 28]).

An interesting aspect of external archives is that they have also serve as inspiration for the development of new MOEAs. The most remarkable example is the Pareto

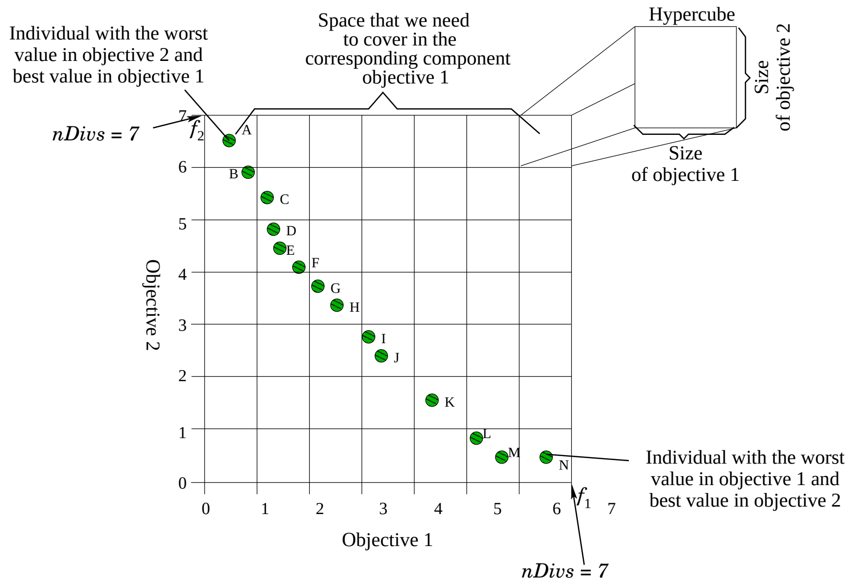

Figure 1: Graphical illustration of the adaptive grid used by PAES.

Archived Evolution Strategy (**PAES**) [48]. This algorithm 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. 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”) as indicated in Figure 1. 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). The adaptive grid of PAES has been adopted (with some variations) by several other modern MOEAs (see for example [12, 11, 7]).

5.3 The NSGA-II

Inspired on some findings from the single-objective optimization literature, some researchers realized that elitism could also be introduced in a MOEA using a plus selection (i.e., to select from the union of parents and offspring). The issue here was how to impose a total order (rather than a partial order) on the population of a MOEA, such that an absolute ranking could be found for the selection of this sort of approach to be effective. Kalyanmoy Deb and his students found a solution in the selection mechanism of the Nondominated Sorting Genetic Algorithm II (**NSGA-II**) [14, 18]. 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).

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. Note however, that the

NSGA-II has some scalability problems (when the number of objectives is increased, its crowding mechanism does not work as well as expected).

5.4 Relaxed Forms of Dominance

More recently, some researchers have proposed the use of relaxed forms of Pareto dominance as a way of regulating convergence of a MOEA [49]. Laumanns et al. [51] proposed a relaxed form of dominance for multi-objective evolutionary algorithms called ε-dominance. This mechanism acts as an archiving strategy to ensure both properties of convergence towards the Pareto-optimal set and properties of diversity among the solutions found. The idea is to use a set of boxes to cover the Pareto front, where the size of such boxes is defined by a user-defined parameter (called ε). Within each box, we only allow a single nondominated solution to be retained (e.g., the one closest to the lower lefthand corner). Thus, by using a large value of ε, the use can accelerate convergence, while sacrificing the quality of the Pareto front obtained. In contrast, if a high-quality of the front is required, then a small value of ε must be adopted. The definition of ε, is then, quite important. However, it's not straightforward to find the most appropriate value of ε when nothing is known in advance about the shape of the Pareto front. Also, to correlate the number of nondominated solutions desired with the value of ε chosen is not easy, and normally some preliminary runs are required in order to estimate the appropriate value. This makes difficult to compare approaches that adopt ε with respect to MOEAs that do not use this concept. Finally, the use of this mechanism naturally eliminates the extreme points of the Pareto front, which may be undesirable in some cases.

Several modern MOEAs have adopted the concept of ε-dominance (see for example [16, 55, 17, 63]). Its limitations have also been recently addressed by some researchers (see for example [75]).

6 The Future of the Field

Many problems still remain to be solved, but they require a higher investment of time than the small problems that were solved during the last ten years. For example, we still do not know what are the actual sources of difficulty that make it hard for a modern MOEA to solve a MOP.

The issue of how to deal with many-objective problems is also worth exploring. Some current research has revealed that a more careful analysis of the Pareto dominance relation is required when dealing with problems that have more than three objectives [59].

New algorithms can be designed, but they require fresh ideas rather than small (and little innovative) changes to existing approaches, which is a common pattern in much of the research that we see nowadays. A few steps in this direction (with nice ideas) have been undertaken with the design of MOEAs that are based on performance measures (see for example [79, 27]). The algorithmic efficiency that share most modern MOEAs should now evolve into new algorithms in which the number of fitness function evaluations is minimized. The use of surrogate techniques is a possible choice (see for example [46, 44, 76]), but it is not the only one (see for example [70]) and much more work in this direction is expected within the next few years.

Parameter control is another issue that certainly deserves attention. Self-adaptation in the context of MOEAs is a very interesting topic that very few researchers have addressed in the specialized literature (see for example [71, 1]).

Twenty years after its inception, evolutionary multi-objective optimization still looks like a healthy research area. However, things are different now of what used to be ten or fifteen years ago. Today, we can say that there is a certain *establishment* that makes necessary to know more about performance measures, test functions and parameters fine-tuning. So, the field may seem less friendly to newcomers than in the old days, but that is only a sign of certain maturity. Ironically, more than ever, we need this “new blood” to challenge the establishment and propose new ideas that can keep this field alive for a long time.

7 Conclusions

This paper does not intend to serve as a survey,[2] but more like a summary of achievements that have shaped this field. Due to space limitations, other interesting topics, such as applications of MOEAs were omitted (see [6], for more on this topic). Also, we didn't mention anything about the (now so popular) multi-objective extensions of other meta-heuristics such as artificial immune systems [5, 43], the ant colony [21, 36], scatter search, and particle swarm optimization [8, 54].

This paper, however, aims to motivate other researchers to get interested in the field. Hopefully, these newcomers will bring their diverse backgrounds into this area,

[2]Surveys are available in [3, 73, 2].

proposing new ideas, relating our ideas to concepts in other fields, and challenging what we believe to be the foundations of this area.

As more and more papers get published in this field,[3] we see more work done by analogy, and less new ideas. We need more significant contributions and more challenging ideas that can constitute trends for others to follow.

Acknowledgements

The author acknowledges support from CONACyT through project 45683-Y.

References

[1] Dirk Büche, Sibylle Müller, and Petro Koumoutsakos. Self-Adaptation for Multi-objective Evolutionary Algorithms. In Carlos M. Fonseca, Peter J. Fleming, Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele, editors, *Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003*, pages 267–281, Faro, Portugal, April 2003. Springer. Lecture Notes in Computer Science. Volume 2632.

[2] 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.

[3] The author maintains the EMOO repository, which currently holds over 2400 bibliographic references. The EMOO repository is located at: `http://delta.cs.cinvestav.mx/˜EMOO`

[3] Carlos A. Coello Coello. Evolutionary multiobjective optimization: A historical view of the field. *IEEE Computational Intelligence Magazine*, 1(1):28–36, February 2006.

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

[5] 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.

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

[7] 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.

[8] 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.

[9] 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.

[10] Carlos Artemio Coello Coello. *An Empirical Study of Evolutionary Techniques for Multiobjective Optimization in Engineering Design*. PhD thesis, Department of Computer Science, Tulane University, New Orleans, LA, April 1996.

[11] 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, Erik D. Goodman, Annie Wu, W.B. Langdon, Hans-Michael Voigt, Mitsuo Gen, Sandip Sen, Marco Dorigo, Shahram Pezeshk, Max H. Garzon, and Edmund Burke, editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'2001)*, pages 283–290, San Francisco, California, 2001. Morgan Kaufmann Publishers.

[12] David W. Corne, Joshua D. Knowles, and Martin J. Oates. The Pareto Envelope-based Selection Algorithm for Multiobjective Optimization. In Marc Schoenauer, Kalyanmoy Deb, Günter Rudolph, Xin Yao, Evelyne Lutton, Juan Julian Merelo, and Hans-Paul Schwefel, editors, *Proceedings of the Parallel Problem Solving from Nature VI Conference*, pages 839–848, Paris, France, 2000. Springer. Lecture Notes in Computer Science No. 1917.

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

[14] 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, Kalyanmoy Deb, Günter Rudolph, Xin Yao,

Evelyne Lutton, Juan Julian Merelo, and Hans-Paul Schwefel, editors, *Proceedings of the Parallel Problem Solving from Nature VI Conference*, pages 849–858, Paris, France, 2000. Springer. Lecture Notes in Computer Science No. 1917.

[15] Kalyanmoy Deb and David E. Goldberg. An Investigation of Niche and Species Formation in Genetic Function Optimization. In J. David Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms*, pages 42–50, San Mateo, California, June 1989. George Mason University, Morgan Kaufmann Publishers.

[16] Kalyanmoy Deb, Manikanth Mohan, and Shikhar Mishra. Towards a Quick Computation of Well-Spread Pareto-Optimal Solutions. In Carlos M. Fonseca, Peter J. Fleming, Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele, editors, *Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003*, pages 222–236, Faro, Portugal, April 2003. Springer. Lecture Notes in Computer Science. Volume 2632.

[17] Kalyanmoy Deb, Manikanth Mohan, and Shikhar Mishra. Evaluating the ϵ-Domination Based Multi-Objective Evolutionary Algorithm for a Quick Computation of Pareto-Optimal Solutions. *Evolutionary Computation*, 13(4):501–525, Winter 2005.

[18] 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.

[19] 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.

[20] 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.

[21] Karl Doerner, Walter J. Gutjahr, Richard F. Hartl, Christine Strauss, and Christian Stummer. Pareto Ant Colony Optimization: A Metaheuristic Approach to Multiobjective Portfolio Selection. *Annals of Operations Research*, 131(1–4):79–99, October 2004.

[22] D. C. Donha, D. S. Desanj, and M. R. Katebi. Genetic Algorithm for Weight Selection in h_∞ Control Design. In Thomas Bäck, editor, *Proceedings of the Seventh International Conference on Genetic Algorithms*, pages 599–606, San Mateo, California, July 1997. Michigan State University, Morgan Kaufmann Publishers.

[23] Francis Ysidro Edgeworth. *Mathematical Psychics*. 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, Anna L. Buczak, Joydeep Ghosh, Mark J. Embrechts, and Okan Ersoy, editors, *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] Michael Emmerich, Nicola Beume, and Boris Naujoks. An EMO Algorithm Using the Hypervolume Measure as Selection Criterion. In Carlos A. Coello Coello, Arturo Hernández Aguirre, and Eckart Zitzler, editors, *Evolutionary Multi-Criterion Optimization. Third International Conference, EMO 2005*, pages 62–76, Guanajuato, México, March 2005. Springer. Lecture Notes in Computer Science Vol. 3410.

[28] M. Fleischer. The Measure of Pareto Optima. Applications to Multi-objective Metaheuristics. In Carlos M. Fonseca, Peter J. Fleming, Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele, editors, *Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003*, pages 519–533, Faro, Portugal, April 2003. Springer. Lecture Notes in Computer Science. Volume 2632.

[29] David B. Fogel. *Evolutionary Computation. Toward a New Philosophy of Machine Intelligence*. The Institute of Electrical and Electronic Engineers, New York, 1995.

[30] 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.

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

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

[33] 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.

[34] M. P. Fourman. Compaction of Symbolic Layout using Genetic Algorithms. In *Genetic Algorithms and their Applications: Proceedings of the First International Conference on Genetic Algorithms*, pages 141–153. Lawrence Erlbaum, 1985.

[35] 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.

[36] Carlos García-Martínez, Oscar Cordón, and Francisco Herrera. An Empirical Analysis of Multiple Objective Ant Colony Optimization Algorithms for the Bi-criteria TSP. In Marco Dorigo, Mauro Birattari, Christian Blum, Luca M. Gambardella, Francesco Mondada, and Thomas Stützle, editors, *Proceedings of the 4th International Workshop on Ant Colony Optimization and Swarm Intelligence*, pages 61–72. Springer. Lecture Notes in Computer Science Vol. 3172, 2004.

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

[38] 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.

[39] P. Hajela and C. Y. Lin. Genetic search strategies in multicriterion optimal design. *Structural Optimization*, 4:99–107, 1992.

[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] 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.

[42] 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.

[43] Licheng Jiao, Maoguo Gong, Ronghua Shang, Haifeng Du, and Bin Lu. Clonal Selection with Immune Dominance and Anergy Based Multiobjective Optimization. In Carlos A. Coello Coello, Arturo Hernández Aguirre, and Eckart Zitzler, editors, *Evolutionary Multi-Criterion Optimization. Third International Conference, EMO 2005*, pages 474–489, Guanajuato, México, March 2005. Springer. Lecture Notes in Computer Science Vol. 3410.

[44] Joshua Knowles. ParEGO: A Hybrid Algorithm With On-Line Landscape Approximation for Expensive Multiobjective Optimization Problems. *IEEE Transactions on Evolutionary Computation*, 10(1):50–66, February 2006.

[45] 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.

[46] Joshua Knowles and Evan J. Hughes. Multiobjective Optimization on a Budget of 250 Evaluations. In Carlos A. Coello Coello, Arturo Hernández Aguirre, and

Eckart Zitzler, editors, *Evolutionary Multi-Criterion Optimization. Third International Conference, EMO 2005*, pages 176–190, Guanajuato, México, March 2005. Springer. Lecture Notes in Computer Science Vol. 3410.

[47] Joshua D. Knowles. *Local-Search and Hybrid Evolutionary Algorithms for Pareto Optimization*. PhD thesis, The University of Reading, Department of Computer Science, Reading, UK, January 2002.

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

[49] 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.

[50] H. W. Kuhn and A. W. Tucker. Nonlinear programming. In J. Neyman, editor, *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 481–492, Berkeley, CA, 1951. University of California Press.

[51] 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.

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

[53] J.N. Morse. Reducing the size of the nondominated set: Pruning by clustering. *Computers and Operations Research*, 7(1–2):55–66, 1980.

[54] Sanaz Mostaghim. *Multi-Objective Evolutionary Algorithms. Data Structures, Convergence, and Diversity*. PhD thesis, Fakultät Elektrotechnik, Informatik und Mathematik der Universität Paderborn, Paderborn, Germany, November 2004.

[55] 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.

[56] Walid El Moudani, Carlos Alberto Nunes Cosenza, Marc de Coligny, and Félix Mora-Camino. A Bi-Criterion Approach for the Airlines Crew Rostering Problem. In Eckart Zitzler, Kalyanmoy Deb, Lothar Thiele, Carlos A. Coello Coello, and David Corne, editors, *First International Conference on Evolutionary Multi-Criterion Optimization*, pages 486–500. Springer-Verlag. Lecture Notes in Computer Science No. 1993, 2001.

[57] Andrzej Osyczka and Sourav Kundu. A Genetic Algorithm Approach to Multicriteria Network Optimization Problems. In *Proceedings of the 20th International Conference on Computers and Industrial Engineering*, pages 329–332, Kyongju, Korea, October 1996.

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

[59] Robin C. Purshouse and Peter J. Fleming. Conflict, Harmony, and Independence: Relationships in Evolutionary Multi-criterion Optimisation. In Carlos M. Fonseca, Peter J. Fleming, Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele, editors, *Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003*, pages 16–30, Faro, Portugal, April 2003. Springer. Lecture Notes in Computer Science. Volume 2632.

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

[61] 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.

[62] Eric Sandgren. Multicriteria design optimization by goal programming. In Hojjat Adeli, editor, *Advances in Design Optimization*, chapter 23, pages 225–265. Chapman & Hall, London, 1994.

[63] Luis Vicente Santana-Quintero and Carlos A. Coello Coello. An Algorithm Based on Differential Evolution for Multi-Objective Problems. *International Journal of Computational Intelligence Research*, 1(2):151–169, 2005.

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

[65] 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.

[66] Hans-Paul Schwefel. *Evolution and Optimum Seeking*. John Wiley & Sons, Inc., New York, 1995.

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

[68] W. Stadler. Fundamentals of multicriteria optimization. In W. Stadler, editor, *Multicriteria Optimization in Engineering and the Sciences*, pages 1–25. Plenum Press, New York, NY, 1988.

[69] 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.

[70] Gregorio Toscano Pulido. *On the Use of Self-Adaptation and Elitism for Multiobjective Particle Swarm Optimization*. PhD thesis, Computer Science Section, Department of Electrical Engineering, CINVESTAV-IPN, Mexico, September 2005.

[71] Gregorio Toscano Pulido and Carlos A. Coello Coello. The Micro Genetic Algorithm 2: Towards Online Adaptation in Evolutionary Multiobjective Optimization. In Carlos M. Fonseca, Peter J. Fleming, Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele, editors, *Evolutionary Multi-Criterion Optimization. Second Inter-*

national Conference, EMO 2003, pages 252–266, Faro, Portugal, April 2003. Springer. Lecture Notes in Computer Science. Volume 2632.

[72] 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.

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

[74] David A. Van Veldhuizen and Gary B. Lamont. On Measuring Multiobjective Evolutionary Algorithm Performance. In *2000 Congress on Evolutionary Computation*, volume 1, pages 204–211, Piscataway, New Jersey, July 2000. IEEE Service Center.

[75] 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.

[76] Ivan Voutchkov and A.J. Keane. Multiobjective Optimization using Surrogates. In I.C. Parmee, editor, *Adaptive Computing in Design and Manufacture. Proceedings of the Seventh International Conference*, pages 167–175, Bristol, UK, April 2006. The Institute for People-centered Computation (IP-CC).

[77] P. B. Wienke, C. Lucasius, and G. Kateman. Multicriteria target optimization of analytical procedures using a genetic algorithm. *Analytica Chimica Acta*, 265(2):211–225, 1992.

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

[79] Eckart Zitzler and Simon Künzli. Indicator-based Selection in Multiobjective Search. In Xin Yao et al., editor, *Parallel Problem Solving from Nature - PPSN VIII*, pages 832–842, Birmingham, UK, September 2004. Springer-Verlag. Lecture Notes in Computer Science Vol. 3242.

[80] Eckart Zitzler, Marco Laumanns, and Stefan Bleuler. A Tutorial on Evolutionary Multiobjective Optimization. In Xavier Gandibleux, Marc Sevaux, Kenneth Sörensen, and Vincent T'kindt, editors, *Metaheuristics for Multiobjective Optimisation*, pages 3–37, Berlin, 2004. Springer. Lecture Notes in Economics and Mathematical Systems Vol. 535.

[81] 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, E. Cantú-Paz, K. Mathias, R. Roy, D. Davis, R. Poli, K. Balakrishnan, V. Honavar, G. Rudolph, J. Wegener, L. Bull, M.A. Potter, A.C. Schultz, J.F. Miller, E. Burke, and N. Jonoska, editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'2002)*, pages 666–673, San Francisco, California, July 2002. Morgan Kaufmann Publishers.

[82] 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.

[83] 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.

[84] 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.