

Evolutionary Multi-Objective Optimization and its Use in Finance

Carlos A. Coello Coello
CINVESTAV-IPN
Evolutionary Computation Group (EVOCINV)
Sección Computación
Av. IPN No. 2508
Col. San Pedro Zacatenco
México, D.F. 07360
MEXICO
Phone: +52 55 5061 3800 x 6564
Fax: +52 55 5061 3757
email: ccoello@cs.cinvestav.mx

Abstract

This chapter provides with a brief introduction of the use of evolutionary algorithms in the solution of multi-objective optimization problems (an area now called “evolutionary multi-objective optimization”). Besides providing some basic concepts and a brief description of the approaches that are more commonly used nowadays, the chapter also provides some of the current and future research trends in the area. In the final part of the chapter, we provide a short description of the sort of applications that multi-objective evolutionary algorithms have found in finance, identifying some possible paths for future research.

keywords: evolutionary multi-objective optimization, evolutionary algorithms, multiobjective optimization, genetic algorithms, evolution strategies.

Introduction

Many real-world problems have two or more objective functions that we aim to minimize. Such problems are called multi-objective optimization problems, and require of an alternative definition of “optimality”. The most common notion of optimality normally adopted is the so-called Pareto optimality, which indicates that the best possible solutions are those representing the best trade-offs among the objective functions. In other words, the

desirable solutions are those in which one objective cannot be improved without worsening another objective.

Evolutionary algorithms (EAs) are techniques based on the emulation of the mechanism of natural selection, which have been successfully used to solve problems during several years (Fogel, 1999; Goldberg, 1989). One of the problem domains in which EAs have been found to be particularly useful is in multi-objective optimization (Coello Coello, Van Veldhuizen, & Lamont, 2002). EAs are particularly suitable for solving multi-objective optimization problems because they deal simultaneously with a set of possible solutions (the so-called population) which allows us to find several members of the Pareto optimal set (i.e., the best possible trade-offs found) in a single run of the algorithm, instead of having to perform a series of separate runs as in the case of the traditional mathematical programming techniques. Additionally, EAs don't require the derivatives of the objective functions and are less susceptible to any features of the problem (e.g., discontinuities either in decision variable space or in objective function space).

The first documented attempt to solve a multi-objective optimization problem using an evolutionary algorithm dates back to the mid-1980s (Schaffer, 1984, 1985). Since then, a considerable amount of research has been done in this area, now known as evolutionary multi-objective optimization (EMO for short). The growing importance of this field is reflected by a significant increment (mainly during the last ten years) of technical papers in international conferences and peer-reviewed journals, special sessions in international conferences and interest groups on the Internet.¹

Basic Concepts

Definition 1 (Global Minimum): *Given a function $f : \Omega \subseteq \mathcal{R}^n \rightarrow \mathcal{R}$, $\Omega \neq \emptyset$, for $\vec{x} \in \Omega$ the value $f^* \triangleq f(\vec{x}^*) > -\infty$ is called a global minimum if and only if*

$$\forall \vec{x} \in \Omega : f(\vec{x}^*) \leq f(\vec{x}) . \quad (1)$$

Then, \vec{x}^ is the global minimum solution, f is the objective function, and the set Ω is the feasible region ($\Omega \in \mathcal{S}$), where \mathcal{S} represents the whole search space. \square*

Definition 2 (General Multi-objective Optimization Problem (MOP)): *Find the vector $\vec{x}^* = [x_1^*, x_2^*, \dots, x_n^*]^T$ which will satisfy the m inequality constraints:*

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

the p equality constraints

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

and will optimize the vector function

$$\vec{f}(\vec{x}) = [f_1(\vec{x}), f_2(\vec{x}), \dots, f_k(\vec{x})]^T \quad (4)$$

¹The author maintains an EMO repository with over 2100 bibliographical entries at: <http://delta.cs.cinvestav.mx/~ccoello/EMO0>

where $\vec{x} = [x_1, x_2, \dots, x_n]^T$ is the vector of decision variables. \square

Definition 3 (Pareto Optimality): A point $\vec{x}^* \in \Omega$ is **Pareto optimal** if for every $\vec{x} \in \Omega$ and $I = \{1, 2, \dots, k\}$ either,

$$\forall_{i \in I} (f_i(\vec{x}^*) \leq f_i(\vec{x})) \quad (5)$$

and, there is at least one $i \in I$ such that

$$f_i(\vec{x}^*) < f_i(\vec{x}) \quad (6)$$

\square

In words, this definition says that \vec{x}^* is Pareto optimal if there exists no feasible vector \vec{x} which would decrease some criterion without causing a simultaneous increase in at least one other criterion. The phrase “Pareto optimal” is considered to mean with respect to the entire decision variable space unless otherwise specified.

Definition 4 (Pareto Dominance): A vector $\vec{u} = (u_1, \dots, u_k)$ is said to *dominate* $\vec{v} = (v_1, \dots, v_k)$ (denoted by $\vec{u} \preceq \vec{v}$) if and only if u is partially less than v , i.e., $\forall i \in \{1, \dots, k\}, u_i \leq v_i \wedge \exists i \in \{1, \dots, k\} : u_i < v_i$. \square

Definition 5 (Pareto Optimal Set): For a given MOP $\vec{f}(x)$, the Pareto optimal set (\mathcal{P}^*) is defined as:

$$\mathcal{P}^* := \{x \in \Omega \mid \neg \exists x' \in \Omega \vec{f}(x') \preceq \vec{f}(x)\}. \quad (7)$$

\square

Pareto optimal solutions are also termed *non-inferior*, *admissible*, or *efficient* solutions (Horn, 1997); their corresponding vectors are termed *nondominated*.

Definition 6 (Pareto Front): For a given MOP $\vec{f}(x)$ and Pareto optimal set \mathcal{P}^* , the Pareto front (\mathcal{PF}^*) is defined as:

$$\mathcal{PF}^* := \{\vec{u} = \vec{f} = (f_1(x), \dots, f_k(x)) \mid x \in \mathcal{P}^*\}. \quad (8)$$

\square

In the general case, it is impossible to find an analytical expression of the line or surface that contains these points. The normal procedure to generate the Pareto front is to compute the feasible points Ω and their corresponding $f(\Omega)$. When there is a sufficient number of these, it is then possible to determine the nondominated points and to produce the Pareto front.

Origins of Evolutionary Multi-objective Optimization

Traditional evolutionary algorithms cannot properly deal with multi-objective optimization problems because of two main reasons:

1. Due to stochastic noise, evolutionary algorithms tend to converge to a single solution if run for a sufficiently large number of iterations. Thus, it is necessary to block the selection mechanism so that different solutions (which are all nondominated) are preserved in the population of an evolutionary algorithm.

2. It is desirable that all nondominated solutions are sampled at the same rate during the selection stage (i.e., that they all are considered with the same survival probability), since all nondominated solutions are equally good among themselves.

Thus, it is necessary to introduce certain modifications into an evolutionary algorithm in order to make it suitable to solve multi-objective optimization problems.

Over the years, there have been many different proposals to extend evolutionary algorithms to solve multi-objective optimization problems. Historically, it is possible to consider three periods:

1. Origins
2. First Generation
3. Second Generation

In this section, we will focus in the first period, and in the two further sections, we will discuss the others.

The first actual implementation of what it is now called a multi-objective evolutionary algorithm (or MOEA, for short) was Schaffer's *Vector Evaluation Genetic Algorithm* (VEGA), which was introduced in the mid-1980s, mainly aimed for solving problems in machine learning (Schaffer, 1984, 1985; Schaffer & Grefenstette, 1985). So, this work is considered as the origin of research in this area.

VEGA basically consisted of a simple genetic algorithm (GA) with a modified selection mechanism. At each generation, a number of sub-populations are generated by performing proportional selection according to each objective function in turn. Thus, for a problem with k objectives, k sub-populations of size N/k each are generated (assuming a total population size of N). These sub-populations are then shuffled together to obtain a new population of size N , on which the GA applies the crossover and mutation operators in the usual way. Schaffer realized that the solutions generated by his approach were nondominated in a local sense, because their nondominance was limited to the current population, which was obviously not appropriate. Also, he noted something that was called "middling" performance.² An individual which had this problem in all the objectives was, perhaps, a good compromise solution, but could not survive under the selection scheme of VEGA, because it wasn't the best in any particular objective. Thus, this problem opposes to the goal of finding Pareto optimal solutions. Although the middling problem can be dealt with using heuristics or other additional mechanisms, it remained as the main drawback of VEGA.

From the second half of the 1980s up to the first half of the 1990s, few other researchers developed MOEAs. Additionally, most of these approaches were very naive and relied on aggregating functions (linear in most cases) (Syswerda & Palmucci, 1991), lexicographic ordering (Fourman, 1985), and target-vector approaches (Wienke, Lucasius, & Kateman, 1992). All of these approaches were strongly influenced by the work done in the operations research community and in most cases did not require any major modifications to the evolutionary algorithm adopted (except for the definition of the fitness function).

Although most of these early MOEAs are rarely referenced in the current literature, this historical period is of great importance because it provided the first insights into the possibility of using evolutionary algorithms for multi-objective optimization. Over the years,

²By "middling", Schaffer meant an individual with acceptable performance, perhaps above average, but not outstanding for any of the objective functions.

researchers would design more sophisticated MOEAs, giving rise to the two generations that are discussed in the two further sections.

The First Generation

The major step towards the first actual generation of MOEAs was given by David E. Goldberg on pages 199 to 201 of his famous book on genetic algorithms published in 1989 (Goldberg, 1989). In his book, Goldberg analyzes VEGA and proposes a selection scheme based on the concept of Pareto optimality. Goldberg not only suggested what would become the standard first generation MOEA, but also indicated that stochastic noise would make such algorithm useless unless some special mechanism was adopted to block convergence. First generation MOEAs typically adopt niching or fitness sharing (Deb & Goldberg, 1989) for that sake. Three are the most representative algorithms from the first generation:

1. **Nondominated Sorting Genetic Algorithm (NSGA)**: This algorithm was proposed by Srinivas and Deb (Srinivas & Deb, 1994). The NSGA is based on several layers of classifications of the individuals as suggested in (Goldberg, 1989). 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. Stochastic remainder proportionate selection is adopted for this technique. Since individuals in the first front have the maximum fitness value, they always get more copies than the rest of the population. This allows to search for nondominated regions, and results in convergence of the population toward such regions. Sharing, by its part, helps to distribute the population over this region (i.e., the Pareto front of the problem).

2. **Niched-Pareto Genetic Algorithm (NPGA)**: Proposed in (Horn, Nafpliotis, & Goldberg, 1994). The NPGA uses a tournament selection scheme based on Pareto dominance. The basic idea of the algorithm is the following: 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 population) and the other is not, then the nondominated individual wins. When both competitors are either dominated or nondominated (i.e., there is a tie), the result of the tournament is decided through fitness sharing (Goldberg & Richardson, 1987). More recently (Erickson, Mayer, & Horn, 2001), the NPGA 2 was proposed. This algorithm relies on a traditional Pareto ranking approach (similar to Fonseca and Fleming's MOGA (Fonseca & Fleming, 1993)), but it keeps its tournament selection scheme. Ties are solved through fitness sharing as in its predecessor. However, the niche count of the NPGA 2 is computed using individuals from the next partially filled generation using a technique called "continuously updated fitness sharing" (Oei, Goldberg, & Chang, 1991).

3. Multi-Objective Genetic Algorithm (MOGA): Proposed in (Fonseca & Fleming, 1993). In MOGA, the rank of a certain individual corresponds to the number of chromosomes in the current population by which it is dominated. Consider, for example, an individual x_i at generation t , which is dominated by $p_i^{(t)}$ individuals in the current generation.

The rank of an individual is given by (Fonseca & Fleming, 1993):

$$\text{rank}(x_i, t) = 1 + p_i^{(t)} \quad (9)$$

All nondominated individuals are assigned rank 1, while dominated ones are penalized according to the population density of the corresponding region of the trade-off surface.

Fitness assignment is performed in the following way (Fonseca & Fleming, 1993):

1. Sort population according to rank.
2. Assign fitness to individuals by interpolating from the best (rank 1) to the worst (rank $n \leq M$) in the way proposed by Goldberg (1989), according to some function, usually linear, but not necessarily.
3. Average the fitnesses of individuals with the same rank, so that all of them are sampled at the same rate. This procedure keeps the global population fitness constant while maintaining appropriate selective pressure, as defined by the function used.

From these 3 algorithms, a few comparative studies undertaken during the mid and late 1990s, indicated that MOGA was the most effective and efficient approach, followed by the NPGA and by the NSGA (in a distant third place) (Coello Coello, 1996; Van Veldhuizen, 1999). This period was characterized by the use of selection mechanisms based on Pareto ranking and by the use of fitness sharing to maintain diversity. The papers of this period normally rely on visual comparisons of results (little work was done regarding the use of performance measures to allow quantitative comparisons of results before the mid-1990s), and normally incorporate very simple test functions.

The Second Generation

The second generation of MOEAs was born with the introduction of the notion of elitism.³ 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. The use of this external file raises several questions:

- How does the external file interact with the main population (e.g., do we select to the union of the main population and the external file)?
- What do we do when the external file is full (assuming that the capacity of the external file is bounded)?
- Do we impose additional criteria to enter the file instead of just using Pareto dominance (e.g., use the distribution of solutions as an additional criterion)?

³Although there were some early studies that considered the notion of elitism in a multi-objective evolutionary algorithm (see for example (Husbands, 1994; Osyczka & Kundu, 1995)), most authors credit Zitzler with the formal introduction of this concept in a multi-objective evolutionary algorithm, mainly because his SPEA was published in a specialized journal (the *IEEE Transactions on Evolutionary Computation*), (Zitzler & Thiele, 1999) 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 multi-objective evolutionary algorithms.

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. Besides the notion of elitism, efficiency (both at an algorithmic level and at the data structures level) has become a concern for researchers in this area (see for example (Jensen, 2003b; Coello Coello & Toscano Pulido, 2001; Mostaghim, Teich, & Tyagi, 2002)). The second generation is also characterized by the use of performance measures to provide a quantitative (rather than only a qualitative) comparison of results (Zitzler, Deb, & Thiele, 2000; Van Veldhuizen & Lamont, 2000b; Fonseca & Fleming, 1996). However, the several drawbacks of many performance measures developed during the second generation (see for example (J. Knowles & Corne, 2002; Zitzler, Laumanns, Thiele, Fonseca, & Grunert da Fonseca, 2002)) have (ironically) brought back to many researchers to adopt visual comparisons as in the origins of the field.

Some MOEAs that are representative of the research trends of the second generation are the following:

1. Strength Pareto Evolutionary Algorithm (SPEA): This algorithm was introduced in (Zitzler & Thiele, 1999). 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 (Fonseca & Fleming, 1993), 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 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. The approach adopted for this sake was a clustering technique called “average linkage method” (Morse, 1980).

2. Strength Pareto Evolutionary Algorithm 2 (SPEA2): SPEA2 has three main differences with respect to its predecessor (Zitzler, Laumanns, & Thiele, 2001): (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 (J. D. Knowles & Corne, 2000). 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).

4. Nondominated Sorting Genetic Algorithm II (NSGA-II): This approach was introduced in (Deb, Agrawal, Pratab, & Meyarivan, 2000; Deb, Pratap, Agarwal, & Meyarivan, 2002) as an improved version of the NSGA (Srinivas & Deb, 1994).⁴ 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 exist (see for example (Coello Coello & Toscano Pulido, 2001; Corne, Knowles, & Oates, 2000; Corne, Jerram, Knowles, & Oates, 2001; Van Veldhuizen & Lamont, 2000a; Zydallis, Lamont, & Veldhuizen, 2000)). The interested reader should consult additional sources for more details (Coello Coello et al., 2002; Coello Coello, 1999; Deb, 2001; Osyczka, 2002; Collette & Siarry, 2003; Tan, Khor, & Lee, 2005).

⁴Note 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.

Current Research Trends

According to the historical view of evolutionary multi-objective optimization presented at the beginning of this chapter, we are currently living the second generation. So far, researchers haven't produced a breakthrough that is so significant as to redirect most of the research into a new direction. However, there are several interesting ideas that have certainly influenced some of the work being done these days. 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 (Laumanns, Thiele, Deb, & Zitzler, 2002). 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 (Coello Coello, 2000b), 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 (Jensen, 2003a; J. D. Knowles, Watson, & Corne, 2001). This procedure has been found to be helpful in removing local optima from a problem.

- The use of alternative bio-inspired heuristics for multi-objective optimization. The most remarkable examples are particle swarm optimization (Kennedy & Eberhart, 2001) and differential evolution (Price, 1999), whose use has become increasingly popular in multi-objective optimization (see for example (Abbass & Sarker, 2002; Coello Coello, Toscano Pulido, & Salazar Lechuga, 2004)). However, other bio-inspired algorithms such as artificial immune systems have also been used to solve multi-objective optimization problems (Coello Coello & Cruz Cortés, 2005).

Future Research Trends

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

- 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) (Purshouse, 2003).

- How to compare (in a quantitative way) the performance of several MOEAs? Despite the existence of a considerable number of performance measures that intend to compare (in a quantitative way) the performance of several MOEAs, many of them are not appropriate because their definition is not compliant with Pareto dominance (Zitzler,

Thiele, Laumanns, Fonseca, & Fonseca, 2003).

- 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 algorithms to generate nondominated solutions without relying on Pareto ranking?

Some Applications in Finance

Evolutionary Algorithms in general and MOEAs in particular, can be useful in the solution of complex problems for which no efficient deterministic algorithm exists (i.e., there is no deterministic algorithm that can solve them in polynomial time).

It is well known that in finance there are several NP-complete problems for which the use of a heuristic is clearly justified (Schlottmann & Seese, 2004). However, the specialized literature on MOEAs reports few papers that deal with problems in finance. Some examples are the following:

- Solution of portfolio optimization problems, particularly using Markowitz models (see for example (Shoaf & Foster, 1996; Vedarajan, Chan, & Goldberg, 1997; Chang, Meade, & Beasley, 2000; Lin, Wang, & Yan, 2001; Streichert, Ulmer, & Zell, 2004; Doerner, Gutjahr, Hartl, Strauss, & Stummer, 2004; Ehrgott, Klamroth, & Schwehm, 2004)). This has been, by far, the most popular application of MOEAs in finance.
- Time series prediction (Zwir & Ruspini, 1999; Ruspini & Zwir, 1999).
- Risk-return trade-offs for loans (Mukerjee, Biswas, Deb, & Mathur, 2002; Schlottmann & Seese, 2002).

Evidently, it is necessary to identify other types of problems in finance whose complexity justifies the use of a MOEA (see for example (Chen, 2002)). Even the financial problems that have been tackled so far normally require a special treatment and a proper tailoring of the current MOEAs (e.g., regarding the encoding, since portfolio selection problems can be modelled as knapsack problems (Streichert et al., 2004)). Additionally, the decision-making process involved in financial applications is normally very complex and difficult to automate. This presents challenges for the (few) models for incorporation of preferences in current use with MOEAs (Cvetković & Parmee, 2002; Coello Coello, 2000a; Coello Coello et al., 2002). Thus, financial applications present research opportunities both for experts in finance and for researchers working exclusively in the development of multi-objective evolutionary algorithms and associated techniques.

Conclusions

In this chapter, we have presented a brief introduction to evolutionary multi-objective optimization. We have provided some basic concepts, and a historical perspective of the

research that has been done in this area. We have also presented short descriptions of some algorithms that are representative of each historical period under consideration including the current one.

In the last part of the chapter, we have presented some of the current and future research trends in the area, as well as a brief description of the sort of financial applications that have been developed using multi-objective evolutionary algorithms.

The main aim of this chapter is to provide a general overview of the area, identifying some opportunity areas, mainly related to financial applications.

References

- Abbass, H. A., & Sarker, R. (2002). The Pareto Differential Evolution Algorithm. *International Journal on Artificial Intelligence Tools*, 11(4), 531–552.
- Chang, T. J., Meade, N., & Beasley, J. E. (2000). Heuristics for Cardinality Constrained Portfolio Optimization. *Computers and Operations Research*, 27(13), 1271–1302.
- Chen, S.-H. (Ed.). (2002). *Evolutionary Computation in Economics and Finance*. Heidelberg, Germany: Physica-Verlag.
- Coello Coello, C. A. (1996). *An Empirical Study of Evolutionary Techniques for Multiobjective Optimization in Engineering Design*. Unpublished doctoral dissertation, Department of Computer Science, Tulane University, New Orleans, Louisiana, USA.
- Coello Coello, C. A. (1999, August). A Comprehensive Survey of Evolutionary-Based Multiobjective Optimization Techniques. *Knowledge and Information Systems. An International Journal*, 1(3), 269–308.
- Coello Coello, C. A. (2000a, July). Handling Preferences in Evolutionary Multiobjective Optimization: A Survey. In *2000 congress on evolutionary computation* (Vol. 1, pp. 30–37). Piscataway, New Jersey: IEEE Service Center.
- Coello Coello, C. A. (2000b). Treating Constraints as Objectives for Single-Objective Evolutionary Optimization. *Engineering Optimization*, 32(3), 275–308.
- Coello Coello, C. A., & Cruz Cortés, N. (2005, June). Solving Multiobjective Optimization Problems using an Artificial Immune System. *Genetic Programming and Evolvable Machines*, 6(2), 163–190.
- Coello Coello, C. A., & Toscano Pulido, G. (2001). Multiobjective Optimization using a Micro-Genetic Algorithm. In L. Spector et al. (Eds.), *Proceedings of the genetic and evolutionary computation conference (gecco'2001)* (pp. 274–282). San Francisco, California: Morgan Kaufmann Publishers.
- Coello Coello, C. A., Toscano Pulido, G., & Salazar Lechuga, M. (2004, June). Handling Multiple Objectives With Particle Swarm Optimization. *IEEE Transactions on Evolutionary Computation*, 8(3), 256–279.
- Coello Coello, C. A., Van Veldhuizen, D. A., & Lamont, G. B. (2002). *Evolutionary Algorithms for Solving Multi-Objective Problems*. New York: Kluwer Academic Publishers. (ISBN 0-3064-6762-3)
- Collette, Y., & Siarry, P. (2003). *Multiobjective Optimization. Principles and Case Studies*. Springer.
- Corne, D. W., Jerram, N. R., Knowles, J. D., & Oates, M. J. (2001). PESA-II: Region-based Selection in Evolutionary Multiobjective Optimization. In L. Spector et al. (Eds.), *Proceedings of the genetic and evolutionary computation conference (gecco'2001)* (pp. 283–290). San Francisco, California: Morgan Kaufmann Publishers.

- Corne, D. W., Knowles, J. D., & Oates, M. J. (2000). The Pareto Envelope-based Selection Algorithm for Multiobjective Optimization. In M. Schoenauer et al. (Eds.), *Proceedings of the parallel problem solving from nature vi conference* (pp. 839–848). Paris, France: Springer. Lecture Notes in Computer Science No. 1917.
- Cvetković, D., & Parmee, I. C. (2002, February). Preferences and their Application in Evolutionary Multiobjective Optimisation. *IEEE Transactions on Evolutionary Computation*, 6(1), 42–57.
- Deb, K. (2001). *Multi-Objective Optimization using Evolutionary Algorithms*. Chichester, UK: John Wiley & Sons. (ISBN 0-471-87339-X)
- Deb, K., Agrawal, S., Pratap, A., & Meyarivan, T. (2000). A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. In M. Schoenauer et al. (Eds.), *Proceedings of the parallel problem solving from nature vi conference* (pp. 849–858). Paris, France: Springer. Lecture Notes in Computer Science No. 1917.
- Deb, K., & Goldberg, D. E. (1989, June). An Investigation of Niche and Species Formation in Genetic Function Optimization. In J. D. Schaffer (Ed.), *Proceedings of the third international conference on genetic algorithms* (pp. 42–50). San Mateo, California: Morgan Kaufmann Publishers.
- Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002, April). A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2), 182–197.
- Doerner, K., Gutjahr, W. J., Hartl, R. F., Strauss, C., & Stummer, C. (2004, October). Pareto Ant Colony Optimization: A Metaheuristic Approach to Multiobjective Portfolio Selection. *Annals of Operations Research*, 131(1–4), 79–99.
- Ehrgott, M., Klamroth, K., & Schwehm, C. (2004, June). An MCDM approach to portfolio optimization. *European Journal of Operational Research*, 155(3), 752–770.
- Erickson, M., Mayer, A., & Horn, J. (2001). The Niche Pareto Genetic Algorithm 2 Applied to the Design of Groundwater Remediation Systems. In E. Zitzler, K. Deb, L. Thiele, C. A. Coello Coello, & D. Corne (Eds.), *First international conference on evolutionary multi-criterion optimization* (pp. 681–695). Springer-Verlag. Lecture Notes in Computer Science No. 1993.
- Fogel, L. J. (1999). *Artificial intelligence through simulated evolution. forty years of evolutionary programming*. New York, NY: John Wiley & Sons.
- Fonseca, C. M., & Fleming, P. J. (1993). Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. In S. Forrest (Ed.), *Proceedings of the fifth international conference on genetic algorithms* (pp. 416–423). San Mateo, CA: Morgan Kaufmann Publishers.
- Fonseca, C. M., & Fleming, P. J. (1996, September). On the Performance Assessment and Comparison of Stochastic Multiobjective Optimizers. In H.-M. Voigt, W. Ebeling, I. Rechenberg, & H.-P. Schwefel (Eds.), *Parallel problem solving from nature—ppsn iv* (pp. 584–593). Berlin, Germany: Springer-Verlag. Lecture Notes in Computer Science No. 1141.
- Fourman, M. P. (1985). Compaction of Symbolic Layout using Genetic Algorithms. In *Genetic algorithms and their applications: Proceedings of the first international conference on genetic algorithms* (pp. 141–153). Hillsdale, NJ: Lawrence Erlbaum.
- Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Reading, MA: Addison-Wesley Publishing Company.
- Goldberg, D. E., & Richardson, J. (1987). Genetic algorithm with sharing for multimodal function optimization. In J. J. Grefenstette (Ed.), *Genetic algorithms and their applications: Proceedings of the second international conference on genetic algorithms* (pp. 41–49). Hillsdale, NJ: Lawrence Erlbaum.

- Horn, J. (1997). Multicriterion Decision Making. In T. Bäck, D. Fogel, & Z. Michalewicz (Eds.), *Handbook of evolutionary computation* (Vol. 1, p. F1.9:1 - F1.9:15). UK: IOP Publishing Ltd. and Oxford University Press.
- Horn, J., Nafpliotis, N., & Goldberg, D. E. (1994, June). A niched pareto genetic algorithm for multiobjective optimization. In *Proceedings of the first ieee conference on evolutionary computation, ieee world congress on computational intelligence* (Vol. 1, pp. 82–87). Piscataway, NJ: IEEE Service Center.
- Husbands, P. (1994). Distributed Coevolutionary Genetic Algorithms for Multi-Criteria and Multi-Constraint Optimisation. In T. Fogarty (Ed.), *Evolutionary computing* (Vol. 865, pp. 150–165). Springer Verlag.
- Jensen, M. T. (2003a, April). Guiding Single-Objective Optimization Using Multi-objective Methods. In G. R. et al. (Ed.), *Applications of evolutionary computing. evoworkshops 2003: Evobio, evocop, evoiasp, evomusart, evorob, and evostim* (pp. 199–210). Essex, UK: Springer. Lecture Notes in Computer Science Vol. 2611.
- Jensen, M. T. (2003b, October). Reducing the Run-Time Complexity of Multiobjective EAs: The NSGA-II and Other Algorithms. *IEEE Transactions on Evolutionary Computation*, 7(5), 503–515.
- Kennedy, J., & Eberhart, R. C. (2001). *Swarm intelligence*. San Francisco, California: Morgan Kaufmann Publishers.
- Knowles, J., & Corne, D. (2002, May). On Metrics for Comparing Nondominated Sets. In *Congress on evolutionary computation (cec'2002)* (Vol. 1, pp. 711–716). Piscataway, New Jersey: IEEE Service Center.
- Knowles, J. D., & Corne, D. W. (2000). Approximating the nondominated front using the pareto archived evolution strategy. *Evolutionary Computation*, 8(2), 149–172.
- Knowles, J. D., Watson, R. A., & Corne, D. W. (2001). Reducing Local Optima in Single-Objective Problems by Multi-objectivization. In E. Zitzler, K. Deb, L. Thiele, C. A. C. Coello, & D. Corne (Eds.), *First international conference on evolutionary multi-criterion optimization* (pp. 268–282). Zürich, Switzerland: Springer-Verlag. Lecture Notes in Computer Science No. 1993.
- Laumanns, M., Thiele, L., Deb, K., & Zitzler, E. (2002, Fall). Combining Convergence and Diversity in Evolutionary Multi-objective Optimization. *Evolutionary Computation*, 10(3), 263–282.
- Lin, D., Wang, S., & Yan, H. (2001). *A multiobjective genetic algorithm for portfolio selection*. Working Paper, Institute of Systems Science, Academy of Mathematics and Systems Science Chinese Academy of Sciences, Beijing, China.
- Morse, J. (1980). Reducing the size of the nondominated set: Pruning by clustering. *Computers and Operations Research*, 7(1–2), 55–66.
- Mostaghim, S., Teich, J., & Tyagi, A. (2002, May). Comparison of Data Structures for Storing Pareto-sets in MOEAs. In *Congress on evolutionary computation (cec'2002)* (Vol. 1, pp. 843–848). Piscataway, New Jersey: IEEE Service Center.
- Mukerjee, A., Biswas, R., Deb, K., & Mathur, A. P. (2002, September). Multi-objective evolutionary algorithms for the risk-return trade-off in bank-load management. *International Transactions in Operational Research*, 9(5), 583–597.
- Oei, C. K., Goldberg, D. E., & Chang, S.-J. (1991, December). *Tournament Selection, Niching, and the Preservation of Diversity* (Tech. Rep. No. 91011). Urbana, Illinois: Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana-Champaign.

- Osyczka, A. (2002). *Evolutionary Algorithms for Single and Multicriteria Design Optimization*. Germany: Physica Verlag. (ISBN 3-7908-1418-0)
- Osyczka, A., & Kundu, S. (1995). A new method to solve generalized multicriteria optimization problems using the simple genetic algorithm. *Structural Optimization*, 10, 94–99.
- Price, K. V. (1999). An Introduction to Differential Evolution. In D. Corne, M. Dorigo, & F. Glover (Eds.), *New ideas in optimization* (pp. 79–108). UK: Mc Graw-Hill.
- Purshouse, R. C. (2003). *On the Evolutionary Optimisation of Many Objectives*. Unpublished doctoral dissertation, Department of Automatic Control and Systems Engineering, The University of Sheffield, Sheffield, UK.
- Ruspini, E. H., & Zwir, I. S. (1999). Automated Qualitative Description of Measurements. In *Proceedings of the 16th IEEE instrumentation and measurement technology conference*. Venice, Italy: IEEE Press.
- Schaffer, J. D. (1984). *Multiple objective optimization with vector evaluated genetic algorithms*. Unpublished doctoral dissertation, Vanderbilt University, Nashville, TN.
- Schaffer, J. D. (1985). Multiple Objective Optimization with Vector Evaluated Genetic Algorithms. In *Genetic algorithms and their applications: Proceedings of the first international conference on genetic algorithms* (pp. 93–100). Hillsdale, NJ: Lawrence Erlbaum.
- Schaffer, J. D., & Grefenstette, J. J. (1985). Multiobjective learning via genetic algorithms. In *Proceedings of the 9th international joint conference on artificial intelligence (ijcai-85)* (pp. 593–595). Los Angeles, CA: AAAI.
- Schlottmann, F., & Seese, D. (2002, June). Hybrid multi-objective evolutionary computation of constrained downside risk-return efficient sets for credit portfolio. In *Proceedings of the 8th international conference of the society for computational economics. computing in economics and finance*. Aix-en-Provence, France.
- Schlottmann, F., & Seese, D. (2004). Financial Applications of Multi-Objective Evolutionary Algorithms: Recent Developments and Future Research Directions. In C. A. Coello Coello & G. B. Lamont (Eds.), *Applications of multi-objective evolutionary algorithms* (pp. 627–652). Singapore: World Scientific.
- Shoaf, J. S., & Foster, J. A. (1996). A Genetic Algorithm Solution to the Efficient Set Problem: A Technique for Portfolio Selection Based on the Markowitz Model. In *Proceedings of the decision sciences institute annual meeting* (pp. 571–573). Orlando, Florida.
- Srinivas, N., & Deb, K. (1994, Fall). Multiobjective Optimization Using Nondominated Sorting in Genetic Algorithms. *Evolutionary Computation*, 2(3), 221–248.
- Streichert, F., Ulmer, H., & Zell, A. (2004, June). Comparing Discrete and Continuous Genotypes on the Constrained Portfolio Selection Problem. In K. D. et al. (Ed.), *Genetic and evolutionary computation-gecco 2004. proceedings of the genetic and evolutionary computation conference. part ii* (pp. 1239–1250). Seattle, Washington, USA: Springer-Verlag, Lecture Notes in Computer Science Vol. 3103.
- Syswerda, G., & Palmucci, J. (1991). The application of genetic algorithms to resource scheduling. In R. K. Belew & L. B. Booker (Eds.), *Proceedings of the fourth international conference on genetic algorithms* (pp. 502–508). San Mateo, California: Morgan Kaufmann.
- Tan, K., Khor, E., & Lee, T. (2005). *Multiobjective Evolutionary Algorithms and Applications*. London: Springer-Verlag. (ISBN 1-85233-836-9)

- Van Veldhuizen, D. A. (1999). *Multiobjective evolutionary algorithms: Classifications, analyses, and new innovations*. Unpublished doctoral dissertation, Department of Electrical and Computer Engineering. Graduate School of Engineering. Air Force Institute of Technology, Wright-Patterson AFB, OH.
- Van Veldhuizen, D. A., & Lamont, G. B. (2000a). Multiobjective Optimization with Messy Genetic Algorithms. In *Proceedings of the 2000 acm symposium on applied computing* (pp. 470–476). Villa Olmo, Como, Italy: ACM.
- Van Veldhuizen, D. A., & Lamont, G. B. (2000b, July). On Measuring Multiobjective Evolutionary Algorithm Performance. In *2000 congress on evolutionary computation* (Vol. 1, pp. 204–211). Piscataway, NJ: IEEE Service Center.
- Vedarajan, G., Chan, L. C., & Goldberg, D. E. (1997, July). Investment Portfolio Optimization using Genetic Algorithms. In J. R. Koza (Ed.), *Late breaking papers at the genetic programming 1997 conference* (pp. 255–263). Stanford University, California: Stanford Bookstore.
- Wienke, P. B., Lucasius, C., & Kateman, G. (1992). Multicriteria target optimization of analytical procedures using a genetic algorithm. *Analytical Chimica Acta*, 265(2), 211–225.
- Zitzler, E., Deb, K., & Thiele, L. (2000, Summer). Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. *Evolutionary Computation*, 8(2), 173–195.
- Zitzler, E., Laumanns, M., & Thiele, L. (2001, May). *SPEA2: Improving the Strength Pareto Evolutionary Algorithm* (Tech. Rep. No. 103). Gloriastrasse 35, CH-8092 Zurich, Switzerland: Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH) Zurich.
- Zitzler, E., Laumanns, M., Thiele, L., Fonseca, C. M., & Grunert da Fonseca, V. (2002, July). Why Quality Assessment of Multiobjective Optimizers Is Difficult. In W. Langdon et al. (Eds.), *Proceedings of the genetic and evolutionary computation conference (gecco'2002)* (pp. 666–673). San Francisco, California: Morgan Kaufmann Publishers.
- Zitzler, E., & Thiele, L. (1999, November). Multiobjective Evolutionary Algorithms: A Comparative Case Study and the Strength Pareto Approach. *IEEE Transactions on Evolutionary Computation*, 3(4), 257–271.
- Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C. M., & Fonseca, V. G. da. (2003, April). Performance Assessment of Multiobjective Optimizers: An Analysis and Review. *IEEE Transactions on Evolutionary Computation*, 7(2), 117–132.
- Zwir, I. S., & Ruspini, E. H. (1999). Qualitative Object Description: Initial Reports of the Exploration of the Frontier. In *Proceedings of the joint eurofuse—sic'99 international conference*. Budapest, Hungary.
- Zydallis, J. B., Lamont, G. B., & Veldhuizen, D. A. V. (2000, September). Messy Genetic Algorithm Based Multi-Objective Optimization: A Comparative Statistical Analysis. In *Ppsn/sab workshop on multiobjective problem solving from nature (mpsn)*. Paris, France.

Author's Biography

Carlos A. Coello Coello holds a PhD degree in computer science from Tulane University, USA. He is currently an associate professor at CINVESTAV-IPN, in Mexico City, México. Dr. Coello has authored and co-authored over 120 technical papers and several book chapters. He has also co-authored the book *Evolutionary Algorithms for Solving Multi-Objective Problems* (Kluwer Academic Publishers, 2002).