

An Overview of Evolutionary Algorithms in Multiobjective Optimization*

Carlos M. Fonseca[†] and Peter J. Fleming[‡]

The University of Sheffield

Department of Automatic Control and Systems Engineering

Mappin Street

Sheffield S1 3JD, U.K.

Tel.: +44 114 2825236

Fax: +44 114 2731729

April 7, 1995

Abstract

The application of evolutionary algorithms (EAs) in multiobjective optimization is currently receiving growing interest from researchers with various backgrounds. Most research in this area has understandably concentrated on the selection stage of EAs, due to the need to integrate vectorial performance measures with the inherently scalar way in which EAs reward individual performance, i.e., number of offspring.

In this review, current multiobjective evolutionary approaches are discussed, ranging from the conventional analytical aggregation of the different objectives into a single function to a number of population-based approaches and the more recent ranking schemes based on the definition of Pareto-optimality. The sensitivity of different methods to

*To appear in *Evolutionary Computation*, 3(1):1–16, Spring 1995. Final draft.

[†]E-mail: C.Fonseca@shef.ac.uk

[‡]E-mail: P.Fleming@shef.ac.uk

objective scaling and/or possible concavities in the trade-off surface is considered, and related to the (static) fitness landscapes such methods induce on the search space. From the discussion, directions for future research in multiobjective fitness assignment and search strategies are identified, including the incorporation of decision making in the selection procedure, fitness sharing, and adaptive representations.

Keywords: evolutionary algorithms, multiobjective optimization, fitness assignment, search strategies.

1 Introduction

Many real world problems involve multiple measures of performance, or objectives, which should be optimized simultaneously. In certain cases, objective functions may be optimized separately from each other and insight gained concerning the best that can be achieved in each performance dimension. However, suitable solutions to the overall problem can seldom be found in this way. Optimal performance according to one objective, if such an optimum exists, often implies unacceptably low performance in one or more of the other objective dimensions, creating the need for a compromise to be reached. A suitable solution to such problems involving conflicting objectives should offer “acceptable”, though possibly sub-optimal in the single-objective sense, performance in all objective dimensions, where “acceptable” is a problem-dependent and ultimately subjective concept.

The simultaneous optimization of multiple, possibly competing, objective functions deviates from single function optimization in that it seldom admits a single, perfect (or Utopian) solution. Instead, multiobjective optimization (MO) problems tend to be characterized by a family of alternatives which must be considered equivalent in the absence of information concerning the relevance of each objective relative to the others. Multiple solutions, or multimodality, arise even in the simplest non-trivial case of two competing objectives, where both are unimodal and convex functions of the decision variables. As the number of competing objectives increases and less well-behaved objectives are considered, the problem of finding a satisfactory compromise solution rapidly becomes increasingly complex.

Conventional optimization techniques, such as gradient-based and simplex-based methods, and also less conventional ones, such as simulated annealing, are difficult to extend to the *true* multiobjective case, because they were not designed with multiple solutions in mind. In practice, multiobjective problems have to be re-formulated as single-objective prior to optimization, leading to the production of a single solution per run of the optimizer.

Evolutionary algorithms (EAs), however, have been recognized to be possibly well-suited to multiobjective optimization since early in their development. Multiple individuals can search for multiple solutions in parallel, eventually taking advantage of any similarities available in the family of possible solutions to the problem. The ability to handle complex problems, involving features such as discontinuities, multimodality, disjoint feasible spaces

and noisy function evaluations, reinforces the potential effectiveness of EAs in multiobjective search and optimization, which is perhaps a problem area where Evolutionary Computation really distinguishes itself from its competitors.

This paper reviews current evolutionary approaches to multiobjective optimization, discussing their similarities and differences. It also tries to identify some of the main issues raised by multiobjective optimization in the context of evolutionary search, and how the methods discussed address them. From the discussion, directions for future work in multiobjective evolutionary algorithms are identified.

2 Evolutionary approaches to multiobjective optimization

The family of solutions of a multiobjective optimization problem is composed of all those elements of the search space which are such that the components of the corresponding objective vectors cannot be all simultaneously improved. This is known as the concept of Pareto optimality.

A more formal definition of Pareto optimality is as follows: consider, without loss of generality, the minimization of the n components f_k , $k = 1, \dots, n$, of a vector function \mathbf{f} of a vector variable \mathbf{x} in a universe \mathcal{U} , where

$$\mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_n(\mathbf{x})).$$

Then, a decision vector $\mathbf{x}_{\mathbf{u}} \in \mathcal{U}$ is said to be Pareto-optimal if and only if there is no $\mathbf{x}_{\mathbf{v}} \in \mathcal{U}$ for which $\mathbf{v} = \mathbf{f}(\mathbf{x}_{\mathbf{v}}) = (v_1, \dots, v_n)$ *dominates* $\mathbf{u} = \mathbf{f}(\mathbf{x}_{\mathbf{u}}) = (u_1, \dots, u_n)$, i.e., there is no $\mathbf{x}_{\mathbf{v}} \in \mathcal{U}$ such that

$$\forall i \in \{1, \dots, n\}, v_i \leq u_i \quad \wedge \quad \exists i \in \{1, \dots, n\} \mid v_i < u_i.$$

The set of all Pareto-optimal decision vectors is called the Pareto-optimal, efficient, or admissible set of the problem. The corresponding set of objective vectors is called the non-dominated set. In practice, however, it is not unusual for these terms to be used interchangeably to describe solutions of a multiobjective optimization problem.

The notion of Pareto-optimality is only a first step towards the practical

solution of a multiobjective problem, which usually involves the choice a *single* compromise solution from the non-dominated set according to some preference information.

2.1 Plain aggregating approaches

Because evolutionary algorithms require scalar fitness information to work on, a scalarization of the objective vectors is always necessary. In most problems where no global criterion directly emerges from the problem formulation, objectives are often artificially combined, or aggregated, into a scalar function according to some understanding of the problem, and the EA applied. Many such approaches developed for use with conventional optimizers can also be used with EAs.

Optimizing a combination of the objectives has the advantage of producing a single compromise solution, requiring no further interaction with the decision maker (DM). The problem is, if the optimal solution cannot be accepted, either due to the function used excluding aspects of the problem which were unknown prior to optimization or to an inappropriate setting of the coefficients of the combining function, new runs of the optimizer may be required until a suitable solution is found.

Several applications of evolutionary algorithms in the optimization of aggregating functions have been reported in the literature. A number of authors (Syswerda and Palmucci, 1991; Jakob *et al.*, 1992; Jones *et al.*, 1993) provide examples of the use of the popular weighted sum approach. Using target vector optimization, which consists of minimizing the distance in objective space to a given goal vector, Wienke *et al.* (1992) report work on a problem in atomic emission spectroscopy. Goal attainment (Gembicki, 1974), a related technique which seeks to minimize the weighted difference between objective values and the corresponding goals, was used amongst other methods by Wilson and Macleod (1993), who also monitored the population for non-dominated solutions. The use of multiple attribute utility analysis (MAUA) in conjunction with GAs has been suggested by Horn and Nafpliotis (1993), but without experimental results.

Handling constraints with penalty functions (Davis and Steenstrup, 1987; Goldberg, 1989) is yet another example of an additive aggregating function. The fact that penalty functions are generally problem dependent and, as a consequence, difficult to set (Richardson *et al.*, 1989) has prompted the de-

velopment of alternative approaches based on ranking (Powell and Skolnick, 1993).

2.2 Population-based non-Pareto approaches

Schaffer (1985, see also Schaffer and Grefenstette (1985)) was probably the first to recognize the possibility of exploiting EA populations to treat non-commensurable objectives separately and search for multiple non-dominated solutions concurrently in a single EA run. In his approach, known as the Vector Evaluated Genetic Algorithm (VEGA), appropriate fractions of the next generation, or sub-populations, were selected from the whole of the old generation according to each of the objectives, separately. Crossover and mutation were applied as usual after shuffling all the sub-populations together. Non-dominated individuals were identified by monitoring the population as it evolved, but this information was not used by the VEGA itself.

Shuffling and merging all sub-populations corresponds, however, to averaging the *normalized* fitness components associated with each of the objectives. In fact, the expected total number of offspring produced by each parent becomes the sum of the expected numbers of offspring produced by that parent according to each objective. Since Schaffer used proportional fitness assignment, these were in turn, proportional to the objectives themselves. The resulting overall fitness corresponded, therefore, to a linear function of the objectives where the weights depended on the distribution of the population at each generation. This has previously been noted by Richardson *et al.* (1989) and confirmed by Schaffer (1993). As a consequence, different non-dominated individuals were generally assigned different fitness values, in contrast with what the definition of non-dominance would suggest.

The linear combination of the objectives implicitly performed by VEGA explains why the population tended to split into species particularly strong in each of the objectives in the case of concave trade-off surfaces, a phenomenon which Schaffer called *speciation*. In fact, points in concave regions of a trade-off surface cannot be found by optimizing a linear combination of the objectives, for *any* set of weights, as noted in (Fleming and Pashkevich, 1985).

Although VEGA, like the plain weighted-sum approach, is not well suited to address problems with concave trade-off surfaces, the weighting scheme it implicitly implements deserves closer attention. In VEGA, each objective

is effectively weighted proportionally to the size of each sub-population and, more importantly, proportionally to the inverse of the average fitness (in terms of that objective) of the whole population at each generation.

By doing so, and assuming that sub-population sizes remain constant for each objective, VEGA selection adaptively attempts to balance improvement in the several objective dimensions, because more good-performers in one objective cause the corresponding average performance to increase and that objective's weight to decrease accordingly. This is not unlike the way sharing techniques (Goldberg and Richardson, 1987, see below) promote the balanced exploitation of multiple optima in the search space. For the same reason, VEGA can, at least in some cases, maintain different species for many more generations than a GA optimizing a pure weighted sum of the same objectives with fixed weights would, due to genetic drift (Goldberg and Segrest, 1987). Unfortunately, the balance reached necessarily depends on the scaling of the objectives.

Fourman (1985) also addressed multiple objectives in a non-aggregating manner. Selection was performed by comparing pairs of individuals, each pair according to one of the objectives. In a first version of the algorithm, objectives were assigned different priorities by the user and individuals compared according to the objective with the highest priority. If this resulted in a tie, the objective with the second highest priority was used, and so on. This is known as the *lexicographic* ordering (Ben-Tal, 1980).

A second version, reported to work surprisingly well, consisted of randomly selecting the objective to be used in each comparison. Similarly to VEGA, this corresponds to averaging fitness across fitness components, each component being weighted by the probability of each objective being chosen to decide each tournament. However, the use of pairwise comparisons makes it essentially *different* from a linear combination of the objectives, because scale information is ignored. As tournaments constitute stochastic approximations to full ranking, the resulting fitness is closer to the ranking of the population according to each objective separately, and the consequent averaging of each individual's ranks. Thus, the population may still see as convex a trade-off surface actually concave, depending on its current distribution and, of course, on the problem.

Kursawe (1991) formulated a multiobjective version of evolution strategies (ESs). Once again, selection consisted of as many steps as there were objectives. At each step, one objective was selected randomly (with replace-

ment) according to a probability vector, and used to dictate the deletion of an appropriate fraction of the current population. After selection, μ survivors became the parents of the next generation.

While Kursawe's implementation of multiobjective selection possesses a number of similarities to both VEGA and Fourman's second method, individuals in the extremes of the trade-off surface would appear to be likely to be eliminated as soon as any objective at which they performed poorly was selected to dictate deletion, whereas middling individuals seem to be more likely to survive. However, since objectives stood a certain chance of not taking part in selection at each generation, it was possible for some specialists to survive the deletion process and generate offspring, although they might die immediately the generation after.

Kursawe (1991) notes that this deletion of individuals according to randomly chosen objectives creates a non-stationary environment in which the population, instead of converging, must try to adapt to constant change. As hinted above, different choices of objectives could result in significant changes in the cost landscape seen by the ES at each generation. Diploid individuals (Goldberg and Smith, 1987) were used for their improved ability to adapt to sudden environmental changes and, since the population was not expected to converge, a picture of the trade-off surface was produced from the points evaluated during the run.

Finally, and still based on the weighted sum approach, Hajela and Lin (1992) exploited the explicit parallelism provided by a population-based search by explicitly including the weights in the chromosome and promoting their diversity in the population through fitness sharing. As a consequence, one family of individuals evolved for each weight combination, concurrently.

2.3 Pareto-based approaches

The methods of Schaffer, Fourman, Kursawe, and Hajela and Lin, all attempt to promote the generation of multiple non-dominated solutions, a goal at which they reportedly achieved a reasonable degree of success. However, none makes *direct* use of the actual definition of Pareto-optimality. At most, the population is monitored for non-dominated solutions, as in Schaffer (1985) and Kursawe (1991).

Pareto-based fitness assignment was first proposed by Goldberg (1989), as a means of assigning equal probability of reproduction to all non-dominated

individuals in the population. The method consisted of assigning rank 1 to the non-dominated individuals and removing them from contention, then finding a new set of non-dominated individuals, ranked 2, and so forth.

Fonseca and Fleming (1993) have proposed a slightly different scheme, whereby an individual's rank corresponds to the number of individuals in the current population by which it is dominated. Non-dominated individuals are, therefore, all assigned the same rank, while dominated ones are penalized according to the population density in the corresponding region of the trade-off surface. The algorithm proceeds by sorting the population according to the multiobjective ranks previously determined. Fitness is assigned by interpolating, e.g., linearly, from the best to the worst individuals in the population, and then averaging it between individuals with the same multiobjective rank. Selection is performed with Baker's (1987) Stochastic Universal Sampling (SUS) algorithm. (Srinivas and Deb (1994) have implemented a similar sorting and fitness assignment procedure, but based on Goldberg's version of Pareto-ranking.)

By combining Pareto dominance with partial preference information in the form of a goal vector, they have also provided a means of evolving only a given region of the trade-off surface. While the basic ranking scheme remains unaltered, the now Pareto-like comparison of the individuals selectively excludes those objectives which already satisfy their goals. Specifying fully unattainable goals causes objectives never to be excluded from comparison, which corresponds to the original Pareto ranking. Changing the goal values during the search alters the fitness landscape accordingly and allows the decision maker to direct the population to zoom in on a particular region of the trade-off surface.

Tournament selection based on Pareto dominance has also been proposed by Horn and Nafpliotis (1993, see also Horn *et al.* (1994)). In addition to the two individuals competing in each tournament, a number of other individuals in the population was used to help determine whether the competitors were dominated or not. In the case where both competitors were either dominated or non-dominated, the result of the tournament was decided through sharing (see below).

Cieniawski (1993) and Ritzel *et al.* (1994) have implemented tournament selection based on Goldberg's Pareto-ranking scheme. In their approach, individual ranks were used to decide the winner of binary tournaments, which is in fact a stochastic approximation to the full sorting of the population, as




Figure 1: The concavity of the trade-off set is related to how the objectives are scaled.

performed by Fonseca and Fleming (1993) and Srinivas and Deb (1994).

The convexity of the trade-off surface depends on how the objectives are scaled. Non-linearly rescaling the objective values may convert a concave surface into a convex one, and vice-versa, as illustrated in Figure 1. The darker surface is the original, concave trade-off surface, corresponding to plotting $f_1(\mathbf{x})$ against $f_2(\mathbf{x})$, where \mathbf{x} denotes the vector of free variables. The lighter surfaces correspond to plotting $[f_1(\mathbf{x})]^\alpha$ against $[f_2(\mathbf{x})]^\alpha$, for $\alpha = 5$ and $\alpha = 9$, the latter being clearly convex. Nevertheless, all are formulations of the same minimization problem which admit exactly the same solution set in phenotypic space.

Since order is preserved by monotonic transformations such as these, Pareto-ranking is blind to the convexity or the non-convexity of the trade-off surface. This is not to say that Pareto-ranking always precludes speciation. Speciation can still occur if certain regions of the trade-off are simply easier to find than others, but Pareto-ranking does eliminate sensitivity to the possible non-convexity of the trade-off.

A second possible advantage of Pareto-ranking, is that, because it rewards good performance in any objective dimension *regardless* of the others,

solutions which exhibit good performance in many, if not all, objective dimensions are more likely to be produced by recombination. This argument also applies to an extent to the population-based methods described in the previous subsection, although they do not necessarily treat all non-dominated individuals equally. The argument assumes some degree of independence between objectives, and was already hinted at by Schaffer in his VEGA work, and has been noted in more detail by Louis and Rawlins (1993). While Pareto-based selection may help find Utopian solutions if they exist, that is rarely the case in multiobjective optimization. Also, the assumption of loosely coupled objectives is less likely to hold near the admissible region, but the argument may still be valid in the initial stages of the search.

2.4 Niche induction techniques

Pareto-based ranking correctly assigns all non-dominated individuals the same fitness, but that, on its own, does not guarantee that the Pareto set be uniformly sampled. When presented with multiple equivalent optima, finite populations tend to converge to only one of them, due to stochastic errors in the selection process. This phenomenon, known as genetic drift (Goldberg and Segrest, 1987), has been observed in natural as well as in artificial evolution, and can also occur in Pareto-based evolutionary optimization.

The additional use of fitness sharing (Goldberg and Richardson, 1987; Deb and Goldberg, 1989) was proposed by Goldberg (1989) to prevent genetic drift and to promote the sampling of the whole Pareto set by the population. Fonseca and Fleming (1993) implemented fitness sharing in the objective domain and provided theory for estimating the necessary niche sizes based on the properties of the Pareto set. Horn and Nafpliotis (1993) also arrived at a form of fitness sharing in the objective domain. In addition, they suggested the use of a metric combining both the objective and the decision variable domains, leading to what they called nested sharing. Cieniawski (1993) performed sharing on a single objective dimension, that in which diversity appeared to be more important. Srinivas and Deb (1994) performed sharing in the decision variable domain.

Although sharing has mainly been used together with Pareto ranking (Fonseca and Fleming, 1993; Cieniawski, 1993; Srinivas and Deb, 1994) and Pareto tournaments (Horn and Nafpliotis, 1993; Horn *et al.*, 1994), it should be noted that Hajela and Lin (1992) had already implemented a form of shar-

ing to stabilize the population around given regions of the trade-off surface. VEGA’s selection has also been noted earlier in this work to implement a sort of sharing mechanism, well before “sharing” as such was introduced to GAs by Goldberg and Richardson (1987).

The viability of mating is another aspect which becomes relevant as the population distributes itself around multiple regions of optimality. Different regions of the trade-off surface may generally have very different genetic representations, which, to ensure viability, requires mating to happen only locally (Goldberg, 1989). So far, mating restriction has only been implemented based on the distance between individuals in the objective domain, either directly, by Fonseca and Fleming (1993), or indirectly, by Hajela and Lin (1992). Nevertheless, the use of mating restriction in multiobjective EAs does not appear to be widespread.

Both sharing and mating restriction in the objective domain necessarily combine objectives to produce a distance measure, which may appear to be in contradiction with the philosophy behind Pareto-based selection. However, the *uniform* sampling of the whole Pareto set is only a meaningful requirement for a given scaling of the objectives. Sharing in the phenotypic domain abandons this requirement and replaces it by the uniform sampling of the admissible set.

Sharing and Pareto-selection should, ideally, have orthogonal effects: while Pareto-selection promotes improvement by exerting a scale-independent selective pressure on the population in a direction normal to the trade-off surface, sharing should attempt to balance the distribution of the population along the front by applying a, possibly scale-dependent, selective pressure tangentially to that surface.

Unfortunately, the possibility that sharing in the objective domain may, by concentrating search effort in some regions of the trade-off surface, favour improvement in those regions to the detriment of others, cannot be discarded. Performing fitness sharing in decision variable space (Srinivas and Deb, 1994) would provide a selection mechanism truly independent from objective scaling, as long as guidelines for the setting of the sharing parameters in that domain in the multiobjective case could be developed.

Fortunately, such guidelines may be already available, although outside the EA community. In fact, if share count calculation in sharing is recognized to be no more than a form of *kernel density estimation* (Silverman, 1986) in n dimensions, well studied heuristics for the setting of the corresponding

smoothing parameter (read niche size) can suddenly be used. More advanced methods of density estimation, such as adaptive smoothing, also become available. Since those heuristics are based on the dimensionality of the space in which sharing is to be performed and on population statistics such as the sample covariance matrix, but not on the function to be optimized (such a function is outside the density estimation problem itself), they may well provide a much more general approach to niche size setting than the current one (Deb and Goldberg, 1989).

3 Discussion

The handling of multiple objectives strongly interacts with evolutionary computation on many fronts, raising issues which can generally be accommodated in one of two broad classes, fitness assignment and search strategies.

3.1 Fitness assignment

The extension of evolutionary algorithms to the multiple objective case has mainly been concerned with multiobjective fitness assignment. According to how much preference information is incorporated in the fitness function, approaches range from complete preference information given, as when combining objective functions directly or prioritizing them, to no preference information given, as with Pareto-based ranking, and include the case where partial information is provided in order to restrict the search to only part of the Pareto set. Progressive refinement of partial preferences is also possible with EAs.

Independently of how much preference information is provided, the assigned fitness reflects a decision maker's understanding of the quality, or *utility*, of the points under assessment. Each selection step of an EA can be seen as a decision making problem involving as many alternatives as there are individuals in the population.

The fitness landscape associated with a multiobjective problem clearly depends on the fitness assignment strategy used. Consider the simple bi-objective problem of simultaneously minimizing

$$f_1(x_1, x_2) = 1 - \exp\left(-(x_1 - 1)^2 - (x_2 + 1)^2\right)$$




Figure 2: Surface plots of functions f_1 and f_2

$$f_2(x_1, x_2) = 1 - \exp\left(- (x_1 + 1)^2 - (x_2 - 1)^2\right)$$

Surface plots of these two objectives are shown in Figure 2. Note that the z -axis is inverted to facilitate the visualization. The corresponding trade-off surface is the one shown earlier in Figure 1 for $\alpha = 1$.

If individuals are ranked according to how many members of the population outperform them (Fonseca and Fleming, 1993), the ranking of a large, uniformly distributed population, normalized by the population size, can be interpreted as an estimate of the fraction of the search space which outperforms each particular point considered. (Global optima should be ranked zero.) This applies equally to single-objective ranking.

Plotting the normalized ranks against the decision variables, x_1 and x_2 in this case, produces an anti-fitness, or cost, landscape, from which the actual fitness landscape can be inferred. Clearly, as the population evolves, its distribution is no longer uniform and the cost landscape it induces will change dynamically. Nevertheless, the “static” landscapes considered here do provide insight into the different *selection mechanisms*. Such surfaces may also help *explain* the behaviour of EAs based on those selection mechanisms, but they cannot be expected to be predictive of EA performance when considered in isolation.

Static cost landscapes for the example above are shown in Figures 4 to 7, corresponding to four different fitness assignment strategies based on ranking. The cost landscape induced by ranking each objective separately is shown in Figure 3.




Figure 3: The cost landscapes defined by ranking each objective separately (and contour plots of the corresponding objective functions)

Figure 4 illustrates the single-objective ranking of the sum of the two objectives. The two peaks arise due to the problem exhibiting a concave trade-off surface. More importantly, these peaks would remain present (but would be no longer symmetric) if the objectives were weighted differently.

Although the surface in Figure 4 can only be seen as a (scaled) representation of the cost landscape induced by VEGA-selection on a uniformly distributed population (since, in this case, the average performance of the population would be the same for both objectives, f_1 and f_2), it clearly illustrates how trade-off surface concavities lead to peaks in the cost surface obtained by linearly combining the objectives. Speciation in VEGA corresponds to the population distributing itself by these “persistent” peaks, in a balanced way: objectives corresponding to highly populated peaks are weighted less (as performance in terms of the corresponding objective increases), causing the population to shift to other peaks until an equilibrium is reached. As a result, genetic drift can be controlled, and different species maintained on each peak in the long run.

In Figure 5, the average of the ranks computed according to each of the two objectives is shown. In this case, a single peak is located towards the middle of the Pareto-optimal set, and the concavity of the trade-off surface is no longer apparent. Binary tournaments according to one objective drawn at random, as in Fourman (1985), can be expected to define a similar landscape.

Figure 6 shows the cost landscape for the ranking of the maximum of the two objectives: a simple case of goal programming. The single-peak is




Figure 4: The cost landscape defined by ranking the sum of the objectives (The contour plots are those of the individual objective functions f_1 and f_2)




Figure 5: The cost landscape defined by ranking objectives separately and averaging the ranks

located on a non-smooth ridge, which makes direct gradient-based optimization difficult. For this reason, alternative formulations are usually preferred to this approach. For example, the goal attainment method as proposed by Gembicki (1974) avoids the problem by introducing an auxiliary scalar parameter λ and solving:

$$\min_{\lambda, \mathbf{x} \in \mathcal{U}} \lambda$$

subject to

$$f_i - w_i \lambda \leq g_i$$

where g_i are goals for the design objectives f_i , and $w_i \geq 0$ are weights which must be specified beforehand by the designer.

Finally, in Figure 7, Pareto-ranking is used. Note how the Pareto-optimal set defines a ridge-shaped plateau in the cost landscape. As desired, this plateau includes all admissible solutions and, thus, all possible optima produced by any coordinatewise monotonic function of the objectives (Steuer, 1986), of which the methods in Figures 4 to 6 are just examples.




Figure 6: The cost landscape defined by ranking the maximum of the two objectives




Figure 7: The cost landscape defined by Pareto-ranking

3.2 Search strategies

The ridges defined in the fitness landscape by Pareto-ranking and/or minimax approaches may not be parallel to any of the decision variable axes, or even follow a straight line. Although ridges, or equivalently, valleys, need not occur in single-objective optimization (Mühlenbein and Schlierkamp-Voosen, 1993), they do appear in this context, and can certainly be expected in almost any multiobjective problem.

Ridge-shaped plateaus raise two problems already encountered with other types of multimodality. Firstly, genetic drift may lead to the poor sampling of the solution set. Fitness sharing has proved useful in addressing this problem, although it requires that a good closeness measure be found. Secondly, mating of well-performing individuals very different from one another may not be viable, i.e., lead to the production of unfit offspring. Mating restriction in the objective domain, or the absence of mating altogether, interprets the individuals populating the Pareto-front as a continuum of species. It seeks to reduce the formation of lethals by encouraging the formation of offspring similar to their parents, which means a less exploratory search. This is the non-random mating strategy adopted by Hajela and Lin (1992) and Fonseca and Fleming (1993).

The alternative interpretation of the Pareto-set as a genetically similar and, therefore, reproductively viable family of points would require the search

for a suitable genetic representation in addition to the solution itself, because the location of the optima is not known prior to optimization. A fixed genetic representation also produces a reproductively viable family of points, but it does not necessarily correspond to the Pareto-set.

Ridges impose a second type of difficulty. Theoretical results by Wagner (1988) show that, under biologically reasonable assumptions, the rate of progression of unconstrained phenotypes on certain types of ridge-shaped landscapes is bounded, in which case it decreases rapidly as the number of decision variables increases. Fast progression cannot be achieved unless the genetic operators tend to produce individuals which stay inside the corridor. The self-adaptation of mutation variances and correlated mutations (Bäck *et al.*, 1991), as implemented in evolution strategies, addresses this same problem, but has not yet been tried in Pareto-based search. Binary mutation, as usually implemented in genetic algorithms, can be particularly destructive if the ridge expresses a strong correlation between a large number of decision variables. The same applies to the discrete recombination of decision variables, since it can only produce offspring at vertices of the hypercube defined by the mating parents. Similarly, single and two-point crossover of concatenated binary strings will change at most one or two decision variables. Uniform crossover (Syswerda, 1989) and shuffle crossover (Caruana *et al.*, 1989) are less biased in this respect, in that the value of all decision variables may be altered in a single recombination step.

Finally, multiobjective fitness landscapes become non-stationary once the DM is allowed to interact with the search process and change the current preferences, even if the objective functions themselves remain unchanged. Diploidy has already revealed its importance in handling non-stationary environments (Goldberg and Smith, 1987). Other relevant work is the combination of evolutionary and pure random search proposed by Grefenstette (1992).

4 Future perspectives

As discussed in the previous section, the EA can be seen as a sequence of decision making problems, each involving a finite number of alternatives. Current decision making theory, therefore, can certainly provide many answers on how to perform multiobjective selection in the context of EAs.

On the other hand, progress in decision making has always been strongly dependent on the power of the numerical techniques available to support it. Certain decision models, although simple to formulate, do not necessarily lead to numerically easy optimization problems (Dinkelbach, 1980). By easing the numerical difficulties inherent to other optimization methods, evolutionary algorithms open the way to the development of simpler, if not new, decision making approaches.

A very attractive aspect of the multiobjective evolutionary approach is the production of useful intermediate information which can be used by an intelligent DM to refine preferences and terminate the search upon satisfaction. In fact, the DM is not only asked to assess individual performance, but also to adjust the current preferences in the search for a compromise between the ideal and the possible in a limited amount of time. Goal setting, for example, is itself the object of study (Shi and Yu, 1989). This is an area where combinations of EAs and other learning paradigms may be particularly appropriate.

As far as the search strategy is concerned, much work has certainly yet to be done. In particular the emergence of niches in structured populations (Davidor, 1991) suggests the study of such models in the multiobjective case. The development of adaptive representations capable of capturing and exploiting directional trends in the fitness landscape, well advanced in the context of ESs, and/or the corresponding operators, is another important avenue for research. Combinations of genetic search and local optimization resulting in either Lamarckian or developmental Baldwin learning (Gruau and Whitley, 1993) may also provide a means of addressing the difficulties imposed by ridge-shaped landscapes.

The question of which fitness assignment method is better remains largely open, although Pareto-based methods seem more promising for their lack of sensitivity to the possible concavity of the trade-off surface. In the few comparative studies of multiobjective EAs available to date (Wilson and Macleod, 1993; Cieniawski, 1993; Ritzel *et al.*, 1994; Srinivas and Deb, 1994), VEGA has understandably been a strong point of reference, but the comparison has remained largely qualitative. No extensive, quantitative comparison of multiobjective EAs has been reported in the literature so far, which is, however, hardly surprising. Ideally, the quality of every point of the trade-off surface produced should be assessed, meaning that the performance of multiobjective EAs is itself a vector quantity. So, how should the trade-off surfaces

produced by *sets of runs* of different EAs be compared, in a meaningful and, preferably, statistically sound way? Should the scaling of the objectives affect the comparison? These questions need yet to be answered.

In any case, the time may be right for EA users and implementors to consider experimenting with some of the available multiobjective EA techniques on their real-world problems, while not losing sight of any alternative approaches. However, a cautionary word is due here. As noted independently by Horn and Nafpliotis (1993) and Fonseca and Fleming (1993), *pure* Pareto-EAs cannot be expected to perform well on problems involving many competing objectives and may simply fail to produce satisfactory solutions due to the large dimensionality and size of the trade-off surface. As the number of actually competing objectives increases, more and more of the search space can be expected to conform to the definition of Pareto optimality, which makes the “theoretical” problem of finding non-dominated solutions easier! Unfortunately, in the total absence of preference information, the EA will face the “impossible” task of finding a satisfactory compromise in the dark, which can only occur by pure chance. It was the observation of this fact on real-world, engineering problems that prompted Fonseca and Fleming (1993) to combine preference articulation and Pareto-ranking.

Finally, a theory of multiobjective EAs is much needed, ideally incorporating single-objective EAs as a particular case. The study of the fitness assigned to large populations as proposed in the previous section, but considering also non-uniform distributions for the population, may well prove useful in understanding how different selection mechanisms work, and indeed, how EAs based on them may behave, provided that the effect of mutation, recombination, and any other operators used, on the distribution of the population can be modelled as well.

Acknowledgement

The first author gratefully acknowledges support by Programa CIENCIA, Junta Nacional de Investigação Científica e Tecnológica, Portugal (Grant BD/1595/91-IA). The authors also wish to acknowledge the support of the UK Engineering and Physical Sciences Research Council (Grant GR/J70857) in completion of this work. The valuable comments and advice provided by the anonymous reviewers in the preparation of the (final) manuscript is greatly appreciated.

References

- Bäck, T., Hoffmeister, F., and Schwefel, H.-P. (1991). A survey of evolution strategies. In (Belew and Booker, 1991), pages 2–9.
- Baker, J. E. (1987). Reducing bias and inefficiency in the selection algorithm. In (Grefenstette, 1987), pages 14–21.
- Belew, R. K. and Booker, L. B., editors (1991). *Genetic Algorithms: Proceedings of the Fourth International Conference*. Morgan Kaufmann, San Mateo, CA.
- Ben-Tal, A. (1980). Characterization of Pareto and lexicographic optimal solutions. In (Fandel and Gal, 1980), pages 1–11.
- Caruana, R. A., Eshelman, L. J., and Schaffer, J. D. (1989). Representation and hidden bias II: Eliminating defining length bias in genetic search via shuffle crossover. In Sridharan, N. S., editor, *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 750–755. Morgan Kaufmann.
- Cieniawski, S. E. (1993). An investigation of the ability of genetic algorithms to generate the tradeoff curve of a multi-objective groundwater monitoring problem. Master’s thesis, University of Illinois at Urbana-Champaign, Urbana, Illinois.
- Davidor, Y. (1991). A naturally occurring niche and species phenomenon: The model and first results. In (Belew and Booker, 1991), pages 257–263.
- Davis, L. and Steenstrup, M. (1987). Genetic algorithms and simulated annealing: An overview. In Davis, L., editor, *Genetic Algorithms and Simulated Annealing*, Research Notes in Artificial Intelligence, chapter 1, pages 1–11. Pitman, London.
- Deb, K. and Goldberg, D. E. (1989). An investigation of niche and species formation in genetic function optimization. In (Schaffer, 1989), pages 42–50.
- Dinkelbach, W. (1980). Multicriteria decision models with specified goal levels. In (Fandel and Gal, 1980), pages 52–59.

- Fandel, G. and Gal, T., editors (1980). *Multiple Criteria Decision Making Theory and Application*, volume 177 of *Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, Berlin.
- Fleming, P. J. and Pashkevich, A. P. (1985). Computer aided control system design using a multiobjective optimization approach. In *Proc. IEE Control'85 Conference*, pages 174–179, Cambridge, U.K.
- Fonseca, C. M. and Fleming, P. J. (1993). Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. In (Forrest, 1993), pages 416–423.
- Forrest, S., editor (1993). *Genetic Algorithms: Proceedings of the Fifth International Conference*. Morgan Kaufmann, San Mateo, CA.
- Fourman, M. P. (1985). Compaction of symbolic layout using genetic algorithms. In (Grefenstette, 1985), pages 141–153.
- Gembicki, F. W. (1974). *Vector Optimization for Control with Performance and Parameter Sensitivity Indices*. PhD thesis, Case Western Reserve University, Cleveland, Ohio, USA.
- Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, Massachusetts.
- Goldberg, D. E. and Richardson, J. (1987). Genetic algorithms with sharing for multimodal function optimization. In (Grefenstette, 1987), pages 41–49.
- Goldberg, D. E. and Segrest, P. (1987). Finite markov chain analysis of genetic algorithms. In (Grefenstette, 1987), pages 1–8.
- Goldberg, D. E. and Smith, R. E. (1987). Nonstationary function optimization using genetic algorithms with dominance and diploidy. In (Grefenstette, 1987), pages 59–68.
- Grefenstette, J. J., editor (1985). *Genetic Algorithms and Their Applications: Proceedings of the First International Conference on Genetic Algorithms*. Lawrence Erlbaum.

- Grefenstette, J. J., editor (1987). *Genetic Algorithms and Their Applications: Proceedings of the Second International Conference on Genetic Algorithms*. Lawrence Erlbaum.
- Grefenstette, J. J. (1992). Genetic algorithms for changing environments. In (Männer and Manderick, 1992), pages 137–144.
- Gruau, F. and Whitley, D. (1993). Adding learning to the cellular development of neural networks: Evolution and the baldwin effect. *Evolutionary Computation*, 1(3):213–233.
- Hajela, P. and Lin, C.-Y. (1992). Genetic search strategies in multicriterion optimal design. *Structural Optimization*, 4:99–107.
- Horn, J. and Nafpliotis, N. (1993). Multiobjective optimization using the niched Pareto genetic algorithm. IlliGAL Report 93005, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA.
- Horn, J., Nafpliotis, N., and Goldberg, D. E. (1994). 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.
- Jakob, W., Gorges-Schleuter, M., and Blume, C. (1992). Application of genetic algorithms to task planning and learning. In (Männer and Manderick, 1992), pages 291–300.
- Jones, G., Brown, R. D., Clark, D. E., Willet, P., and Glen, R. C. (1993). Searching databases of two-dimensional and three-dimensional chemical structures using genetic algorithms. In (Forrest, 1993), pages 597–602.
- Kursawe, F. (1991). A variant of evolution strategies for vector optimization. In Schwefel, H.-P. and Männer, R., editors, *Parallel Problem Solving from Nature, 1st Workshop, Proceedings*, volume 496 of *Lecture Notes in Computer Science*, pages 193–197. Springer-Verlag, Berlin.
- Louis, S. J. and Rawlins, G. J. E. (1993). Pareto optimality, GA-easiness and deception. In (Forrest, 1993), pages 118–123.

- Männer, R. and Manderick, B., editors (1992). *Parallel Problem Solving from Nature, 2*. North-Holland, Amsterdam.
- Mühlenbein, H. and Schlierkamp-Voosen, D. (1993). Predictive models for the breeder genetic algorithm i. continuous parameter optimization. *Evolutionary Computation*, 1(1):25–49.
- Powell, D. and Skolnick, M. M. (1993). Using genetic algorithms in engineering design optimization with non-linear constraints. In (Forrest, 1993), pages 424–431.
- Richardson, J. T., Palmer, M. R., Liepins, G., and Hilliard, M. (1989). Some guidelines for genetic algorithms with penalty functions. In (Schaffer, 1989), pages 191–197.
- Ritzel, B. J., Eheart, J. W., and Ranjithan, S. (1994). Using genetic algorithms to solve a multiple objective groundwater pollution containment problem. *Water Resources Research*, 30(5):1589–1603.
- Schaffer, J. D. (1985). Multiple objective optimization with vector evaluated genetic algorithms. In (Grefenstette, 1985), pages 93–100.
- Schaffer, J. D., editor (1989). *Proceedings of the Third International Conference on Genetic Algorithms*. Morgan Kaufmann, San Mateo, CA.
- Schaffer, J. D. (1993). Personal communication.
- Schaffer, J. D. and Grefenstette, J. J. (1985). Multi-objective learning via genetic algorithms. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 593–595. Morgan Kaufmann.
- Shi, Y. and Yu, P. L. (1989). Goal setting and compromise solutions. In Karpak, B. and Zionts, S., editors, *Multiple Criteria Decision Making and Risk Analysis Using Microcomputers*, volume 56 of *NATO ASI Series F: Computer and Systems Sciences*, pages 165–204. Springer-Verlag, Berlin.
- Silverman, B. W. (1986). *Density Estimation for Statistics and Data Analysis*, volume 26 of *Monographs on Statistics and Applied Probability*. Chapman and Hall, London.

- Srinivas, N. and Deb, K. (1994). Multiobjective optimization using nondominated sorting in genetic algorithms. *Evolutionary Computation*, 2(3). To appear.
- Steuer, R. E. (1986). *Multiple Criteria Optimization: Theory, Computation, and Application*. Wiley Series in Probability and Mathematical Statistics. John Wiley and Sons, New York.
- Syswerda, G. (1989). Uniform crossover in genetic algorithms. In (Schaffer, 1989), pages 2–9.
- Syswerda, G. and Palmucci, J. (1991). The application of genetic algorithms to resource scheduling. In (Belew and Booker, 1991), pages 502–508.
- Wagner, G. P. (1988). The influence of variation and of developmental constraints on the rate of multivariate phenotypic evolution. *Journal of Evolutionary Biology*, 1(1):45–66.
- Wienke, D., Lucasius, C., and Kateman, G. (1992). Multicriteria target vector optimization of analytical procedures using a genetic algorithm. Part I. Theory, numerical simulations and application to atomic emission spectroscopy. *Analytica Chimica Acta*, 265(2):211–225.
- Wilson, P. B. and Macleod, M. D. (1993). Low implementation cost IIR digital filter design using genetic algorithms. In *IEE/IEEE Workshop on Natural Algorithms in Signal Processing*, volume 1, pages 4/1–4/8, Chelmsford, U.K.