Multi-Objective Optimization Using Metaheuristics

Carlos A. Coello Coello

ccoello@cs.cinvestav.mx
CINVESTAV-IPN
Evolutionary Computation Group (EVOCINV)
Computer Science Department
Av. IPN No. 2508, Col. San Pedro Zacatenco
México, D.F. 07360, MEXICO

Eindhoven University of Technology
Eindhoven, The Netherlands July 2023

LECTURE 3

Techniques to Maintain Diversity

Diversity

Since the early days of evolutionary computation, researchers working in this field, realized that evolutionary algorithms tend to converge to a single solution because of stochastic noise (digital computers can only provide pseudo-random numbers). Since then, maintaining diversity in the population of an evolutionary algorithm has been a fundamental research topic in evolutionary computation.

Techniques to Maintain Diversity

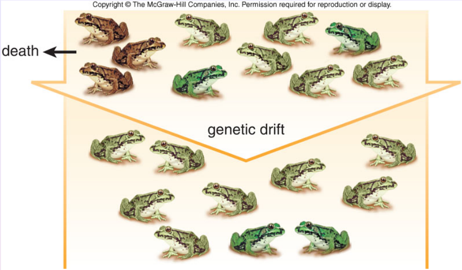

Genetic Drift

Loss of diversity is a pheonomenon that also occurs in nature and it is called **genetic drift**.

This problem has been studied for a long time, and several mechanisms to preserve diversity in the population of an evolutionary algorithm have been proposed.

Techniques to Maintain Diversity

Crowding

Holland [1975] proposed an operator called **crowding** which was able to identify situations in which more individuals dominate a certain ecological niche, since it is precisely in those cases in which the competition for limited resources rapidly increases, giving rise to a lower life expectancy and to a lower birth rate.

Techniques to Maintain Diversity

Crowding

De Jong [1975] experimented with Holland's crowding operator. For this sake, De Jong used a non-generational genetic algorithm that worked in the following way: a fraction of the population (defined by a parameter called **generational gap (GG)** was selected using proportional selection, so that it was subject to crossover and mutation.

Techniques to Maintain Diversity

Crowding

After applying crossover and mutation, $GG \times n$ individuals were selected from the population to die (i.e., they were replaced by the new offpsring). Each offspring finds the individual that it will replace by taking a random sample of CF individuals in the population (CF is the so-called **crowding factor**). Each offspring replaces the individual which is most similar to it in the population.

Techniques to Maintain Diversity

Crowding

A value of $CF = 1$ indicates that no crowding will take place. As the value of CF increases, it is more likely that similar individuals replace among themselves.

Similarity is measured, in this case, using Hamming distances of the individual's genotypes.

Techniques to Maintain Diversity

Cavicchio's Work

Cavicchio [1970] proposed several pre-selection schemes from which one was oriented to preserve diversity.

Cavicchio's idea was that if an offspring had a higher fitness than the worst parent, then this offspring would replace that worst parent.

D.J. Cavicchio, "**Adaptive search using simulated evolution**", PhD thesis, University of Michigan, Ann Arbor, USA, 1970.

Techniques to Maintain Diversity

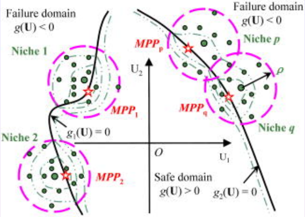

Fitness Sharing

Fitness sharing was originally proposed by Goldberg and Richardson [1987]. In this case, the population is subdivided into several subpopulations based on the similarity of the individuals in either of two possible spaces: **phenotypic** (i.e., decoded parameters) or **genotypic** (binary encoding).

D.E. Goldberg and J. Richardson, "**Genetic algorithms with sharing for multimodal function optimization**", in *Proceedings of the Second International Conference on Genetic algorithms and their Application*, pp. 41–49, Lawrence Erlbaum Associates Inc., Hillsdale, NJ, USA, 1987.

Techniques to Maintain Diversity

Fitness Sharing

Fitness sharing is defined in the following way:

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

where: $\alpha = 1$, d_{ij} is a metric that indicates the distance between solutions i and j and σ_{share} is the sharing parameter (or threshold) that controls the size of the niche.

Techniques to Maintain Diversity

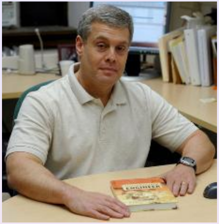

Fitness Sharing

Using the sharing parameter, the fitness of an individual i is modified using:

$$f_{s_i} = \frac{f_i}{\sum_{j=1}^{M} \phi(d_{ij})} \tag{2}$$

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

Techniques to Maintain Diversity

Fitness Sharing

Deb & Goldberg [1989] proposed a methodology to compute σ_{share}. In phenotypic space, they adopt an Euclidean distance in a p-dimensional space, where p refers to the number of decision variables encoded in the evolutionary algorithm.

Kalyanmoy Deb and David E. Goldberg, "**An investigation of niche and species formation in genetic function optimization**", in J. David Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms*, pp. 42–50, Morgan Kaufmann Publishers, San Mateo, California, USA, June 1989.

Techniques to Maintain Diversity

In Deb and Goldberg [1989], the value of d_{ij} is computed using:

$$d_{ij} = \sqrt{\sum_{k=1}^{p} (x_{k,i} - x_{k,j})^2} \tag{3}$$

where $x_{1,i}, x_{2,i}, \ldots, x_{p,i}$ and $x_{1,j}, x_{2,j}, \ldots, x_{p,j}$ are the decoded variables.

Techniques to Maintain Diversity

Fitness Sharing

In order to estimate the value of σ_{share}, Deb and Goldberg [1989] proposed the following expression:

$$\sigma_{share} = \frac{r}{\sqrt[p]{q}} = \frac{\sqrt{\sum_{k=1}^{p}(x_{k,max} - x_{k,min})^2}}{\sqrt[p]{2q}} \tag{4}$$

where r is the volume of a p-dimensional sphere of radius σ_{share} and q is the number of niches that the evolutionary algorithm pretends to find.

Techniques to Maintain Diversity

Fitness Sharing

In genotypic fitness sharing, d_i is defined as the Hamming distance between the strings and σ_{share} is the maximum number of different bits that are allowed between strings to form separate niches in the population.

Techniques to Maintain Diversity

Fitness Sharing

The expression that they suggest to use is:

$$\frac{1}{2^l}\sum_{i=0}^{k}\binom{l}{i}=\frac{1}{q} \tag{5}$$

where: l is the length of the chromosomic string (in bits), k is the maximum difference (in bits) allowed between sub-strings to produce q sub-divisions of the solutions space (in other words, $\sigma_{share}=k$).

Techniques to Maintain Diversity

Fitness Sharing

For large values of l, they suggest to use:

$$\sigma_{share} = \frac{1}{2}(l + z^* \sqrt{l}) \tag{6}$$

where: z^* is the normalized difference (in bits) corresponding to $\frac{1}{q}$ of the total probability space.

The value of z^* can be obtained from a cumulative normal distribution chart.

Techniques to Maintain Diversity

Fitness Sharing

The experiments conducted by Deb and Goldberg [1989] indicated that fitness sharing was better than crowding.

Additionally, these experiments indicated that fitness sharing worked better in phenotypic space than in genotypic space.

Techniques to Maintain Diversity

Fitness Sharing

Some researchers within evolutionary multi-objective optimization have also proposed their own techniques to compute σ_{share}.

For example, Fonseca & Fleming [1993] proposed:

$$N = \frac{\prod_{i=1}^{k}(\Delta_i + \sigma_{share}) - \prod_{i=1}^{k}\Delta_i}{\sigma_{share}^{k}}, \tag{7}$$

where N is the population size, Δ_i is the difference between the maximum and minimum objective function values in dimension i and k is the number of objectives of the problem.

Techniques to Maintain Diversity

Deterministic Crowding

It is worth noting, however, that Mahfoud [1992] showed that the conclusions from Deb and Goldberg [1989] are no longer correct if a few minor changes are introduced to the crowding scheme analyzed in their paper.

Techniques to Maintain Diversity

Deterministic Crowding

Mahfoud [1992] proposed a crowding algorithm that tends to minimize the replacement errors from De Jong's proposal (which was one of its main drawbacks).

S.W. Mahfoud, "**Crowding and preselection revisited**", in Reinhard Männer and Bernard Manderick (Editors), *Parallel Problem Solving from Nature 2*, pp. 27–36, North-Holland, Amsterdam, 1992.

Techniques to Maintain Diversity

Deterministic Crowding

This algorithm was called **deterministic crowding** and its main feature was that introduces a competition between parents and offspring that are located in identical niches.

After performing crossover and mutation, each offspring replaces its nearest parent if the offspring has a better fitness than its parent.

Samir W. Mahfoud, "**Niching Methods for Genetic Algorithms**", PhD thesis, University of Illinois at Urbana-Champaign, 1995.

Techniques to Maintain Diversity

Deterministic Crowding

Mahfoud’s algorithm starts by grouping all the elements from the population into $N/2$ pairs.

Then, these pairs are recombined and the resulting offspring are mutated. For each pair of offspring, two sets of tournaments parent-offspring are possible: (1) parent 1 vs. offspring 1, (2) parent 2 vs. offspring 2, (3) parent 1 vs. offspring 2 and (4) parent 2 vs. offspring 1.

Determinitic crowding performs these tournaments in such a way that the most similar individuals are forced to compete between themselves.

The algorithmic complexity of deterministic crowding is $O(N)$.

Techniques to Maintain Diversity

Probabilistic Crowding

Mengshoel and Goldberg [2008] proposed a niching technique called **probabilistic crowding**.

Ole J. Mengshoel and David E. Goldberg, “**The Crowding Approach to Niching in Genetic Algorithms**”, *Evolutionary Computation*, Vol. 16, No. 3, pp. 315–354, Fall 2008.

Techniques to Maintain Diversity

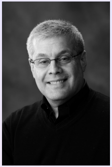

Probabilistic Crowding

This approach is similar to Mahfoud's deterministic crowding, but now, a probabilistic acceptance function is adopted (in the original approach, the fittest individuals always won the tournaments).

Mengshoel and Goldberg [2008] showed that this approach produces more stable sub-populations than those obtained with deterministic crowding.

Techniques to Maintain Diversity

Cedeño [1995] and Harik [1995] proposed algorithms with a similar behavior as Mahfoud's approach.

In the case of Cedeño et al. [1995], they suggest the use of phenotypic crossover and specialized operators to reduce the replacement error of the technique.

Walter Cedeño, "**The Multi-Niche Crowding Genetic Algorithm: Analysis and Applications**", PhD thesis, University of California at Davis, USA, 1995.

Techniques to Maintain Diversity

Restricted Tournament Selection

Harik [1995] proposed an approach called **Restricted Tournament Selection** (RTS) for multimodal optimization.

Georges R. Harik, “**Finding Multimodal Solutions Using Restricted Tournament Selection**”, in *Proceedings of the 6th International Conference on Genetic Algorithms*, pp. 24–31, Morgan Kaufmann Publishers, San Mateo, California, USA, 1995.

Techniques to Maintain Diversity

Restricted Tournament Selection

In this case, two elements from the population are selected first, and they are recombined and mutated.

After recombination, a random sample of CF individuals are taken from the population, as in the original crowding algorithm.

Each offspring competes with the nearest individual in the sample. The winners are inserted into the population. This procedure is repeated $N/2$ times.

This algorithm has a complexity $O(CF \cdot N)$ and can vary from $O(N)$ to $O(N^2)$.

Techniques to Maintain Diversity

Clearing

Pétrowski [1995] proposed an approach called **clearing**, which is similar to fitness sharing, but it's based on the notion of limited existing resources in the environment.

In this case, instead of sharing the resources among all the individuals in a sub-population, they are shared only among the best of them.

Alain Pétrowski, "**A Clearing Procedure as a Niching Method for Genetic Algorithms**", *Proceedings of 1996 IEEE International Conference on Evolutionary Computation (ICEC'96)*, pp. 798–803, IEEE Press, Nagoya, Japan, 1996.

Techniques to Maintain Diversity

Clearing

In practice, this approach preserves the fitness of the *k* best individuals (the dominant ones) of the niche and resets the fitness of the others that belong to the same population (the dominated ones).

Techniques to Maintain Diversity

Clearing

As in fitness sharing, individuals belonging to the same niche (or sub-population) are penalized if their distance in the search space is less than a certain similarity threshold σ_s (this is called the **clearing radius**).

This technique can be coupled with an elitist strategy to preserve the best elements in a niche throughout generations.

This algorithm has a complexity $O(qN)$, where q is the number of niches mantained during the search process.

Techniques to Maintain Diversity

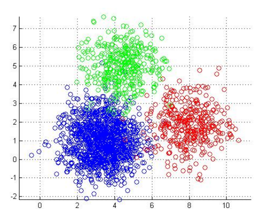

Clustering

Clustering algorithms have also been frequently used to implement niching approaches.

Yin and Germay [1993] proposed a framework for implementing niches using clustering algorithms.

A clustering algorithm such as $K - Means$ [Haykin, 1999] first divides the population into several clusters and then considers the centroids of the recently partitioned sub-populations.

Techniques to Maintain Diversity

Clustering

Let's make d_{ic} denotes the distance between individual i and its *centroid*, and that f_i denotes the original fitness (*raw fitness*) of individual i. Let's assume that there are n_c individuals in the same niche of individual i. Its fitness is then defined as:

$$f_i^{Clustering} = \frac{f_i}{n_c \cdot (1 - (d_{ic}/2d_{max})^{\alpha}} \tag{8}$$

where d_{max} is the maximum distance allowed between an individual and the centroid of its niche, and α is a user-defined parameter.

Xiaodong Yin and Noël Germay, "**A fast genetic algorithm with sharing scheme using cluster analysis methods in multimodal function optimization**", in R.F. Albrecht, C.R. Reeves and N.C. Steele (Editors), *Proceedings of the International Conference on Artificial Neural Nets and Genetic Algorithms*, p. 450–457, Springer, 1993.

Techniques to Maintain Diversity

Thermodynamical Genetic Algorithm

Kita et al. [1996] proposed the **Thermodynamical Genetic Algorithm** (TDGA) which adopts Pareto ranking, combined with the principle of minimal free energy that is adopted in simulated annealing.

Hajime Kita, Yasuyuki Yabumoto, Naoki Mori and Yoshikazu Nishikawa, "**Multi-Objective Optimization by Means of the Thermodynamical Genetic Algorithm**", in Hans-Michael Voigt et al. (Editors), *Parallel Problem Solving from Nature–PPSN IV*, Springer, Lecture Notes in Computer Science, pp. 504–512, Berlin, Germany, September 1996.

Techniques to Maintain Diversity

Thermodynamical Genetic Algorithm

The core idea in the TDGA is to select the individuals of the next generation in such a way that the free energy is minimized.

The free energy is given by:

$$F = < E > - HT \tag{9}$$

where: $< E >$ is the average energy in the system, H is the entropy and T is the temperature.

Techniques to Maintain Diversity

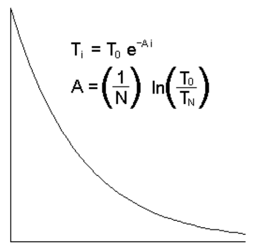

Thermodynamical Genetic Algorithm

Diversity of the population is controlled by adjusting *T* according to certain (given) cooling schedule.

According to Kita et al. [1996], *T* is less sensitive to the population size and to the size of the feasible region than traditional fitness sharing functions.

Techniques to Maintain Diversity

Emphasis on Efficiency

Over the years, a great emphasis to efficiency has been given to the research on niching.

For example, Yin and Germay [1993] proposed to alternate the use of clustering with niching.

They indicated that this sort of scheme would reduce the algorithmic complexity of the approach from $O(n^2)$ to $O(nq)$, where q is the number of niches.

Techniques to Maintain Diversity

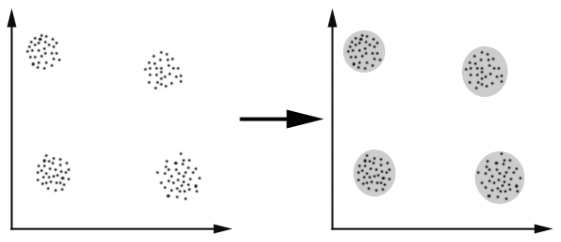

Emphasis on Efficiency

It is worth noting, however, that fitness sharing algorithms can work with an algorithmic complexity $O(nq)$ when the shared fitness values are sampled from a sub-population of size $O(q)$ [Oei et al., 1991].

Techniques to Maintain Diversity

Dynamic Niching

There have also been proposals of dynamic niching schemes. For example, Miller and Shaw [1996] proposed the use of a fixed number of dynamic niches with fixed radii and niche centers. Such values would be determined through a sorting procedure applied to the entire population.

Brad L. Miller and Michael J. Shaw, "**Genetic Algorithms with Dynamic Niche Sharing for Multimodal Function Optimization**", in *Proceedings of the 1996 IEEE International Conference on Evolutionary Computation (ICEC'96)*", pp. 786–791, IEEE Press, May 1996.

Techniques to Maintain Diversity

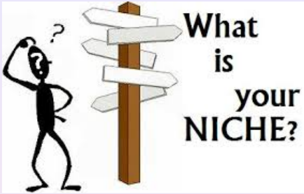

Dynamic Niching

For those individuals that are not located within any of these niches, conventional fitness sharing is applied.

This reduces the computational cost because it is no longer required to re-compute the niches except when the individuals cannot be placed in any of the previously defined niches.

Techniques to Maintain Diversity

Coevolutionary Niching

Goldberg and Wang [1998] proposed a coevolutionary scheme to maintain diversity.

This scheme is inspired on a monopolic competition as modelled by economists. The algorithm uses two populations: one containing business men and another one containing customers.

These populations are set up in such a way that the location of the business men corresponds to the locations of the niches and the locations of the customers corresponds to the solutions.

David E. Goldberg and Liwei Wang, "**Adaptive Niching via Coevolutionary Sharing**", in D. Quagliarella et al. (eds.), *Genetic Algorithms and Evolution Strategies in Engineering and Computer Science. Recent Advances and Industrial Applications*, Chapter 2, pp. 21–38, John Wiley & Sons, Chichester, UK, 1998.

Techniques to Maintain Diversity

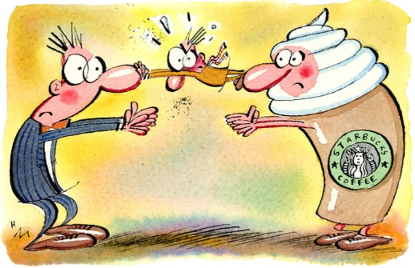

Coevolutionary Niching

In this proposal, the location and radii of the niches automatically adapt to the specific features of the fitness landscape at the current generation.

Techniques to Maintain Diversity

Tagging

Spears [1994] proposed a mechanism called **tagging**, which aims to improve the performance of the techniques to maintain diversity that are based on distances (e.g., fitness sharing) through a labeling system for the individuals.

W.M. Spears, “**Simple subpopulation schemes**”, in A.V. Sebald and L.J. Fogel (Editors), *Third Annual Conference on Evolutionary Programming*, IEEE Press, pp. 296–307, San Diego, California, USA, 1994.

Techniques to Maintain Diversity

Tagging

The idea of this approach is that, instead of computing distances, labels are adopted to identify the sub-population to which they belong. This enforces a mating restriction, and gives rise to a fitness sharing scheme.

Techniques to Maintain Diversity

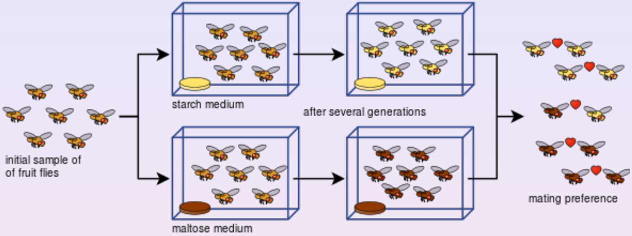

Tagging

This is a very interesting concept, which simplifies the process of classifying individuals and, evidently, reduces the computational cost.

In the process, this approach constitutes a new bio-inspired technique for niche formation: under this scheme, individuals belong to a certain species because their parents were members of that species, and not because they are adjacent to an individual that is located in the peak of the function being optimized.

Techniques to Maintain Diversity

Multinational Evolutionary Algorithm

Ursem [1999] proposed a complex model for differentiating sub-populations, which was called **multinational evolutionary algorithm**.

Rasmus K. Ursem, “**Multinational evolutionary algorithms**”, in *Proceedings of the 1999 IEEE Congress on Evolutionary Computation (CEC’99)*, pp. 1633-1640, IEEE Press, July 1999.

Techniques to Maintain Diversity

Multinational Evolutionary Algorithm

This approach considers a world of “nations”, “governments” and “politicians”, with dynamics dictated by the migration of individuals, the union of sub-populations and the selection scheme.

Additionally, it introduces an auxiliary sampling mechanism based on a topology, which detects if a feasible solution shares the same attraction point.

Although interesting, the main drawback of this approach is that it becomes inefficient in high-dimensional problems.

Techniques to Maintain Diversity

Elitist Generational Genetic Chromodynamics Algorithm

Stoean et al. [2005] proposed the **elitist generational genetic chromodynamics algorithm**, which is a technique based on radii.

The core idea of this approach is to define a mating region (with its corresponding mating radii, replacing radii and union radii), which dictates the dynamic of the genetic operators.

Catalin Stoean, Mike Preuss, Ruxandra Gorunescu and D. Dumitrescu, “**Elitist generational genetic chromodynamics - a new radii-based evolutionary algorithm for multimodal optimization**”, in *2005 IEEE Congress on Evolutionary Computation (CEC’2005)*, pp. 1839–1846, IEEE Press, September 2005.

Techniques to Maintain Diversity

S-Energy

The *s*-energy indicator was proposed by Hardin and Saff [2004] and is defined as follows:

$$E_s(\mathcal{A}) = \sum_{i \neq j} \left\| \vec{a}_i - \vec{a}_j \right\|^{-s} \qquad (10)$$

where $\mathcal{A} = \{\vec{a}_1, \ldots, \vec{a}_{|\mathcal{A}|}\}$, $\vec{a}_i \in \mathbb{R}^m$, and $s > 0$ is a fixed parameter. The aim is to minimize its value because it produces a uniform distribution of points. Normally, $s = m$ (where m is the number of objectives).

D. P. Hardin and E. B. Saff, "**Discretizing Manifolds via Minimum Energy Points**", *Notices of the AMS*, Vol. 51, No. 10, pp. 1186–1194, November 2004.

Techniques to Maintain Diversity

Moment of Inertia

Morrison and De Jong [2001] proposed a new measure of population diversity which is based on an extension of the concept of inertia for measurement of mass distribution into arbitrarily high dimensionality spaces. Extending into n-space, the coordinates of the centroid of P equally weighted points in n-space, $C = (c_1, c_2, c_3, \ldots c_n)$ are computed using:

$$c_i = \frac{\sum_{j=1}^{j=P} x_{ij}}{P} \tag{11}$$

where $x_{ij} \in \mathbf{R}$ and c_i is the i^{th} coordinate of the centroid.

Techniques to Maintain Diversity

Moment of Inertia

Continuing with P equally-weighted points in n-space, the moment-of-inertia based measure of diversity of these points about their centroid is:

$$I = \sum_{i=1}^{i=n} \sum_{j=1}^{j=P} (x_{ij} - c_i)^2 \tag{12}$$

According to its authors, the main advantage of this measure of diversity is that, in comparison with traditional methods of computing pair-wise population diversity, which are quadratic on population size (P), this method is linear in P (it requires a total of $4(nP) + n$ calculations).

Ronald W. Morrison and Kenneth A. De Jong, "**Measurement of Population Diversity**", in Pierre Collet et al. (Editors) *Artificial Evolution, 5th International Conference, Evolution Artificielle, EA 2001*, pp, 31–41, Springer, Lecture Notes in Computer Science Vol. 2310, Le Creusot, France, October 29-31, 2001.

Techniques to Maintain Diversity

Ectropy

Lacevic and Almadi [2011] point out that measures based on the sum (average) of pairwise distances between points in $\mathbf{R}^m$ have substantial shortcomings. For instance, these measures reach their maximum value when the population consists of very few (sometimes only two) mutually distant clusters of collocated points. They propose a notion called *ectropy* (which is the opposite from *entropy*, which can be seen as the measure of disorder in the population). This concept aims for reaching a near-maximum value for a population that has many collocated points (e.g., consists of several clusters of collocated points). This property is clearly in contrast with the (more intuitive) notion of diversity measure, because the collocation of points is rewarded instead of being penalized.

Bakir Lacevic and Edoardo Amaldi, "**Ectropy of diversity measures for populations in Euclidean space**", *Information Sciences*, Vol. 181, No. 11, pp. 2316–2339, June 2011.

Techniques to Maintain Diversity

Mating Restrictions

The idea of using mating restrictions is not new. Goldberg [1989] mentions its use in single-objective optimization as a means to prevent or to minimize the appearence of the so-called “lethals” (offspring with a low fitness value).

In other words, mating restrictions bias the way in which individuals mate for recombination purposes. The aim is to increase the effectiveness and the efficiency of the evolutionary algorithm.

Goldberg [1989] presents an example using genotypic similarities as a mating criterion.

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

Techniques to Maintain Diversity

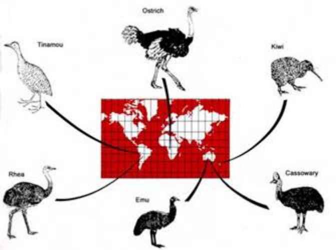

Mating Restrictions

Biologically, mating restrictions are equivalent to geographical isolation, or to setting up a barrier to restrict the flow of genes. This allows to complete the speciation stage.

Geographical isolation is a key component of evolution, since it creates the basic conditions for speciation to occur (i.e., for new species to appear). Therefore, it should not be surprising that several niching techniques rely on mating restrictions.

Techniques to Maintain Diversity

Mating Restrictions

Goldberg and Deb [1989] suggested the use of mating restriction based on phenotypic distances.

The idea is to allow that two individuals recombine only if they are very similar (i.e., if their phenotypic distance is less than a certain user-defined threshold, called σ_{mate}, which is measured with some metric).

Techniques to Maintain Diversity

Mating Restrictions

This intends to produce different “species” (mating groups) in the population [Mitchell, 1996].

Parallel genetic algorithms with an island model also use mating restrictions which are defined in a geographical sense. In this case, an individual can only recombine with its neighborhood (i.e., the other individuals in its own sub-population) according to the interconnection topology.

Melanie Mitchell, **An Introduction to Genetic Algorithms**, The MIT Press, Cambridge, Massachusetts, USA, 1996.

Techniques to Maintain Diversity

Mating Restrictions

Some researchers have indicated that mating restrictions should motivate the recombination of different individuals with the aim of preventing the generation of “lethals”.

Regardless of the restriction criterion adopted, severam MOEAs incorporate mating restrictions whose aim is to reduce the number of dominated solutions in the population (this is the case, for example, of MOGA [Fonseca and Fleming, 1993]).

Techniques to Maintain Diversity

Mating Restrictions

For example, Baita [1995] and Loughlin & Ranjithan [1997] placed the solutions in a grid and restrict the area within which individuals can recombine.

Lis and Eiben [1996] only allow the recombination of individuals that have a different “gender”.

Jakob et al. [1992] implemented a rather atypical mating restriction mechanism based on the values of the weights of each solution (they adopt a linear aggregating function in this case).

Techniques to Maintain Diversity

Mating Restrictions

Although in many MOEAs, it is common practice to adopt $\sigma_{mate} = \sigma_{share}$, there are no studies that indicate that this is the most appropriate setting for the mating threshold.

In fact, most EMOO researchers who adopt mating restrictions don't normally justify the incorporation of this sort of mechanism and in most cases, no empirical evidence of their effectiveness is provided.

Techniques to Maintain Diversity

Mating Restrictions

There are, indeed few studies that provide empirical evidence of the usefulness of mating restrictions in evolutionary multi-objective optimization.

Zitzler and Thiele [1998] showed that, for the different values of σ_{mate} with which they experimented, no actual improvement on the performance of their MOEA (SPEA) was detected, when solving a certain set of test problems.

In this case, the presence or absence of mating restrictions didn't produce any statistically significant difference.

Eckart Zitzler and Lothar Thiele, "**Multiobjective Optimization Using Evolutionary Algorithms–A Comparative Study**", in A.E. Eiben (Ed.), *Parallel Problem Solving from Nature V*, Springer-Verlag, pp. 292–301, Amsterdam, September 1998.

Techniques to Maintain Diversity

Mating Restrictions

Shaw and Fleming [1996] reported results similar to those of Zitzler.

Horn et al. [1994] offered empirical evidence that directly contradicts the foundational principles for using mating restrictions.

K.J. Shaw and P.J. Fleming, “**An Initial Study of Practical Multi-Objective Production Scheduling using Genetic Algorithms**, in *Proceedings of the International Conference on Control’96*, University of Exeter, UK, September 1996.

Techniques to Maintain Diversity

Mating Restrictions

Horn [1994] indicated that recombining solutions whose associated vectors are in different portions of $PF_{known}(t)$ can INDEED produce offspring whose vectors are in $PF_{known}(t+1)$.

However, he also indicated that these offspring will be located between their parents.

Jeffrey Horn, Nicholas Nafpliotis and David E. Goldberg, “**A Niched Pareto Genetic Algorithm for Multiobjective Optimization**”, in *Proceedings of the First IEEE Conference on Evolutionary Computation*, Vol. 1, pp. 82–87, IEEE Press, Piscataway, New Jersey, USA, June 1994.

Techniques to Maintain Diversity

Mating Restrictions

Horn [1994] also indicates that for a given multi-objective optimization problem, the constant re-generation of vectors through the recombination of “dissimilar” parents, maintains solutions in PF_{known}.

He also claims that most of the recombinations in P_{known} produce solutions that are also in P_{known}.

Techniques to Maintain Diversity

Mating Restrictions

Based on what we have seen about mating restrictions, we can conclude that, as in single-objective optimization, it is unclear if they are at all beneficial.

One would expect this mechanism to be useful in certain situations, but it is still unclear when.

In any case, mating restriction are another technique to maintain diversity which, however, is rarely used these days.

Most modern MOEAs don’t adopt any sort of mating restriction mechanism.

Techniques to Maintain Diversity

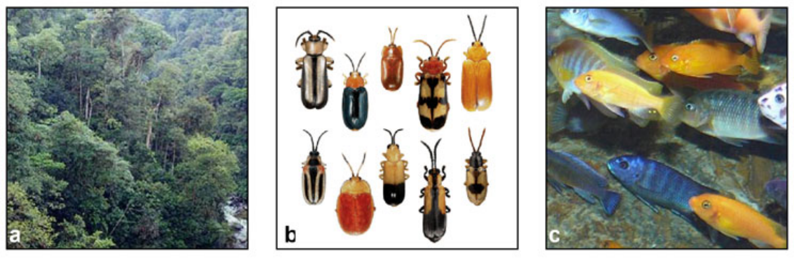

Many other techniques to maintain diversity exist (see for example [Mahfoud, 1995], but only a small percentage of them have been used in the context of evolutionary multi-objective optimization.

Samir W. Mahfoud, "**Niching Methods for Genetic Algorithms**", PhD thesis, University of Illinois at Urbana-Champaign, 1995.

Techniques to Maintain Diversity

There is no solid evidence regarding the advantages or disadvantages of some particular technique to maintain diversity over the others in the context of evolutionary multi-objective optimization.

The information that we have in this regard is fairly limited. For example, we know that fitness sharing can be used with any number of objectives, unlike other mechanisms (for example the crowded-comparison operator of NSGA-II or the adaptive grid from PAES).

Techniques to Maintain Diversity

It is also unclear the space in which fitness sharing should be applied.

Horn et al. [1993] wisely indicated that fitness sharing should be applied in the space that we “care the most”.

In evolutionary multi-objective optimization, it is very common to apply it in phenotypic space, since the main concern is normally a uniform distribution of the Pareto front.

However, in Operations Research, it is common to aim for a uniform distribution in decision variables space [Benson & Sayin, 1997].

Techniques to Maintain Diversity

For more information on this topic, see:

- Ofer M. Shir, **Niching in Evolutionary Algorithms**, in G. Rozenberg et al. (Editors), *Handbook of Natural Computing*, Chapter 32, pp. 1035–1069, Springer-Verlag, Berlin, Germany, 2012.
- Bruno Sareni and Laurent Krähenbühl, **Fitness Sharing and Niching Methods Revisited**, *IEEE Transactions on Evolutionary Computation*, Vol. 2, No. 3, pp. 97–106, September 1998.
- N.N. Glibovets and N.M. Gulayeva, "**A Review of Niching Genetic Algorithms for Multimodal Function Optimization**", *Cybernetics and Systems Analysis*, Vol. 49, No. 6, pp. 815–820, November 2013.

Secondary Populations

Horn [1997] indicates that every implementation of a MOEA should adopt a **secondary population** (also called **external archive**).

A secondary population is required to preserve the nondominated solutions generated by a MOEA at each iteration. Without it, such solutions could be lost when we apply the genetic operators to them. Secondary populations are also required for theoretical reasons (i.e., to guarantee convergence [Rudolph and Agapie, 2000]).

Jeffrey Horn, “**Multicriterion Decision Making**, in Thomas Bäck, David Fogel and Zbigniew Michalewicz (editors), *Handbook of Evolutionary Computation*, pp. F1.9:1 - F1.9:15, Vol. 1, IOP Publishing Ltd. and Oxford University Press, 1997.

Günter Rudolph and Alexandru Agapie, “**Convergence Properties of Some Multi-Objective Evolutionary Algorithms**”, in *Proceedings of the 2000 IEEE Conference on Evolutionary Computation*, Vol. 2, pp. 1010–1016, IEEE Press, Piscataway, New Jersey, USA, July 2000.

Secondary Populations

Nowadays, nobody questions the use of a secondary population in a MOEA. However, their use raises several issues. For example:

- Should the secondary population participate in the selection process of the MOEA (as done in SPEA)?
- What sort of filtering scheme should we use to bound the size of the secondary population (if the size of the secondary population grows too much, and it is used in the selection process, then the selection pressure will quickly dilute)?
- Can we use the secondary population as a density estimator?

The reason for which a secondary population is required is because a MOEA is trying to construct a discrete image of a Pareto front that is probably continuous. Therefore, it is recommended to keep as many nondominated solutions as possible, at least at the beginning of the search. This will allow us to produce Pareto fronts with well-distributed solutions.

Secondary Populations

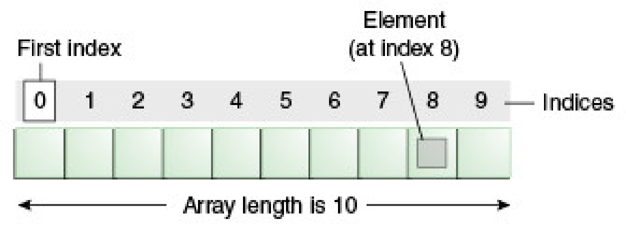

Although early MOEAs used secondary populations that were implemented as linear lists to store nondominated solutions, over the years, several researchers have proposed different types of data structures for this sake.

Secondary Populations

For example, Mostaghim et al. [2002] proposed to use quadtrees [Finkel and Bentley, 1974] in SPEA. This work is based on a paper from Habenicht [1983] which introduced the use of quadtrees to identify nondominated solutions. There is also a follow-up of this work, by Sun and Steuer [1996] in which Habenicht's algorithm is improved.

Sanaz Mostaghim, Jürgen Teich and Ambrish Tyagi, "**Comparison of Data Structures for Storing Pareto-sets in MOEAs**, in *2002 IEEE Congress on Evolutionary Computation (CEC'2002)*, Vol. 1, pp. 843–848, IEEE Service Center, Piscataway, New Jersey, USA, May 2002.

Secondary Populations

Mostaghim et al. [2002] compared the arrays (or linear lists) with three types of quadtrees: **quadtree1** (Habenicht's algorithm [1983]), **quadtree2** (a version proposed by Mostaghim et al. [2002] which uses flags to delete dominated solutions) and **quadtree3** (the algorithm of Sun and Steuer [1996]).

This study shows that the use of quadtrees is beneficial for archives with at least 5,000 solutions, because with smaller archive sizes, the use of an array is much more efficient.

Something interesting is that Mostaghim et al. [2002] report that **quadtree1** is the best algorithm from the three used in the comparison for archives with 5,000 or more solutions. This is remarkable if we consider that the other two algorithms are supposed to be improved versions of this one.

Secondary Populations

Fieldsend et al. [2003] indicate that bounding the size of the external population can produce a “shrinkage” and an “oscillation” phenomena in the Pareto front approximations that we produce.

On the other hand, they recognize that the main problem related to the use of unbounded archives is their high computational cost.

Thus, they proposed the use of two new data structures (the **nondominated trees** and the **PQRS trees**) for storing and retrieving solutions from an unbounded archive in an efficient manner.

Secondary Populations

Something interesting is that they also proposed the use of a secondary population for defining a stopping criterion for a MOEA.

Jonathan E. Fieldsend, Richard M. Everson and Sameer Singh, "**Using Unconstrained Elite Archives for Multiobjective Optimization**", *IEEE Transactions on Evolutionary Computation*, Vol. 7, No. 3, pp. 305–323, June 2003.

Secondary Populations

Laumanns et al. [2002] proposed a relaxed form of Pareto dominance called ε-dominance, which is used as an archiving technique that allows to filter out solutions generated by a MOEA.

The core idea is to define a set of boxes of size ε and to allow only one solution in each of these boxes (e.g., from the nondominated solutions inside a box, we only keep the one that is closest to the low lefthand corner).

Marco Laumanns, Lothar Thiele, Kalyanmoy Deb and Eckart Zitzler, “**Combining Convergence and Diversity in Evolutionary Multi-objective Optimization**”, *Evolutionary Computation*, Vol. 10, No. 3, pp. 263–282, Fall 2002.

Secondary Populations

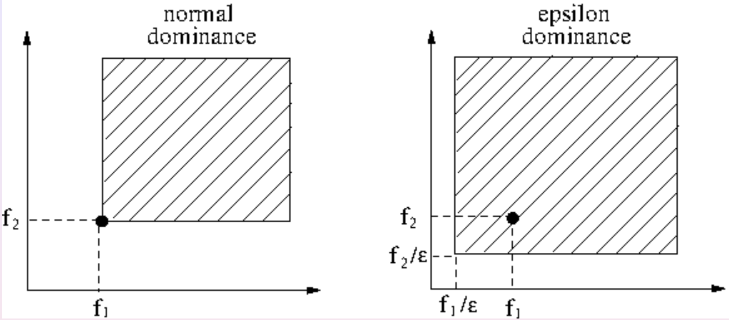

To the left, we can see the area that is dominated by a certain solution. To the right, we graphically show the concept of ε-dominance. In this case, the dominated area has been extended in a magnitude that is proportional to the (user-defined) parameter ε.

Secondary Populations

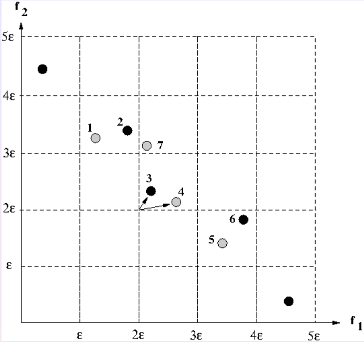

An example of the use of ε-dominance in an external archive. Solution 1 dominates solution 2, and it's therefore, preferred. Solutions 3 and 4 are incomparable. However, we prefer solution 3, because it is closer to the lower lefthand corner. Solution 5 dominates solution 6, so it's preferred. Solution 7 is not accepted because its box (represented by the point (2ε,3ε) is dominated by the box represented by the point (2ε,2ε).

Secondary Populations

ε-dominance has some limitations, including the following:

1. We can lose a high number of nondominated solutions if the geometrical shape of the true Pareto front is unknown or is not taken into consideration by the decision maker.
2. The extreme portions of the Pareto front are normally lost. Additionally, nondominated points that are located in segments of the Pareto front that are almost horizontal or almost vertical will also be lost.
3. The upper bound on the number of points allowed in the external archive may be difficult to reach. For a non-adaptive scheme (i.e., when ε is kept at a fixed value), this upper bound can only be reached when the true Pareto front is linear.

Secondary Populations

With the aim of addressing these limitations, Hernández-Díaz et al. [2007] proposed the so-called **Pareto-adaptive ε-dominance**.

In this proposal, different ε-dominated regions are considered, depending on the characteristics of the true Pareto front. In order to obtain such regions, each Pareto front is associated to a curve from the family:

$$\left\{x^p + y^p = 1 : 0 \leq x, y \leq 1, 0 < p < \infty\right\}.$$

for bi-objective problems, or

$$\left\{x^p + y^p + z^p = 1 : 0 \leq x, y, z \leq 1, 0 < p < \infty\right\}$$

for problems with three objective functions.

Alfredo G. Hernández-Díaz, Luis V. Santana-Quintero, Carlos A. Coello Coello and Julián Molina, “**Pareto-Adaptive epsilon-Dominance**”, *Evolutionary Computation*, Vol. 15, No. 4, pp. 493–517, Winter 2007.

Secondary Populations

With this proposal, it is possible to obtain more solutions than when adopting the original ε-dominance, because the size of the boxes is adjusted based on the geometrical characteristics of the Pareto front. This scheme also provides a more uniform distribution of solutions.

Its main drawback is that it requires an initial sampling to estimate the geometrical shape of the Pareto front (it is worth noticing, however, that this is also a requirement of the original ε-dominance).

Secondary Populations

Schütze et al. [2010] proposed two different archiving strategies for a stochastic multi-objective optimizer. The authors prove convergence with probability one to gap free (and thus 'tight') Pareto front approximations.

The limit set of the first strategy is a tight ϵ-approximate Pareto set which provides a guaranteed uniformity level, while the limit set of the second strategy forms a tight ϵ-Pareto set, which, however, lacks the uniformity.

Oliver Schuetze, Marco Laumanns, Emilia Tantar, Carlos A. Coello Coello and El-Ghazali Talbi, "**Computing Gap Free Pareto Front Approximations with Stochastic Search Algorithms**", *Evolutionary Computation*, Vol. 18, No. 1, pp. 65–96, Spring, 2010.

Secondary Populations

Knowles and Corne [2003] analyzed three convergent archiving algorithms:

- Unbounded archiving
- Simple bounded archiving (similar to Rudolph and Agapie [2000]).
- $\mathcal{S}$ Metric Archiving

These three archiving strategies are able to converge to the true Pareto front under certain conditions, but are not able to maintain a well-distributed set of points along the Pareto front.

Joshua Knowles and David Corne, "**Properties of an Adaptive Archiving Algorithm for Storing Nondominated Vectors**", *IEEE Transactions on Evolutionary Computation*, Vol. 7, No. 2, pp. 100–116, April 2003.

Secondary Populations

Knowles and Corne [2003] proposed the Adaptive Grid Archiving (AGA) algorithm.

This approach is shown to have a low computational cost, it adapts itself to the values of points in objective space and maintains a nondominated set archive which uses “crowding” (not to be confused with NSGA-II’s crowded comparison operator) to encourage an even distribution of points. This approach is based on PAES’ archiving method.

For this algorithm, the authors show that, although convergence to a subset of the true Pareto front is not guaranteed, under certain conditions, the grid boundaries do converge. When this occurs, certain grid regions will become constantly occupied, guaranteeing a certain minimum quality of points in the archive. This also encourages diversity.

Finally, the authors also discuss the conditions under which the AGA’s upper grid boundaries do not converge, which is basically when the Pareto front has a small extent than the whole objective space.

Secondary Populations

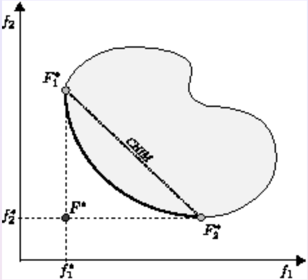

Zapotecas and Coello [2010] proposed the use of the *Convex Hull of Individual Minima (CHIM)* for maintaining well-distributed solutions in the secondary population of a MOEA.

Saúl Zapotecas Martínez and Carlos A. Coello Coello, "**An Archiving Strategy Based on the Convex Hull of Individual Minima for MOEAs**", in *2010 IEEE Congress on Evolutionary Computation (CEC'2010)*, pp. 912–919, IEEE Press, Barcelona, Spain, July 18-23, 2010.

Secondary Populations

Zhao and Suganthan [2010] proprosed an *ensemble* (set) of ε values and an *ensemble* of external archives for a MOPSO.

The idea is to avoid performing a pre-sampling for estimating the proper value of ε.

Shi-Zheng Zhao and Ponnuthurai Nagaratnam Suganthan, "**Multi-Objective Evolutionary Algorithm with Ensemble of External Archives**", *International Journal of Innovative Computing Information and Control*, Vol. 6, No. 4, pp. 1713–1726, April, 2010.

Secondary Populations

Suggested Readings

- Günter Rudolph and Alexandru Agapie, “**Convergence Properties of Some Multi-Objective Evolutionary Algorithms**”, in *Proceedings of the 2000 IEEE Conference on Evolutionary Computation*, Vol. 2, pp. 1010–1016, IEEE Press, Piscataway, New Jersey, USA, July 2000.
- Sanaz Mostaghim, Jürgen Teich and Ambrish Tyagi, “**Comparison of Data Structures for Storing Pareto-sets in MOEAs**, in *2002 IEEE Congress on Evolutionary Computation (CEC’2002)*, Vol. 1, pp. 843–848, IEEE Service Center, Piscataway, New Jersey, USA, May 2002.
- Jonathan E. Fieldsend, Richard M. Everson and Sameer Singh, “**Using Unconstrained Elite Archives for Multiobjective Optimization**”, *IEEE Transactions on Evolutionary Computation*, Vol. 7, No. 3, pp. 305–323, June 2003.
- Shi-Zheng Zhao and P. N. Suganthan, “**Multi-Objective Evolutionary Algorithm with Ensemble of External Archives**”, *International Journal of Innovative Computing Information and Control*, **6**:4, pp. 1713–1726, 2010.

Secondary Populations

Suggested Readings

- Marco Laumanns, Lothar Thiele, Kalyanmoy Deb and Eckart Zitzler, "**Combining Convergence and Diversity in Evolutionary Multi-objective Optimization**", *Evolutionary Computation*, Vol. 10, No. 3, pp. 263–282, Fall 2002.
- Alfredo G. Hernández-Díaz, Luis V. Santana-Quintero, Carlos A. Coello Coello and Julián Molina, "**Pareto-Adaptive epsilon-Dominance**", *Evolutionary Computation*, Vol. 15, No. 4, pp. 493–517, Winter 2007.
- Alfredo G. Hernandez Diaz, Luis V. Santana Quintero, Carlos A. Coello Coello, Julian Molina and Rafael Caballero, "**Improving the Efficiency of Epsilon-Dominance Based Grids**", *Information Sciences*, Vol. 181, No. 15, pp. 3101–3129, August 1, 2011.
- Oliver Schuetze, Marco Laumanns, Emilia Tantar, Carlos A. Coello Coello and El-Ghazali Talbi, "**Computing Gap Free Pareto Front Approximations with Stochastic Search Algorithms**", *Evolutionary Computation*, Vol. 18, No. 1, pp. 65–96, Spring, 2010.

Secondary Populations

Suggested Readings

- Saúl Zapotecas Martínez and Carlos A. Coello Coello, “**An Archiving Strategy Based on the Convex Hull of Individual Minima for MOEAs**”, in *2010 IEEE Congress on Evolutionary Computation (CEC’2010)*, pp. 912–919, IEEE Press, Barcelona, Spain, July 18-23, 2010.
- Joseph Yuen et al., “**An adaptive data structure for evolutionary multi-objective algorithms with unbounded archives**, in *2012 IEEE Congress on Evolutionary Computation (CEC’2012)*, pp. 1004–1011, IEEE Press, Brisbane, Australia, June 10-15, 2012.
- Karl Bringmann, Tobias Friedrich, Frank Neumann and Markus Wagner, “**Approximation-Guided Evolutionary Multi-Objective Optimization**”, in *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2011)*, pp. 1198–1203, AAAI Press, Barcelona, Spain, 16-22 July 2011.
- Joshua Knowles and David Corne, “**Properties of an Adaptive Archiving Algorithm for Storing Nondominated Vectors**”, *IEEE Transactions on Evolutionary Computation*, Vol. 7, No. 2, pp. 100–116, April 2003.