Multi-objective Optimization for Paroxysmal Atrial Fibrillation Diagnosis

Francisco de Toro[1], Eduardo Ros[2], Sonia Mota[2], Julio Ortega[2]

[1] Departamento de Ingeniería Electrónica, Sistemas Informáticos y Automática, Universidad de Huelva, Spain, ftoro@uhu.es
[2] Departamento de Arquitectura y Tecnología de Computadores, Universidad de Granada, Spain, {eduardo,sonia,julio}@atc.ugr.es

Abstract. This paper deals with the application of multi-objective optimization to the diagnosis of Paroxysmal Atrial Fibrillation (PAF). The automatic diagnosis of patients that suffer PAF is done by analysing Electrocardiogram (ECG) traces with no explicit fibrillation episode. This task presents difficult problems to solve, and, although it has been addressed by several authors, none of them has obtained definitive results. A recent international initiative to study the viability of such an automatic diagnosis application has concluded that it can be achieved, with a reasonable efficiency. Furthermore, such an application is clinically important because it is based on a non-invasive examination and can be used to decide whether more specific and complex diagnosis testing is required. In this paper we have formulated the problem in order to be approached by a multi-objective optimisation algorithm, providing good results through this alternative.

1 Introduction

Most real-world optimisation problems are multi-objective in nature, since they normally have several (usually conflicting) objectives that must be satisfied at the same time. These problems are known as MOP (*Multi-objective Optimisation Problems*) [1]. The notion of *optimum* has to be re-defined in this context, as we no longer aim to find a single solution; a procedure for solving MOP should determine a set of good compromises or *trade-off* solutions, generally known as *Pareto optimal solutions* from which the decision maker will select one. These solutions are optimal in the wider sense that no other solution in the search space is superior when all objectives are considered.

Evolutionary Algorithms (EAs) have the potential to find multiple Pareto optimal solutions in a single run and have been widely used in this area [2,3]**.** After the first studies on evolutionary multi-objective optimisation (EMO) in the mid-1980s, a number of Pareto-based techniques were proposed, e.g. MOGA [4], NSGA [5] and NPGA [6]. These approaches did not explicitly incorporate elitism: recently, however, the importance of this concept in multi-objective searching has been recognized and supported experimentally [2,7,8]. In this sense, the present work continues exploring the benefits of elitism by presenting a new elitist EMO algorithm that uses the first

front of non-dominated individuals in the current population to update an external set of best current solutions. Periodically, a filtering function is used to preserve diversity in the population. A selection procedure is applied to obtain the mating pool by randomly choosing candidate solutions from the external set and the current population. The usual variation operators are then applied to produce an offspring set which becomes the new population.

The Atrial Fibrillation is the heart arrhythmia that causes most frequently embolic events that may generate cerebrovascular accidents. The NSFGA and some others MO evolutionary algorithms are applied to Paroxysmal Atrial Fibrillation (PAF) diagnosis. The automatic diagnosis of patients that suffer Paroxysmal Atrial Fibrillation is done by analysing Electrocardiogram (ECG) traces with no explicit fibrillation episode. This task presents difficult problems to solve, and, although it has been addressed by several authors[9], none of them has obtained definitive results. Thus the automatic PAF diagnosis remains an open problem. A recent international initiative to study the viability of such an automatic diagnosis application has concluded that it can be achieved, with a reasonable efficiency [9,10]. Furthermore, such an application is clinically important because it is based on a non-invasive examination and can be used to decide whether more specific and complex diagnosis testing is required. In this paper we have formulated the problem in order to be approached by a multi-objective optimisation algorithm, providing good results through this alternative.

In this paper, Section 2 introduces the MOPs and provides a brief survey about the evolutionary multi-objective techniques previously proposed. Section 3 presents the new single front genetic algorithm for multi-objective optimisation (NSFGA) while Section 4 describes the way PAF diagnosis has been solved by using a multi-objective optimisation approach. Finally, experimental results corresponding to the comparison of the different multi-objective optimization algorithms are given in Section 5 and the concluding remarks are summarized in Section 6.

2 Evolutionary algorithms for multi-objective optimization

A multi-objective optimisation problem (MOP) can be defined [1] as one of finding a vector of *decision variables* that satisfies a set of constraints and optimises a vector function whose elements represent the objectives. These functions form a mathematical description of performance criteria, and are usually in conflict with each other.

The problem can be formally stated as finding the vector $\mathbf{x^*} = [x_1^*, x_2^*, \ldots, x_n^*]$ which satisfies the *m* inequality constraints

$$g_i(\mathbf{x}) \geq 0 \qquad i=1,2,\ldots,m \tag{1}$$

the *p* equality constraints

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

and optimizes the vector function

$$\mathbf{f(x)} = [f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_k(\mathbf{x})]^T \tag{3}$$

where $\mathbf{x} = [x_1, x_2, ..., x_n]^T$ is a vector of *decision variables.*

The constraints given by (1) and (2) define the *feasible region* Φ: any point $\mathbf{x}$ in Φ is a *feasible solution.*

The vector function f ($\mathbf{x}$) is one that maps the set Φ in the set χ of all possible values of the objective functions. The k components of the vector $\mathbf{f}$ ($\mathbf{x}$) represent the *non-commensurable criteria* to be considered. The constraints g_i ($\mathbf{x}$) and h_i (x) represent the restriction imposed on the decision variables. The vector $\mathbf{x^*}$ is reserved to denote the optimal solutions (normally there will be more than one).

The meaning of *optimum* is not well defined in this context, since in these problems it is difficult to have an $\mathbf{x^*}$ where all the components of f_i ($\mathbf{x}$) have a minimum in Φ. Thus, a point $\mathbf{x^*} \in \Phi$ is defined as *Pareto Optimal*

$$\text{If } \forall \mathbf{x} \in \Phi, \forall i := 1..k\ ,\ (f_i(\mathbf{x}) = f_i(\mathbf{x}^*)) \quad or \quad f_i(\mathbf{x}^*) < f_i(\mathbf{x}) \tag{4}$$

The inequality equation in (4) must be fulfilled by at least one component i.

This means that $\mathbf{x^*}$ is Pareto optimal if there exists no feasible vector $\mathbf{x}$ which would decrease one criterion without causing a simultaneous increase in at least one of the others. The notion of Pareto optimum almost always gives, not a single solution, but rather a set of solutions called non-inferior or *non-dominated* solutions (Figure1). This set of non-dominated solutions is known as the *Pareto front*. As, in general, it is not easy to find an analytical expression for the *Pareto front*, the usual procedure is to determine a set of *Pareto optimal* points that provide a good approximate description of the *Pareto front* .

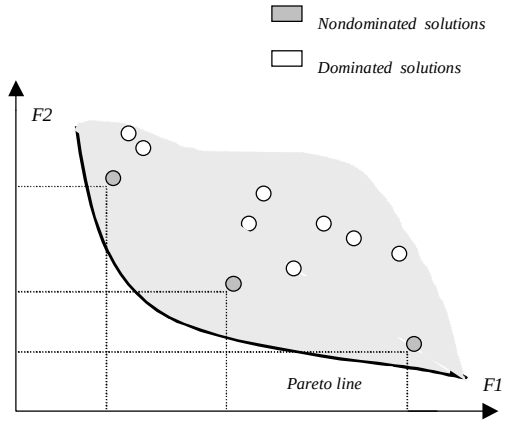

Fig. 1. Two-objective space

Evolutionary algorithms seem to be especially suited to multi-objective optimisation because they are able to capture multiple Pareto-optimal solutions in a single run, and may exploit similarities of solutions by recombination. Indeed, some research suggests that multi-objective optimisation might be an area where EAs perform better than other search strategies. The considerable amount of research

related to MOEAs currently reported in the literature is evidence of present interest in this subject.

Pareto-based fitness assignment in a genetic algorithm (GA) was first proposed by Goldberg [11]. The basic idea is to find a set of Pareto non-dominated individuals in the population. These individuals are then assigned the highest rank and eliminated from further competition. Then, another set of Pareto non-dominated individuals are determined from the remaining population and are assigned the next highest rank. This process continues until the whole population is suitably ranked. Goldberg also suggested the use of a *niching technique* [12] to preserve diversity in the population, in order to converge to different solutions**.**

The *Non-dominated Sorting Genetic Algorithm (NSGA)* [5] uses several layers of ranked individuals. Before selection is performed, the population is ranked on the basis of non-domination: all non-dominated individuals are classified into one category (with a dummy fitness value, which is proportional to the population size, to provide an equal reproductive potential for these individuals). To maintain the diversity of the population, those so classified are shared with their dummy fitness values. Then this group of classified individuals is ignored and another layer of non-dominated individuals is considered. The process continues until all individuals in the population have been classified. Then a stochastic remainder proportionate selection is used, followed by the usual cross and mutation operators. An improved version of NSGA (called NSGA-II) that uses less parameter has been proposed more recently [8].

Fonseca and Fleming [4] have proposed an algorithm called *Multiple Objective Genetic Algorithm (MOGA)* where the rank of each individual is obtained from the number of individuals in the current population that dominate it. Thus, if at generation t, an individual x_i is dominated by p_i (t) individuals, its current rank can be given by:

$$Rank\ (x_i, t)=1+p_i\ (t)$$

All non-dominated individuals are assigned rank 1, while dominated ones are penalized according to the population density of the corresponding region of the trade-off surface. In this way, the fitness assignment is performed by the following steps:

1. - Sort population according to the rank of the individuals.
2. - Assign fitness to individuals by interpolating from the best (rank 1) to the worst (rank n<=N), according to a function (not necessarily linear)
3. - Average the fitness of individuals with the same rank, so that all of them will be sampled at the same rate.

In their *Niched Pareto Genetic Algorithm (NPGA),* Horn and Nafpliotis [6] proposed a tournament selection scheme based on Pareto dominance. Instead of limiting the comparison to two individuals, a number of other individuals (usually about 10) in the population are used to help determine dominance. Whether the competitors are dominated or non-dominated, the result is decided through fitness sharing.

Zitzler and Thiele [7] suggested an elitist multi-criterion EA with the concept of non-domination in their *Strength Pareto Evolutionary Algorithm (SPEA)*. They suggested maintaining an external population at every generation storing all non-dominated solutions discovered so far beginning from the initial population. This external population participates in genetic operations. At each generation, a combined population with the external and the current population is first constructed. All non-dominated solutions in the combined population are assigned a fitness based on the

number of solutions they dominate, and dominated solutions are assigned fitness worse than the worst fitness of any non-dominated solution. This fitness assignment assures that the search is directed towards the non-dominated solutions. A deterministic clustering technique is also used to maintain diversity among non-dominated solutions. Although the implementation suggested is O (mN^3), with appropriate bookkeeping the complexity of SPEA can be reduced to O (mN^2). An improved version of SPEA known as SPEA2 [13] has recently been proposed. This incorporates additionally fine-grained fitness assignment strategy, a density estimation technique, and an enhanced archive truncation method.

Knowles and Corne [14] suggested a simple MOEA using an evolutionary strategy (ES). In their *Pareto-Archived ES (PAES)*, a parent and a child are compared. If the child dominates the parent, the child is accepted as the next parent and the iteration continues. On the other hand, if the parent dominates the child, the child is discarded and a new mutated solution (a new child) is found. However, if the child and the parent do not dominate each other, the choice between the child and the parent considers the second objective of maintaining diversity among obtained solutions. To achieve this diversity, an archive of non-dominated solutions is created. The child is compared with the archive to determine whether it dominates any member of the archive. If so, the child is accepted as the new parent and the dominated solution is eliminated from the archive. If the child does not dominate any member of the archive, both parent and child are examined for their proximity to the solutions of the archive. If the child resides in an uncrowded region in the parameter space among the members of the archive, it is accepted as a parent and a copy is added to the archive. Knowles and Corne later, suggested a *multiparent PAES* with similar principles to the above. The authors have calculated the worst case complexity of PAES for N evaluations as O(amN) where a is the archive length. Since the archive size is usually chosen proportional to the population size N, the overall complexity of the algorithm is O(mN^2).

3 The New Single Front Genetic Algorithm

The Single Front Genetic Algorithm [15], previously proposed by some of the authors, implements a super elitist procedure in which only the non-dominated (and well-diversified) individuals in the current population are copied to the mating pool for recombination purposes and all non-dominated individuals in the current population are copied to the next population (see Figure 2.a). The rest of the individuals required to complete the population are obtained by recombination and mutation of these non-dominated individuals. The preservation of diversity in the population is ensured by means of a filtering function, which prevents the crowding of individuals by removing individuals according to a given *grid* in the *objective* space. The filtering function uses the distance evaluated in the objective space. This approach has proved very effective when applied to Zitzler test functions in comparison to other similar algorithms that use a more complex selection scheme to produce the mating pool

In the new single front genetic algorithm (Figure 2.b), some features have been added to the original SFGA. Firstly an external archive keeps track of the best current solutions found during the running of the algorithm. A selection procedure produces a

mating pool of size *S* by randomly choosing individuals from the external set and the filtered current population. The variation operators produce the offspring set that is copied to the next population. The updating procedure adds the first front of non-dominated individuals to the current population and deletes the dominated individuals from the archive.

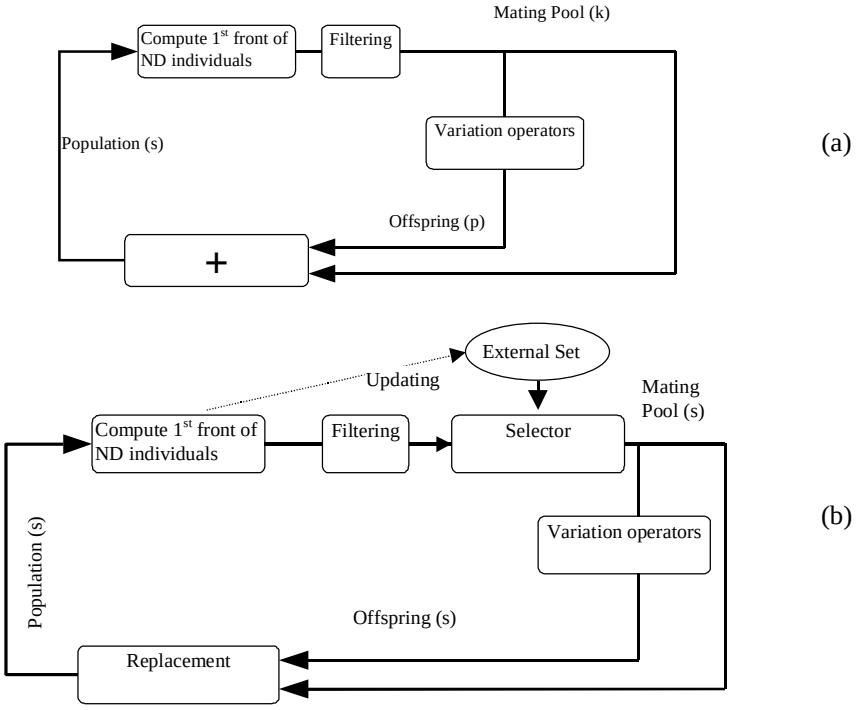

Fig. 2.: Description of SFGA (a) and NSFGA (b).

4 The problem of PAF Diagnosis

A public database for PAF diagnosis applications is available [16], comprising registers obtained from 25 healthy individuals and 25 patients diagnosed with PAF. In this paper, multi-objective optimisation is applied to improve the ability to automatically discriminate registers of these two groups with a certain degree of accuracy. For this purpose, 48 parameters were extracted from each ECG register [17] to obtain a 48 component vector that characterizes each subject ($p_1 \ldots p_{48}$). We have implemented a diagnosis scheme based on weighted threshold dependent decision rules that the multi-objective optimisation procedure will determine. For each parameter we can apply 4 different decision rules:

If $p_i<U_{i(Low\_1)}$ then $C_{PAF}=C_{PAF}+W_{i1}$
If $p_i<U_{i(Low\_2)}$ then $C_{PAF}=C_{PAF}-W_{i2}$
If $p_i>U_{i(High\_1)}$ then $C_{PAF}=C_{PAF}+W_{i3}$
If $p_i>U_{i(High\_2)}$ then $C_{PAF}=C_{PAF}-W_{i4}$

Where U represents different thresholds, C_{PAF} is a level that will determine the final diagnosis and the weights (W_{ij}) are constrained in the interval [0,1]. Thus, the first and third rules increment C_{PAF} (leading to a positive PAF diagnosis) while the second and fourth rules decrement C_{PAF} (leading to a negative PAF diagnosis).

In the diagnosis procedure, the C_{PAF} level is finally compared with a security interval [-F, F]. If C_{PAF} remains within this interval we consider there is not enough certainty about the diagnosis and we leave this case undiagnosed. If $C_{PAF}>F$ the subject is diagnosed positive (as PAF patient), while if $C_{PAF}<F$ it is diagnosed negative.

The multi-objective optimisation procedure uses two objectives to be applied to this diagnosis application: the Classification Rate (CR) and the Coverage Level (CL), defined in (5) and (6), respectively.

$$CR = \frac{Correct\ Diagnosed\ Cases}{Number\ of\ Diagnosed\ Cases} \quad (\mathbf{5})$$

$$CL = \frac{Number\ of\ Diagnosed\ Cases}{Total\ Number\ of\ Cases} \quad (\mathbf{6})$$

High levels of F increment CR because the certainty of the diagnosis rises but it leads to a low CL because it leaves more undiagnosed cases.

This approach could take into consideration 48 parameters with 4 decision rules associated with each of them. In each decision rule the variables to be fixed are the threshold (U) and the weight (W). Therefore we can configure a chromosome with 192 weights and 192 thresholds to be optimised.

In order to reduce the complexity of the optimisation problem, and after a statistical study, an expert selected a subset of 32 rules and their associated 32 thresholds that apparently maximize the discrimination power. In this way, the chromosome length is reduced to 32 (weights) if the expert thresholds are adopted or to 64 if weights and thresholds are optimised.

5 Experimental Results

For performance comparison, the hypervolume metric [7] for *maximization* problems has been used (see Table1). For the sake of simplicity just S metric is shown. S(A) is the volume of the space that is *dominated* by the pareto optimal solution set A. All of the algorithms were executed with the same initial population. The filter parameter, ft, is set to 0.01, the mutation probability per gene is 0.01 and the crossover probability is 0.6. Each Algorithm is executed for 1000 iterations and the population size is set to 200. In Fig 3-5 are shown the solutions in the objective space: CR (Vertical Axis) *versus* CL (horizontal axis) for all SFGA, NSFGA and SPEA [7]. We have considered two cases:

(1) *PAF diagnosis: optimisation of weights of decision rules given by an expert.* In this case the chromosome length is 32, and only the weights associated with the decision rules given by an expert are optimised (see Figures 3a-5a)
(2) *PAF diagnosis: optimisation of threshold and weights of the decision rules given by an expert.* In this case the chromosome length is 64. The parameters to be optimised are the 32 weights and 32 thresholds for the rules selected by an expert (see Fig. 3b-5b). Solutions found by an expert are available for F=1 and included in Fig.4b.

6 Concluding Remarks

This paper presents the *New Single Front Genetic Algorithm (NSFGA).* The NSFGA, SFGA and SPEA are applied to the PAF diagnosis problem. All algorithms show a similar performance. The experiments show that NSFGA and SPEA provide very similar results, and that better results are obtained when both threshold and weights are optimised ($SPEA_{64}$, $NSFGA_{64}$ in Table 1). Moreover, our procedure has improved SPEA results in some cases (threshold given by an expert with F=1.5 and threshold not given by an expert and F=1). The application of multi-objective optimisation to PAF diagnosis is a new approach to tackle this problem. This procedure achieves classification results that are similar to classic schemes but searches optimised values for weights and thresholds that would otherwise have to be given by an expert. In this way, the integration of new rules is made automatic by applying EAs. Furthermore, these MO optimisation schemes lead to multiple solutions that can be of interest for certain patients who suffer from other disorders. In these cases some decision rules may become unreliable and certain solutions are more suitable than others.

Acknowledgements. This paper has been supported by the Spanish *Ministerio de Ciencia y Tecnología* under grant TIC2000-1348.

Table 1. Performance of the algorithms for the 32-gene and 64-gene problems

F	$SFGA_{32}$	$NSFGA_{32}$	$SPEA_{32}$	$SFGA_{64}$	$NSFGA_{64}$	$SPEA_{64}$
0.5	0.82016	0.81849	**0.83010**	0.92555	0.93026	**0.96605**
1	0.80538	0.79240	**0.81015**	0.94943	**0.96698**	0.95733
1.5	0.77278	**0.78852**	0.78331	0.86934	0.89811	**0.92828**

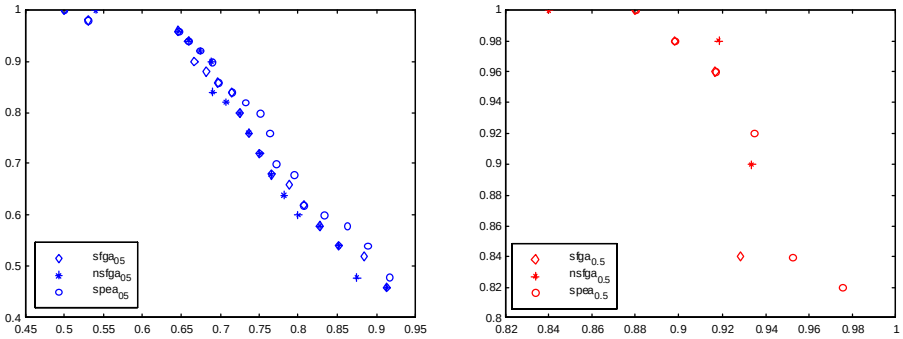

a) b)

Fig. 3. Pareto optimal solutions for F=0.5

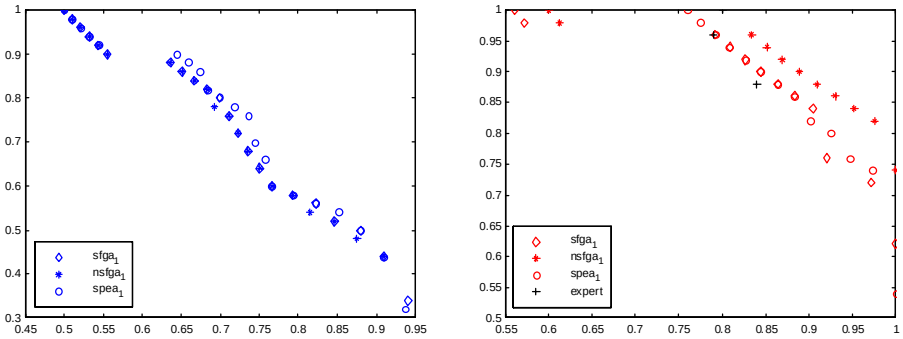

a) b)

Fig 4. Pareto optimal solutions for F=1

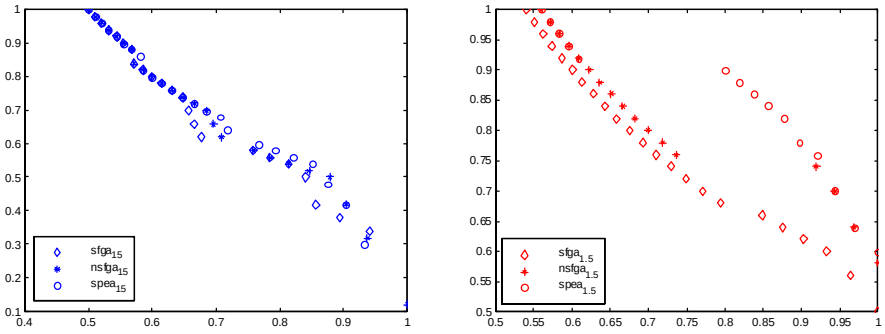

a) b)

Fig 5. Pareto optimal solutions for F=1.5

References

1. Carlos A. Coello Coello. An Updated Survey of GA-Based Multiobjective Optimization Techniques, Technical Report Lania-RD-98-08, Laboratorio Nacional de Informática Avanzada (LANIA), 1998.
2. Parks, G.T. and I. Miller.*"Selective breeding in a multiobjective genetic algorithm"*. In A.E. Eiben, T. Bäck, M. Schoenauer, and H.-P. Schwefel (Editos).5th International Conference on Parallel Problem Solving from Nature (PPSN-V), Berlin, Germany, pp.250-259. Springer.
3. Fonseca, C.M. and P.J. Fleming. Multiobjective optimization and multiple constraint handling with evolutionary algorithms-part i: A unified formulation. IEEE Transactions on Systems, Man, and Cybernetics 28(1), 38-47.
4. Carlos M. Fonseca and Peter J. Fleming. *Genetic Algorithms for Multiobjective Optimization: Formulation, Discussion and Generalization*, In Stephanie Forrest, editor, Proceedings of the Fifth International Conference on Genetic Algorithms, pages 416-423, San Mateo, California, 1993. University of Illinois at Urbana-Champaign, Morgan Kauffman Publishers.
5. N.Srinivas and Kalyanmoy Deb, *"Multiobjective Optimization Using Nondominated Sorting in Genetic Algorithms"*. Evolutionary Computation, Vol. 2, No. 3, pages 221-248
6. Jeffrey Horn and Nicholas Nafpliotis.*Multiobjetive Optimization Using the Niched Pareto Genetic Algorithm.* Proceedings of the first IEEE Conference on Evolutionary Computation. Vol 1, 1994.
7. Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. *Comparison of Multiobjective Evolutionary Algorithms: Empirical Results*, Technical Report 70, Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH) Zurich, Gloriastrasse 35, CH-8092 Zurich, Switzerland, December 1999.
8. Kalyanmoy Deb, Samir Agrawal, Amrit Pratab, and T. Meyarivan. A Fast Elitist Non-Dominated Sorting Genetic Algorithm for Multi-Objective Optimization: NSGA-II, KanGAL report 200001, Indian Institute of Technology, Kanpur, India, 2000.
9. http://www.cinc.org/LocalHost/CIC2001_1.htm
10. Goldberger AL, Amaral LAN, Glass L, Hausdorff JM, Ivanov PCh, Mark RG, Mietus JE, Moody GB, Peng CK, Stanley HE. PhysioBank, PhysioToolkit, and Physionet: Components of a New Research Resource for Complex Physiologic Signals. *Circulation* 101(23):e215-e220 [Circulation Electronic Pages; http://circ.ahajournals.org/cgi/content/full/101/ 23/ e215]; 2000 (June 13).
11. D.E. Goldberg, "*Genetic Algorithms in Search, Optimization and Machine Learning*", New York: Addison Wesley, 1989.
12. B.Sareni and L. Krähenbühl, *"Fitness Sharing and Niching Methods Revisited"*. IEEE Transaction on Evolutionary Computation, Vol 2, No. 3, 1998.
13. Eckart Zitzler, M. Laumannas; L.Thiele. SPEA2: Improving the Strengh Pareto Evolutionary Algorithm. TIK – Report 103, May 2001.
14 Corne, D.W., Knowles, J.D., and Oates, M.J. (2000). The Pareto envelope-based selection algorithm for multiobjective optimization. Proceedings of the Parallel Problem Solving from Nature VI Conference, pp. 839-848.
15. F. de Toro; J.Ortega.;J.Fernández; A.F.Díaz. PSFGA: A parallel Genetic Algorithm for Multiobjective Optimization. 10th Euromicro Workshop on Parallel and Distributed Processing. Gran Canaria, January 2002
16. http://physionet.cps.unizar.es/physiobank/database/afpdb/
17. Mota S., Ros E., Fernández F.J., Díaz A.F., Prieto, A.: ECG Parameter Characterization of Paroxysmal Atrial Fibrillation. Submitted to the 4th International Workshop on Biosignal Interpretation (BSI2002), 24th -26th June , 2002, Como, Italy.