

A Genetic Algorithm for Feature Selection in Data-Mining for Genetics

*Laetitia Jourdan**

*Clarisse Dhaenens**

*El-Ghazali Talbi**

* LIFL, University of Lille, Bât M3 Cité Scientifique
59655 Villeneuve d'Ascq cedex France
Email: {jourdan, dhaenens, talbi }@lifl.fr

1 Introduction

In this paper, we are interested in discovering genetic features and environmental factors that are involved in multifactorial diseases such as (obesity, diabetes). Therefore, experiments have been achieved by the Biological Institute of Lille and lot of data has been generated. To exploit this data, data mining tools are required and we propose a 2-phase approach using a specific genetic algorithm.

This heuristic approach has been chosen as the number of features to consider is large (up to 3654 for biological data we are studying). Data provided indicates for pairs of affected individuals of a same family their similarity at given points (locus) of their chromosomes. This is represented in a matrix where each locus is represented by a column and each pairs of individuals considered by a row. The objective is first to isolate the most relevant associations of features, and then to class individuals that have the considered disease according to these associations.

For the first phase, the feature selection problem, we use a genetic algorithm (GA). To deal with this very specific problem, some advanced mechanisms have been introduced in the genetic algorithm such as sharing, random immigrant, dedicated genetic operators and a particular distance operator has been defined. Then, the second phase, a clustering based on the features selected during the previous phase, will use the clustering algorithm K -means, which is very popular in clustering.

2 A GA for a feature selection problem

The first phase of our algorithm deals with isolating the very few relevant features from the large set. This is not exactly the classical feature selection problem known in Data mining as, in [9], for example around 50% of features are selected. Here, we have the idea that less than 5% of the features have to be selected. But this problem is close from the classical feature selection problem, and we will use a genetic algorithm as we saw they are well adapted for problems with a large number of features [6, 7, 8]. We present here the main characteristics and adaptations we made to deal with this particular feature selection problem. Our genetic algorithm has different phases. It proceeds for a fixed number of generations. A chromosome, here, is a string of bits whose size corresponds to the number of features. A 0 or 1, at position i , indicates whether the feature i is selected (1) or not (0).

2.1 The genetic operators

These operators allow GAs to explore the search space. However, operators typically have destructive as well as constructive effects. They must be adapted to the problem.

Crossover

We use a Subset Size-Oriented Common Feature Crossover Operator (SSOCF)[2] which keeps useful informative blocks and produces offsprings which have the same distribution than the parents. Offsprings are kept, only if they fit better than the least good individual of the population.

Features shared by the 2 parents are kept by offsprings and the non shared features are inherited by offsprings corresponding to the i^{th} parent with the probability $(n_i - n_c/n_u)$ where n_i is the number of selected features of the i^{th} parent, n_c is the number of commonly selected features across both mating partners and n_u is the number of non-shared selected features.

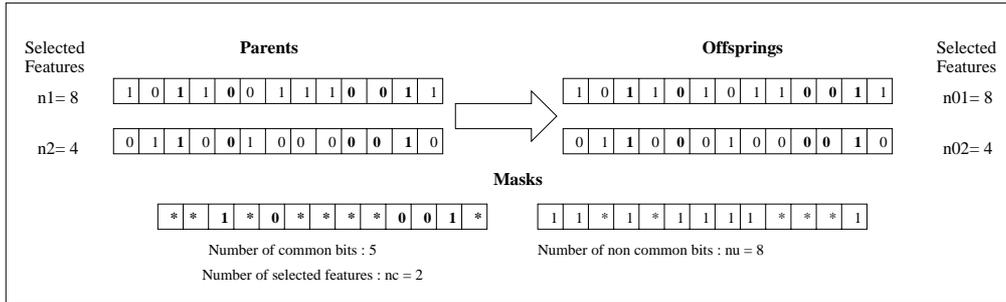


Figure 1: The SSOCF crossover Operator.

Mutation

The mutation is an operator which allows diversity. During the mutation stage, a chromosome has a probability p_{mut} to mutate. If a chromosome is selected to mutate, we choose randomly a number n of bits to be flipped then n bits are chosen randomly and flipped. In order to create a large diversity, we set p_{mut} around 10% and $n \in [1, 5]$.

Selection

We implement a probabilistic binary tournament selection. Tournament selection holds n tournaments to choose n individuals. Each tournament consists of sampling 2 elements of the population and choosing the best one with a probability $p \in [0.5, 1]$.

2.2 Specific adaptations and mechanisms

2.2.1 The chromosomal distance (A distance adapted to the problem)

The biologist experts indicate that a gene is correlated with its neighbors situated on the same chromosome at a distance smaller than σ equals to 20 CMorgan (a measure unit). So in order to compare two individuals, we create a specific distance which is a kind of bit to bit distance where not a single bit i is considered but the whole window $(i - \sigma, i + \sigma)$ of the two individuals are compared. If one and only one individual has a selected feature in this window, the distance is increased by one.

2.2.2 The fitness function

The fitness function we developed refers to the support notion, for an association, which, in data mining, denotes the number of times an association is met over the number of times at least one of the members of the association is met.

The function is composed of two parts. The first one favours for a small support a small number of selected features because biologists have in mind that associations will be composed of few features and if an association has a bad support, it is better to consider less features (to have opportunity to increase

the support). The second part, the most important (multiplied by 2), favours for a large support a large number of features because if an association has a good support, it is generally composed of few features and then we must try to add other features in order to have a more complete association. What is expected is to favour good associations (in term of support) with as much as features as possible. This expression may be simplified, but we let it in this form in order to identify the two terms.

$$F = \left((1 - S) \times \frac{\frac{T}{10} - 10 \times SF}{T} \right) + 2 \times \left(S \times \frac{\frac{T}{10} + 10 \times SF}{T} \right)$$

$$\text{Where : } \begin{cases} \text{Support : } S = \frac{|A \cap B \cap C \dots|}{|A \cup B \cup C \dots|} \text{ where A, B, C ... are the selected features,} \\ T = \text{Total Number of features.} \\ SF = \text{Number of selected significant features (selected features that are not} \\ \text{too close in term of the chromosomal distance).} \end{cases}$$

2.2.3 Sharing

To avoid premature convergence and to discover different good solutions (different relevant associations of features), we use a niching mechanism. A comparison of such mechanisms has been done in [4]. Both crowding and sharing give good results and we choose to implement the fitness sharing [3]. The objective is to boost the selection chance of individuals that lie in less crowded area of the search space. We use a niche count that measures of how crowded the neighborhood of a solution is. The distance D is the chromosomal distance adapted to our problem presented before. The fitness of individuals situating in high concentrated search space regions is degraded and a new fitness value is calculated and used, in place of the initial value of the fitness, for the selection.

The sharing fitness $f_{sh}(i)$ of an individual i , where n is the size of the population, $\alpha_{sh}=1$ and $\sigma_{sh}=3$), is:

$$f_{sh}(i) = \frac{F(i)}{\sum_{j=1}^n Sh(D(I_i, I_j))} \text{ where : } Sh(D(I_i, I_j)) = \begin{cases} 1 - \left(\frac{D(I_i, I_j)}{\sigma_{sh}} \right)^{\alpha_{sh}} & \text{if } D(I_i, I_j) < \sigma_{sh} \\ 0 & \text{else} \end{cases}$$

2.2.4 Random Immigrant

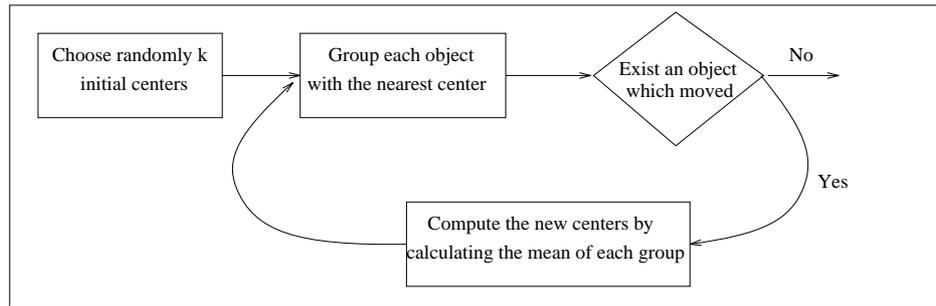
Random Immigrant is a method that helps to maintain diversity in the population. It should also help to avoid premature convergence [1].

We use random immigrant as follows : if the best individual is the same during N generations, each individual of the population, whose fitness is under the mean, is replaced by a new randomly generated individual. When random immigrant is done, we add an extra step in our algorithm.

3 The clustering phase : use of K -Means algorithm

The k -means algorithm is an iterative procedure for clustering which requires an initial classification of the data. The k -means algorithm proceeds as follows : it computes the center of each cluster, then computes new partitions by assigning every object to the cluster whose center is the closest (in term of the Hamming distance) to that object. This cycle is repeated during a given number of iterations or until the assignment has not changed during one iteration [5].

Since the number of features is now very small, we implement a classical k -means algorithm widely used in clustering, and to initialise the procedure we randomly select initial centers.

Figure 2: The k -means algorithm : the states.

4 Experiments

4.1 Validation on an artificial database

Experiments have been first executed on an artificial database in order to validate the method. This database was constructed for the workshop challenge GAW11 (Genetic Analysis Workshop) which was hold in 1998 and was organised by Dr David A. Greenberg¹ to test different data mining methods. This database is a public one. This is an artificial database constructed to be close to real problems and we know, by construction, the relevant associations of features which can influence the disease. Results to obtain are associations A+B+D and C+E1. This test base is composed of 491 features and 165 pairs of individuals.

For ten runs, we wanted to know how many times associations were discovered by the GA. We noted the following results :

Association	A + B	A + D	B + D	C + E1
	100%	50%	20%	10%

The first phase was able to discover real interactions of locus. Some of them are more difficult than other to find.

Then, we ran the k -means algorithm with the results of the GA. We gave 11 features selected by the GA, instead of the initial 491 to the k -means algorithm.

The k -means algorithm helps us to discover associations genes-genes and genes-environmental factors. We have experimented the classical k -means algorithm without any feature selection. The execution time was very large (over 7500 minutes) and results can not be interpreted (we didn't know which were the features involved in the disease) so the feature selection phase is required. With the feature selection, the time of execution of k -means had decreased to 1 minute and the results are exploitable. We present here clusters obtained with $k=2$ and their number of occurrences :

(A B D)	(A B D)	(A B)	(A B)	(E1 A D B)	(A B D)
(E1 C)	(E1 C D)	(E1 D C)	(E1)	(E1 C B)	(E1 D)
4	1	1	2	1	1

This table shows that the k -means algorithm using results of the GA, is able to construct clusters very closely related to the solution presented in results of the workshop. Moreover this solution has been exactly found 4 times over 10 of executions.

¹Dr David A. Greenberg, Department of Psychiatry, Box 1229, Mt Sinai Medical center, 1 Gustave Levy Place, New York 10029

4.2 Experiments on real data

Now, experiments are executed on real data provided by the Biological Institute of Lille (B.I.L) for the study of diabetes. The dataset is composed of 1079 pairs of individuals who have diabetes. The biologists take 3652 points of comparison and 2 covariables (age at onset, the age of the individual when diabetes was diagnosed and BMI Body Mass Index, which is a measure of obesity). The data are confidential and we can not give any biological results here, but here is some aspects :

First, we tested the performances of the method in term of size of problems it can deal with. It appeared that the execution time grows linearly with the number of features and the number of pairs. So the method is able to deal with very large size problem.

Then, we ran several times the algorithm. The genetic algorithm managed to select interesting features and the k -means algorithm was able to class pairs of individuals according to these features and to confirm interesting associations of features.

5 Conclusion

We present here a genetic algorithm dedicated for a particular feature selection problem encountered in genetic analysis of different diseases. The specificities of this problem is that we are not looking for single feature but for several associations of features that may be involved in the studied disease. Results are promising for biologists as the algorithm seems to be robust and to be able to isolate interesting associations. Those associations, found thanks to the method, have now to be studied and validated by the biologists.

References

- [1] C. Bates Congdon. *A comparison of genetic algorithm and other machine learning systems on a complex classification task from common disease research*. PhD thesis, University of Michigan, 1995.
- [2] C. Emmanouilidis, A. Hunter, and J. MacIntyre. A multiobjective evolutionary setting for feature selection and a commonality-based crossover operator. In *Congress on Evolutionary Computing 2000*, volume 2, pages 309–316. CEC, 2000.
- [3] J. Horn, D.E. Goldberg, and K. Deb. Implicit niching in a learning classifier system : Nature's way. *Evolutionary Computation*, 2(1):37–66, 1994.
- [4] S.W. Mahfoud. *Niching Methods for Genetic Algorithms*. PhD thesis, University of Illinois, 1995.
- [5] N. Monmarché, M. Slimane, and G. Venturini. Antclass : discovery of cluster in numeric data by an hybridization of an ant colony with the kmeans algorithm. Technical Report 213, Ecole d'Ingénieurs en Informatique pour l'Industrie (E3i), Université de Tours, Jan. 1999.
- [6] M. Pei, E.D. Goodman, and W.F. Punch. Feature extraction using genetic algorithms. Technical report, Michigan State University : GARAGE, June 1997.
- [7] M. Pei, E.D. Goodman, W.F. Punch, and Y. Ding. Genetic algorithms for classification and feature extraction. In *Annual Meeting : Classification Society of North America*, June 1995.
- [8] M. Pei, M. Goodman, and W.F. Punch. Pattern discovery from data using genetic algorithm. In *Proc of the first Pacific-Asia Conference on Knowledge Discovery and Data Mining*, Feb. 1997.
- [9] J. Yang and V. Honoavar. *Feature Extraction, Construction and Selection : A data Mining Perspective*, chapter 1: Feature Subset Selection Using a Genetic Algorithm, pages 117–136. H. Liu and H. Motoda Eds, massachussetts : kluwer academic publishers edition, 1998.

