A greedy randomized adaptive search procedure applied to the clustering problem as an initialization process using K-Means as a local search procedure

J.R. Cano[a], O. Cordón[b], F. Herrera[b] and L. Sánchez[c]

[a]*Department of Software Engineering, University of Huelva, 21071 La Rabida (Huelva), Spain*
E-mail: Jose.cano@diesia.uhu.es
[b]*Department of Computer Science and A.I., University of Granada, 18071-Granada, Spain*
E-mail: ocordon,herrera@decsai.ugr.es
[c]*Department of Computer Science, University of Oviedo, Oviedo, Spain*
E-mail: luciano@lsi.uniovi.es

Abstract. We present a new approach for Cluster Analysis based on a Greedy Randomized Adaptive Search Procedure (GRASP), with the objective of overcoming the convergence to a local solution. It uses a probabilistic greedy Kaufman initialization to get initial solutions and K-Means as a local search algorithm. The approach is a new initialization one for K-Means. Hence, we compare it with some typical initialization methods: Random, Forgy, Macqueen and Kaufman. Our empirical results suggest that the hybrid GRASP – K-Means with probabilistic greedy Kaufman initialization performs better than the other methods with improved results. The new approach obtains high quality solutions for eight benchmark problems.

Keywords: Clustering, greedy randomized adaptive search procedure, K-Means

1. Introduction

Clustering is a basic process to human understanding. The grouping of related objects can be found such diverse fields as statistics, economics, physics, psychology, biology, pattern recognition, engineering, and marketing [6,14].

The Clustering problem involves partitioning a set of entities into a given number of subsets and finding the location of a centre for each subset in such a way that a dissimilarity measure between entities and centres is minimized.

K-Means is a popular clustering algorithms [2]. Although it is known for its robustness, it can fall in local optimal solutions under certain conditions [13]. It is also widely reported that the K-Means algorithm suf-fers from its dependence on initialization conditions (initial clustering and instance order) as shown in [11].

Since heuristic approaches are good for avoiding convergence to a locally optimal solution, they could be used to find a globally optimal solution. Heuristic approaches used in clustering include those based on simulated annealing [4,7,12], tabu search [1], evolution strategies [3] and genetic algorithms [8–10,15].

In this paper, we focus on the problem suffered by the K-Means due to the initial conditions, proposing a heuristic procedure to this problem. We propose a Greedy Randomized Adaptive Search Procedure (GRASP) [5] applied to the Clustering problem as an initialization process, using K-Means as a local search procedure.

Our approach tries to eliminate the classical problem of the K-Means algorithm, by permitting a higher ex-

ploration and exploitation of the search space, with a low computational cost.

A GRASP is an iterative process with each iteration consisting of two phases, a construction phase and a local search phase:

– The construction phase is based on a probabilistic greedy Kaufman initialization to get initial solutions. It encourages search space exploration.
– The local search phase uses the K-Means algorithm to exploit the search space. We use K-Means as a local search technique because it offers better solutions than other classical algorithms with a low computational cost.

The construction phase builds a feasible solution, whose neighborhood is explored by a local search. The best solution over all GRASP iterations is returned as the result.

In order to introduce this new approach, this paper is organized as follows. Section 2 introduces a background on clustering, K-means and initialization approaches. Section 3 presents the GRASP approach to clustering. Section 4 shows experiments and their analysis, and finally, Section 5 points out some concluding remarks.

2. Background

A common problem in cluster analysis is partitioning objects into a fixed number of groups to optimize an objective function-based clustering. These objects are measured according to several features that characterize them. Patterns can be viewed as vectors in a high dimensional space, where each dimension corresponds to one feature.

In this section we introduce the formalization of clustering, the K-Means algorithms, and four initialization approaches.

2.1. Clustering problem

The clustering problem can be formalized as follows [2]: Considering N entities e_i, each with an associated weight w_i $(i = 1, \ldots, N)$, we search for k centres $c_j (j = 1, \ldots, k)$ minimizing:

$$f(c_1, \ldots, c_k) = \sum_{i=1}^{N} \min_j (w_i d(e_i, c_j))$$

where $d(e_i, c_j)$ measures the dissimilarity between e_i and c_j. In our case, where the entities are described

Table 1
Iris results, $k = 3$

	$K = 3$		
	Arithmetic mean	Standard deviation	Best solution
Random	97.339	0.010	97.326
Forgy	97.326	0	97.326
Macq.	97.326	0	97.326
Kaufm.	97.326	0	97.326
GRASP	97.326	0	97.326

Table 2
Iris results, k = 4

	$K = 4$		
	Arithmetic mean	Standard deviation	Best solution
Random	83.769	0.026	83.729
Forgy	83.749	0.027	83.729
Macq.	83.750	0.031	83.729
Kaufm.	83.786	0	83.786
GRASP	83.729	0	83.729

by their co-ordinates in $\Re^m$, $d(e_i, c_j)$ is the Euclidean distance.

Basically, clustering can be considered as a combinatorial optimization problem, formulated as follow:

Let Q be set containing all objects to be clustered,
C be the set of all feasible clustering of Q,
J: C → ℜ be the internal clustering criterion,
then
Minimize J(c) subject to c ∈ C.

The complexity of the clustering problem is due to different factors:

1. Clustering is an NP-HARD problem when it is considered as the optimization of a certain measure. Therefore, an exhaustive approach is not practicable due to the exponential number of potential partitions of the input data. The number of possible partitions of N elements into k clusters is given by

$$\prod(k, N) = \frac{1}{k!} \sum_{j=1}^{k} (-1)^{k-j} \binom{k}{j} (j)^N.$$

2. Clustering complexity grows if the number of groups is unknown. In such a case he number of solutions becomes:

$$\prod(k, N) = \sum_{i=1}^{k} \frac{1}{i!} \sum_{j=1}^{i} (-1)^{i-j} \binom{i}{j} (j)^N.$$

Due to these reasons, trying to get a global optimal solution by means of an efficient and robust method is difficult. Thus, there is a considerable interest in the design of new heuristics to solve large-sized practical clustering problems.

Table 3
Ruspini results, $k = 4$

| | $K = 4$ | | |
	Arithmetic mean	Standard deviation	Best solution
Random	720.142	0	720.142
Forgy	720.142	0	720.142
Macq.	720.142	0	720.142
Kaufm.	720.142	0	720.142
GRASP	720.142	0	720.142

Table 4
Ruspini results, $k = 5$

| | $K = 5$ | | |
	Arithmetic mean	Standard deviation	Best solution
Random	654.918	12.193	647.672
Forgy	654.074	7.320	647.672
Macq.	655.437	6.692	647.672
Kaufm.	647.672	0	647.672
GRASP	647.672	0	647.672

2.2. K-Means algorithm

K-means evaluates a set of k selected objects, which are considered representatives for the k clusters to be found within the source set of N objects. Given the set of representative objects, the remaining objects are assigned to the nearest representative one, using a chosen distance measure.

The philosophy is that a better set of clusters is obtained when the k representative objects are more centrally located in the cluster they define. For this reason, a suitable objective function to be minimized is the sum of the distances from the respective centers to all the other objects of the same cluster. The function value depends on the current partition of the database $\{C_1, \ldots, C_k\}$:

$$J : \prod_k (\Omega) \to \Re$$

with $\pi_k(\Omega)$ being the set of all partitions of the database $\Omega = \{e_1, \ldots, e_N\}$ in k non-empty clusters. Each instance e_i of the N instances in the database Ω is an m-dimensional vector.

The K-Means algorithm finds locally optimal solutions using the Euclidean distance in the clustering criterion. This criterion is sometimes referred to as square-error criterion. Therefore, it follows that:

$$J(\{C_1, \ldots, C_k\}) = \sum_{i=1}^{k} \sum_{j=1}^{k_i} ||e_{ij} - c_i||$$

where k is the number of clusters, k_i the number of objects of the cluster i, e_{ij} is the j-th object of the i-th cluster and c_i is the centroid of the i-th cluster defined

Table 5
Glass results, $k = 2$

| | $K = 2$ | | |
	Arithmetic mean	Standard deviation	Best solution
Random	323.323	0	323.323
Forgy	323.323	0	323.323
Macq.	323.323	0	323.323
Kaufm.	323.323	0	323.323
GRASP	323.323	0	323.323

as:

$$c_i = \frac{1}{k_i} \sum_{j=1}^{k} e_{ij}, i = 1, \ldots, k$$

The pseudo-code for this algorithm is:

1. Select an initial partition of the database in k clusters $\{C_1, \ldots, C_k\}$
2. Calculate cluster centroids, using the expression of its definition.
3. For every ei in the database and following the instance order DO

   – Reassign instance e_i to its closest cluster centroid. Hence, $e_i \in C_s$ is moved to C_t if $||e_i - e_t|| \leqslant ||e_i - c_s||$ for all $t = 1, \ldots, k, t \neq s$.
   – Recalculate centroids for those clusters.

4. If cluster membership is stabilized then stop else go to step 3.

The K-Means algorithm has the following drawbacks:

– It assumes that the number of clusters k in the database is known, which is not necessarily true.
– It is especially sensitive to initial conditions (initial clusters and instance order).
– It converges finitely to a local minimum, defined by a deterministic mapping from the initial conditions to the final solution.

2.3. Initialization approaches

The second problem, sensitivity to initial conditions, may be mitigated using different values of k, some instance orders and different initialization methods.

In this section we describe four initialization approaches analyzed in [11]. Each one generates k initial clusters following some kind of heuristic and produces a different K-Means response.

These four classical methods are:

Table 6
Glass results, $k = 7$

	$K = 7$		
	Arithmetic mean	Standard deviation	Best solution
Random	206.632	0.704	205.307
Forgy	206.493	0.723	205.889
Macq.	206.635	0.676	205.889
Kaufm.	211.658	0	211.158
GRASP	205.239	0.486	204.992

Table 7
Glass results, $k = 10$

	$K = 10$		
	Arithmetic mean	Standard deviation	Best solution
Random	179.042	2.670	175.945
Forgy	176.710	1.205	176.044
Macq.	177.144	1.389	176.044
Kaufm.	188.691	0	188.691
GRASP	176.058	0.089	175.945

- *Random:* Divides the database into a partition of K randomly selected clusters.
- *Forgy:* k instances of the database (seeds) are chosen at random and the rest of the instances are assigned to the cluster represented by the nearest seed.
- *Macqueen:* k instances of the database (seeds) are chosen at random. Following the instance order, the rest of the instances are assigned to the cluster with the nearest centroid. After each assignment, a recalculation of the centroid has to be carried out.
- *Kaufman:* The initial clustering is obtained by the successive selection of k representative instances. The first representative is chosen to be the most centrally located instance in the database. The rest of representative instances are selected according to the heuristic rule of choosing the instances that promise to have around them a higher number of the rest of instances and that are located as far as possible from the previously fixed ones.

Differences between these initialization methods are:

- Random and Forgy generate an initial partition independently of the instance order.
- Macqueen generates an initial partition that depends on the instance order.
- Kaufman is the only deterministic initialization, based on a greedy approach.

Table 8
Titanic results, $k = 2$

	$K = 2$		
	Arithmetic mean	Standard deviation	Best solution
Random	1617.563	0	1617.563
Forgy	1617.564	0	1617.564
Macq.	1617.564	0	1617.564
Kaufm.	1617.564	0.00066	1617.563
GRASP	1617.562	0	1617.562

3. GRASP approach to the clustering problem

In this section we introduce the GRASP approach and present its application to clustering.

3.1. Greedy randomized adaptive search procedure (GRASP)

A generic GRASP pseudo-code is shown as follows [5]:

```
Procedure grasp( )
InputInstance();
For GRASP stopping criterion not satisfied
ConstructGreedyRandomizedSolution (Solution);
LocalSearch(Solution);
UpdateSolution(Solution,BestSolutionFound);
Rof
Return(BestSolutionFound);
End grasp;
```

In the construction phase, a feasible solution is iteratively constructed, choosing one element at a time. At each construction algorithm iteration, the choice of the next element to be added is determined by ordering all elements in a candidate list with respect to a greedy selection function. This function measures the benefit of selecting each element.

The probabilistic component of a GRASP is characterized by randomly choosing one of the best candidates in the list, but not necessarily the top one. The list of the best candidates is called the restricted candidate list (RCL) and has dimension l. This choice technique allows different solutions to be obtained at each GRASP iteration, but does not necessarily compromise the power of the adaptive greedy component of the method.

The GRASP construction phase pseudo-code is:

```
Procedure ConstructGreedyRandomizedSolution
(Solution)
Solution= {};
For Solution construction not done
MakeRCL(RCL);
s = SelectElementAtRandom(RCL);
```

Table 9
Titanic results, $k = 4$

	$K = 4$		
	Arithmetic mean	Standard deviation	Best solution
Random	1110.034	19.662	1080.426
Forgy	1437.635	297.603	1282.327
Macq.	1442.327	301.718	1282.327
Kaufm.	1170.012	0	1170.012
GRASP	1071.458	1.368	1070.661

Table 10
Titanic results, $k = 10$

	$K = 10$		
	Arithmetic mean	Standard deviation	Best solution
Random	880.530	71.877	670.166
Forgy	992.681	82.042	930.914
Macq.	991.914	84.564	930.914
Kaufm.	340.696	0	340.696
GRASP	329.234	4.196	327.234

Table 11
Image segmentation results, $k = 6$

	$K = 6$		
	Arithmetic mean	Standard deviation	Best solution
Random	158581.932	470.346	158141.047
Forgy	158261.497	626.579	157923.297
Macq.	158334.78	321.169	157923.297
Kaufm.	159233.335	0.0178	159233.297
GRASP	158246.504	312.505	157917.547

Table 12
Image segmentation results, $k = 12$

	$K = 12$		
	Arithmetic mean	Standard deviation	Best solution
Random	117819.845	3263.117	114229.797
Forgy	114244.441	7.570	114237.953
Macq.	115645.357	8.724	114237.842
Kaufm.	114248.172	0	114248.172
GRASP	114221.817	1.066	114220.75

 Solution = Solution ∪ s;
 AdaptGreedyFunction(s);
 Rof
 End ConstructGreedyRandomizedSolution;

The solutions generated by a GRASP construction algorithm are not guaranteed to be locally optimal with respect to simple neighborhood solutions. Hence, it is almost always useful to apply a local search to attempt to improve each constructed solution.

A local search algorithm works in an iterative fashion by successively replacing the current solution by a better solution in the neighborhood of the current one. The key to success for a local search algorithm involves a suitable choice for a neighborhood structure, efficient neighborhood search techniques, and the starting solution.

Finally, the GRASP local search phase pseudo-code is:

 Procedure Local (P,N(P),s)
 For s not locally optimal
 Find a better solution $t \in N(s)$;
 Let $s = t$;
 Rof
 Return (s as local optimal for P)
 End local;

with P being a problem (clustering in our case), s and t being a solution for P, and $N(s)$ being a mechanism to obtain neighbors for s.

3.2. Using GRASP for the clustering problem

Following the generic GRASP structure, it is easy to adapt the algorithm to the clustering problem using the K-means algorithm and greedy Kaufman initialization. The stopping criterion is defined in terms of the maximum number of iterations, while the K-Means is introduced as a local search algorithm. To construct the greedy randomized solution, the Kaufman initialization [6] is taken as a base, because it is a greedy deterministic initialization algorithm. Using the Kaufman criterion, the RCL list is generated by the most promising objects for each center of the solution and one of those candidates is randomly selected.

A generic pseudo-code of the GRASP construction phase for clustering is:

Step 1. Select the most centrally located instance as the first seed.

Step 2. FOR every non selected instance e_i DO

 Step 2.1. FOR every non selected instance e_j DO
 Calc $C_{ji} = \max(D_j - d_{ji}, 0)$
 where $d_{ji} = ||e_i - e_j||$
 and $D_j = \min s d s j$ being s one of the selected seeds

 Step 2.2. Calculate the gain of selecting e_i by $\sum_j C_{ji}$

Step 3. MakeRCL(RCL) by selecting the l instances ei which maximizes $\sum_j C_{ji}$

Step 4. SelectElementAtRandom(RCL)

Step 5. If there are k selected seeds THEN stop ELSE go to Step 2.

Step 6. For having a clustering assign each non-selected instance to the cluster represented by the nearest seed.

Table 13
Texture results, $k = 6$

	$K = 6$		
	Arithmetic mean	Standard deviation	Best solution
Random	406.635	0	406.635
Forgy	406.635	0	406.635
Macq.	406.635	0	406.635
Kaufm.	406.640	0	406.640
GRASP	406.635	0	406.635

Table 14
Texture results, $k = 12$

	$K = 12$		
	Arithmetic mean	Standard deviation	Best solution
Random	361.731	0.782	360.172
Forgy	361.361	1.032	360.468
Macq.	361.491	0.964	360.468
Kaufm.	360.354	0	360.354
GRASP	359.989	0.135	359.803

Table 15
Vehicle results, $k = 4$

	$K = 4$		
	Arithmetic mean	Standard deviation	Best solution
Random	427.314	0	427.314
Forgy	427.351	0.124	427.314
Macq.	427.635	0.116	427.314
Kaufm.	427.314	0	427.314
GRASP	427.314	0	427.314

When the k objects have been taken, the local search (K-Means in this case) is applied taking the clustering obtained as initialization. Then, we compare the local solution cost with the best solution cost found so far, and we take the best. This process continues until all the iterations have been done.

The use of K-Means offers an efficient and low computational cost method to obtain relatively good solutions, but converges to a local minimum. The GRASP construction phase corrects this problem, performing a wide space exploration search. Note that our algorithm constitutes a new initialization method for K-Means, which better explores the search space.

4. Experimental results

Next, we present the sample process, the results obtained in our experiments, and their analysis.

4.1. Sampling process

The performance of our algorithm is studied with various instance sets, trying to get conclusions independent of the problem. Eight real-world databases are considered:

- Iris, which has 150 instances, 4 attributes and 3 classes.
- Ruspini , which has 75 instances, 2 attributes and 4 classes.
- Glass, which has 214 instances, 9 attributes and 7 clusters that can be grouped in 2 bigger classes.

- Titanic, which has 2200 instances, 4 attributes and 2 classes.
- Image Segmentation, which has 2310 instances, 19 attributes and 7 classes.
- Textures, which has 1024 instances, 32 attributes and 6 classes.
- Vehicle, which has 846 instances, 18 attributes and 4 classes.
- Pima, which has 768 instances, 8 attributes and 2 classes.

Since the K-Means algorithm strongly depends on initial conditions, this problem is mitigated using different values of k and some instance orders. In this case, the following initial number of clusters have been considered:

- Iris: $k = 3, 4$
- Ruspini: $k = 4, 5$
- Glass: $k = 2, 7, 10$
- Titanic: $k = 2, 4, 10$
- Textures: $K = 6, 12$
- Image Segmentation: $k = 6, 12$
- Vehicle: $K = 4, 8, 12$
- Pima: $k = 6, 10$

First, we applied three K-Means variants, each one with its own initialization: Random, Forgy and Macqueen.

The sampling process followed is based on the combination of four initial partitions and four instance orders (see [11]), running the K-Means over every one of them, and taking the best result of these sixteen runs.

This sampling process is repeated ten times and the following three values are taken:

- the arithmetic mean of the objective function, square-error criterion ($J(\cdot)$), of the best solution obtained in each one of the runs of the ten sampling processes, denoted by "*Arithmetic mean*".
- the standard deviation, and
- the best solution (objective function value) obtained from the run of the ten sampling processes.

Table 16
Vehicle results, $k = 8$

	$K = 8$		
	Arithmetic mean	Standard deviation	Best solution
Random	370.256	0.79	369.768
Forgy	370.344	0.754	369.703
Macq.	369.945	0.464	369.703
Kaufm.	371.934	0	371.934
GRASP	369.7	0.001	369.699

Table 17
Vehicle results, $k = 12$

	$K = 12$		
	Arithmetic mean	Standard deviation	Best solution
Random	335.661	1.585	333.045
Forgy	334.343	0.94	333.267
Macq.	334.671	1.163	333.301
Kaufm.	336.131	0	336.131
GRASP	333.313	0.675	332.845

Table 18
Pima results, $k = 2$

	$K = 2$		
	Arithmetic mean	Standard deviation	Best solution
Random	52072.211	0.005	52072.203
Forgy	52072.211	0.005	52072.203
Macq.	52072.211	0.005	52072.203
Kaufm.	52072.233	0.010	52072.219
GRASP	52072.233	0.009	52072.219

Table 19
Pima results, $k = 6$

	$K = 6$		
	Arithmetic mean	Standard deviation	Best solution
Random	29415.066	15.681	29403.762
Forgy	29426.369	14.519	29403.762
Macq.	29428.447	15.117	29403.762
Kaufm.	30729.2	0.013	30729.195
GRASP	29403.763	0.001	29403.762

On the other hand, the evaluation of K-Means with Kaufman initialization has been done using 10 random instance orders of each problem. Although Kaufman initialization seems to be a completely deterministic initialization method, this is not completely true. When one instance has the same Euclidean distance to two cluster centers, this instance will be associated to the first of them. If we have some instances in this situation, their order will modify the evolution of the k centers.

The proposed GRASP algorithm is studied using sixteen iterations in each execution (running sixteen K-means). Ten executions of the GRASP algorithm are made, getting the arithmetic mean, the standard deviation and the best solution. The RCL size used is fifteen ($l = 15$), which is flexible enough to obtain optimal solutions due the database sizes. Therefore, in order to compare GRASP evaluation with the remaining K-Means initialization methods, we use the same number of K-Means runs as GRASP iterations for every execution of Random, Forgy and Macqueen initialization methods.

4.2. Results

Experimental results are presented in Tables 1–20.

4.3. Analysis

Comparing results, we notice that Forgy and Macqueen usually give very similar results. They seem to have the same performance, keeping the same or similar local minima. Their results improve those given by Random initialization, with any number of clusters, except in the Titanic database. Thus, both Forgy and Macqueen initializations respond better than Random, with this difference more significant when k grows.

It is clear that Kaufman initialization by itself does not obtain the best results compared to the classical initialization methods. We conclude that the heuristic used in Kaufman initialization needs a more flexible centroid selection to reach the global optimum.

Finally, we compare the results obtained from GRASP with the remaining K-Means initialization methods. As said, this comparison is feasible because we are comparing ten GRASP runs versus ten sampling processes runs (16 K-Means runs with different initialization approaches). This comparison is based on effectiveness and robustness:

– If we study each individual table, we notice that GRASP gets the best results in every problem for the majority of executions.
  Table 21 shows the percentage of times that GRASP executions are the best, equal or worse than the remaining ones (names in the fist row) with respect to the "Arithmetic mean column" and "Best Solution column" in the 20 tables of results, denoting them as AM and BS respectively.
– On the other hand, GRASP presents similar robustness (small values for the standard deviation) to Kaufman K-Means in most of the cases, and better robustness than the remaining approaches, Random, Forgy and Macqueen.

Although GRASP is computationally a little more demanding than the remaining K-Means initializations with the same number of runs/iterations (due to the

Table 20
Pima results, $k = 10$

	$K = 10$		
	Arithmetic mean	Standard deviation	Best solution
Random	24287.956	134.991	24120.801
Forgy	24136.125	26.718	24117.748
Macq.	24137.881	23.687	24117.748
Kaufm.	24340.811	0	24340.811
GRASP	24127.770	7.929	24117.748

Table 21
Percentage of the performance improvement of GRASP versus the remaining initializations

	Random		Forgy		Macq.		Kaufm.	
	AM	BS	AM	BS	AM	BS	AM	BS
GRASP best	75%	50%	75%	50%	75%	50%	70%	70%
GRASP equal	20%	45%	20%	45%	20%	45%	25%	25%
GRASP worst	5%	5%	5%	5%	5%	5%	5%	5%

computational time needed by the construction phase), GRASP induces the K-Means algorithm to present a better performance and a more robust behaviour.

5. Concluding remarks

The K-Means algorithm suffers from its dependence on initial conditions. We present a GRASP algorithm to the clustering problem based on the greedy Kaufman initialization to avoid this drawback. The GRASP algorithm has been empirically compared with four classical initialization methods (Random, Forgy, Macqueen, and Kaufman), and it proved to be the most effective algorithm.

References

[1] K.S. Alsultan, A Tabu Search Approach to the Clustering Problem, *Pattern Recognition* **28** (1995), 1443–1451.

[2] M.R. Anderberg, Cluster Analysis and Applications, Academic Press, 1973.

[3] G.P. Babu and M.N. Murty, Clustering with Evolution Strategies, *Pattern Recognition* **27** (1994), 321–329.

[4] D.E. Brown and C.L. Hutley, A Practical Application of Simulated Annealing to Clustering, *Pattern Recognition* **25** (1992), 401–412.

[5] T.A. Feo and M.G.C. Resende, Greedy Randomized Adaptive Search Procedure, *Journal of Global Optimization* **2** (1995), 1–27.

[6] L. Kaufman and P.J. Rousseeuw, *Finding Groups in Data*, An Introduction to Cluster Analysis, Wiley, 1990.

[7] R.W. Klein and R.C. Dubes, Experiments in Projection and Clustering by Simulated Annealing, *Pattern Recognition* **22** (1989), 213–220.

[8] K. Krishna and M.N. Murty, Genetic K-Means Algorithm. IEEE Trans. on Systems, *Man and Cybernetics – Part B: Cybernetics* **29** (1999), 433–439.

[9] U. Maulik and S. Bandyopadhyay, Genetic Algorithm-Based Clustering Technique, *Pattern Recognition* **33** (2000), 1455–1465.

[10] J.A. Lozano and P. Larrañaga, Applying Genetic Algorithms to Search for the Hierarchical Clustering of Data sets, *Pattern Recognition Letters* **20** (1999), 911–918.

[11] J.M. Peña, J.A. Lozano and P. Larrañaga, An Empirical Comparison of Four Initialization Methods for the K-Means Algorithm, *Pattern Recognition Letters* **20** (1999), 1027–1040.

[12] S.Z. Selim and K. Alsultan, A Simulated Annealing Algorithm for the Clustering Problem, *Pattern Recognition* **24** (1991), 1003–1008.

[13] S.Z. Selim and M.A. Ismail, K-Means-Type Algorithm: Generalized Convergence Theorem and Characterization of Local Optimality, *IEEE Trans. Pattern Anal. Machine Intelligence* **6** (1984), 81–87.

[14] S. Theodoridis and K. Koutroumbas, *Pattern Recognition*, Academic Press, 1999.

[15] L.Y. Tseng and S.B. Yang, A Genetic Clustering Algorithm for Data with Non-Spherical-Shape Clusters, *Pattern Recognition* **33** (2000), 1251–1259.