UNIVERSITY OF JOENSUU
DEPARTMENT OF COMPUTER SCIENCE
Report Series A

Stepwise Clustering Algorithm for Unknown Number of Clusters

Ismo Kärkkäinen and Pasi Fränti

Report A-2002-5

ACM I.4.2, I.5.3
UDK 519.712
ISSN 0789-7316
ISBN 952-458-129-9

Stepwise clustering algorithm for unknown number of clusters

Ismo Kärkkäinen and Pasi Fränti
University of Joensuu, Department of Computer Science,
P.O. Box 111, FIN-80101 Joensuu, FINLAND

Email: ismo.karkkainen@cs.joensuu.fi, pasi.franti@cs.joensuu.fi.

We consider a criterion-based approach to solve dynamic clustering problems for all potential number of clusters, and use a given evaluation criterion to determine the correct one. We propose two novel improvements for reducing the work load to find the correct clustering. The first idea is to use stepwise clustering algorithm that utilizes the previous solution when solving the next clustering with different (one more or one less) number of clusters. The second idea is to use a heuristic stopping criterion in the algorithm that solves a single clustering. The stopping criterion estimates potential further improvement on the basis of the improvement achieved so far. The iteration will automatically stop if the estimated improvement stays below a predefined threshold value. By experiments we have found out that any threshold value 0.1 % or less results in the correct clustering with a confidence of more than 99%. Their effect on the run time and clustering quality is studied.

Keywords: pattern recognition, algorithms, cluster analysis, number of clusters.

1 Introduction

Clustering is an important problem that must often be solved as a part of more complicated tasks in pattern recognition, image analysis and other fields of science and engineering [1, 2, 3]. The goal is to partition a given set of N data vectors into M groups so that similar data vectors are grouped together and dissimilar data vectors to different groups. The clustering task can be formalized as a combinatorial optimization problem, in which the goal is to find the partition that minimizes a given cost function. Clustering is also needed for designing a *codebook* in vector quantization [4].

Clustering aims at solving two main problems: *how many clusters* there are in the data set and *where* they are located. We define the problem here as *static clustering* if the number of clusters is known beforehand, and as *dynamic clustering* if the number of clusters must also be solved.

The static clustering problem can be solved by methods such as *Generalized Lloyd algorithm* (GLA) [5], *simulated annealing* [6], *deterministic annealing* [7,8], *genetic algorithm* [9, 10], *agglomerative methods* [11], and *local search* [12, 13] among many others. The *Randomized Local Search* (*RLS*) [12] is a simple and effective method for solving the problem iteratively. It is based on trial-and-error approach, where the location of the clusters are tentatively changed using *random swapping*. Thus, it corrects the location of misplaced cluster one by one. The ease of implementation makes it possible to tailor the RLS algorithm for various clustering applications with different distance metrics and evaluation criteria. For example, *stochastic complexity* [14, 15] has been successfully applied with the RLS method as the evaluation function for the classification of bacteria in [16].

The dynamic clustering problem can be solved using heuristic methods to determine the number of clusters. For example, *competitive agglomeration* [17] decreases the number of clusters until there are no clusters smaller than a predefined threshold value. The drawback is that the threshold value must be experimentally determined. Divisive clustering, such as the *X-means* in [18], uses an opposite, top-down approach for generating the clustering. The method starts with a single cluster, and new clusters are then created by dividing existing clusters. The splitting continues until a predefined stopping criterion is met. The divisive approach typically requires much less computation than the agglomerative clustering methods but are far less used because of inferior clustering results.

Criterion-based approach, on the other hand, divides the problem in two parts. First a suitable *evaluation function* is defined, which includes the number of clusters as parameter. Static clustering problems are then solved for all reasonable numbers of clusters by using any existing clustering algorithm. The resulting solutions are compared using the evaluation function, and the final result is the clustering that minimizes the evaluation function. Criteria such as *Davies-Bouldin index* [19], *variance-ratio F-test* [20], *stochastic complexity* [15], or *minimum description length* [21] can be applied as the evaluation function among many others [22]. The advantage of the criterion-based approach is that the existing solutions for the static clustering can be utilized as such.

In this paper, we consider the criterion-based approach, and generalize the RLS method to the dynamic clustering problem where the number of clusters is unknown. We propose two variants: (1) *brute force algorithm*, and (2) *stepwise algorithm*. The brute force algorithm applies the RLS method independently to every reasonable number of clusters. It guarantees that the optimal solution will be found if the evaluation criterion is properly designed, and if the RLS method is iterated sufficiently long.

The brute force algorithm is general but not very efficient. For example, we can solve static clustering problem for a data set with 5000 data vectors and 32 clusters in about 40-50 seconds on a 500 MHz Pentium-III machine. However, to solve the dynamic clustering problem when the number of clusters is varied from 1 to 50, the total run time would be about 40 minutes. We proposed two methods to speed-up the brute force method.

The first speed-up idea is the stepwise algorithm. It is based on the idea that if we have already found a solution for m clusters, then this solution can be utilized when searching for the solution of $m + 1$ clusters. It is expected that these two solutions do not differ greatly from each other and, therefore, fewer iterations would be needed. The proposed method starts by finding solution for $m = 2$, and then repeating the process until a predefined number $m=M_{\max}$. At each step, one new cluster is added and the RLS algorithm is then applied. In principle, any other iterative clustering algorithm could also be applied within the method. We consider three possible algorithms: the GLA, LBG-U, and the RLS.

The second speed-up idea idea is to use a heuristic stopping criterion instead of a fixed number of iterations in the RLS algorithm. The stopping criterion estimates the potential further improvement on the basis of the improvement achieved so far. The iteration will stop if the estimated improvement stays below a predefined threshold value (e.g. 0.1 %). With the brute force algorithm, the RLS is iterated about 179-209 times, on average, and with the stepwise method only about 64-84 iterations. More importantly, the method removes the need for manual tuning of the number of iterations.

The rest of the paper is organized as follows. The static clustering problem with fixed number of clusters is defined and the relevant clustering algorithms are recalled in Section 2. The dynamic clustering problem is the considered in Section 3. Various evaluation criteria are first discussed in Section 3.1, and the brute force algorithm is given in Section 3.2. The stepwise approach is then studied in Section 3.3, and three stopping criterion are proposed in Section 3.4. Extensive experiments are then carried out in Section 4, and conclusions are drawn in Section 5.

2 Static clustering problem

Static clustering problem can be defined as follows. Given a data set X of N vectors x_i, partition the data set into M disjoint classes so that similar vectors are grouped together and dissimilar vectors belong to different classes. Partition P defines the classification by giving each vector an index p_i of the class to which the vector is assigned to. Each class is represented by its representative vector c_j, which is here defined as the centroid of the cluster.

We assume that the vectors of the data set are normalized and that they are in metric space, so that we can use Euclidean distance to estimate the distances between the vectors. This allows us to estimate the goodness of solution of M clusters by calculating the distances from vectors to their cluster centroids. We also assume that all clusters are spherical, so all we need to do is to measure the distance to the centroid when calculating the distance to a cluster.

Given the above assumptions, we can use the well-known error measure, that is mean square error (MSE):

$$MSE = \frac{1}{N}\sum_{i=1}^{N} d(x_i, c_{p_i})^2 \qquad (1)$$

Here d is the distance function between the two vectors. Fig. 1 shows a synthetic two-dimensional data set with fifteen clusters, their centroids, and the Voronoi diagram demonstrating the boundaries of the clusters. The small black dots are the data vectors and white dots are the cluster centroids.

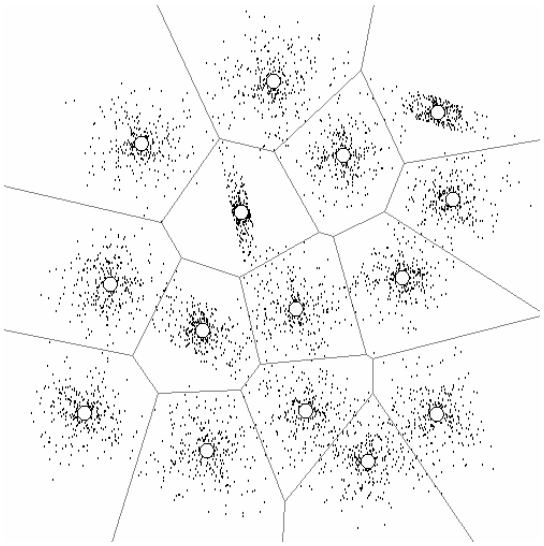

Fig. 1: An example of data set with centroids and cluster boundaries.

A well-known clustering algorithm for minimizing (1) is the Generalized Lloyd algorithm [6], also known the LBG, see Fig. 2. It is simple to implement and it can be applied to any initial solution. It proceeds by changing the representation between centroid and mapping descriptions of the clustering using two optimization steps until the solution does not improve anymore. It has been shown that the algorithm converges to a local or global optimum.

Another virtually as simple clustering method is the randomized local search (RLS), which has been shown to outperform most of the comparative algorithms [12]. The RLS proceeds from a given initial solution by replacing a randomly selected centroid with a random data vector, and then fine-tuning the solution using the GLA. The initial solution is a random set of centroids. Our experiments indicate that it is better to apply only a few of GLA-iterations than to allow the GLA to converge to the local optimum [23]. Pseudocode for the RLS is presented in Fig. 3.

GLA(*X*, *C, Iterations*) return *C*, *P*
REPEAT *Iterations* TIMES
- Generate partition *P* by mapping all data vectors to cluster represented by the nearest centroid vector
- Generate *C* by calculating mean vector for each partition

Return *C*, *P*;

Fig. 2: Pseudocode for GLA.

RLS(*X*, *C*) return *C*, *P*
$C_{best} \leftarrow C$;
Generate partition P_{best} using C_{best} by mapping all data vectors to cluster represented by the nearest centroid vector
REPEAT *Iterations* TIMES
- $C \leftarrow C_{best}$;
- Swap randomly chosen centroid to randomly chosen data vector
- Generate partition *P* using *C* by mapping all data vectors to cluster represented by the nearest centroid vector
- *C*, *P* ← GLA(*X*, *C*, 2);
- Update the best solution if new solution is better:
  IF MSE(*X*, *C*, *P*) < MSE(*X*, C_{best}, P_{best}) THEN
  $C_{best} \leftarrow C$; $P_{best} \leftarrow P$;

Return *C*, *P*;

Fig. 3: Pseudocode for randomized local search.

3 Dynamic clustering problem

We consider the dynamic clustering problem, where also the number of clusters must be solved. We recall first two suitable criteria for evaluating the clustering in Section 3.1. The *brute force algorithm* (*BF*) for finding the best clustering is then introduces in Section 3.2, and the *stepwise clustering algorithm* (*Step*) in Section 3.3. Three stopping criterion are then proposed in Section 3.4.

3.1 Evaluation criteria

MSE has the flaw that it does not take into account the number of clusters. It has the property that the greater the number of clusters is, the lower values we get; the limit being the case where the number of clusters equals the number of data vectors and the MSE-value will be zero. Therefore, we cannot use MSE to determine the number of clusters.

We consider two criteria: *Davies-Bouldin index* (Davies and Bouldin [19]), which has been used in [22, 23, 24], and *variance-ratio F-test* [20]. Davies-Bouldin index (DBI) measures for each cluster the ratio of the intracluster distortion relative to the inter cluster distance of the nearest cluster. This is denoted here as the *mixture* of the two clusters *j* and *k*, and is calculated as:

$$R_{j,k} = \frac{MAE_j + MAE_k}{d(c_j, c_k)} \tag{2}$$

Here *d* is the distance between the cluster centroids, and MAE_j and MAE_k are the mean absolute errors within the clusters *j* and *k* respectively. The higher the intracluster distortion and the closer their centroids, the higher is the index *R* measuring the mixture. The *MAE* for cluster *j* is calculated as:

$$MAE_j = \frac{1}{n_j} \sum_{p_i = j} \left| d(x_i, c_j) \right| \tag{3}$$

The *mixture* of a cluster *j* is defined as the maximum mixture between cluster *j* and all other clusters. The overall DBI-value is then calculated as the average mixtures of the clusters:

$$DBI = \frac{1}{M}\sum_{j=1}^{M} \max_{j \neq k} R_{j,k} \tag{4}$$

The second criterion we use is based on a statistical ANOVA test procedure [20], which we have modified slightly for use with multi-dimensional data. We also omit checking the obtained values against *F*-distribution and use the values directly instead. First we calculate the *squared error* of the given clustering as the sum of squared distances of the data vectors to their cluster centroid:

$$SE = \sum_{i=1}^{N} d(x_i, c_{p_i})^2 \tag{5}$$

Then we estimate the diversity of the clustering by calculating the weighted squared distance from the cluster centroids to the centroid of the data set:

$$Q_B = \sum_{i=1}^{M} n_i d\left(c_i, \bar{x}\right)^2 \tag{6}$$

where n_i is the number of data vectors in the cluster *i*. The centroid of the data set is calculated as:

$$\bar{x} = \frac{1}{N}\sum_{i=1}^{N} x_i \tag{7}$$

The evaluation function is now defined as the ratio of the cluster diversity (Q_B) divided by intra-cluster diversity (*SE*):

$$F = \frac{Q_B \cdot (N - M)}{SE \cdot (M - 1)} \tag{8}$$

In brief, the idea is to evaluate all potential clusterings by an evaluation criteria, and determine the clustering with provides the minimum value. This procedure is demonstrated for the data set 3 (see Section 4) using the DBI in Fig. 4 as a function of the number of clusters, and for the F-test in Fig. 5. In these examples, both evaluation criterion finds the number of clusters ($m = 15$) as there are clear downwards peaks in the graphs of both criteria. The behavior of the F-test is also somewhat different than that of the DBI; the F-value tends to have a bend in the graph after the minimum point instead of a symmetric minimum point. In general, we have found out that the variance-ratio F-test is better for determining the number of clusters than the DBI using the minimum point detection. It is also noted, that the individual clustering obtained might be suboptimal and, therefore, it may happen that wrong number of cluster would be selected despite the evaluation criterion.

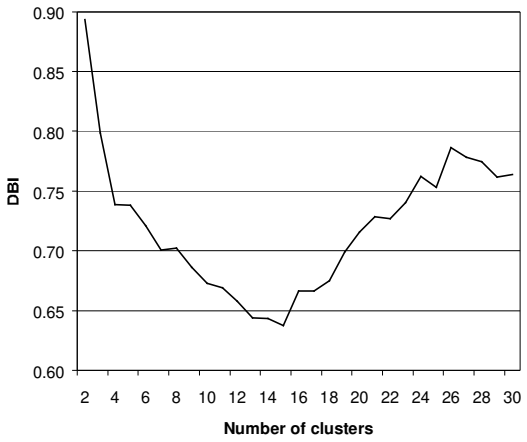

Fig. 4: values of DBI as a function of the number of clusters.

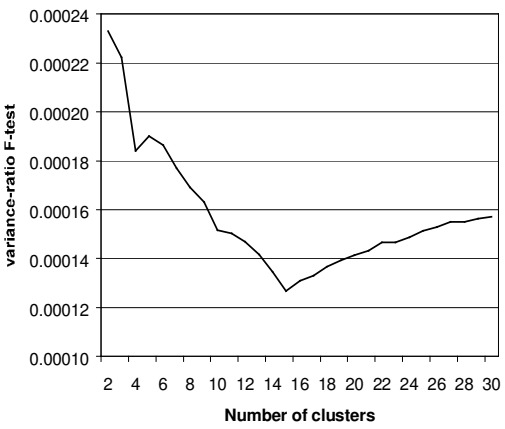

Fig. 5: values of variance-ratio F-test as a function of the number of clusters.

3.2 The Brute Force (BF) algorithm

The simplest search strategy is to go through a specified range of number of clusters, and find the best solution for every possible number. The solutions are then evaluated and the one that minimizes the chosen evaluation function is the final result of the clustering. This search strategy can use any clustering algorithm to find the individual solutions but we will use the RLS due to its benefits discussed in Section 2.

Pseudocode of the brute force algorithm using the RLS is given in Fig. 6. For each number of clusters, the algorithm first generates a random solution, which is then improved by the RLS algorithm. The algorithm is expected to find the correct number of clusters, provided that the chosen evaluation function can distinguish it. The main drawback of the algorithm is that it is time-consuming to generate all possible clusterings in the range $[M_{\min}, M_{\max}]$.

```
BruteForce(X, Mmin, Mmax) return C, P
Generate random solution Cbest with Mmin clusters.
Generate Cbest, Pbest with RLS(X, Cbest);
FOR m ← Mmin + 1 TO Mmax DO
    • Generate random solution C with m clusters
    • Generate C, P with RLS(X, C);
    • If new solution is better than the best one found so far,
      replace the best solution by the new one:
      IF f(X, C, P) < f(X, Cbest, Pbest) THEN
          Cbest ← C; Pbest ← P;
END FOR
Return Cbest, Pbest;
```

Fig. 6: Pseudocode for the brute force algorithm.

3.3 Stepwise algorithm

We next describe how to speed up the brute force algorithm without sacrificing its thoroughness. Since we can freely decide the initial solution that is passed on to the RLS, we can utilize the previous solution. If, for example, we have gained a good solution with m clusters, then this solution can be used as a starting point for finding the solution with m+1

clusters. We perform this by adding a new cluster centroid, performing repartitioning of the data vectors, and then starting the RLS iterations from this solution.

We expect that the clusters in the solution with m+1 clusters are located mainly in the same places as in the solution with m clusters. This is demonstrated in Fig. 7 as a sequence of clusterings with different number of groups. Cluster centroids are shown as white circles. The clustering results are similar for the three middlemost images. Thus, the hypothesis is that fewer number of iterations are needed for optimizing the solution because most of the clusters are already in their correct location.

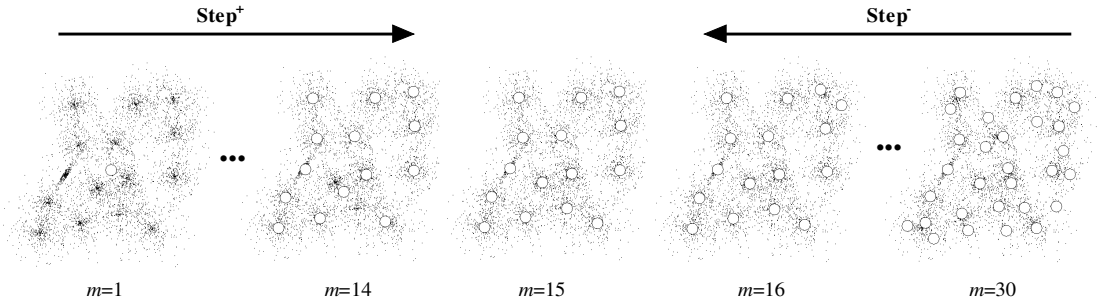

Fig. 7: Location of different number centroids (white dots) in the data.

The stepwise algorithm involves two main design questions: whether we should add or remove clusters, and how the clusters are modified. The first design question can be divided into two approaches. In the first approach (denoted as $Step^+$), we start with a solution of $M_{\min}$ clusters and add new clusters one at a time in the main loop. In the second approach (denoted as $Step^-$), we start with a solution of $M_{\max}$ clusters and remove the clusters one at a time. With brute force algorithm, it does not matter whether we increase or decrease the number of clusters in the loop; the result will be the same as the solutions are generated independently from each other. In the stepwise approach, however, the direction can have an effect on the expected number of iterations needed.

The second design question is to define how we add or remove the clusters. We derive the method directly from the original RLS algorithm and perform the change randomly. This is argued by the fact that the correct place of the clusters must be optimized by a series of random swap operations anyway. Thus, the $Step^+$ algorithm inserts new cluster into a random location, and by performing repartition in respect to the newly created cluster. The $Step^-$ algorithm, on the other hand, removes a randomly chosen cluster, and repartition the data vectors from the eliminated cluster. Pseudocode for the $Step^+$ and $Step^-$ algorithms are shown in Figures 8 and 9.

Step$^+$(X, M_{min}, M_{max}) return C, P

Generate random solution C with M_{min} clusters
Generate C_{best}, P_{best} with RLS(X, C);
Set C to C_{best}
FOR $m \leftarrow M_{min}$ + 1 TO M_{max} DO
- Increase size of C by adding a random centroid
- Generate C, P with RLS(X, C)
- If new solution is better than best one found so far, replace the best solution:
  IF $f(X, C, P) < f(X, C_{best}, P_{best})$ THEN
  $C_{best} \leftarrow C$; $P_{best} \leftarrow P$;

END FOR
Return C_{best}, P_{best};

Fig. 8: Stepwise algorithm Step$^+$ for increasing number of clusters.

Step$^-$(X, M_{min}, M_{max}) return C, P

Generate random solution with M_{max} clusters
Generate C_{best}, P_{best} with RLS(X, C);
Set C to C_{best}
FOR $m \leftarrow M_{max}$ - 1 DOWNTO M_{min} DO
- Decrease size of C by removing random centroid
- Generate C, P with RLS(X, C)
- If new solution is better than the best one found so far, replace the best solution:
  IF $f(X, C, P) < f(X, C_{best}, P_{best})$ THEN
  $C_{best} \leftarrow C$; $P_{best} \leftarrow P$;

END FOR
return C_{best}, P_{best};

Fig. 9: Stepwise algorithm Step$^-$ for decreasing number of clusters.

3.4 Stopping criterion

In general, it is not clear how many iterations we need to perform the RLS in order to find the correct clustering. This question arises both in the brute force and in the stepwise-algorithms. In the stepwise-algorithm the question is even more important because the aim is to decrease the number of iterations and yet obtain the correct number of clusters. We consider next few simple heuristics that can be used for estimating whether it is still worthwhile to continue improving the solution by the RLS. If we had an algorithm that always improves the solution in each iteration this would not be needed as we could check against the previous value and stop if difference were small enough. Since this is not the case (most likely the best solution remains the same) we can take this into acocunt in the first criterion.

The first heuristic is called as *T-times*, as it stops after the solution has been iterated T times since the last improvement to solution occurred. For example, if T is 3 and the previous improvement occurred in the iteration number 100, we will stop in iteration 300 if there will appear no further improvement. We also need a minimum number of iterations that must be performed so that we do not stop right at the beginning. Let f_i be the value of the solution after iteration i, i.e. the f-value for the currently best C and P, and let k be the current iteration number. Stopping condition is:

$$k > T_{\min} \ \wedge \ \left(f_{k/T} - f_k = 0\right) \tag{9}$$

Second heuristic is called as *50-50 ratio*, as it stops when the improvement for the latter half ($f_{k/2}$ - f_k) divided by the improvement that occurred in the first half (f_1 - $f_{k/2}$) falls below certain limit. Taking first and second halves is a consequence of the fact that the solution doesn't improve in evevry iteration. The following formula gives the appearance of smoothly decreasing error:

$$k > T_{\min} \ \wedge \ L > \frac{f_{k/2} - f_k}{f_1 - f_{k/2}} \tag{10}$$

Here L is the predefined threshold value. As with the first heuristic, we need to specify the minimum number of iterations, otherwise we might stop immediately. Fig. 10 shows how the values relate to the number of iterations.

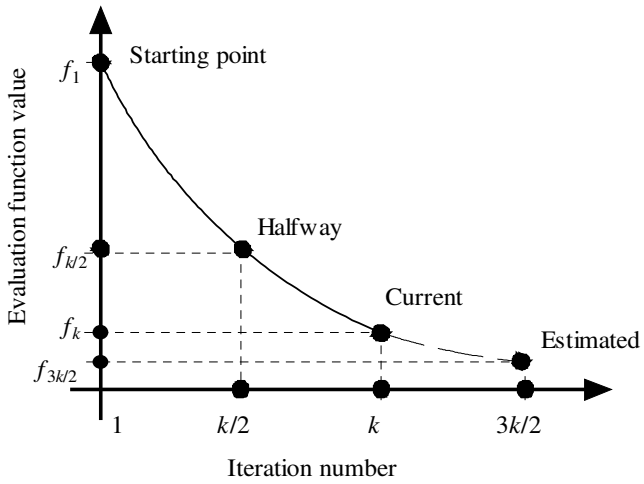

Fig. 10: Values used with stopping criteria.

Third heuristic is called as *Estimate*, as it tries to estimate the future improvement of *f* relative to the improvement so far. The estimated value is shown in Fig. 10 as $f_{3k/2}$. We calculate the estimated value by assuming that the relative slowdown of the improvement of the solution from k iterations to 3k/2 iterations when compared to the slowdown from k/2 iterations to k iterations is the same as that for ranges from k/2 to k and 1 to k/2. The estimated value is:

$$f_{3k/2} = (f_{k/2} - f_k) \cdot \frac{f_{k/2} - f_k}{f_1 - f_{k/2}} = \frac{(f_{k/2} - f_k)^2}{f_1 - f_{k/2}} \tag{11}$$

We calculate the ratio of the estimated future improvement and the total improvement (f_1 - f_k) in order to make the heuristic independent of the range of values of the criterion function *f*. Once again, we must perform some minimum number of iterations in order not to stop right away. The heuristic is:

$$k > T_{\min} \ \wedge \ L > \frac{f_{3k/2}}{(f_1 - f_k)} \tag{12}$$

The numerator estimates how much the solution would improve if number of iterations would increase by half, presuming that the rate of improvement decreases at the same rate as it decreased between first half and second half. The denominator merely normalizes the value so that we get the relative improvement estimation in respect to the total improvement obtained so far. Thus we do not need to take into account the actual range of values of the evaluation function.

4 Test results

We use the four data sets shown in Fig. 11 with varying complexity in terms of spatial data distributions. The test sets have 5000 data vectors randomly scattered around 15 predefined clusters with a varying degree of overlap between the clusters. The overlap is smallest in the set 1 and largest in the set 4. The proposed method is then compared with existing approaches using these four data sets, and also a series of data sets with varying dimensionality (from 2 to 9). Results are also given for the Iris data set [25], which is known to include three clusters.

We first study the parameter setup of the proposed method in Section 4.1 by finding out how many iterations the proposed algorithms need to be run in order to find out the correct number of clusters. This describes the minimum amount of work that must be performed. We perform clustering for all four data sets using the different clustering algorithms, evaluation and stopping criteria. We repeat the clustering process several times for each parameter combination. The successfulness of the clustering is measured by counting how many times the algorithm finds the original number of clusters (15). The goal of the study is to find out the best algorithm and parameter combination that consistently finds out the correct clustering with least amount of work. We then study the effect of the various stopping criteria on the Brute Force and Stepwise algorithms in Section 4.2. The proposed algorithm is then compared with existing approaches in Section 4.3.

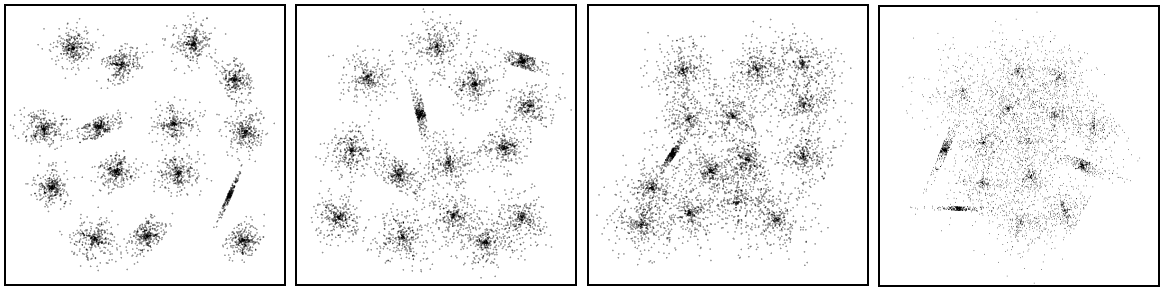

Fig. 11: Two-dimensional data sets used in the experiments.

4.1 Fixed number of iterations

In order to estimate the minimum work load, we first study the work load required by the straightforward brute force algorithm (BF) when manually fixing the number of iterations. We performed the clustering process for a wide range of the number of iterations. Each run (with a fixed number of iterations) was performed independently and repeated 25 times. The process was then repeated for the stepwise algorithms, too.

The results are summarized in Figures 12 and 13 for the data set 3 by showing the number of times (%) the algorithm found out the correct number of clusters. In the case of F-test, the brute force algorithm requires approximately 100 iterations to consistently reach the correct answer (Fig. 12). In the case of the DBI, however, the algorithms cannot distinguish very well between the correct number of clusters. Consequently, the required number of iterations is much greater (Fig. 12). The stepwise algorithm performs slightly better than the BF when a small number of iterations is used. The correct number of clusters, however, is not always found in these cases and the improvement is therefore not significant.

The stepwise algorithm do not start from scratch as BF but it uses the previous solution as a starting point. The quality of the result may therefore depend on the search range of the number of clusters applied. This might have an effect on the quality of the result in case if only a small number of iterations is applied, and when the search range is limited. However, Fig. 13 indicates that the results are virtually the same regardless of the range and direction of the search. For example, the choice between $Step^+$ and $Step^-$ do not have any significant effect on the result. To sum up, there are no significant differences between the three algorithms (BF, $Step^+$ and $Step^-$) when a fixed number of iterations were applied.

The minimum work load of the algorithms is estimated as the smallest number of iterations for which the correct answer is found every time. The numbers are reported in Table 1. In the case of the F-test, the minimum work load is relatively constant over the different algorithms and data sets. In the case of the DBI, on the other hand, the work load is much greater for the more complex data sets. The extreme case is the data set 4, for which the correct answer was not found by the DBI. Instead, the DBI tends to find the optimum results with 14 clusters.

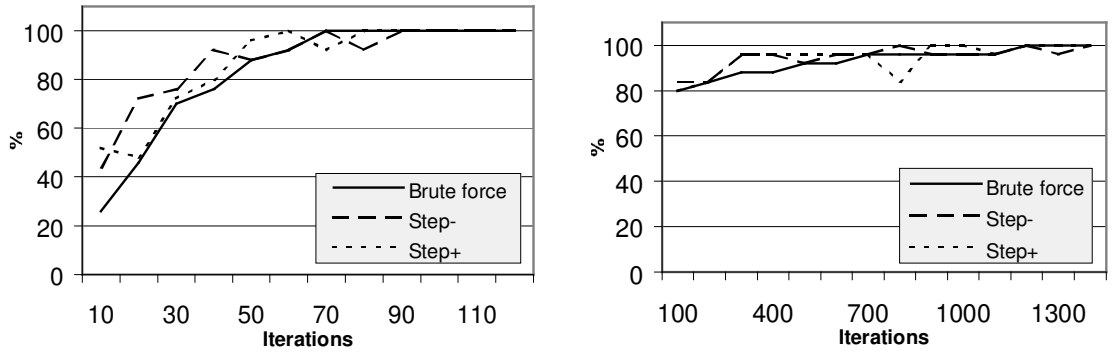

Fig. 12: The number of times (%) the algorithms found out the correct answer, as a function of the number of iterations. The results are for the data set 3 using F-test (left) and DBI (right). Note that the graph of the DBI (right) has much larger scale than that of the F-test (left).

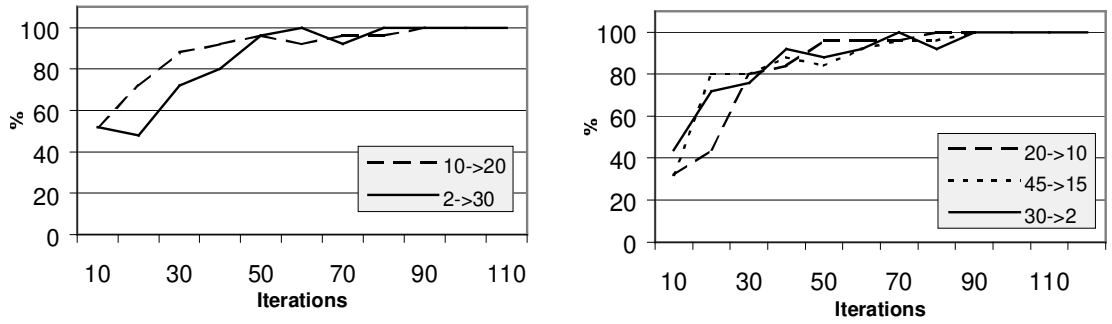

Fig. 13: The number of times (%) the algorithms found out the correct answer, as a function of the number of iterations. The results are for the $Step^+$ (left) and $Step^-$ (right) algorithms with different search ranges using F-test with data set 3.

Table 1: The minimum number of iterations required to find out the correct answer with 100% of the cases.

	F-test			DBI		
	BF	$STEP^+$	$STEP^-$	BF	$STEP^+$	$STEP^-$
Data set 2	130	60	60	90	100	70
Data set 3	70	60	70	1200	900	800
Data set 4	70	80	70	N/A	N/A	N/A

4.2 Using stopping criteria with brute force

In the previous chapter we studied the minimum work load required by the different algorithms. In practice, the correct clustering is not known beforehand and the number of iterations cannot be manually tuned. The use of a proper stopping criterion is therefore important for minimizing the work load. The aim is to get the correct clustering with the least number of iterations as possible. The results in Table 1 can be used for estimating the lower limit of this work load.

The brute force and stepwise algorithms were tested using the three stopping criteria presented in Section 3.4 (*T-times*, *50-50 ratio*, *Estimate*). The threshold parameter was experimentally set to T=3 for *T-times*, and L=0.001 for the other two criteria. These criteria require also some minimum number of iterations before the stopping criterion is allowed to take effect. We studied the performance of the criterion by varying the minimum number of iterations from 10 to 90. The results are summarized in Table 2 as the average number iterations applied with the first parameter setup that were able to produce the correct clustering in all runs.

The results show that the number of iterations remains reasonably low when *F-test* is used as the evaluation function. For example, the BF was iterated only 122-162 times (on average) in the case of the *50-50 ratio test*, and 179-209 times in the case of the *Estimate* in order to get the correct clustering. The corresponding numbers for the Stepwise are 88-183 (with *50-50 ratio test*) and 81-121 (with *Estimate*). The average number of iterations behaves basically in the same way with the stepwise algorithms as with the brute force algorithm. The *T-times* performed less consistent in these tests as it keeps iterating unnecessarily long with the more difficult test sets (Set 3 and 4).

With *DBI* the results are not as good as with the *F-test*. The correct results is found only in the case of the data set 2, and even then, the algorithm sometimes iterates unnecessarily long. The average value of the stepwise algorithms is affected by a single large iteration count, where for 28 clusters the RLS algorithm performed 244,053 iterations. As the threshold value was set to T=3, it means that the last improvement has occurred at the 81,351 iteration, after which there has been no improvement whatsoever. In the case of the data set 3, on the other hand, the methods were unable to keep the algorithm iterating long enough in order to find the correct number of clusters for the data set 3.

Table 2: Average number of iterations applied when using a stopping criterion. The numbers shown in parentheses gives the fixed lower limit for the number of iterations.

	F-test			DBI		
Brute force	T-times	50-50 ratio	Estimate	T-times	50-50 ratio	Estimate
Data set 2	180 (80)	143 (130)	179 (80)	982 (50)	136 (80)	191 (70)
Data set 3	881 (60)	122 (90)	209 (60)	N/A	N/A	N/A
Data set 4	743 (50)	162 (100)	208 (50)	N/A	N/A	N/A
Step$^+$						
Data set 2	208 (50)	133 (80)	81 (50)	1935 (60)	210 (80)	N/A
Data set 3	282 (30)	88 (40)	92 (50)	N/A	N/A	N/A
Data set 4	416 (60)	141 (60)	121 (70)	N/A	N/A	N/A
Step$^-$						
Data set 2	177 (30)	88 (50)	85 (70)	2923 (50)	284 (50)	112 (70)
Data set 3	497 (60)	107 (40)	65 (50)	N/A	N/A	N/A
Data set 4	405 (40)	183 (70)	84 (60)	N/A	N/A	N/A

4.3 Behavior of the stopping criteria

Despite the fact that the minimum number of iterations in Table 2 may be half of what is required in comparison to Table 1, the average number of iterations may yet be much higher. The stopping criteria try to determine when the solution has ceased to improve, not when enough work has been done to find the correct number of clusters. Another reason is that not all cluster counts require the same amount of iterations. The less there are clusters, the less iterations are needed. After the correct number of clusters has been passed, the number of iterations increases noticeably. This is illustrated in Fig. 14. For *T-times* there is a general rising trend after the correct number of clusters has been exceeded. Similar trend can also be observed for *Estimate*.

The graphs in Fig. 14 show marked jaggedness. This is due to the random nature of the underlying algorithm. Since the RLS does not necessarily improve the solution in every iteration, it is possible that a noticeable drop happens just before the stopping criterion would stop the iterations and the limit is increased. With *T-times* this means that the number of iterations can grow to be extremely large. Fig. 15 shows an example of how the improvements occur during a single run of the RLS. This Fig. also offers an explanation why the setup of the threshold value for the *50-50 ratio* and the *Estimate* is not so vital; if there is a sufficiently long time period without an improvement, these stopping criteria will halt regardless of the threshold since the numerator becomes zero. Thus, these criteria behave as the *T-times* with the threshold value T=2 with the exception that they can stop sooner.

Iterations are also not distributed evenly over the specified range of clusters. It appears that finding a good clustering for the correct number of clusters can be an easier task than for an incorrect number of clusters. Fig. 16 illustrates how the number of iterations are divided on the average over the range from 3 to 30. It shows that the stopping criteria have the effect of distributing the work so that it is spent more on the incorrect numbers of clusters and less around the correct number of clusters. Another effect with Step$^+$ is that once the correct number of clusters (15) has been exceeded, the number of iterations tends to grow with the number of clusters.

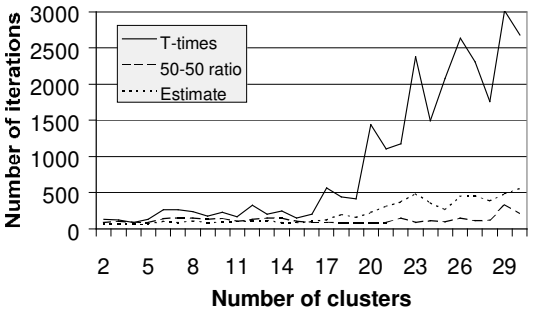

Fig. 14: Average number of iterations as function of the number of clusters.

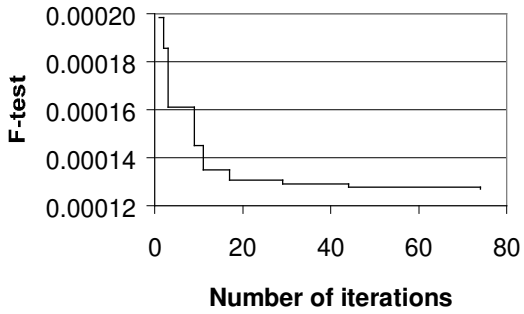

Fig. 15: An example of how improvements in solution occur when using RLS.

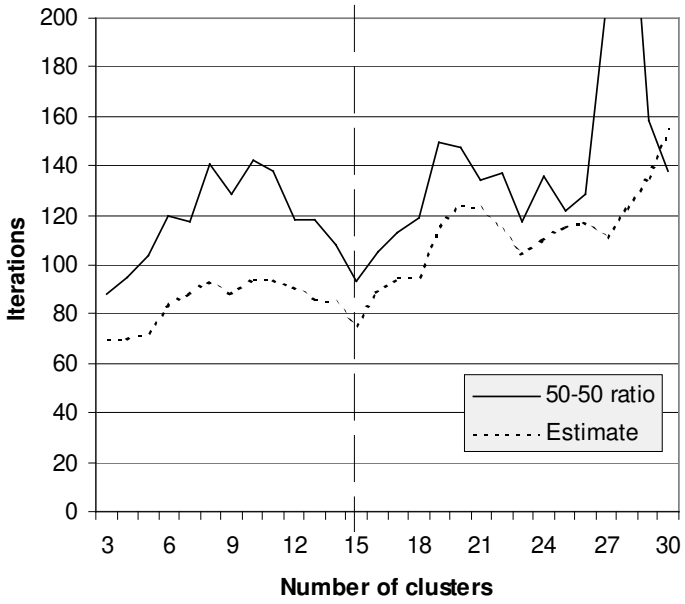

Fig. 16: Distribution of iterations over the range from 3 to 30 for the data set 2. F-test is used as the evaluation function and 50-50 ratio and Estimate as the stopping criteria.

The effect on the quality of the solutions found is shown in Table 3. It shows the average values of the error function for the correct number of clusters. The results show that actual differences in solutions obtained are small despite the differences in the number of iterations used. In all cases the solutions are equal to three significant digits. We can therefore conclude that finding the optimal solution and the correct number of clusters are practically the same thing. In other words, if we iterate long enough to find the correct number of clusters, we have also found the optimal solution. Therefore, further iteration for the solution with the correct number of clusters does not seem to be necessary.

Table 3: Comparison of the average values of clustering criteria.

Brute force	F-test ($\times 10^{-3}$)				DBI			
	Fixed	T-times	50-50 ratio	Estimate	Fixed	T-times	50-50 ratio	Estimate
Data set 2	0.074071	0.074071	0.074071	0.074075	0.465317	0.465003	0.465381	0.465363
Data set 3	0.126805	0.126811	0.126819	0.126840	0.638389	N/A	N/A	N/A
Data set 4	0.161385	0.161276	0.161373	0.161361	N/A	N/A	N/A	N/A

Step$^+$	T-times	50-50 ratio	Estimate	T-times	50-50 ratio	Estimate
Data set 2	0.0740624	0.0740812	0.0740717	0.4650288	0.4653249	N/A
Data set 3	0.126843	0.126872	0.126957	N/A	N/A	N/A
Data set 4	0.161265	0.161371	0.161506	N/A	N/A	N/A
Step$^-$						
Data set 2	0.0740650	0.0740610	0.0740698	0.465056	0.465421	0.465317
Data set 3	0.126832	0.126940	0.126887	N/A	N/A	N/A
Data set 4	0.161325	0.161297	0.161347	N/A	N/A	N/A

4.4 Comparison with the existing algorithms

We next compare the following approaches:

- Stepwise with GLA
- Stepwise with LBG-U
- Stepwise with RLS (the proposed method)
- Competitive agglomeration (CA)
- Brute Force (BF)

The Stepwise GLA combines the *Step*- approach with the GLA [5] as follows. It generates $M_{\max}$=30 random clusters, and then proceeds to decrease the number of clusters by removing each cluster in turn. There are M different choices for the cluster removal, and every of them are iterated by the GLA until convergence. As a result, we obtain a set of candidate solutions, of which we keep the best one for the initial solution for next cluster removal.

The Stepwise with LBG-U is implemented in the same manner as the Stepwise GLA but the LBG-U algorithm [13] is used instead. The LBG-U is also similar to that of the RLS but it uses deterministic choice for replacing the location of the cluster, whereas the RLS uses random choice. This also means that we do not need to repeat the RLS for every possible cluster removal but we can simply start from random starting point as the algorithm will anyway optimize the cluster locations. With the GLA and LBG-U, however, the repeats are done as there are no randomness in this part and all possibilities can be explored.

The competitive agglomeration [17] uses competitive learning for optimising the location of the clusters. The algorithm starts with a predefined number of initial clusters, and then removes all clusters that are smaller than a predefined threshold value ε. The process continues until there are no more clusters to be removed. The initial number of clusters is set to 100, and the threshold value was experimentally determined by varying it from 10^{-6} to 10^{-1}.

The results for the data sets 1 to 4, and for the Iris data set are summarized in Table 4. When F-test is used as the criterion, the proposed method (Stepwise RLS) finds the correct clustering almost in all cases. The Stepwise with LBG-U works fine most of the times whereas the Stepwise GLA gives significantly worse success rate with all data sets. When DBI is used as the criterion, the results are worse with all variants. The relative performance of the different

methods is similar to that of the F-test. The CA, on the other hand, consistently fails to find the correct clustering except in the case of easiest test set. It tends to remove too many clusters no matter of the parameter setup.

The methods are also compared using a series of data sets with varying dimensionality. The spatial complexity of these sets is slightly easier than that of Data sets 1 to 4 but the dimensions are higher (varying from 2 to 9). The results are summarized in Fig. 17. The main observation is that the choice of the static clustering algorithm is crucial: the performance of the Stepwise RLS is much better and more consistent than that of the GLA and LBG-U. The second observation is that the CA performs reasonably well independent of the dimensionality, as long as the clusters are spatially clearly separable.

Table 4: The number of times the correct clustering was found (among 100 repeats) by the three Stepwise variants, Brute force algorithm, and by Competitive agglomeration.

	F-test				DBI				CA
	Stepwise GLA	Stepwise LBG-U	Stepwise RLS	Brute Force	Stepwise GLA	Stepwise LBG-U	Stepwise RLS	Brute Force	
Data set 1	22 %	94 %	98 %	98 %	14 %	94 %	98 %	95 %	70 %
Data set 2	30 %	74 %	100 %	98 %	19 %	53 %	96 %	100 %	8 %
Data set 3	39 %	80 %	100 %	100 %	28 %	72 %	76 %	82 %	0 %
Data set 4	51 %	92 %	100 %	100 %	31 %	36 %	33 %	52 %	0 %
Iris data	100 %	100 %	100 %	100 %	0 %	0 %	0 %	0 %	5 %

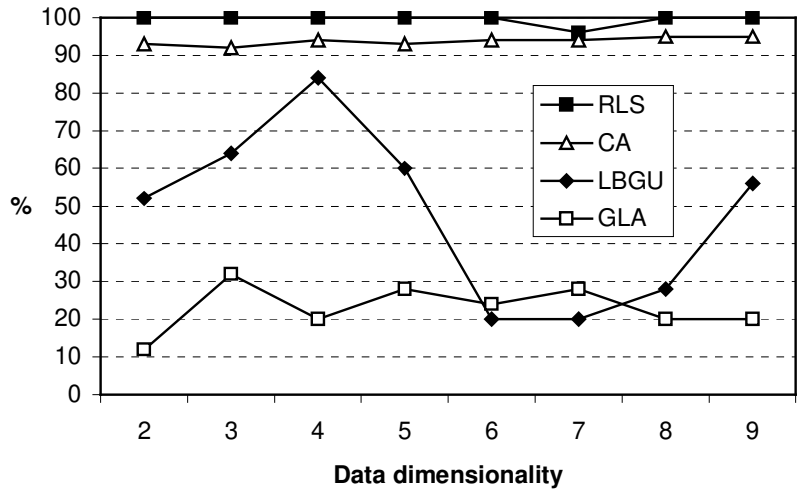

Fig. 17: Comparison of the algorithms as a function of the data set dimensions (2-9).

5 Conclusions

We have introduced stepwise clustering algorithm for finding the correct clustering in a case when the number of groups is unknown. The method generates solutions for all cluster counts within a given search range. According to the experiments, the work load can be reduced to about 200 iterations per cluster count with the brute force approach. With the stepwise algorithm (decreasing variant), the algorithm iterates only about 65-85 times on average. The results indicate that there is obvious dependency between the number of iterations and how complex is the data set (how much the clusters overlap). The choice of the evaluation function seems also to be important.

References

[1] Everitt BS, Cluster Analysis, 3rd Edition. Edward Arnold / Halsted Press, London, 1992.

[2] Kaufman L, Rousseeuw PJ, Finding Groups in Data: An Introduction to Cluster Analysis. John Wiley Sons, New York, 1990.

[3] Dubes R, Jain A, Algorithms that Cluster Data. Prentice-Hall, Englewood Cliffs, NJ, 1987.

[4] Reeves C, Modern Heuristic Techniques for Combinatorical Optimization Problems. McGraw - Hill, 1995.

[5] Linde Y, Buzo A, Gray RM, An algorithm for vector quantizer design. IEEE Transactions on Communications 1980; 28(1): 84-95.

[6] Zeger K, Gersho A, Stochastic relaxation algorithm for improved vector quantiser design. Electronics Letters 1989; 25(14): 896-898.

[7] Rose K, Gurewitz E, Fox G, A deterministic annealing approach to clustering. Pattern Recognition Letters 1990; 11: 589-594.

[8] Hoffmann T, Buhmann J, Pairwise Data Clustering by Deterministic Annealing. IEEE Transactions on Pattern Analysis and Machine Intelligence 1997; 19(1): 1-14

[9] Fränti P, Kivijärvi J, Kaukoranta T, Nevalainen O, Genetic algorithms for large scale clustering problems. The Computer Journal 1997; 40(9): 547-554.

[10] Fränti P, Genetic algorithm with deterministic crossover for vector quantization. Pattern Recognition Letters 2000; 21(1): 61-68.

[11] Ward JH, Hierarchical grouping to optimize an objective function. J. Amer. Statist.Assoc. 1963; 58: 236-244.

[12] Fränti P, Kivijärvi J, Randomized local search algorithm for the clustering problem. Pattern Analysis and Applications 2000; 3(4): 358-369.

[13] Fritzke B, The LBG-U method for vector quantization – an improvement over LBG inspired from neural networks. Neural Processing Letters 1997; 5(1): 35-45.

[14] Rissanen J, Stochastic Complexity. Journal of Royal Statistical Society B 1987; 49(3): 223-239.

[15] Gyllenberg M, Koski T, Verlaan M, Classification of binary vectors by stochastic complexity. Journal of Multivariate Analysis 1997; 63(1): 47-72.

[16] Fränti P, Gyllenberg HH, Gyllenberg M, Kivijärvi J, Koski T, Lund T, Nevalainen O, Minimizing stochastic complexity using local search and GLA and with applications to classification of bacteria. Biosystems 2000; 57(1): 37-48.

[17] Frigui H, Krishnapuram R, Clustering by Competitive Agglomeration. Pattern Recognition 1997; 30(7): 1109-1119.

[18] Pelleg D, Moore A, X-means: Extending K-means with Efficient estimation of the Number of Clusters. Proc. 17th International Conf. on Machine Learning. Morgan Kaufmann, San Francisco CA 2000; pp. 727-734.

[19] Davies DL, Bouldin DW, A cluster separation measure. IEEE Transactions on Pattern Analysis and Machine Intelligence 1979; 1(2): 224-227.

[20] Ito PK, Robustness of ANOVA and MANOVA Test Procedures. In: Krishnaiah PR (ed). Handbook of Statistics 1: Analysis of Variance. North-Holland Publishing Company, 1980, pp 199-236.

[21] Bischof H, Leonardis A, Selb A, MDL Principle for robust vector quantization. Pattern Analysis and Applications 1999; 2(1): 59-72.

[22] Bezdek JC, Pal NR, Some new indexes of cluster validity. IEEE Transactions on Systems, Man and Cybernetics 1998; 28(3): 302-315.

[23] Kärkkäinen I, Fränti P, Minimization of the value of Davies-Bouldin index. In Proceedings of the IASTED International Conference on Signal Processing and Communications (SPC'2000). IASTED/ACTA Press, 2000, pp 426-432.

[24] Sarkar M, Yegnanarayana B, Khemani D, A clustering algorithm using an evolutionary programming-based approach. Pattern Recognition Letters 1997; 18(10): 975-986.

[25] Duda RO, Hart PE, Pattern Classification and Scene Analysis. (Q327.D83) John Wiley & Sons, New York, page 218, 1973.