

Vector Quantization by Lazy Pairwise Nearest Neighbor Method*

(published in *Optical Engineering*, 38 (11), 1862-1868, November 1999)

Timo Kaukoranta¹, Pasi Fränti² and Olli Nevalainen¹

¹ Turku Centre for Computer Science (TUCS)
Department of Computer Science
University of Turku
Lemminkäisenkatu 14A, FIN-20520 Turku
FINLAND

² Department of Computer Science
University of Joensuu
P.O. Box 111, FIN-80101 Joensuu
FINLAND

Abstract: The Pairwise Nearest Neighbor (PNN) algorithm is a well-known method for the codebook construction in vector quantization, and for the clustering of data sets. The algorithm has simple structure and it gives high quality solutions. A drawback of the method is the large running time of the original (exact) implementation. In this paper we prove the monotony of the merge costs of the PNN. The monotony property is utilized for speeding-up an existing PNN variant. The idea is to postpone a number of distance calculations. In this way we can reduce the computation by about 35% while preserving the exactness of the PNN.

Keywords: Vector Quantization, Codebook Generation, Clustering Algorithms, Pairwise Nearest Neighbor Algorithm.

1. Introduction

We study the problem of generating a *codebook* for a *vector quantizer* (VQ). The aim is to find M *code vectors* (*codebook*) for a given set of N *training vectors* (*training set*) by minimizing the average pairwise distance between the training vectors and their representative code vectors. The problem of generating an optimal codebook is a combinatorial optimization problem and it is NP-complete [1]. In other words, there is no known polynomial time algorithm for finding the globally optimal solution. However, reasonable suboptimal solutions are typically obtained by heuristic algorithms [2-7]. The most cited and widely used algorithm is the *generalized Lloyd algorithm* (GLA) [2, 3]. It starts with an initial solution, which is iteratively improved using two optimality criteria in turn until a local minimum has been reached.

A different approach is to build the codebook hierarchically. The *pairwise nearest neighbor* (PNN) algorithm [4] is a member of agglomerative clustering methods [8]. The PNN starts by constructing an initial codebook in which each training vector is considered as its own code vector. Two nearest code vectors are merged at each step of the algorithm and the process is repeated until the desired size of the codebook has been reached. The algorithm is straightforward to implement in its basic form and in comparison to the GLA [3] it gives good results. The PNN has also the advantage that

* Preliminary version of the paper was presented at *GraphiCon'98: The 8-th International Conference on Computer Graphics and Visualization*.

the hierarchical approach produces codebooks of differing sizes as a side-product. Thus, the PNN can be suited for joint minimization of distortion and entropy, see entropy-constrained PNN [9]. The algorithm can also be used to produce an initial codebook for another algorithm, such as the GLA, or it can be embedded into hybrid methods such as a genetic algorithm [6], or the iterative split-and-merge method [7].

A drawback of the PNN is the relatively high running time in its exact form. There are a large number of steps because typically $M \ll N$, and at each step, all pairwise distances are calculated for finding the pair of vectors to be merged. This is very slow for large training sets. Most of the computation originates from the calculation of the pairwise distances. However, only two code vectors are changed at each step of the PNN and therefore most of the distance calculations are unnecessary. A fast and space efficient implementation of the PNN has been recently given independently by [10] and [11]. The idea is to maintain for each cluster a pointer to its nearest neighbor and in this way, avoid unnecessary distance calculations. After the merge operation, the pointers must be updated only for clusters whose nearest neighbor is one of the merged clusters.

In this paper we propose an improved version of the above *nearest neighbor variant* [10]. The main idea is to reduce the distance calculations further by delaying the updates. We can do this because, as we will show, the cost function is monotonically increasing as a function of time. This means that the cost of merging any cluster with its nearest neighbor does not decrease during the algorithm. It is therefore sufficient to update a cost value only when it becomes the minimum one. The new method, referred as *Lazy-PNN*, reduces the number of updates considerably. Empirical tests show that in comparison to the PNN algorithm of [10] the Lazy-PNN improves the running time by 35% on average.

The rest of the paper is organized as follows. The problem formulation and the structure of the PNN are given in Section 2. The Lazy-PNN is then introduced in Section 3. Simulation results for various training sets are shown in Section 4, and conclusions are finally drawn in Section 5.

2. The PNN method

We next use the following notations:

- T Set of N training vectors $T = \{T_1, T_2, \dots, T_N\}$.
- C Codebook of m code vectors $C = \{C_1, C_2, \dots, C_m\}$.
- M Size of the final codebook.
- K The dimension of the vectors.
- S_i Cluster (set) of n_i training vectors.
- NN_i Index of the nearest neighbor of the cluster S_i .
- d_i Increase of the distortion if the clusters S_i and S_{NN_i} are merged.
- R_i Validity indicator; $R_i = \text{true}$ if and only if d_i is valid.

We consider a set of N training vectors in a K -dimensional Euclidean space. The task of the codebook construction is to find a set of M code vectors (i.e. a codebook) by

minimizing the average squared distance D between the training vectors T_i and their representative code vectors C_j :

$$D = \frac{1}{N} \sum_{j=1}^M \sum_{T_i \in S_j} \|T_i - C_j\|^2. \quad (1)$$

Here $S = \{S_1, \dots, S_M\}$ defines the clustering of the training set T . For a given codebook C , the optimal clustering can be constructed by assigning each training vector T_i to the cluster j_0 for which:

$$\|T_i - C_{j_0}\|^2 = \min_{j=1, \dots, M} \|T_i - C_j\|^2 \quad (2)$$

The basic structure of the PNN is shown in Fig. 1. The method starts by initializing each training vector T_i as its own cluster S_i . At each step of the algorithm, two nearest clusters (S_a and S_b) are searched and merged. The distance (*merge cost*) d between two clusters is defined as the increase in the distortion of the codebook if the clusters are merged. It is calculated as the squared Euclidean distance of the cluster centroids (code vectors) weighted by the number of vectors in the two clusters [4]:

$$d(S_a, S_b) = \frac{n_a n_b}{n_a + n_b} \cdot \|C_a - C_b\|^2 \quad (3)$$

The chosen clusters S_a and S_b are then merged. The size of the combined cluster S_{a+b} is $n_{a+b} = n_a + n_b$, and the corresponding code vector is the centroid of the training vectors in the cluster. It can be calculated as the weighted average of C_a and C_b :

$$C_{a+b} = \frac{n_a C_a + n_b C_b}{n_a + n_b} \quad (4)$$

It is thus sufficient to maintain only the cluster centroids (C_i) and the sizes of the clusters (n_i) in the implementation of the algorithm. The merge process is repeated until the codebook reaches the size M .

Let each training vector be a code vector ($m=N$).
Repeat
 Find two nearest clusters S_a and S_b to be merged.
 Merge S_a and S_b ; $m \leftarrow m-1$.
 Update data structures.
Until $m=M$.

Figure 1. Structure of the PNN method.

The exact PNN applies local optimization strategy where all possible cluster pairs are considered and the one increasing the distortion least (smallest cost function value) is

chosen for merge. Straightforward implementation [4] recalculates all distances at each step of the algorithm. No additional data structures are needed but the algorithm takes $O(N^3K)$ time because there are $O(N)$ steps in total, and there are $O(N^2)$ cluster pairs to be checked at each step.

Most of the computation of the PNN originates from the calculation of the pairwise distances. Since only two code vectors are changed at each step, most of the distance calculations are unnecessary. To reduce the number of distance calculations, previous pairwise cluster distances can be stored in an $N \times N$ matrix. The minimum cluster distance is searched from the matrix and the corresponding cluster pair is merged. New distances are then calculated between the new cluster and remaining clusters only. The algorithm runs in $O(N^2K + N^3)$ time where the former term originates from the distance calculations and the latter from the search for the minimum [12]. The disadvantages of this approach are cubic running time and quadratic memory consumption.

Kurita's method [13] stores all pairwise distances into a matrix, as above, but it utilizes a heap structure for searching the minimum distance. The merged clusters can be found by popping the smallest element from the top of the heap in $O(\log N)$ time. Only $O(N)$ distance updates are needed after each merge step; each of these updates takes $O(K + \log N)$ time because of the distance calculation and the heap operation. The method thus runs in $O(N^2K + N^2 \log N)$ time. The method still requires $O(N^2)$ memory, which is impractical for large training sets.

Another approach (τ -PNN) has been recently studied by Fränti and Kaukoranta in [10], and by Shen and Chang in [11]. The main idea is to maintain only a nearest neighbor pointer for each cluster. The index of the nearest cluster (NN_i) and the corresponding cost function value d_i are stored in the *nearest neighbor table*. The optimal cluster pair (S_a, S_b) to be merged can be found by a linear search among the d_i -values. After the merge operation, the nearest neighbor pointers must be updated for those clusters for which $NN_i = a$ or $NN_i = b$. Fortunately, in practice, there are only a small number (denoted by τ) of pointers to be updated on average. The method thus takes $O(\tau N^2K)$ time in total. In addition to that, the memory requirement of this approach is only $O(N)$.

Several approximate variants of the PNN have also been considered in the literature. Equitz uses K - d tree for localizing the search for the code vectors, and the algorithm merges several vector pairs at the same time [4]. Another possibility is to generate a preliminary codebook of size M_0 ($N > M_0 > M$) using the GLA and then apply the exact PNN until the codebook reaches its final size M [14]. These variants decrease the running time at the cost of increased distortion. In the following, we focus on the exact PNN variants.

3. Lazy-PNN algorithm

We propose next an improved version of the nearest neighbor variant of the PNN [10]. The idea is to reduce the distance calculations even further. Although the total number of updates (τ) is rather small on average, the search of the nearest neighbor is still an expensive $O(NK)$ time operation and it dominates the running time of the algorithm.

However, some of the distance calculations can be delayed and therefore a remarkable number of updates may be avoided. The new method is referred here as *Lazy-PNN*.

3.1. Monotony property

The application of delayed distance calculations is based on the *monotony property* of the nearest neighbor distances of the clusters, which is defined as follows. Suppose that at a certain moment the minimal merge cost is $d(S_a, S_b)$ and the clusters S_a and S_b are merged. It is possible that the centroid of the merged cluster (C_{a+b}) becomes closer to the centroid of a third cluster S_c than C_c was in respect to the original cluster centroids (C_a and C_b), see Fig. 2. However, the merge cost $d(S_{a+b}, S_c)$ is never smaller than $\min\{d(S_a, S_c), d(S_b, S_c)\}$ because the large size of the merged cluster compensates the potential decrease of the Euclidean distance. The cost function d is therefore monotonically increasing as a function of time. This is formalized in the following lemma:

Lemma 1. Consider the clusters S_a, S_b, S_c with centroids C_a, C_b, C_c , and frequencies n_a, n_b, n_c . Assume that $d(S_a, S_b) \leq d(S_a, S_c) \leq d(S_b, S_c)$ and $n_a, n_b, n_c \geq 1$. Then it holds that $d(S_a, S_c) \leq d(S_{a+b}, S_c)$.

Proof. We have the following relationships between the distances of the cluster centroids:

$$\|C_a - C_c\|^2 = \|C_a - C_b\|^2 + \|C_c - C_b\|^2 - 2(C_a - C_b) \cdot (C_c - C_b) \quad (5)$$

and

$$\|C_{a+b} - C_c\|^2 = \left(\frac{n_a}{n_a + n_b}\right)^2 \|C_a - C_b\|^2 + \|C_c - C_b\|^2 - 2\left(\frac{n_a}{n_a + n_b}(C_a - C_b)\right) \cdot (C_c - C_b). \quad (6)$$

We can thus write (6) in the form:

$$\begin{aligned} \|C_{a+b} - C_c\|^2 &= \left(\frac{n_a}{n_a + n_b}\right)^2 \|C_a - C_b\|^2 + \|C_c - C_b\|^2 \\ &\quad - \frac{n_a}{n_a + n_b} \left(\|C_a - C_b\|^2 + \|C_b - C_c\|^2 - \|C_a - C_c\|^2 \right) \\ &= \frac{n_a}{n_a + n_b} \|C_a - C_c\|^2 + \frac{n_b}{n_a + n_b} \|C_c - C_b\|^2 - \frac{n_a n_b}{(n_a + n_b)^2} \|C_a - C_b\|^2 \end{aligned}$$

This gives the condition

$$(n_a + n_b) \|C_{a+b} - C_c\|^2 = n_a \|C_a - C_c\|^2 + n_b \|C_c - C_b\|^2 - d(S_a, S_b).$$

Now, we can write the difference of the merge costs in form:

$$\begin{aligned}
d(S_{a+b}, S_c) - d(S_a, S_c) &= \frac{n_c(n_a + n_b)}{n_a + n_b + n_c} \|C_{a+b} - C_c\|^2 - d(S_a, S_c) \\
&= \frac{n_a n_c}{n_a + n_b + n_c} \|C_a - C_c\|^2 + \frac{n_b n_c}{n_a + n_b + n_c} \|C_c - C_b\|^2 - \frac{n_c}{n_a + n_b + n_c} d(S_a, S_b) - d(S_a, S_c) \\
&= \frac{n_a + n_c}{n_a + n_b + n_c} \cdot \frac{n_a n_c}{n_a + n_c} \|C_a - C_c\|^2 + \frac{n_b + n_c}{n_a + n_b + n_c} \cdot \frac{n_b n_c}{n_b + n_c} \|C_c - C_b\|^2 \\
&\quad - \frac{n_c}{n_a + n_b + n_c} d(S_a, S_b) - d(S_a, S_c) \\
&= \frac{n_a + n_c}{n_a + n_b + n_c} \cdot d(S_a, S_c) + \frac{n_b + n_c}{n_a + n_b + n_c} \cdot d(S_b, S_c) - \frac{n_c}{n_a + n_b + n_c} d(S_a, S_b) - d(S_a, S_c) \\
&= \frac{n_b + n_c}{n_a + n_b + n_c} \cdot d(S_b, S_c) - \frac{n_c}{n_a + n_b + n_c} d(S_a, S_b) - \frac{n_b}{n_a + n_b + n_c} d(S_a, S_c)
\end{aligned}$$

This gives us the formula

$$d(S_{a+b}, S_c) - d(S_a, S_c) = \frac{n_b [d(S_b, S_c) - d(S_a, S_c)] + n_c [d(S_b, S_c) - d(S_a, S_b)]}{n_a + n_b + n_c}. \quad (7)$$

The value of (7) is clearly now positive due to the assumptions made in lemma. This proves the lemma. ■

3.2. Lazy update

Because of the monotony property, we know that a given distance value never decreases due to its update. We may therefore delay the distance calculations until the old cost function value becomes a candidate for being the smallest distance. We accomplish this by marking each value whether it is up to date or not. The optimal cluster pair to be merged can now be searched as before with only one difference; when an out-dated distance value is found to be minimal it is recalculated. This practice does not compromise the exactness of the algorithm but it may remarkably reduce the number of expensive distance calculations.

The lazy processing can be applied to the nearest neighbor method as such. We take one step further and maintain a min-heap of the distance values (d_i) and the corresponding nearest neighbor pointers (NN_i). The heap elements contain an additional flag (R_i), which indicates whether the distance value is up-to-date or not. The difference to Kurita's method is that we only store one element per cluster whereas Kurita stores all distances. The use of the heap has no asymptotic influence on the running time of the Lazy-PNN because the recalculation of the distance values still dominates the running time. The heap, however, may speed-up the practical implementation because we need to consider only the root of the heap and therefore may potentially avoid a number of distance recalculations.




Figure 2. Illustration of the clusters in 2-dimensional space. C_a and C_b are the centroids of the two clusters to be merged, C_{a+b} is the centroid of the merged cluster, and C_c is the centroid of any other cluster.

3.3. Implementation

The pseudo code of the Lazy-PNN algorithm is presented in Fig. 3. The algorithm starts by initializing each cluster (or code vector) with one training vector. For each cluster S_i , the nearest neighbor according to (3) is searched among the other clusters S_j ($i \neq j$). Information of the nearest neighbor is stored and marked valid. The nearest neighbor distances d_i of all clusters are inserted in the minimum heap H . The algorithm is then iterated until the size of the codebook reduces to M . At each step, the cluster S_a with the smallest d -value is deleted from H . If the d -value of S_a is out-of-date ($R_a = \text{False}$) its nearest neighbor is recalculated and reinserted in the heap. The process is repeated until a valid minimal distance is obtained from the top of H .

The cluster S_a and its nearest neighbor S_b ($b = NN_a$) are merged according to (4) and the number of vectors in the new cluster is calculated. The nearest neighbor of the new cluster is then determined and inserted in the heap according to its d -value. The information of the merged cluster is stored into the place of S_a leaving S_b unused. The non-existing cluster S_b is removed from the heap. This can be done in $O(\log N)$ time by maintaining a position index to the heap for each cluster. To avoid gaps in the indexing, the last cluster S_m replaces S_b . All the nearest neighbor pointers NN_i to S_m are reassigned to S_b . Finally, the size of the intermediate codebook m is subtracted by one.

```

Input: Training set  $\{T_i\}$  and final codebook size  $M$ .
Output:  $\{C_i\}$ .
Main program
 $H \leftarrow \emptyset$ ;
 $m \leftarrow N$ ;
for  $\forall i \in [1, m]$ :  $C_i \leftarrow T_i$ ;  $n_i \leftarrow 1$ ;
for  $\forall i \in [1, m]$ : UpdateNNpointer( $H, i$ );
repeat
    PickPair( $H, a, b$ );
    MergeClusters( $a, b$ );
    UpdateNNpointer( $H, a$ );
until  $m = M$ ;

Procedure PickPair( $H, a, b$ ):
 $a \leftarrow \text{DeleteMin}(H)$ ;
while  $R_a = \text{False}$ 
    UpdateNNpointer( $H, a$ );
     $a \leftarrow \text{DeleteMin}(H)$ ;
 $b \leftarrow NN_a$ ; Remove( $H, b$ );

Procedure MergeClusters( $a, b$ ):
 $C_a \leftarrow (n_a C_a + n_b C_b) / (n_a + n_b)$ ;
 $n_a \leftarrow n_a + n_b$ ;
for  $\forall i \in [1, m]$ : if  $NN_i = a \vee NN_i = b$  then  $R_i \leftarrow \text{False}$ ;
 $C_b \leftarrow C_m$ ;  $n_b \leftarrow n_m$ ;  $NN_b \leftarrow NN_m$ ;  $d_b \leftarrow d_m$ ;  $R_b \leftarrow R_m$ ;
for  $\forall i \in [1, m-1]$ : if  $NN_i = m$  then  $NN_i \leftarrow b$ ;
 $m \leftarrow m-1$ ;

Procedure UpdateNNpointer( $H, a$ )
 $NN_a \leftarrow \text{FindNearestCluster}(a)$ ;
 $d_a \leftarrow d(a, NN_a)$ ;
 $R_a \leftarrow \text{True}$ ;
Insert( $H, a$ );

```

Figure 3. Pseudo code of the Lazy-PNN.

4. Practical results

We generated training sets from six different images: *Bridge*, *Camera*, *Miss America*, *Table tennis*, *Airplane* and *House*, see Fig. 4. The vectors in the first two sets (*Bridge*, *Camera*) are 4×4 pixel blocks from the image. The third and fourth sets (*Miss America*, *Table Tennis*) have been obtained by subtracting two subsequent image frames of the original video image sequences, and then constructing 4×4 spatial pixel blocks from the residuals. Only the first two frames have been used. The fifth and sixth data sets (*Airplane*, *House*) consist of color values of the *RGB* images, prequantized to 5 bits per color component. Applications of this kind of data sets is found in image and video image coding (*Bridge*, *Camera*, *Miss America*, *Table tennis*), and in color image quantization (*Airplane*, *House*).



Figure 4. Sources of the training sets. *The training set *Table tennis* is constructed by random sampling only every fourth block. **The images *Airplane* and *House* are prequantized by 5 bits per color component.

Properties of the compared PNN methods are presented in Table 1. Table 2 shows a summary of the test results for three main variants. Kurita’s method was not applied because its memory consumption is too high for these training sets. The size of the codebook was fixed to $M=256$. Both nearest neighbor variants (τ -PNN and Lazy-PNN) are clearly superior to the original PNN being about 100 to 500 times faster. From these two variants, the Lazy-PNN is about 35% faster. The speed-up originates mainly from the decreased number of distance recalculations; the average number of updates (τ) varied from 4.4 to 5.6 in the τ -PNN, and from 3.0 to 3.8 in the Lazy-PNN. Small improvement is also due to the use of the heap structure.

In order to compare the nearest neighbor variants with the Kurita’s method we generated subsets from *Bridge* and *House* by random sampling. The smaller training sets are of size $N=(128, 256, 384, 512, 640, 768, 896, 1024)$. In these tests the codebook size was set to $M=1$ for getting the maximal number of iterations. The results for the two cases are illustrated in Fig. 5. The Lazy-PNN is comparable to the Kurita’s method in speed but the Lazy-PNN has the benefit of smaller memory consumption. The actual running times are virtually the same for training sets with large vector dimensions (e.g. *Bridge*, $K=16$), whereas for training sets with smaller dimensions (e.g. *House*, $K=3$) the Lazy-PNN is faster.

The exact PNN is compared in Fig. 6 with two approximative variants of the PNN. The *Fast-PNN* refers to the $O(N \log N)$ time variant of the PNN given by Equitz [4]. The figure shows the (running time, MSE) pairs for the Fast-PNN, when the maximum bucket size of the K -d tree ranges from 8 to 160. The method is very fast with small bucket sizes and gets slower as the maximum bucket size increases. Another

approximative PNN variant, the GLA-PNN by deGarrido et al. [14], starts with an initial codebook of size M_0 ($M < M_0 < N$). The initial codebook is first generated by the GLA and then reduced to the final size using the exact PNN (Lazy-PNN). The method is a compromise between higher speed of the GLA and better quality of the PNN, and is parameterized by the choice of M_0 . From the two approximative methods, the GLA-PNN has a better time-distortion performance than Equitz's K - d tree variant, see Fig. 6. The exact PNN, on the other hand, always produces the lowest MSE values.

Table 1. Summary of the compared exact PNN methods.

Method	Time complexity	Extra data	Space complexity
Original PNN [4]	$O(N^3K)$	-	$O(N)$
Kurita's method [11]	$O(N^2K + N^2 \log N)$	Distance matrix	$O(N^2)$
τ -PNN [8]	$O(\tau N^2K)$	Nearest neighbor table	$O(N)$
Lazy-PNN	$O(\tau N^2K)$	Nearest neighbor table	$O(N)$

Table 2. Summary of the running times (in seconds).

Training set	Original PNN	τ -PNN	Lazy-PNN	Time saved by Lazy-PNN
<i>Bridge</i>	73385	334	224	32.8%
<i>Camera</i>	73797	303	213	29.8%
<i>Miss America</i>	292351	874	569	34.9%
<i>Table tennis</i>	177019	650	429	34.0%
<i>Airplane</i>	5123	51	32	38.1%
<i>House</i>	2514	29	19	33.8%




Figure 5. Comparison of the fast exact PNN methods for subsets of *House* ($K=3$) and *Bridge* ($K=16$) when $M=1$. The size of the training set is limited to $N=1024$ due to the memory requirements of Kurita's method.

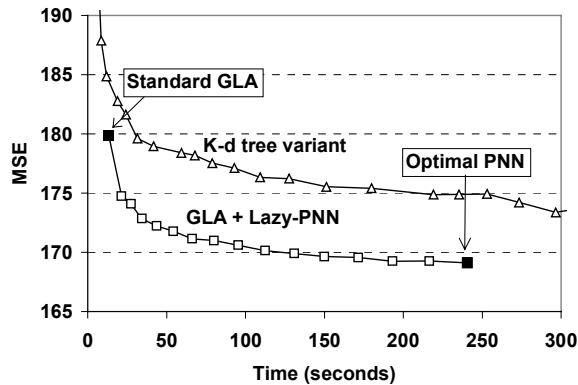


Fig. 6. Comparison of the suboptimal PNN variants (for *Bridge*). The size of the initial codebook of GLA-PNN varies from 256 (standard GLA) to 4096 (optimal PNN). The *K-d tree* variant is parameterized by varying the maximum bucket size from 8 to 160.

5. Conclusion

A fast variant of the exact PNN algorithm was introduced. The main idea of the algorithm is to maintain a table of nearest neighbors as in the τ -PNN algorithm. In addition to that we postpone the updating of the closest distance information to the moment when the (old) distance becomes the new tentative minimum among the cluster distances. This action is possible due to the monotony of the cluster distances. The monotony property is utilized further by using a heap structure as a priority queue to maintain the set of cluster distances.

Our practical tests indicate that the Lazy-PNN is about 100 to 500 times faster than the original PNN. The new method is comparable to the Kurita's algorithm in speed but it has the benefit of smaller memory by factor N . In the comparison to τ -PNN the number of updated cluster distances was observed to reduce by 35% on average. The proposed method is rather simple to implement and practical because no distance matrix is needed for storing the pairwise distances.

We also gave a proof of the monotony property for the vectors in Euclidean space. It is an open question whether the result generalizes to other cluster distances. This would expand the usefulness of the new algorithm to the general clustering problem.

Acknowledgements

The work of Pasi Fränti was supported by a grant from the Academy of Finland.

References

- [1] M.R. Garey, D.S. Johnson, H.S. Witsenhausen, "The Complexity of the Generalized Lloyd-Max Problem". *IEEE Transactions on Information Theory*, Vol.28 (2), pp.255-256, March 1982.

- [2] A. Gersho and R.M. Gray, *Vector Quantization and Signal Compression*. Kluwer Academic Publishers, Boston 1992.
- [3] Y. Linde, A. Buzo and R.M. Gray, "An Algorithm for Vector Quantizer Design". *IEEE Transactions on Communications*, **28** (1), pp. 84-95, January 1980.
- [4] W.H. Equitz, "A new vector quantization clustering algorithm", *IEEE Transactions on Acoustics, Speech, and Signal Processing* **37** (10), pp. 1568-1575, October 1989.
- [5] P. Fränti, T. Kaukoranta and O. Nevalainen, "On the splitting method for VQ codebook generation", *Optical Engineering*, **36** (11), pp. 3043-3051, November 1997.
- [6] P. Fränti, J. Kivijärvi, T. Kaukoranta and O. Nevalainen, "Genetic algorithms for large scale clustering problem", *The Computer Journal*, **40** (9), pp. 547-554, 1997.
- [7] T. Kaukoranta, P. Fränti and O. Nevalainen, "Iterative split-and-merge algorithm for VQ codebook generation", *Optical Engineering*, **37** (10), pp. 2726-2732, October 1998.
- [8] G.N. Lance and W.T. Williams, "A general theory of classificatory sorting strategies 1. Hierarchical systems", *Computer Journal*, **9**, pp. 373-380, 1967.
- [9] F. Kossentini and M.J.T. Smith, "A fast PNN design algorithm for entropy-constrained residual vector quantization", *IEEE Transactions on Image Processing*, **7** (5), pp. 1045-1050, July 1998.
- [10] P. Fränti and T. Kaukoranta, "Fast implementation of the optimal PNN method", *Proc. IEEE Int. Conf. on Image Processing (ICIP)*, Chicago, Illinois, October 1998.
- [11] D.-F. Shen and K.-S. Chang, "Fast PNN algorithm for design of VQ initial codebook", *Proc. SPIE 3309, Visual Communications and Image Processing '98*, San Jose, California, pp. 842-850, 1998.
- [12] J. Shanbehzadeh and P.O. Ogunbona, "On the computational complexity of the LBG and PNN algorithms". *IEEE Transactions on Image Processing* **6** (4), pp. 614-616, April 1997.
- [13] T. Kurita, "An efficient agglomerative clustering algorithm using a heap". *Pattern Recognition*, **24** (3), pp. 205-209, March 1991.
- [14] D.P. de Garrido, W.A. Pearlman and W.A. Finamore, "A clustering algorithm for entropy-constrained vector quantizer design with applications in coding image pyramids". *IEEE Transactions on Circuits and Systems for Video Technology* **5** (2), pp. 83-95, April 1995.