Classification of Medical Images Using Local Representations

Roberto Paredes, Daniel Keysers[1], Thomas M. Lehmann[2], Berthold Wein[3], Hermann Ney[1], and Enrique Vidal

Instituto Tecnológico de Informática,
Universidad Politécnica de Valencia, Camino de Vera s/n, E-46022, Valencia, Spain.
rparedes@iti.upv.es
[1]Lehrstuhl für Informatik VI, RWTH Aachen, D-52056 Aachen, Germany
[2]Institute of Medical Informatics, RWTH Aachen, D-52057 Aachen, Germany
[3]Department of Diagnostic Radiology, RWTH Aachen, D-52057 Aachen, Germany

Abstract In medical image retrieval, the images are usually subject to a large range of variability. In order to classify medical images, we therefore propose the use of local representations, which are small square windows taken from the images. This approach is combined with a fast approximate k-nearest neighbor technique and yields state-of-the-art results on a medical image database of 1617 images.

1 Introduction

Recently, research within the field of content-based medical image retrieval has attracted a lot of attention. The concept for content-based image retrieval in medical applications (IRMA) is based on the separation of the following steps: categorization of the entire image, registration with respect to prototypes, extraction and query-dependent selection of local features, hierarchical blob representation including object identification and finally, image retrieval performed on a data-reduced abstract image representation [1].

This work presents a new approach to the first step of this paradigm –categorization of the entire image– using local representations. Medical images are usually complex, subject to high variability and composed of different smaller structures. Therefore, we propose to use local representations of the images for classification, which is an approach that has been applied very successfully to the task of face recognition [2].

2 Methods

In a classical classifier [3], each object for training and test is represented by a feature vector, and a discrimination rule is applied to classify a test vector. In the image classification problem, this feature vector is usually obtained from the entire image, using the appearance-based approach (each pixel corresponds to one feature) or some type of feature extraction. Due to the complexity and variability of medical images, a local representation scheme is adopted. Many local representations have been proposed, mainly

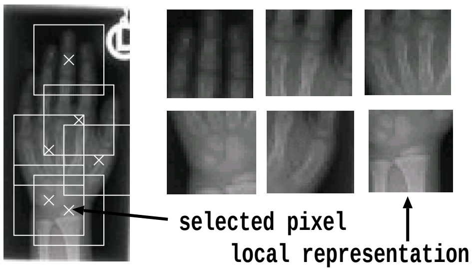

Figure 1. Example of extracted local representations from selected pixels.

in the image database retrieval literature [4,5,6]. In the present work, each image is represented by several (possibly overlapping) square windows of size $W \times W$, which correspond to a set of "local appearances" (Figure 1).

To obtain the local feature vectors from an image, a selection of windows with highly relevant and discriminative content is needed. Although a number of methods exist to detect such windows [7], most of them are not appropriate for medical imaging or they are computationally too expensive. In this work, the local variance in a small window around each pixel is used as measure of information. Those pixels having local variance above a certain global threshold are selected and the surrounding window is used as a representation of the whole image. Figure 1 shows an example of possible local representations.

In the classification procedure, each test image is also represented using the local features selected. To approximate the posterior probability that a certain local feature belongs to an image of a given class, the k-nearest neighbor algorithm is used. This posterior probability is then used to obtain a combined decision for the set of local features.

Representing objects by several local features involves a computational problem if the number of local features for representing an object is very large. The k-nearest neighbor algorithm needs to compare every local feature of a test object with every local feature of every training object. This high computational cost is considerably reduced by using a fast approximate k-nearest neighbor search technique. This technique uses a kd-tree structure to store the set of local features from the training images. In a kd-tree, the search of the nearest neighbor of a test point is performed starting from the root, which represents the whole space, and choosing at each node the sub-tree that represents the region of the space containing the test point. When a leaf is reached, an exhaustive search of the b prototypes residing in the associated region is performed. Since the closest point may also be a member of some other region the algorithm needs to backtrack until all possible regions are checked.

If a guaranteed exact solution is not needed, as can be assumed in our case, the backtracking process can be aborted as soon as a certain criterion is met by the current best solution. In [8], the concept of $(1 + \epsilon)$-approximate nearest neighbor query is introduced. A point p is a $(1 + \epsilon)$-approximate nearest neighbor of q if the distance from p to q is less than $1 + \epsilon$ times the distance from p to its nearest neighbor. This concept is

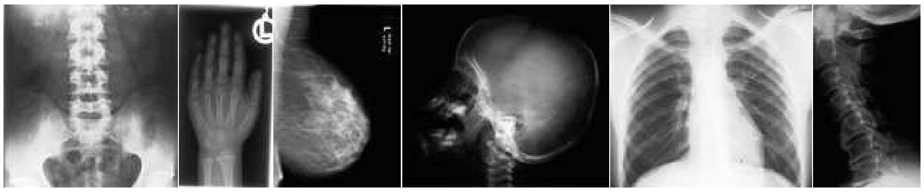

Figure 2. Example images (IRMA). Left to right: abdomen, limbs, breast, skull, chest, spine.

used here to obtain an efficient approximate search that can easily cope with very large sets of reference vectors at significantly lower runtime.

Moreover, the possibly high dimensionality of $W \times W$ vector components is reduced using a principal components analysis on the set of all local features extracted from the training set.

For each local feature of the test image, the k-nearest neighbor algorithm gives a fraction of votes to each class, which is an approximation of the posterior probability of each local feature to belong to each class. As each of the vectors obtained from the test image can be classified into a different class, a joint decision scheme is required to finally decide on a single class for the entire test image. The probabilities obtained from each local feature are combined using the sum rule to obtain the overall posterior probability for the entire image for each class. The test image is assigned to the class with highest posterior probability. In the context of the sun rule and the k-nearest neighbor approximation for the posterior probability, this decision corresponds to the most voted class counting all votes from all local features of the test image [2].

3 Experimental Results

The proposed approach was applied to a database of 1617 medical medical radiographs from the RWTH Aachen University of Technology IRMA project [9]. Figure 2 shows example images from the six different categories.

A leaving-one-out procedure was carried out to estimate the classification error rate. This means that each image of the corpus is classified using the remaining 1616 images.

Using a window size $W = 19$ for the local features, a 10.6% error rate estimation was obtained. Table 1 shows the results of the proposed method in comparison with other techniques [10]. Using a thresholding distance to avoid great differences between the distances of two image pixels, the classification error rate of the proposed local feature approach was reduced from 10.6% to 9.6%. The same technique reduced the error rate of the approach based on distorted tangent distance from 10.4% to 8.0% [10].

4 Conclusion

A novel approach is proposed for medical image classification, which combines a simple local representation method with a direct voting scheme based on the sum of votes from the k-nearest neighbor approximation for each local feature. This approach rises computational problems that can be effectively solved using a fast approximate k-nearest neighbor search technique. One strong advantage of the presented method is

method	error [%]
cooccurrence matrices	29.0
square images, 1-nearest neighbor	18.1
square images, kernel densities	16.4
square images, kernel densities, aspect ratio feature	14.9
kernel densities, Mahalanobis distance	14.0
kernel densities, tangent distance	13.3
kernel densities, image distortion model	12.1
local features	**10.6**
kernel densities, distorted tangent distance	10.4

Table 1. Performance of several methods on the IRMA data set.

that it uses very few parameters that need to be tuned to the task at hand and works equally well for a large number of image classification tasks such as medical image classification, face recognition [2] or handwritten character recognition. Experimental results with the IRMA medical image database are reported showing the effectiveness of the proposed approach.

Current work is under way to test the proposed approach on other databases. We are also interested in studying other voting schemes and approximations to the posterior probabilities of each local feature instead of the k-nearest neighbor.

References

1. Lehmann T, Wein B, Dahmen J, et al.: Content-based image retrieval in medical applications: a novel multi-step approach. Procs SPIE 3972(32):312–320, February 2000.
2. Paredes R, Perez-Cortes JC, Juan A, et al.: Local representations and a direct voting scheme for face recognition. Workshop on Pattern Recognition in Information Systems, Setúbal, Portugal, July 2001.
3. Duda R and Hart P: Pattern Recognition and Scene Analysis. John Wiley, New York, 1973.
4. Mohr R, Picard S, and Schmid C. Bayesian decision versus voting for image retrieval. Procs CAIP-97, 1997.
5. Schmid C and Mohr R: Local grayvalue invariants for image retrieval. IEEE TPAMI 19(5):530–535, 1997.
6. Shyu C, Brodley C, Kak A, et al.: Local versus global features for content-based image retrieval. Procs IEEE Workshop on Content-Based Access of Image and Video Libraries 30–34, June 1998.
7. Deriche R and Giraudon G: A computational approach to corner and vertex detection. Int. Journal Computer Vision 10:101–124, 1993.
8. Arya S, Mount D, Netanyahu N, et al.: An optimal algorithm for approximate nearest neighbor searching fixed dimensions. Journal of the ACM 45:891–923, 1998.
9. Dahmen J, Keysers D, Motter M, et al.: An automatic approach to invariant radiograph classification. Procs BVM-01:337–341, Lübeck, Germany, March 2001.
10. Dahmen J, Keysers D, Ney H, et al.: Statistical image object recognition using mixture densities. Journal of Mathematical Imaging and Vision 14(3):285–296, May 2001.