AN ARCHITECTURE FOR THE RECOGNITION AND CLASSIFICATION OF MULTIPLE SCLEROSIS LESIONS IN MR IMAGES

E. Ardizzone *†, R. Pirrone *†

* *DIAI – Dipartimento di Ingegneria Automatica e Informatica, University of Palermo Italy*
† *CERE – CEntro di studio sulle Reti di Elaboratori, CNR, Palermo Italy*

Abstract
A software architecture is presented that is able to perform classification of focal lesions due to the multiple sclerosis disease in MR images of the brain. The methodology proceeds through four main steps: tissue segmentation, re-clustering and tissue classification, lesion localization and lesion classification. Images are first segmented using the FCM algorithm; then the images of each cluster are processed in order to classify and label non-pathologic tissues making use of simple decision algorithms based on suitable numerical indices related to tissue morphology. All possible candidates to be sclerosis lesions are then located by means of morphological operations applied to binary images of single tissues. Finally the classification step is performed together with an estimate of the position and the shape for each lesion. Each candidate has been characterized by means of a set of measurements related to its shape, position and brightness as a way to code the clinical knowledge about the disease under investigation. Classification is implemented using both an algorithmic classifier and a multi-layer perceptron trained using a back-propagation scheme, and the performances of the two approaches are compared. The outline of the whole architecture is presented and the experimental results are reported.

1. Introduction

During last decades, a lot of work has been spent to the aim of segmenting brain tissues and lesions in MRI. This task is of enormous relevance in the diagnosis and therapy of all the afflictions of the brain. An efficient and reliable segmentation tool allows the doctor to determine more precisely the disease entity and to better follow its evolution.

The separation of the brain from skull, eyes and other tissues and the classification of multiple sclerosis lesions are the two main activities in the research work related to this paper. However, the general framework of the present work is the development of a design methodology for tools acting as a support to the diagnosis process in all those affections that require medical imaging.

Such a tool presents to the clinician both a qualitative and a quantitative description of the disease and allows him to formulate more detailed diagnoses. Besides, it's possible to better follow the evolution of the particular syndrome investigated because comparisons can be easily carried out between successive tests for the same patient.

The realization of general-purpose tools of this kind is very hard or impossible, due to the nature of the problem under consideration. On the contrary, the design of tools devoted to analyze a single disease may be supported by the well defined nature of the necessary clinical knowledge, that can be described as a set of features e.g. shape or brightness measures in the regions of interest extracted from medical images.

The clinical picture can be submitted to several specialized modules concurring to provide a pool of possible diagnoses each with its degree of confidence and a set of quantitative measures. The doctor can evaluate the system responses and assess the more likely ones in order to make his diagnosis.

In these cases a multi-module approach as opposed to the classic AI diagnostic systems can result in an overall increase of diagnosis accuracy [13]. In particular, a software system able to perform segmentation and classification of tissues in MR images of the brain is presented. The system is aimed to locate and recognize focal lesions deriving from multiple sclerosis disease.

The presented system can classify and label pathologic and non-pathologic tissues in the analyzed images, and give a numeric estimate of the shape and the position of each lesion.

The architecture is made up by two main stages: first the segmentation and labeling of brain tissues is performed and then the potential sclerosis objects are located and classified.

The tissue segmentation is based on the well-known fuzzy c-means (FCM) algorithm [6,7,8] that has already been successfully applied to the segmentation of MR images representing tumor lesions of the brain [1,2,4,5]. The application of fuzzy c-means gives a correct distinction between different tissues, but it often results in an over-segmentation in which a single tissue is split into more than one cluster. A re-clustering phase follows which uses a decision algorithm to merge the over-segmented clusters and to label the different tissues in the RM slices.

Sclerosis candidate objects are located through a morphological analysis of the binary images representing the re-clustered tissues.
In the classification step we have defined a set of measures that characterize a possible lesion object. We've put them together with the information derived in the tissues labeling phase in order to obtain a feature vector representing each potential lesion.
The architecture of the classifier is a multi-layer perceptron trained using the back-propagation algorithm. Due to the characteristics of the problem, a connectionist approach may be regarded as the more appropriate to accomplish this task.
Sclerosis segmentation and classification is a very hard goal. In this task, the anatomical knowledge of the disease is involved, which has to be used to perform a fine discrimination between white matter (WM), gray matter (GM), peripheral and ventricular cerebro-spinal fluid (CSF) in order to locate all possible sites of the lesions.
Several approaches have been proposed in the last years: Zijdenbos [15] proposes a neural approach to segment brain in five classes: background, white matter, gray matter, CSF and white matter lesions (WML) by means of a multi-layer perceptron trained with the back-propagation algorithm. Kamber [11] and Johnston [10] propose stochastic approaches to the 3D segmentation of the brain. In [11] a probabilistic model providing the a priori probability of the tissues distribution is derived from the analysis of several volunteers' slices, and Bayesian classification is performed on 3D data sets that have been previously affine-transformed in order to normalize them. In [10] a stochastic relaxation approach to 3D analysis is proposed that is based on the use of the iterated conditional modes (ICM) algorithm. This algorithm produces a map of voxels, each of them with an estimate of its percentage composition in terms of the various tissues and/or lesions. The method works separately on conspicuous and subtle lesions. Udupa [14] makes use of an original interactive procedure that is based on the concept of fuzzy-connectedness. This is a re-formulation of the topological connectedness with variable degrees of strength. WM, GM and CSF fuzzy-connected 3D objects are defined starting from some regions manually outlined by the operator. The holes existing between these objects are located as possible sites of conspicuous lesions, then the operator interacts with the system to accept or reject its results. Subtle lesions are segmented starting from a weaker definition of fuzzy-connectedness in the white matter object. Krishnan [12] proposes a method based solely on image analysis techniques. The approach starts from the computation of the ratio image between the PD and the T2 signal. First, a manual thresholding of the ratio image is used to obtain the white matter, CSF and lesions masks. These masks are applied to the PD and T2 images and the related histograms are analyzed in order to obtain four threshold values that are used to actually segment the T2, PD images. In this way the CSF is removed while conspicuous and subtle lesions are obtained.
Most of the previously described approaches require an initial manual segmentation and need a comparison with the work of an expert. On the contrary, our approach is fully automated as regards segmentation and re-clustering in which we use both data information and clinical knowledge. Besides, we encode the expert assessment in the learning phase of the classifier: so this component is trained to provide estimates that are close to those of the expert both for lesions and artifacts.
In the present paper we start from some suggestions from the works of Clark [1], Hall [5], Liang [3], and Bezdek [4] on the segmentation of MR images in order to classify pathological and non-pathological tissues. We derive a quantitative description of the position and the shape of focal lesions.
The rest of the paper is arranged as follows. In the following section we'll explain in detail the segmentation and re-clustering steps in order to achieve a correct classification of tissues. Next, the procedure for candidate sclerosis location and building of the feature vectors for classification will be described. Finally, a comparison between the performances of the two classifiers we adopted will be carried out, along with the analysis of the experimental results.

2. Tissue segmentation

In this section an outline of the segmentation, re-clustering and tissue labeling stages is explained: further details can be found in [9].
The tissue segmentation subsystem is based on the well-known FCM approach to MR images segmentation.
Formally, a clustering problem can be expressed as follows: given a set $X=\{\boldsymbol{x_1}, \boldsymbol{x_2}, \ldots \boldsymbol{x_n}\}$, where $\boldsymbol{x_i} \in R^P$ denotes an object characterized by p attributes, subdivide the n elements of X into $c \in [2, n)$ clusters.
To solve this problem it is convenient to represent each cluster by means of its center $\boldsymbol{v_i} \in R^P$ and to assign an object to a cluster according to the distance of the object from the cluster center, which can be measured by means of Euclidean norm.
Let define the c x n matrix $\boldsymbol{U} = [\mu_{ik}]$ where μ_{ik} denotes the membership degree of $\boldsymbol{x_k}$ to the cluster S_i. If the following conditions are satisfied, $\boldsymbol{U}$ is called a fuzzy c-partition [1].

$$\mu_{ik} \in [0,1]\ ,\ i \in [1,c]\ ,\ k \in [1,n]\ ;\ \sum_{i=1}^{c} \mu_{ik} = 1\ ,\ k \in [1,n];\ \sum_{k=1}^{n} \mu_{ik} \in (0,n)\ \ ,\ i \in [1,c].$$

M_{fc} denotes the set of matrices satisfying the above conditions. Now, both the c x n values of μ_{ik} and the cluster centers $\boldsymbol{v_i}$ are unknown. A criterion frequently used for clustering is the *variance criterion* according to which the following problem is solved, that is to find a fuzzy c-partition in order to minimize the index $J(\boldsymbol{U}, \boldsymbol{v_1}, \ldots \boldsymbol{v_n})$ given by:

$$J(\boldsymbol{U}, \boldsymbol{v}_1, \cdots \boldsymbol{v}_n) = \sum_{k=1}^{n} \sum_{i=1}^{c} \mu_{ik}^{m} \left\| \boldsymbol{x}_i - \boldsymbol{v}_i \right\|_{\boldsymbol{G}}^{2} \tag{2.1}$$

Here *m* is a suitable weight exponent and

$$\left\| \boldsymbol{x}_i - \boldsymbol{v}_i \right\|_{\boldsymbol{G}}^{2} = (\boldsymbol{x}_i - \boldsymbol{v}_i)^T \boldsymbol{G} (\boldsymbol{x}_i - \boldsymbol{v}_i)\ ,\ \boldsymbol{G} > \boldsymbol{0}$$

is the *inner product norm*, where $\boldsymbol{G}$ is a p x p positive definite matrix; if $\boldsymbol{G} \equiv \boldsymbol{I}$ we obtain the Euclidean norm. Summations in (2.1) are with respect to μ_{ik} and $\boldsymbol{v_i}$, $i \in [1, c]$, $k \in [1, n]$ subject to $\boldsymbol{U} \in M_{fc}$ and $\boldsymbol{v_i} \in R^p$. The FCM algorithm provides an iterative solution to the problem stated above starting from a suitable initial choice $\boldsymbol{U}^0$ for the fuzzy c-partition [1].

In our implementation, the segmentation tool classifies the pixels in a MR slice in 10 different clusters (c = 10). Each pixel is characterized by a 2-dimensional feature vector (p = 2) made up by the gray levels of the pixel in the T2-weighted and in the Proton Density image of the slice. We don't take into account the T1-weighted image because it is of minor interest in the multiple sclerosis diagnosis, due to the poor contrast between focal lesions and surrounding tissue. The choice of c has been made on the basis of a trial and error approach. It has been noted that for c = 10 most pixels are assigned to clusters well representative of tissues, whereas for c > 10 meaningless clusters are obtained. The results of the segmentation procedure described above are displayed in figure 2.1 in which the two original MR images are shown together with the clustered image.

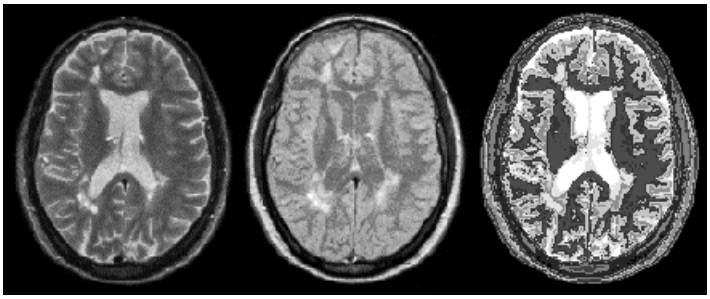

Fig. 2.1 from left to right: a T2-weighted MR slice, the corresponding PD image and the segmentation results using the FCM algorithm.

3. Re-clustering and tissue labeling

The application of FCM gives a correct distinction between white matter, gray matter, CSF, brain cortex and skull tissues, but it results in an over-segmentation in which a single tissue is split in more than one cluster.

The re-clustering procedure is based on the fact that a single tissue is split into clusters that are adjacent in the feature-space. Our approach is inspired to the work presented in [1,2], where the main goal was the detection of abnormal shapes with respect to some templates for qualitative tissue models derived from statistical analysis of the data images. The system described in [1,2] makes use of a hierarchical classification frame where, at first, air and skull tissues are located and discarded; then white matter is recognized and searched in, for abnormal shapes. If white matter shape matches to its template, then the process is iterated on CSF and, if even CSF is normal, on gray matter.

The most evident difference between our approach and the work in [1,2] is that the latter is aimed to the search of very large lesions, like those arising in tumor afflictions, and not to exact tissue labeling. On the contrary, our approach is aimed to provide a complete description of the anatomical structure under investigation, both for good tissues and for sick ones, and to locate subtle lesions.

We have carried out some heuristic measures involving the number and the distribution of pixels in the binary images representing single clusters as they result from the FCM algorithm. Moreover some measures have been performed on the mutual position of the clusters in the feature space.

The main steps of the re-clustering algorithm are the following:

1. Discarding air and skull tissues;

2. Locating and labeling white matter;
3. Locating and labeling other tissues: at this step sclerosis objects are merged with brain cortex.

In order to perform separation between extraneous tissues and the brain we can use anatomical knowledge. In fact all tissues that are to be discarded, form a band around the brain and there's a gap between them and the brain itself. First, we discard the first two clusters in T2 space (0 and 1) without further investigation because of the experimental evidence that they belong to air and skull tissues in all the processed images. A binary image is built setting to 1 all the bits corresponding to cluster from 2 to 9 and the center of the resulting area is computed as first order moment. We use a contour points search procedure (CPS) that is the same as in [1,2] in order to determine the extreme points of the vertical and horizontal axes of the brain. Let l_1 and l_2 be their respective lengths.
We know that the regions of no interest for our analysis are placed completely outside of the brain area and that they encompass it: so we use a rough procedure to analyze clusters 2, 3 and 4 and to discard those ones that contain only air and skull tissues. A rectangle is drawn on the image, which has the same center as the brain region, with dimension V x H where $V = 0.6 \cdot l_2$ and $H = 0.7 \cdot l_1$. We found experimentally that all the clusters representing brain tissues have more than 600 points inside the rectangle area even if they contain spurious points, while the clusters of no interest have no more than 250 points in it. So we use a threshold of 400 to discriminate between clusters, discarding those ones with less than 400 points in the rectangle area. This procedure creates some holes in the binary image of the brain: so we locate and eliminate them using an algorithm to search simply connected regions made up by background pixels.
A more fine procedure is used at this point to merge over-segmented clusters. First, the CPS procedure is applied to approximate with a polygon the brain. Then all clusters are analyzed to compute how many points are inside the polygon. Again we use a threshold of 400 points. All those clusters that have less than 400 points inside the polygon are merged with the next cluster in the feature space, while their external points are discarded. This procedure isn't applied to clusters 8 and 9 because they are always representatives of CSF and brain cortex mixed with sclerosis objects: the anatomical structure of these tissues doesn't allow obtaining more than 400 pixels in the slice image. Finally, we obtain the vector of adjacent clusters in the feature space that is used to perform tissue labeling.
The algorithm for locating and labeling white matter starts from the experimental evidence that all points of the first cluster in the vector belong to white matter: sometimes white matter is split in the first two clusters. We'll refer to the first cluster in the vector as the *known cluster*, while the second will be the *unknown cluster*: now, the problem is to determine if the unknown cluster is representative of white or gray matter.
To accomplish our task we implemented a simple decision algorithm that relies on the computation of four numerical indices derived from topological measures on the points of the first two classes in the vector:

- boundary density;
- compactness;
- width variation;
- slope variation.

Morphological operators have been heavily used on the binary images representing the various clusters in order to derive these indices. Here follows a brief outline of their definition: implementation details are explained in [9].
The boundary density d_b indicates how much points of the unknown cluster are lying near the brain contour: in fact, gray matter is mainly placed in this region, while white matter is not. The boundary density is obtained as:

$$d_b = \frac{u_b}{b} \cdot 100 \tag{3.1}$$

Here b is the total number of points in a 7 pixels-wide boundary region of the brain, not including discarded tissues, while u_b is the number of pixels belonging to the unknown cluster in the same region.
Compactness is used in order to discriminate between white and gray matter because of the experimental evidence that gray matter forms more sparse clusters than white matter. Compactness is defined as:

$$c = \frac{N_f}{N} \tag{3.2}$$

Here N_f is the pixels number after a suitable morphological filtering [9] of the cluster image, while N is the original number. In order to obtain a compactness measure relative to the known cluster, we derive

a constant value k for each image, so that the normalized value c_{kn} of the known cluster compactness is always set to a fixed value. For our convenience we have chosen:

$$c_{kn} = 100 \cdot k \cdot c_k = 16$$
$$c_{un} = 100 \cdot k \cdot c_u \tag{3.3}$$

If white matter is split in two clusters, it can be noted that the unknown cluster is mainly placed around the known one and results "larger": this is particularly true in the median region of the brain.
Following this idea, we have derived experimentally a numerical index called width variation Δw. A symmetric horizontal strip is considered in the image whose width is set to 30% of the length of the vertical brain axis. In this strip we find the couple of points with the maximum left and right distance from the vertical axis of the brain: we refer to the horizontal displacement w between them as the strip's width. Finally we compute:

$$\Delta w = w_u - w_k \tag{3.4}$$

The last numerical index used to discriminate among white and gray matter is the slope variation of the segments connecting centers of clusters in the T2-PD space (see fig. 3.1).

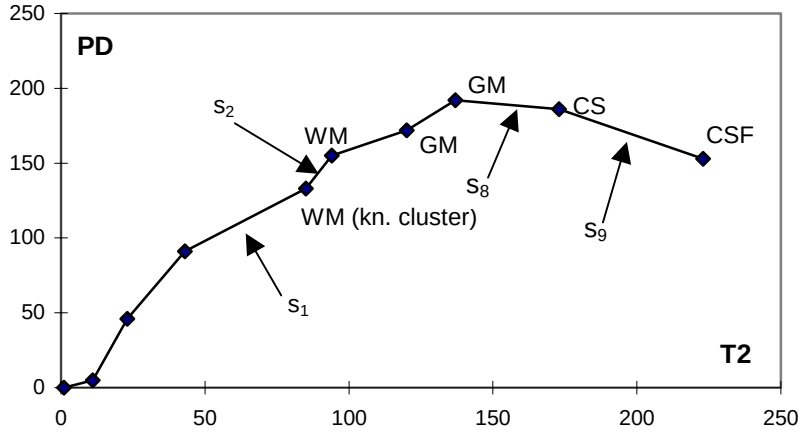

Fig. 3.1 Clusters distribution in the feature space; slope variation is computed between the two segments connected to the known cluster. In this example, white matter is split in two clusters, thus obtaining a positive slope variation.

In particular, if we consider the two segments connected to the center of the known cluster, slope variation can be computed as:

$$\Delta s = s_2 - s_1 \tag{3.5}$$

Slope analysis allows relating brightness information for each tissue both in the T2 and the PD image. From the direct analysis of MR images, experimental evidence arises that PD images exhibit strong brightness variation between white matter and gray one: the former results darker while the latter is brighter. On the contrary, white and gray matter are almost the same in T2 images, even if gray matter is a little brighter.
On the basis of the values of the numeric indices described above the following decision algorithm has been set up in order to merge the over-segmented clusters and to label the different tissues in the MR slices.

Label the first cluster in the vector of adjacent clusters (known cluster) as WM

IF *((d_b > 26)* ***OR*** *(c_{un} < 10)* ***OR*** *(Δs > 15))* ***THEN***
Label the unknown cluster as GM
ELSE
IF *((d_b < 10)* ***OR*** *(Δw > 17)* ***OR*** *(Δs < 0))* ***THEN***
Label unknown cluster as WM
ELSE
Split the unknown cluster into WM and GM

(α)

In order to label other tissues we make use of our experimental knowledge about the results of the FCM algorithm applied to all the images in the database. It is to be noted that CSF is always enclosed in a single cluster, while sclerosis objects are grouped with the brain cortex or the CSF itself. Moreover, CSF and brain cortex exhibit the higher brightness both in PD and T2 images, so they are always located in the last two clusters obtained from FCM (cluster 8 and 9). In many cases, cluster 8 corresponds to brain cortex mixed with sclerosis (CS) while cluster 9 is representative of CSF, but they can result occasionally swapped.

Tissue labeling is accomplished through a decision algorithm similar to algorithm α that makes again use of numerical indices. Particularly, we have used slopes s_8 and s_9 of the segments connected to cluster 8 and the difference of the PD signal computed in each cluster center with respect to the average PD value of white matter:

$$\begin{aligned} \Delta PD_8 &= PD_8 - \overline{PD}_{WM} \\ \Delta PD_9 &= PD_9 - \overline{PD}_{WM} \end{aligned} \tag{3.6}$$

The use of white matter is due to the observation that CSF has a PD signal close to white matter itself, while brain cortex signal is much higher. In this way it is always possible to determine the CSF as the cluster with the lower ΔPD value.

```
IF (ΔPD_9 ≤ ΔPD_8) OR (s_9 ≤ 0) THEN        /* CSF is located in cluster 9 */
 {
    IF (-0.8 ≤ s_9 ≤ 0) THEN                /* CSF contains CS points */
       Label points in cluster 9 with ΔPD >60 as CS
       Label points in cluster 9 with ΔPD ≤ 60 as CSF
    ELSE
       Label cluster 9 as CSF

    Label cluster 8 as CS
 }
ELSE                                        /* CSF is located in cluster 8 */
 {
    IF ( 0 ≤ s_9 ≤ 1) THEN                  /* CSF contains CS points */
       Label points in cluster 8 with ΔPD >60 as CS
       Label points in cluster 8 with ΔPD ≤ 60 as CSF
    ELSE
       Label cluster 8 as CSF

    Label cluster 9 as CS
 }

Label all the remaining clusters that are not WM as GM
                                                              ( β )
```

The application of α and β algorithms allows us to reduce the original 10 clusters to only 6, each representative of a particular tissue: the only exception is the CS class where are grouped brain cortex and lesions. Figure 3.2 shows the result of FCM segmentation and re-clustering on a MR brain slice.

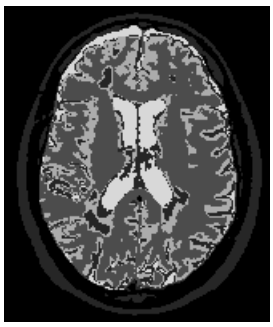

Fig. 3.2 Slice of image 2.1 re-clustered: the original ten classes are reduced to only six; here are depicted even discarded tissues.

4. Location of possible lesions

The algorithm to locate sclerosis objects is driven by clinical knowledge about this disease and experimental observations on the lesions' signal intensity in both PD and T2 images.
In detail, the whole procedure is based upon the following statements:

1. lesions have an almost rounded shape;
2. multiple sclerosis disease afflicts only white matter;
3. only some parts of the brain are affected by multiple sclerosis, particularly all the border areas between white matter and CSF and between white and gray matter;
4. lesions' signal is, in general, higher than other tissues both in PD and T2 images, but subtle lesions can be confused with gray matter.

The use of suitable thresholds on signal intensity doesn't suffice to the correct detection of all the sclerosis objects, while a deeper analysis of the anatomical structure of the lesions provides more satisfactory results. Using the first three above statements, we focused our search on the clusters that are representative of white and gray matter and of CSF. In particular, for each re-clustered image, we built three binary images:

- the image of the WM pixels (WM image);
- the image of the WM and GM pixels (WM+GM image);
- the image of WM, GM and CSF pixels (WM+GM+CSF image).

In these three images we searched for the "holes", or better, all the simply connected background regions (see fig. 4.1).

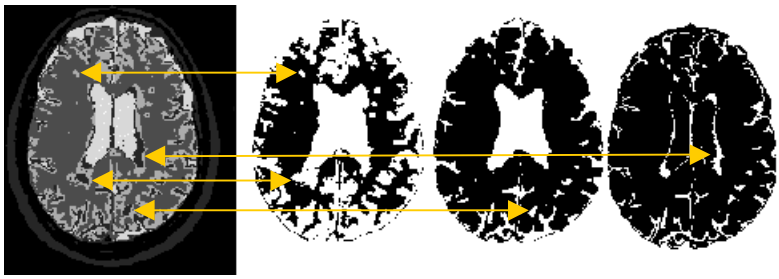

Fig. 4.1 from left to right: the re-clustered image, WM image, WM+GM image and WM+GM+CSF image. Arrows indicate the position of some lesions and the corresponding holes.

In all the observed slices, lesions are located in regions that correspond to holes in the images previously described. All the three kinds of images have been processed in order to classify all the potential sclerosis objects as small, medium or large. Five different labels have been derived for these objects depending on their size and their location between or inside the WM, GM and CSF clusters. For all the candidate objects a normalized value of surface S_N is computed. First, the WM binary image is processed with a simple 3x3-median filter in order to eliminate very small holes that are irrelevant for our analysis. All candidates are found using an algorithm for searching simply connected background regions and a new binary image is built where all the pixels belonging to the holes are set to 1, while the rest is set to 0. All the objects whose S_N value is under a suitable threshold are collected together and labeled as S_w if they don't contain pixels belonging to the CS class and S_{wc} in the opposite case. We have already noticed that, according to our experimental observations, most of sclerosis objects are in the CS class. Thus S_{wc} objects are more likely to be actual lesions than S_w ones and we'll weight them differently in the classification step.
All the remaining objects inside WM and made up by CS or CSF pixels are collected in a new binary image and labeled as B_w because they are potential big lesions. These object are smaller than white matter holes because of the presence of some gray matter pixels around them: these pixels can be discarded because they are never part of a lesion. Again, we pick CSF and CS points because of the experimental evidence that they contain sclerosis objects.
As regards holes in the WM+GM images the procedure is slightly different. In some cases, sclerosis objects result linked with pixels belonging to CSF and brain cortex because of the fact that the lesion is adjacent to these tissues. In this case, the regularization phase is accomplished using a 3x3-dilation followed by a 3x3-erosion filter in order to divide abnormal holes. As in the previous case, only the objects inside the holes and made by CS or CSF pixels are selected as the candidate sclerosis lesions and are labeled as B_{wg} objects.

Finally the WM+GM+CSF images are analyzed because some lesions are placed near the ventricular system: so they appear connected to the CSF points even if they don't belong to CSF class. The image is regularized using a 3x3 median filter, then all the holes are collected together and labeled as B_{wgc}. At this point all the candidate sclerosis objects have been selected from the re-clustered image and it is possible to go on with the classification step in order to determine which objects are actual lesions and which ones can be discarded.

5. Neural classification

The output of the potential sclerosis localization subsystem is a first labeling of the candidate objects on the basis of their size and cluster membership.
In the classification step, we use medical knowledge about the sclerosis affliction, together with some heuristics derived from our experimental investigations to obtain a set of measurable features characterizing each object that can be processed by a neural architecture.
All the information can be expressed in terms of collection of position, shape and signal intensity measures for each candidate object. The 16 measures we've used for classification are grouped in four categories:

- surface measure and type label deriving from the locating step;
- differences between object mean and maximum brightness and WM or GM mean brightness value in both T2 and PD images;
- shape measures: regularity coefficient and axes ratio;
- position measures: distances between the object center and the contour, the median axis or the center of brain.

In what follows these measures are described in detail. The surface measure is taken as the value of S_N, and it's re-scaled in [0,1]. Objects' labels have been put in direct correspondence to suitable numerical values in the range [0,1] with a 0.2 step. These values reflect the heuristic knowledge about lesions: subtle lesions are more likely to contain points of the CS class ($S_{wc} > S_w$). Large objects at the borders between white and gray matter or between white matter and CSF are very likely to be lesions (highest values for B_{wg} and B_{wgc}). Typically, large objects inside white matter (B_w label) are not sclerosis lesions: in fact, the largest one is the region corresponding to the entire CSF.
For signal intensity features, we have derived eight different measures. At a first glance, the mean value of one object's signal intensity in both PD and T2 images could seem be sufficient to characterize the object. Therefore, one might think that objects of the same class (that is with the same label) should appear similar in different slices. On the contrary, we have experienced that the mean value of signal intensity can be quite different from one image to another, depending on the nature and settings of the digitizing equipment. Images can be obtained from the direct digital output of MR equipment, or from the digital scanning of MR film. Even in the case of direct digital output, there may be little differences in the intensity of the magnetic field from one MR device to another or from one test to another in the same device.
On the other hand, differences in signal intensity between lesions and the other tissues are more stable even if we compare MR slices acquired in distinct ways.
The previous considerations led us to take into consideration the following values:

$$\begin{aligned}
\Delta PD_{OW,m} &= PD_{O,m} - PD_{W,m} \\
\Delta PD_{OW,M} &= PD_{O,M} - PD_{W,m} \\
\Delta T2_{OW,m} &= T2_{O,m} - T2_{W,m} \\
\Delta T2_{OW,M} &= T2_{O,M} - T2_{W,m} \\
\Delta PD_{OG,m} &= PD_{O,m} - PD_{G,m} \\
\Delta PD_{OG,M} &= PD_{O,M} - PD_{G,m} \\
\Delta T2_{OG,m} &= T2_{O,m} - T2_{G,m} \\
\Delta T2_{OG,M} &= T2_{O,M} - T2_{G,m}
\end{aligned} \qquad (5.1)$$

Here, the subscript *m* identifies the mean value, while the subscript *M* identifies the maximum value. For each object, differences are computed between its mean or maximum value in a PD or T2 image and the mean value of white or gray matter in the same image. These values can vary in the range [-255,255].
The computation of shape features starts from the observation that an almost rounded shape characterizes multiple sclerosis lesions: they're regular and not too elongated. We derived two coefficients, the regularity measure *R* and the axes ratio *A* defined as follows:

$$R = \frac{S_O}{S_R} \cdot 100$$
$$A = \frac{L_a}{L_A} \cdot 100 \qquad (5.2)$$

Here S_O is the object surface in pixels, while S_R is the surface of the smallest rectangle enclosing the object itself. In the second formula L_a is the length of the minor axis of the object, while L_A is the length of the major one. Higher R and A values indicate that the object shape is regular and not elongated: they range in [0,100].

The last set of four measures is related to the object's position inside the brain:

- distance from the center of the brain d_c;
- distance from the major axis of the brain d_a;
- distance from the boundary of the brain area d_b;
- distance from the ventricular area or from the line separating brain hemispheres d_v.

In order to obtain the values for d_c and d_a, we consider the binary image made up by all the brain tissues of interest that is white and gray matter, brain cortex, CSF and all candidate lesions. We compute the position of the center of mass of this image and take this point as the geometrical center of the brain, while the major axis will be the vertical line passing through it. The two Euclidean distances are computed in pixels from the center of mass of each candidate object.

The computation of d_b and d_v provides us with information about candidate objects position inside the brain, with respect to the borders with CSF: we don't take into account CSF because the sclerosis never affects it. d_b gives us the distance to the border with the external CSF region, while d_v is a measure of the distance to the boundary of the ventricular area which is the most inner region containing CSF.

The whole set of measures derived so far constitutes a feature vector that fully describes a candidate sclerosis object to the extent of classification.

The classifier has been set up as a multi-layer perceptron trained with the back-propagation algorithm. The choice for this architecture follows from various considerations.

First, we have to represent the medical knowledge on the sclerosis affection. We have a series of potential lesions together with a doctor assessment on each of them. It's not so easy to derive a lesion prototype because lesions can be small or very big, close to brain borders or nearer to CSF, some lesions have strong brightness in both T2 and PD images, while others are more vivid only in T2 images. Only a doctor can actually evaluate them. So, we're facing a learning-by-example problem with a weak object prototyping.

Besides, we have made up a first trial classifier based on the maximization of a parametric cost function:

$$\mathrm{C}(\boldsymbol{x}) = \sum_i w_i \sum_j w_{ij}\, \mathrm{f}_j(x_j) \qquad (5.3)$$

Here, $\boldsymbol{x}$ is the feature vector, the terms $\mathrm{f}_j()$ represent some suitable input processing functions, while w_i and w_{ij} are tunable weights.

For area, label and brightness features we've adopted an input function of the form:

$$\mathrm{h}(x_j) = k\frac{x_j}{m_j} \qquad (5.4)$$

Here m_i is the mean value of x_i over the whole training set, while k is a suitable constant. For the distance features we adopted the Gaussian function whose mean value and standard deviation were computed for each x_i over the entire training set. There is no theoretical foundation for this choice, except that we have found a relevant improvement in classification performance with respect to simply summing up the inputs.

Experiments have been carried out on a set of 70 MR slices from eight patients. Using a training database extracted from 50 image couples and made up by 4557 objects with 212 sclerosis lesions (as they have been assessed by a doctor) we have obtained a satisfactory classification score (table 1). In this former set-up, the weights were updated through a pseudo exhaustive search in order to maximize $C(\boldsymbol{x})$ to the value 100 for the sclerosis objects while forcing it to 0 for non-sclerosis objects. An object was recognized as lesion if it scored more than 60.

The previous reasons led us to don't take into account some classical connectionist classifiers such as LVQ or SOM: the former because of prototyping weakness, and the latter because it performs an unsupervised learning. Moreover the mathematical form of (5.3) is very close to the multi-layered network architecture with linear activation functions in hidden and output units.

We tried several architectures with both one and two hidden layers. All the networks had sixteen input units and one output unit, and were fully connected. All units used the logistic activation function scaled in the range [0, 1]; activation near 1 stands for a lesion classification, while activation close to 0 indicates a non-lesion object. The threshold for classification was set experimentally at 0.25, the learning rate was fixed to η = 0.1 and the networks were trained for 21000 iterations with the same training set of the algorithmic classifier mentioned above. Tests were carried out using 20 image couples not presented during the learning phase. The test set contains 1902 candidates and in particular 84 sclerosis lesions and 1818 non-sclerosis objects.
In figure 5.1 are depicted the learning curves for the various architectures. The first trial was a network with only one hidden layer containing four neurons in order to follow the architecture of the algorithmic classifier. As it can be noted from figures this network resulted in a very poor performance with respect to the two hidden layers ones. For the other architectures we chose a first hidden layer made up by sixteen units, while varying the number of the second hidden layer. While the networks with the higher number of units in the second layer have a very good performance in learning the training set, they exhibit an over-fitting of data in the test phase. On the contrary, the better trade off between learning and testing was obtained using the network with only four hidden units in the second hidden layer.
Table 1 reports the detail of training and test errors of the two classifiers, respectively for sclerosis objects and non-sclerosis objects while figure 5.2 shows some experimental classification results.

	Train Score		**Test Score**	
	Quantity	Percentage	Quantity	Percentage
Sclerosis	194/212	91.5%	76/84	90.4%
Non-Sclerosis	4338/4345	99.8%	1813/1818	99.7%
Total	4532/4557	99.5%	1889/1902	99.3%

	Train Score		**Test Score**	
	Quantity	Percentage	Quantity	Percentage
Sclerosis	212/212	100%	79/84	94%
Non-Sclerosis	4345/4345	100%	1810/1818	99.6%
Total	4557/4557	100%	1889/1902	99.3%

Table 1 Classification score for the algorithmic (top) and neural (bottom) classifier.

Early results are satisfactory: the percentage classification error is strongly enhanced with respect to the non-neural classifier in both the training and test phase. Besides, the network has performed a correct generalization. The lower percentage score for the sclerosis objects classification is due to the difficult detection of the smallest lesions that are hard to discern from the rest of the image.

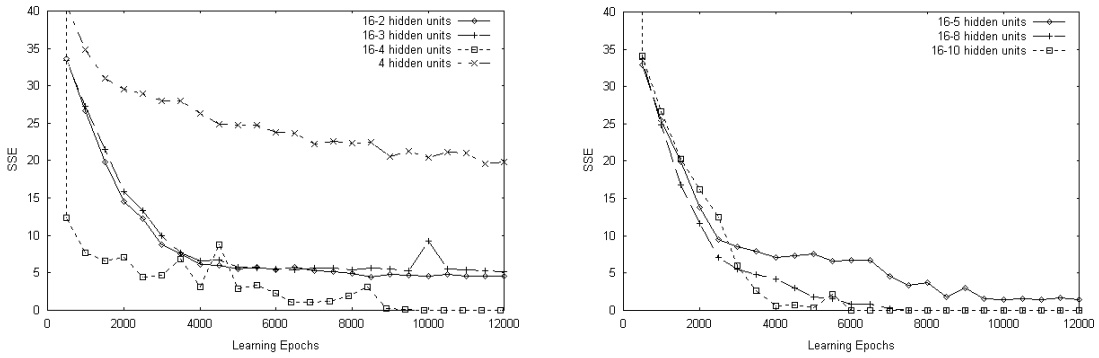

Fig. 5.1 Learning curves for the various neural networks we have tested.

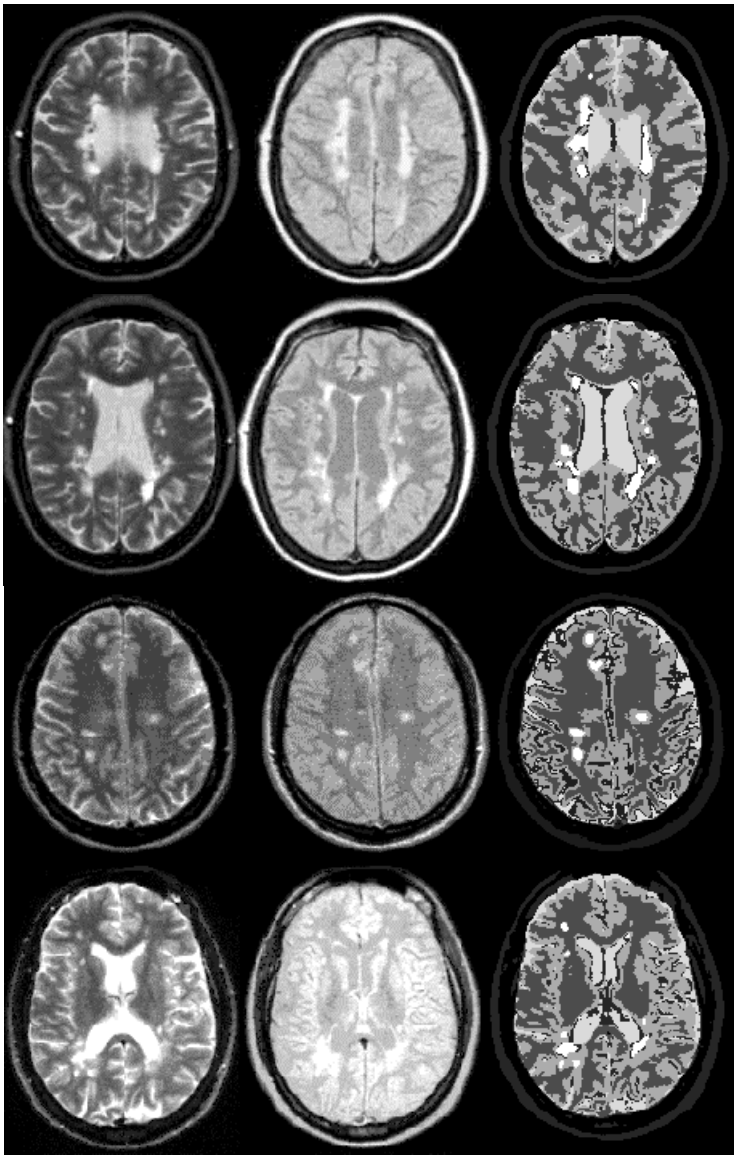

Fig. 5.2 Experimental results: in all the four series the leftmost image is the T2 image, next there is the PD image and the segmented one with the sclerosis objects outlined as the brighter blobs.

6. Conclusions

We have presented a neural classification system able to locate multiple sclerosis lesions in MR slices of the brain and to derive a numerical estimate of their position and shape.

The segmentation and classification algorithms presented in the paper have been specifically tailored to the kind of images under consideration, that is PD and T2-weighted MR slices of the brain. Moreover they are strongly grounded on the clinical knowledge about the problem faced with, that is the multiple sclerosis disease. For example, the knowledge about the possible localization of lesions within brain tissues has guided both the segmentation and the classification step.

However, the methodology we've developed is applicable to all classes of syndromes that need investigation by means of medical imaging tests. In fact we can say that all clinical knowledge to assess this kind of afflictions in a medical image, can be encoded in a set of measures regarding shape, position, brightness and so on. These form a suitable basis for a connectionist classification.

The paradigm preprocessing-segmentation-classification can be enriched by a 3D-reconstruction step able to provide a more detailed information to the clinician involved in the diagnosis process. From this point of view, we are now extending the present work to the treatment of the entire brain volume.

Neural networks ensure an on-line performance not comparable with classical AI systems, which have to derive inferences on a set of assertions regarding the pathology under investigation and need complex formalisms to represent medical knowledge.

The design of a family of specialized tools, each aimed to evaluate data in order to discover a particular disease, results quite simple and allows setting up multi-module concurrent systems providing the clinician with several responses each with its different degree of confidence. The doctor can then assess the more likely outputs in order to formulate the diagnosis.
Finally, the modular approach is to be preferred to the use of expert systems because the simplicity of the knowledge representation paradigm allows the acquisition of information about new afflictions without interfering with the existing one. On the other hand, to add new knowledge about an illness we have to re-train the neural set-up.

Acknowledgements
This work has been partially supported by Italian Government through "Programma Biotecnologie L.95/95 (MURST 5%)" and "Programma Galileo '99".

References

[1] Clark, M.C., Hall, L.O., Goldof, D.B., Clarke, L.P., Velthuizen, R.P., Silbiger, M.S., *MRI segmentation using fuzzy clustering techniques*, IEEE Engineering in Medicine and Biology, pp. 730/742, November 1994.
[2] Li, C., Goldof, D.B., Hall, L.O., *Knowledge-based classification and tissue labeling of MR images of human brain*, IEEE Transactions on Medical Imaging, vol. 12, no. 4, pp. 740/750, December 1993.
[3] Liang, Z., *Tissue classification and segmentation of MR images*, IEEE Engineering in Medicine and Biology, pp. 81/85, March 1993.
[4] Bezdek, J.C., Hall, L.O., Clarke, L.P., *Review of MR image segmentation techniques using pattern recognition*, Medical Physics, 20(4), pp. 1033/1048, 1993
[5] Hall, L.O., Bensaid, A.M., Clarke, L.P., Velthuizen, R.P., Silbiger, M.S., Bezdek, J.C., *A comparison of neural network and fuzzy clustering techniques in segmenting magnetic resonance images of the brain*, IEEE Transactions on Neural Networks, vol. 3, no. 5, pp. 672/682, September 1992.
[6] Bezdek, J.C., *Fuzzy models for pattern recognition*, pp. 1/33, 1992.
[7] Bobrowski, L., Bezdek, J.C., *C-means clustering with the l1 and l∞ norms*, IEEE Transactions on Systems, Man, Cybernetics, vol. 21, no. 3, pp. 545/554, May/June 1991.
[8] Cannon, R.L., Dave, J.V., Bezdek, J.C., *Efficient implementation of the fuzzy c-means clustering algoritm*, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 8, no. 2, pp. 248/255, March 1986.
[9] Alonge, F., Ardizzone, E., Pirrone, R., *Recognition of multiple sclerosis lesions in MR images*, Technical Report DIAI, University of Palermo, submitted to Artificial Intelligence in Medicine, February 1999.
[10] Johnston, B., Atkins, M.S., Mackiewich, B., Anderson, M., *Segmentation of Multiple Sclerosis Lesions in Intensity Corrected Multispectral MRI*, IEEE Transactions on Medical Imaging, vol. 15, no. 2, pp. 154/169, April 1996.
[11] Kamber, M., Shinghal, R., Collins, D.L., Francis, G.S., Evans, A.C., *Model-Based 3-D Segmentation of Multiple Sclerosis Lesions in Magnetic Resonance Brain Images*, IEEE Transactions on Medical Imaging, vol. 14, no. 3, pp. 442/453, September 1995.
[12] Krishnan, K., Atkins, M.S., *Segmentation of Multiple Sclerosis Lesions in MRI – An Image Analysis Approach*, Proceedings of the SPIE-Medical Imaging 1998, 3338, pp. 1106/1116, February 1998.
[13] Russel, S., Norvig, P., *Artificial Intelligence – A Modern Approach*, Prentice-Hall, 1995.
[14] Udupa, J.K., Wei, L., Samarasekera, S., Miki, Y., van Buchem, M.A., Grossman, R.I., *Multiple Sclerosis Lesion Quantification Using Fuzzy-Connectedness Principles*, IEEE Transactions on Medical Imaging, vol. 16, no. 5, pp. 598/609, October 1997.
[15] Zijdenbos, A.P., Dawant, B.M., Margolin, R.A., Palmer, A.C., *Morphometric Analysis of White Matter Lesions in MR Images: Method and Validation*, IEEE Transactions on Medical Imaging, vol. 13, no. 4, pp. 716/724, December 1994.