Segmentation of Multispectral MR Images Using a Hierarchical Self–Organizing Map

Suchendra M. Bhandarkar Premini Nammalwar
Department of Computer Science
The University of Georgia
Athens, Georgia 30602–7404, USA
Email: suchi@cs.uga.edu, pnamm@yahoo.com

Abstract

The application of a hierarchical self–organizing map (HSOM) to the problem of segmentation of multispectral magnetic resonance (MR) images is investigated. The HSOM is composed of several layers of the self–organizing map (SOM) organized in a pyramidal fashion. The SOM has been used for the segmentation of multispectral MR images but the results often suffer from undersegmentation and oversegmentation. By combining the concepts of self–organization and topographic mapping with multiscale image segmentation, the HSOM is seen to overcome the major drawbacks of the SOM. The segmentation results of the HSOM are compared with those of the SOM and the k–means clustering algorithm on multispectral MR images of the human brain representing both, normal conditions and pathological conditions such as multiple sclerosis. The multiscale segmentation results of the HSOM are shown to have interesting consequences from the viewpoint of clinical diagnosis of pathological conditions.

1 Introduction

Magnetic resonance imaging (MRI) was introduced in clinical medicine in 1981 and is regarded as one of the most significant advancements in medical imaging since the discovery of X-rays a century ago. During the past decade, conventional MRI has had a tremendous impact on the diagnostic imaging of most human organ systems and has made strides in several clinically useful directions. The fundamental physical principle underlying MRI is the interaction of proton spins with a strong static magnetic field and carefully designed dynamic magnetic fields that vary as precisely defined functions of both space and time [7].

Image contrast in MRI depends primarily on three intrinsic parameters: the proton density (PD) which quantifies the number of hydrogen nuclei (protons) within the sample volume, the longitudinal relaxation time (T1) which is the time taken for the spins in a sample volume to return from an excited state to a ground state (spin-lattice relaxation) and the transverse relaxation time (T2) which is the time taken for the spins in the excited state to lose phase coherence (spin–spin relaxation). An MR image is acquired by a pulse sequence which renders it sensitive or weighted with respect to one of these intrinsic parameters. Thus, for any given field of view (FOV), slice thickness and voxel resolution, a band of images can be obtained where each image corresponds to a given weighted factor or intrinsic parameter. This band of images constitutes a multispectral or multichannel MR image. The most commonly used multispectral data sets in MRI consist of T1-weighted, T2-weighted and PD-weighted images. Each pixel value in such a multispectral MR image is denoted by a triple (a, b, c) which denotes the intensity in each of the T1-, T2- and PD-weighted images.

MR image segmentation has significant ramifications in clinical medicine since it enables tissue classification which is critical to the study of the morphology of an anatomical structure [3, 5]. Image segmentation is the process of dividing an image I into regions, S_i, $i = 1, 2, ..., m$, such that for given a logical predicate P that represents a homogeneity criterion, the following four conditions are met: (1) $I = \bigcup_{i=1}^{m} S_i$; (2) $S_i \cap S_j = \phi, \forall i \neq j$, where $1 \leq i, j \leq m$; (3) $P(S_i) = true$, where $1 \leq i \leq m$; and

Proceedings of the Fourteenth IEEE Symposium on Computer-Based Medical Systems (CBMS'01)
1063-7125/01 $10.00 © 2001 IEEE

Figure 1: The network structure of the SOM

(4) $P(S_i \cup S_j) = \mathit{false}$ for all adjacent regions S_i and S_j, where $i \neq j$ and $1 \leq i, j \leq m$. Ideally, each region in a segmented image should represent a semantically meaningful entity in the input image. The goal of segmentation, is to present the relevant information in an image in a manner such that it can be easily accessed and interpreted.

A major goal of MR image segmentation is tissue classification where a distinction is made between normal and pathologic tissues. For example, in cerebral MR images some of the tissue classes of interest are white matter (WM), gray matter (GM), cerebrospinal fluid (CSF) and abnormal (pathologic) tissues [5]. Multispectral segmentation serves to differentiate between tissue classes having similar characteristics in a single imaging modality such as edema and necrotic (scar) tissue and is of great significance in the evaluation of response to therapy [6]. Multispectral segmentation techniques have been used successfully in the case of MR images of the human brain, especially in the study of neurodegenerative diseases such as Alzheimer's and multiple sclerosis (MS) and in the detection of brain tumors [1, 6]. On account of their ready availability, the MR images used in this work are those of the human brain. The complex structure of the human brain makes it an interesting and challenging subject for MR image acquisition and analysis [1, 7].

2 The Self–Organizing Map (SOM)

The Self–Organizing Map (SOM), originally developed by Kohonen [4], is a spatially interacting artificial neural network (ANN). The SOM embodies a nonparametric regression whereby a number of ordered discrete reference (or weight) vectors are fitted to the distribution of the input vectors. A weight vector is associated with each connection between an input vector and a neural unit (Figure 1). Each neural unit in the SOM forms lateral (excitatory or inhibitory) connections with its local neighbors. Each neural unit in the 2-D SOM array receives the same input information. The learning procedure in the SOM proceeds in two phases, the competitive phase and the cooperative phase. In the competitive phase, the distances between each input vector and all the weight vectors are determined using a distance measure, the weight vector which is at the least distance from the input vector is determined and the corresponding neural unit is designated a winner. In the cooperative phase, the winner and its neighboring neural units are updated so as to increase their sensitivities to the input vector. The SOM iterates through these two learning phases until all the input vectors are exhausted.

The most important property of the SOM is that it represents a *topology-preserving* or *topographic* mapping which preserves topological relations in the input space while simultaneously performing a dimensionality reduction (projection) of the input space onto the 2-D mesh of neural units in the competitive layer. The SOM models a vector quantization procedure which has proved to be very effective for image segmentation. Vector quantization is a process of partitioning an n-dimensional vector space into M regions so as to optimize a criterion function when all the points in each region are approximated (or represented) by the representation vector associated with that region. The goal is to find an appropriate partition and an appropriate set of representation vectors that would result in an optimum of the criterion function. The vector quantization procedure is one of unsupervised clustering where the number of clusters are known a priori and the desired clusters are

Proceedings of the Fourteenth IEEE Symposium on Computer-Based Medical Systems (CBMS'01)
1063-7125/01 $10.00 © 2001 **IEEE**

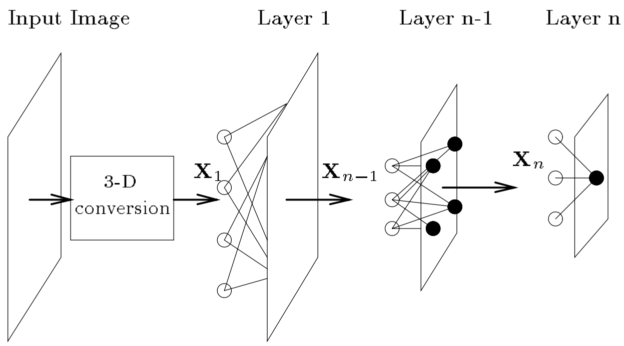

Figure 2: The structure of the HSOM

those which optimize the criterion function. However, the SOM has certain fundamental limitations in the context of image segmentation. The number of neural units in the SOM must correspond to the number of segmentable regions in the input image. This initial determination of the number of neural units in the SOM is very subjective. A gross mismatch between the number of regions into which the image can be partitioned and the number of neural units in the SOM will result in an undersegmented or oversegmented image [2].

3 The Hierarchical Self–Organizing Map (HSOM)

The hierarchical SOM (HSOM) seeks to remedy the shortcomings of the SOM. The HSOM (Figure 2) combines the concepts of self–organization and topographic mapping with multiscale segmentation. A multiscale image representation is one where different features in the image are explicitly represented at different scales. Using a single scale for an entire image may not be appropriate since the scale at which to represent or analyze a particular feature may depend strongly on the presence of other features of different sizes in its spatial proximity. The HSOM achieves multiscale segmentation of an input image i.e., the segmented result may be viewed at different levels of abstraction of the feature space. This would enable an anatomical entity to be potentially represented at a level of abstraction that is best suited for it. In the context of MRI of the brain, this means that gray matter will be represented at one level of abstraction, white matter at another and so forth which would make identification of potential pathological situations easier. This property potentially renders the HSOM more useful than the SOM from a clinical viewpoint.

The HSOM, as the name implies, is a network composed of successive SOM layers arranged in a pyramidal fashion (Figure 2). Input data arrives at the lowest layer and information flows to the higher layers in a strictly feed-forward manner. The number of neural units in a SOM layer decreases as one proceeds from the base to the apex of the HSOM. Since, the number of representation vectors generated in a given layer is typically proportional to the number of neural units in that layer, in a higher layer of the HSOM which contains fewer neural units, each weight vector represents a larger cluster in input vector space i.e., a higher level of abstraction of the input image.

Each SOM layer within the HSOM implements competitive learning. The input layer receives input from the external world and propagates the input to all the neural units in the first layer. On completion of competitive learning in the first layer, the resulting weight vectors are converted and propagated to the next layer as input to that layer. The above process is repeated until the top layer in the HSOM is reached. The input vector to the HSOM is a 5–tuple of the form $(T1(x,y), T2(x,y), PD(x,y), x, y)$ where (x,y) denotes the coordinates of the pixel in the image and $T1(x,y)$, $T2(x,y)$ and $PD(x,y)$ denote the intensity values of the pixel in the T1-weighted, T2-weighted and PD-weighted images, respectively. The coordinate values (x,y) are intended to enforce spatial connectivity of the resulting clusters whereas $T1(x,y)$, $T2(x,y)$ and $PD(x,y)$ ensure spectral homogeneity of the clusters in each spectral band. The learning procedure in each HSOM layer consists of the same two phases as the learning procedure in the SOM. The competitive phase of the learning process determines the winning neural unit for weight adjustment. The cooperative phase adjusts the weight vectors. The distance between the input vector $(T1(x,y), T2(x,y), PD(x,y), x, y)$

Proceedings of the Fourteenth IEEE Symposium on Computer-Based Medical Systems (CBMS'01)
1063-7125/01 $10.00 © 2001 IEEE

Figure 3: (a) A 3-level abstraction tree, (b) A tree to corresponding the segmented image

and a weight vector $\mathbf{w} = (w_{T1}, w_{T2}, w_{PD}, w_x, w_y)$ is computed using a weighted Euclidean distance measure [2]. The learning parameters such as learning rate, the initial and final lateral width of the neighborhood and the convergence criterion were empirically determined. The same set of parameters were used for all the MR images in the experiments.

4 The Abstraction Tree

The winning neural units (and their corresponding weight vectors) from each layer result in a hierarchical structure termed as an *abstraction tree* (Figure 3(a)). Note that each node in the abstraction tree represents a region in the image at a specified level of abstraction. A segmented image is generated on demand by traversing the abstraction tree in a breadth-first manner starting from the root node until some stopping criterion is met. A node in the abstraction tree is expanded if the sum of the variances of the feature attribute values of its pixels is larger than a prespecified threshold value. Otherwise, the node is labeled as a closed node and none of its descendants are visited. Regions corresponding to the *closed* nodes constitute a segmented image. Note that the resulting segmented image usually contains regions from different abstraction levels (Figure 3(b)). Also note that the abstraction tree is not limited to the traversal criterion mentioned in this paper. One could very well use a traversal criterion based on statistical hypothesis testing and/or local surface fitting.

One of the drawbacks of the SOM is that the resulting image segmentation is sensitive to the neighborhood or interconnection topology of the SOM array of neural units. Conventionally, the 4-, 6- or 8-neighborhood topology is used in the SOM but the final image segmentation has been noted to vary significantly with the choice of neighborhood topology [2]. The segmentation results of the HSOM, on the other hand, are seen to be less sensitive to the choice of neighborhood topology in each of its SOM arrays. The reason for this is that, in the case of the HSOM, some of the neighborhood information is encoded by the abstraction tree in that the child nodes of a parent node have a higher probability of being topological neighbors. The only requirement is that the number of layers in the HSOM is adequate to segment the most complex region in the image. This requirement is satisfied if the number of neural units in the first layer of the HSOM is greater than or equal to the number of visually distinguishable regions in the image. Another advantage of the HSOM is that it also allows one to explore different resolutions of topological ordering [2].

5 Experimental Results

The HSOM was tested with MR images of the human brain. Figure 4(a)–(c) shows the T1-, T2- and PD-weighted images of a human brain. Figure 4(d) depicts the segmented image generated by

Proceedings of the Fourteenth IEEE Symposium on Computer-Based Medical Systems (CBMS'01)
1063-7125/01 $10.00 © 2001 **IEEE**

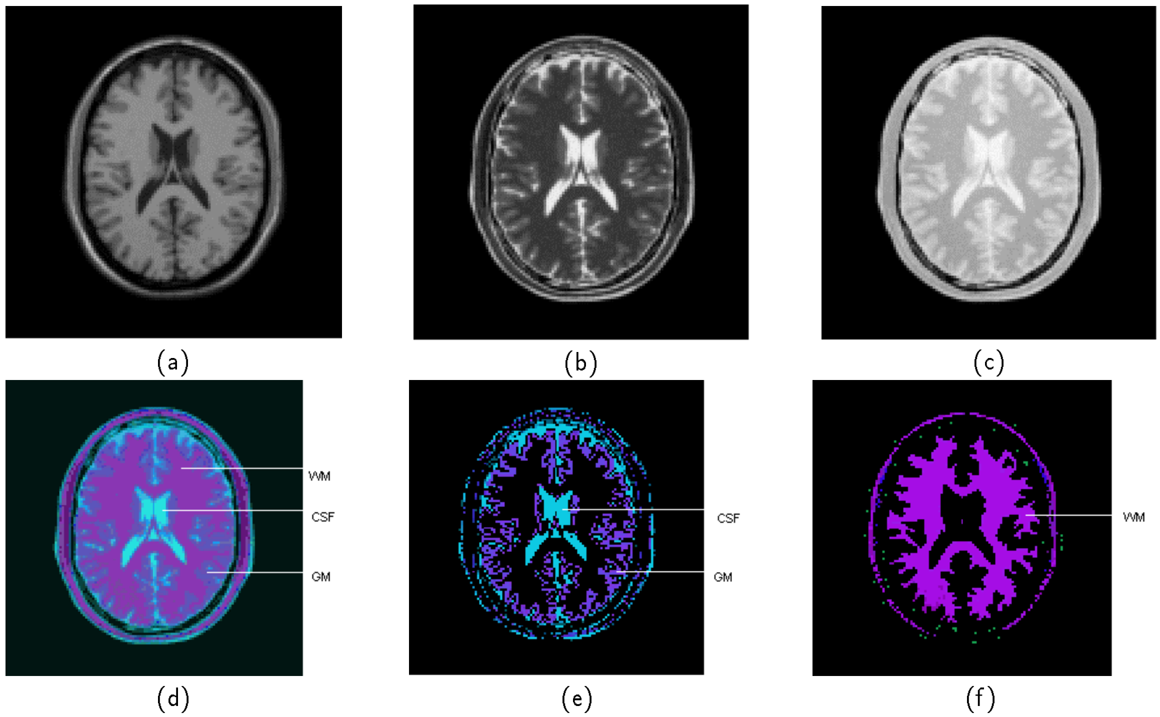

Figure 4: (a) T1-weighted image (b) T2-weighted image (c) PD-weighted image (d) segmented image generated by the HSOM (e) cerebrospinal fluid (CSF) and gray matter (GM) segmented at HSOM level 2 (32×32) (f) white matter (WM) segmented at HSOM level 4 (8×8)

the HSOM in conjunction with the traversal of the abstraction tree. The pixels in the segmented image belonging to a closed node in the abstraction tree are assigned a color in (R, G, B) space based on the (T1, T2, PD) values of the closed node. The segmented image is seen to have a much higher contrast than any of the input T1-, T2- and PD-weighted images since distinct regions of the brain such as cerebrospinal fluid (CSF), gray matter (GM) and white matter (WM) are assigned distinct colors by the HSOM. This is not obvious in the grayscale images shown in the paper but is evident in the color images viewed on the computer screen. This suggests that the segmented image generated by the HSOM can be treated as an enhanced version of the input images which could be viewed better by a radiologist. In fact, the HSOM can be considered as performing an adaptive smoothing of the input data in multi-dimensional space where the extent of smoothing is determined by the local scene content. As shown in Figure 4(e),(f), CSF and GM are segmented at level 2 in the HSOM (with 32×32 neural units) whereas WM is segmented at level 4 (with 8×8 neural units). This shows that the HSOM, unlike the conventional SOM, is able to adapt the level of resolution based on the physiology of the anatomical entity under consideration. For example, in this case, a 32×32 SOM array would have resulted in an oversegmentation of the WM and an 8×8 SOM array in an undersegmentation of the CSF and GM. The HSOM was also used to segment MR images of the human brain with pathalogical conditions such as multiple sclerosis (MS). As can be seen in Figure 5(a)–(f) the HSOM is successful in detecting and delineating the neurodegenerative lesions caused by MS. The conventional SOM and a region growing-based segmentation technique that used the k–means clustering algorithm were not able to detect these lesions.

6 Conclusions and Future Work

In this paper we presented the application of the HSOM to the segmentation of multispectral MR images of the human brain. The HSOM results in an abstraction tree which represents a multiresolution segmentation of the input multispectral MR image. Breadth-first traversal of the abstraction tree results in a more accurate segmentation when compared with conventional tech-

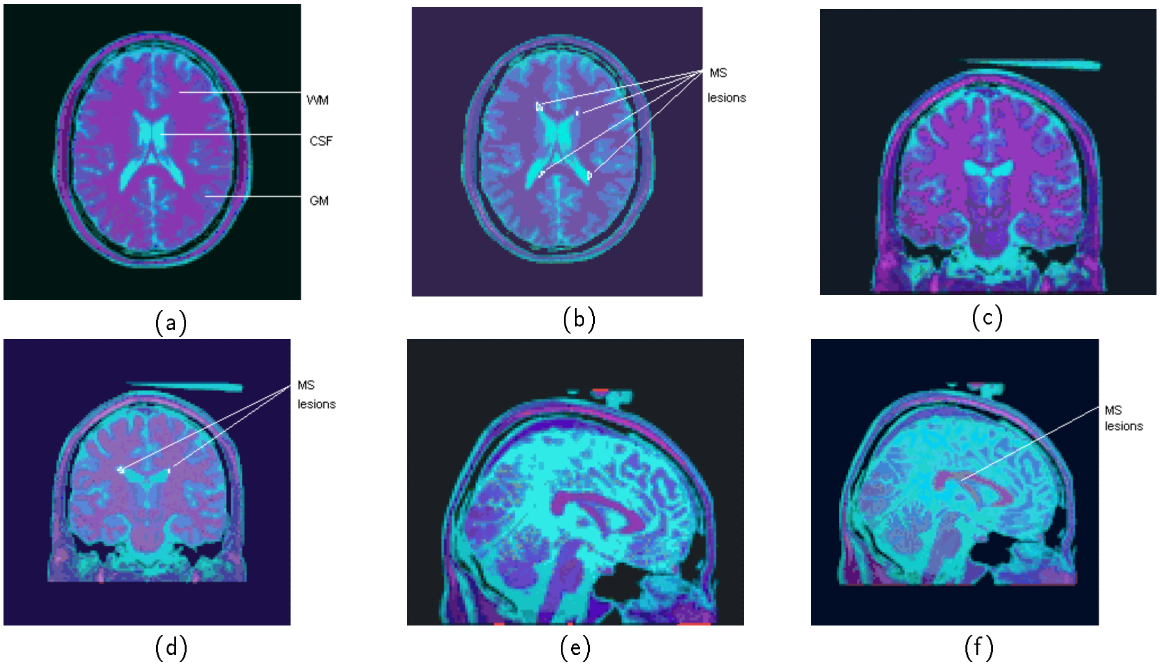

Figure 5: HSOM segmented image of (a) normal brain, (b) brain with multiple sclerosis (MS) lesions, (c) normal brain, (d) brain with MS lesions, (e) normal brain, (f) brain with MS lesions

niques on ground truth data. Results indicate that with an appropriate abstraction tree traversal criterion, it is possible to map different tissue classes to nodes at different levels in the abstraction tree. This underscores the ability of the HSOM to adapt the level of segmentation to the feature complexity and local image content. The HSOM could also be used as an interactive visualization or navigation tool that could enable the user to view or navigate through the MRI data at multiple levels of resolution or abstraction. Future work includes fine-tuning and more exhaustive testing of the HSOM on anatomically more diverse MR image sets. The use of the HSOM to create a hierachical index/hash structure for content-based access to MR image databases is also being explored.

References

[1] J. Alirezaie, M.E. Jernigan and C. Nahmias, Neural Network–based Segmentation of Magnetic Resonance Images of the Brain, *IEEE Trans. Nuclear Science*, Vol. 44(2), pp. 194–198, 1997.

[2] S.M. Bhandarkar, J. Koh and M. Suk, Multiscale Image Segmentation Using a Hierarchical Self-Organizing Map, *Neurocomputing*, Vol. 14(3), pp. 241–272, 1997.

[3] L.P. Clarke, R.P. Velthuizen, M.A. Camacho, J.J. Heine, M. Vaidyanathan, L.O. Hall, R.W. Thatcher and M.S. Silbiger, MRI Segmentation: Methods and Applications, *Magnetic Resonance Imaging*, Vol. 13(3), pp. 343–368, 1995.

[4] T. Kohonen, *Self-Organizing Maps*, Springer–Verlag, Berlin, Germany, 1995.

[5] A. Lundervold and G. Storvik, Segmentation of Brain Parenchyma and Cerebrospinal Fluid in Multispectral Magnetic Resonance Images, *IEEE Trans. Medical Imaging*, Vol. 14(2), pp. 339–349, 1995.

[6] M. Ozkan, B.M. Dawant and R.J. Maciunas, Neural Network–based Segmentation of Multimodal Medical Images: A Comparative and Prospective Study, *IEEE Trans. Medical Imaging*, Vol. 12(3), pp. 534–544, 1993.

[7] T. Taxt and A. Lundervold, Multispectral Analysis of the Brain Using Magnetic Resonance Imaging, *IEEE Trans. Medical Imaging*, Vol. 13(3), pp. 470–81, 1994.