

Context-Dependent Segmentation of Retinal Blood Vessels Using Hidden Markov Models

Amir Pourmorteza¹, Seyed Hamid RezaTofighi¹, Alireza Roodaki¹, Ashkan Yazdani¹,
Hamid Soltanian-Zadeh¹⁻³

¹ Control and Intelligent Processing Center of Excellence, Department of Electrical and
Computer Engineering, University of Tehran, Tehran, Iran
a.pourmorteza@ece.ut.ac.ir, hszadeh@ut.ac.ir

² School of Cognitive Sciences, Institute for Studies in Theoretical Physics and Mathematics
(IPM), Tehran, Iran

³ Image Analysis Lab., Department of Radiology, Henry Ford Hospital,
Detroit, Michigan, USA
hamids@rad.hfh.edu

Abstract. Hidden Markov Models (HMMs) have proven valuable in segmentation of brain MR images. Here, a combination of HMMs-based segmentation and morphological and spatial image processing techniques is proposed for the segmentation of retinal blood vessels in optic fundus images. First the image is smoothed and the result is subtracted from the green channel image to reduce the background variations. After a simple gray-level stretching, aimed to enhance the contrast of the image, the feature vectors are extracted. The feature vector of a pixel is formed from the gray-level intensity of that pixel and those of its neighbors in a predefined neighborhood. The ability of the HHMs to build knowledge about the transitions of the elements of the feature vectors is exploited here for the classification of the vectors. The performance of the algorithm is tested on the DRIVE database and is comparable with those of the previous works.

Keywords: Hidden Markov models (HMM), context-dependent image segmentation, retinal images.

1 Introduction

Early symptoms of systematic diseases can be detected by the assessment of retinal images. Eye is considered as a window to the retinal vasculature, where the influence of the factors that affect the human body vasculature can be observed *in vivo* non-invasively. Furthermore, inspection of optic fundus photographs [1]-[4], and flourocein images [2] may help to diagnose and monitor the progress of general diseases such as diabetes, hypertension, arteriosclerosis, cardiovascular diseases, stroke and eye diseases like retinopathy of prematurity [1]-[3]. Thus, the measurement and analysis of retinal blood vessels is of diagnostic value for a wide range of pathological states. However, manual analysis of the complex retinal blood vessel

trees in fundus images is a tiresome and laborious task, and as the number of images increases, it may even be impossible.

Segmentation of the vessels from the background is an initial requirement for the automatic assessment of retinal images. Many automatic algorithms are proposed for the segmentation of retinal blood vessels. According to [9], methods for detecting blood vessels generally fall into three categories: kernel-based, tracking based, and classifier-based.

In this paper, we will introduce a novel method for the automated segmentation of blood vessels. We start with simple kernel-based preprocessing operations followed by a contrast enhancement step. Then a Hidden Markov Model-based classification is used to assign the pixels to vessel or background classes.

Recently, Hidden Markov Models (HMMs) are used for MRI brain segmentation in [5]. Markov Random Fields (MRF), Hidden Markov Random Fields (HMRF), and Hidden Markov Models, take advantage of the relative information of the vectors and fall under the category of context-dependent classifiers. Kernel-based and thresholding methods, simply assume no relation between different classes, i.e. once a feature vector (corresponding to a pixel) is assigned to a class, the next vector may be assigned to any other class. Whereas in context-dependent classification, the class to which a vector is assigned, depends (a) on its own value, (b) on the value of other feature vectors and (c) on the existing relation among various classes.

MRFs and HMRFs were introduced in several segmentation frameworks [11]-[14]. While these frameworks encode the dependency between the pixel to be segmented and its first-degree neighbors, they are computationally intensive and therefore, they are not welcome in medical environments [5]. In contrast, HMMs have proven valuable in Automatic Speech Recognition (ASR) and MRI brain segmentation tasks [5],[6].

The primary goal of this paper is to report a method which combines kernel-based techniques with state-of-the-art HMM-based segmentation. The algorithm is evaluated using the images obtained from the publicly available DRIVE database [16]. This combined method allows for results comparable with manually segmented images as well as with those reported by other authors. A brief introduction to HMMs and its training algorithms is provided in section 2. Section 3 details the preprocessing and feature extraction algorithm. Section 4 includes the segmentation algorithm and necessary post-processing step to improve the results. Finally, experimental results of the proposed method are presented in section 5.

2 Hidden Markov Models

Similar to the work done in [5], the fundamental idea of this paper relies on the ability of the underlying HMM to encode the knowledge about the input data vectors or sequences that reflect the characteristics of the image to be segmented e.g. intensity information about the pixel and its neighborhood.

HMMs are made of different states statically bound by transition probabilities. An HMM is characterized by a set of internal states, the transition probabilities among the states in response to an input symbol from the sequence, and the emission

probabilities of symbols from the different states[5]. The knowledge is built in the form of the transition and the emission probabilities of the states that are trained during the learning stage. HMMs assume that the states are hidden and cannot be observed at the output stage. Instead, only the output emitted from the states are observable, without knowing which states emitted those outputs. From this point of view, the HMM is regarded as a symbol generating process in which the observations are viewed from the outside without knowing which state emitted them. Here, similar to what has been previously done in [5] HMMs are viewed from a different perspective. In using HMMs for blood vessel segmentation, the goal is to find the best state sequence that might have produced a specific output.

Hence, during the training, the goal of the training algorithm is to increase the output probability of the input sequences representing a class of tissue [5]. The transition and emission probabilities are updated in a manner that maximizes the output probability of a given class. As a result, the relationship between the elements of a sequence (voxels of a neighborhood) is encoded in the transition and emission matrices.

The concept of Minimum Classification Error is used to update the parameters of both discrete and continuous mixture of Gaussian distributions in [5]. In [6], an Expectation Maximization (EM) algorithm is utilized to optimize the parameters of a HMM which has Gaussian pdfs for emission probabilities. Here, the algorithm used in [6] is selected. In our model, the sequence of observations is nothing more than a sequence of pixel intensities. There are many ways to generate this observation sequence. One reasonable way is to consider a 3×3 neighborhood of a certain pixel and put the intensity of the pixel along with the intensities of its neighbors in a 9×1 feature vector. Other choices of neighborhoods are also possible. The increase in size of the neighborhood may adversely affect the accuracy due to smoothing effect of large neighborhoods [5].

3 Preprocessing

The goal of preprocessing phase is to reduce the unwanted effects of noise and background intensity variations, and enhance the contrast between the vessel and non-vessel pixels prior to segmentation. Since the HMMs will have to learn respond to feature vectors from different images, the images must share nearly similar brightness and contrast characteristics. By compensating for these through and across data set variations, one can improve the performance of the segmentation process.

First step is to choose a monochromatic image that shows high contrast between vessel and non-vessel pixels. A Monte Carlo simulation of retinal vessel profiles [15] predicts that light that is observed from a retinal vessel at green band on a RGB color image is predominantly backscattered from the vessel or transmitted once through the vessel [2]. This is confirmed by experimental results [1],[4],[8]. Hence we have selected the green channel representation. In an attempt to compensate for the contrast and brightness variations among the images of the data set, the mean and standard deviation of the histograms of the images are changed so that they nearly match predefined values of mean=0.5 and S=0.075.

To compensate for the background intensity variations, a smoothing kernel is convolved with the green channel. This convolution gives an approximation of the background variations. Next, this approximation is subtracted from the green image. The result is a background normalized image as can be seen in fig.1(c). Size of the smoothing kernel is empirically selected to be 11×11 .

Histogram of the background normalized image resembles a Gaussian distribution with a small variance in which most of the pixels have midrange gray-levels. In an attempt to enhance the global contrast, the histogram is stretched by a sigmoid function. Results of this step are illustrated in fig. 1.




Fig. 1. (a) green channel of original image, (b) image with corrected mean and variance values, (c) background normalized image, and (d) contrast enhanced image.

4 Segmentation and Post-processing

The notion behind HMMs-based segmentation [5] is that, by training a HMM with feature vectors representing a special class of data, the model learns to produce higher outputs when a vector of that class is presented to the model as an input sequence. Hence, the segmentation algorithm is straightforward. A HMM is trained for each class of data, the output probability of every HMM for the feature vector representing the pixel to be classified is then computed using the trained transition matrices and emission vectors that encode the relative dependency of the elements of the feature vector. The segmentation is then done, by simply assigning the pixel to the class associated with the HMM showing the highest output probability.

In this paper, two HMMs are trained. The first model is trained using feature vectors randomly chosen from vessel pixels of manually segmented images. The second HMM is trained using non-vessel feature vectors. Feature vectors corresponding to every non-black pixel of the image is then presented to the HMMs. The output probabilities of the HMMs are represented in two likelihood maps (see figure 2). The maps are then compared pixel by pixel and the pixels whose vessel likelihoods are greater are segmented as vessel. The HMMs do not learn every possible form of blood vessel or background sequences. This is because the training features are chosen randomly.

Therefore, some sequences may never show up in the training step. This fact and the presence of noise and contrast variations, result in falsely detected blood vessels. To eliminate these, a post-processing step is necessary. Morphological operations such as erosion and dilation can be very helpful in this step. First, the image is eroded using a structural element of a certain size. Next, connected components smaller than a specified size are removed. Finally the image is dilated using the same structural element that was used for eroding (see fig. 3.).




Fig. 2. Likelihood maps of (a) vessel and (b) non-vessel classes.

5 Experimental Results

Our proposed method was tested on images of the publicly available DRIVE database [16]. The DRIVE database contains 40 color images of the retina, with 565×584 pixels and 8 bits per color channel, in LZW compressed TIFF format. The images were originally captured from a Canon CR5 nonmydriatic 3 charge-coupled device (CCD) camera and were initially saved in JPEG-format. The database also includes binary images with the results of manual segmentation. These binary images have already been used as ground truth for performance evaluation of several vessel segmentation methods [16]. The 40 images were divided into a training set and a test set by the authors of the database. The results of the manual segmentation are available for all the images of the two sets. For the images of the test set, a second independent manual segmentation also exists.

Segmentation accuracy is selected as performance measure to compare our results with previous retinal vessel segmentation algorithms. The accuracy is estimated by the ratio of the total number of correctly classified points (sum of true positives and true negatives) by the number of non-black points in the image. The ground proof for computing the performance measures was the manual segmentation result. The values for the fraction of pixels erroneously classified as vessel pixels, false positive ratio (FPR), and the percentage of pixels correctly classified as vessel pixels, true positive fraction (TPR), are also reported. Two experiments were carried out to evaluate the performance of the proposed algorithm.




Fig. 3. Result of segmentation step before (a) and after (b) morphological post-processing.

Table 1. Performance of vessel segmentation methods.

Method	Accuracy	TPR	FPR
2 nd human observer	0.9473	0.7761	0.0275
Mendonça et al. [1]	0.9452	0.7344	0.0236
Jiang et al. [4]	0.8911	not reported	
Staal et al. [1]	0.9442	0.7194	0.0227
Niemeijer [1]	0.9417	0.6898	0.0304
Soares et al. [4]	0.9466	not reported	
First Experiment	0.9388	0.7342	0.0332
Second Experiment	0.9401	0.7492	0.0384

In our first experiment, a simple global thresholding algorithm was applied on the test set of the DRIVE database. The images were initially pre-processed according to the algorithm discussed in section 3. Next a global thresholding algorithm was applied on the data set. The threshold level was empirically chosen to be 0.6. Then, we used the first two images of the training set to train two HMMs with 2000 feature vectors randomly selected for each class of pixels using the labels of the results of the manual segmentation. Each HMM had 9 hidden states and the emission probabilities were approximated by mixtures of 15 Gaussian distributions. Next, the images of the test

set were segmented followed by erosion, small connected component removal, and dilation to remove the falsely detected vessel pixels. To further improve the results, the union of the globally thresholded image and the result of the HMM-based segmentation were chosen as the final result, whenever the accuracy of the segmentation step was not acceptable.

The second experiment aimed to decrease the computation time of the segmentation step. The whole algorithm was the same as for the first experiment, except that the output likelihood maps of the HMMs were formed in a 4-connected manner, i.e. starting from the first pixel of the image, its 4-connected neighbors had the same likelihood as the pixel. As a result, the computational time was reduced to 20% of the first experiment. Consequently, once a pixel is chosen to be in class 1, its neighbors are also chosen, resulting in a 4 connected segmented image which in turn, leads to better TPR since the number of small gaps between the connected components of the vessel class is reduced.




Fig. 4. Results from - (a) human observer, (b) our algorithm

6 Discussion and Conclusion

Our method proved the ability of Hidden Markov Models in retinal blood vessel segmentation. The results are particularly comparable to previous works in which classifier-based methods are used. In [4], all 20 images of the training set of the DRIVE database were used for training, with about one million feature vectors. Whereas in our work, only 2000 feature vectors, extracted from 2 images, were used. Thus, the training time was reduced to approximately 1 hour, which is 9 times less than the time reported in [4]. However, the accuracies of these two methods are approximately equal. The simplicity of the feature vectors is also remarkable. Unlike [4] in which feature vectors are formed by Gabor transforms, our feature vectors are merely made of intensities of the pixels in a predefined neighborhood.

The downside to our algorithm is that the information used for training the HMMs is not accurate. In fact, the TPR and FPR for the human observers are 0.7761 and 0.0275, respectively. Therefore, some training vectors are wrongly presented to the HMMs and the models learn the wrong combination of pixel intensities. Another

drawback of the proposed algorithm is that the EM algorithm used for the training of the models, may converge to local minima, resulting in poorly trained HMMs. In [5] the authors claim that under certain conditions, the MCE algorithm may achieve the global minimum with higher probability. The choice of number of hidden states and Gaussian distributions is also open to further investigation.

Acknowledgment

The authors would like to thank the authors of DRIVE database for making their data publicly available. We would also like to thank Kevin Murphy for making his HMM-toolbox for MATLAB available.

References

1. Mendonça, M., Campilho, A.: Segmentation of Retinal Blood Vessels by Combining the Detection of Centerlines and Morphological Reconstruction. *IEEE Trans. Med. Imag.* 147, 1200--1213 (2006)
2. Martinez-Prez, M. E., Hughes, A. D., Thom, S. A., Bharath, A. A., Parker, K. H.: Segmentation of Blood Vessels from Red-Free and Fluorescein Retinal Images. *Med. Image. Anal.* 11, 47--61 (2007)
3. Vermeer, K. A., Vos, F. M., Lemij, H. G., Vossepoel, A. M.: A Model Based Method for Retinal Blood Vessel Detection. *Comput. Biol. Med.* 34, 209--219 (2004)
4. Soares, J. V. B., Leandro, J. J. G., Cesar Jr., R. M., Jelinek, H. F., Cree, M. J.: Retinal Vessel Segmentation Using the 2-D Gabor Wavelet and Supervised Classification. *IEEE Trans. Med. Imag.* 25, 1214--1222 (2006)
5. Ibrahim, M., John, N., Kabuka, M., Younis, A.: Hidden Markov Models-Based 3D MRI Brain Segmentation. *Image. Vision. Comput.* 24, 1065--1079 (2006)
6. Solomon, J., Butamen, J. A., Sood, A.: Segmentation of Brain Tumors in 4D MR Images Using the Hidden Markov Model. *Comput. Meth. Prog. Bio.* 84, 76--85 (2006)
7. Resch, B.: Hidden Markov Models. A Tutorial for the Course Computational Intelligence, [Online]. Available: <http://www.igi.tugraz.at/lehre/CI>
8. Feng, P., Pan, Y., Wei, B., Jin, W., Mi, D.: Enhancing Retinal Images by the Contourlet Transform. *Pattern Recogn. Lett.* 28, 516--522 (2007)
9. Hoover, A., Kouznetsova, V., Goldbaum, M.: Locating Blood Vessels in Retinal Images by Piecewise Threshold Probing of a Matched Filter Response. *IEEE Trans. Medical Imaging.* 19, 203--219 (2000)
10. Chaudhuri, S., Chatterjee, S., Katz, N., Nelson, M., Goldbaum, M.: Detection of Blood Vessels in Retinal Images Using Two-Dimensional Matched Filters. *IEEE Trans. Med. Imag.* 8, 263--269 (1989)
11. Zhang, Y., Brady, M., Smith, S.: Segmentation of Brain MR Images Through a Hidden Markov Random Field Model and the Expectation--Maximization Algorithm. *IEEE Trans. Med. Imag.* 20, 45--57 (2001)
12. Rajapakse, J. C., Piyaratna, J.: Bayesian Approach to Segmentation of Statistical Parametric Maps. *IEEE Trans. Biomed. Eng.* 48, 1186--1194 (2001)
13. Rajapakse, J. C., Kruggel, F.: Segmentation of MR Images with Intensity Inhomogeneities. *Image. Vision. Comput.* 16, 165--180 (1998)
14. Rajapakse, J. C., Giedd, J. N., Rapoport, J. L.: Statistical Approach to Segmentation of Single-Channel Cerebral MR Images. *IEEE Trans. Med. Imag.* 16, 176--186 (1997)
15. Hammer, M., Leistriz, S., Leistriz, L., Schweitzer, D.: Light Paths in Retinal Vessel Oximetry. *IEEE Trans. Biomed. Eng.* 48, 592--598 (2001)
16. Niemeijer, M., van Ginneken, B., (2002) [Online]. Available: <http://www.isi.uu.nl/Reseach/Databases/DRIVE/results.php>