

Detection of Retinal Blood Vessels Using Complex Wavelet Transforms and Random Forest Classification

Reyhaneh Sadeghzadeh
reyhaneh.sadeghzadeh@postgrad.manchester.ac.uk

Michael Berks
michael.berks@manchester.ac.uk

Susan Astley
sue.astley@manchester.ac.uk

Chris Taylor
chris.taylor@manchester.ac.uk

Imaging Science and
Biomedical Engineering
School of Cancer and
Enabling Sciences
The University of Manchester

Abstract

We present a new method for detecting vessels in retinograms. The Dual-tree Complex Wavelet Transform (DT-CWT) [6] is used to provide a rich, multi-scale description of local structure, and a random forest classifier [1] is used to classify pixels as vessel/non-vessel on the basis of their DT-CWT coefficients. The method is tested on retinograms obtained from a publicly available database and our results are compared with previously reported results for the same database. The best method to date achieved an area under the ROC, A_z , of 0.952, using a combination of pixel level and contextual information. We achieve a comparable A_z of 0.944, using only pixel level information.

1 Introduction

Retinograms – optical images of the retina – are an important tool for the early detection of eye disease and, potentially, other health risks. Diabetic retinopathy, the leading cause of adult blindness, has received particular attention [2], though other forms of eye disease are also important, whilst retinal images may, for example, provide a valuable, non-invasive approach to screening for cardiovascular risk [8]. Many developed countries have now introduced a retinal screening programme, based on digital retinography, creating the opportunity to detect disease and monitor progress at a population level. Realistically, methods of quantitative automated analysis will be required to realise this opportunity. An important problem in the analysis of retinograms is detection of the blood vessels that lie on the surface of the retina (see Figure 1). Some forms of disease can be detected directly from changes in the vascular structure [8], whilst the vessel tree always provides an essential anatomical framework for other forms of analysis. Vessel detection is a challenging problem because retinograms are intrinsically noisy and many of the vessels have low contrast. The problem of retinal vessel segmentation has been studied extensively. Staal et al [7] review some of the

most important approaches, and describe a ridge-based analysis method. Niemeijer et al [5] describe an evaluation methodology for retinal vessel segmentation and compare some of the most important approaches to retinal vessel segmentation experimentally, using the publicly available DRIVE database [5]. Other methods have since been evaluated in the same way, and results published on the DRIVE website. These show Staal's method to be the best of those tested. We have used the same database and testing methodology to compare our approach to the state-of-the-art.




Figure 1: Example of retinograms. (a) normal retinogram. (b) abnormal retinogram (abnormal region circled). (c) manual segmentation of (a). (d) manual segmentation of (b) [7]

2 Methodology

2.1 Complex Wavelet Transforms

Wavelet analysis provides a powerful basis for capturing local structure. The discrete wavelet transform (DWT) [3] provides a computationally efficient approach in which the wavelets are discretely sampled and high-pass and low-pass filters are applied to successively down-sampled versions of the original image, giving a set of wavelet coefficients at each pixel which provide a rich, multi-scale description of local structure. A drawback of the DWT is its shift dependence property [3]; another is that it provides very limited information on the orientation of image features [3]. To overcome these problems, a complex wavelet transform can be used [6]. The dual-tree complex wavelet transform (DT-CWT) combines two DWTs, using even and odd wavelets to provide complex coefficients, whilst retaining the efficiency of the DWT approach. In practice, the wavelet analysis is applied in 1-D, along rows and columns, and 6 oriented 2D complex wavelets are constructed from different combinations of the outputs, as shown in Figure 2. This analysis is performed at a series of scales differing by a factor of 2, by successively down-sampling the image. For the coarser scales, a set of responses is obtained for every pixel in the original image by interpolation [4]. The result of applying the DT-CWT is thus a set of complex wavelet coefficients at each pixel for six different orientations (sub-bands) and for each of a number of scales.

2.2 Random Forest Classification

We classify retinogram pixels into two classes – vessel or non-vessel – based on their complex wavelet coefficients, using a random forest classifier [1] – an approach that is well-suited to non-linear classification in a high-dimensional space.

Given a set of training data consisting of N samples each of which is a D -dimensional feature vector labelled as belonging to one of C classes, a random forest comprises a set




Figure 2: The oriented filters of the DT-CWT. Top set: real part. Bottom set: imaginary part.

of tree predictors constructed from the training data. Each tree in the forest is built from a bootstrap sample of the training data (that is, a set of N samples chosen randomly, with replacement, from the original data). The trees are built using a standard classification and regression tree (CART) algorithm; however, rather than assessing all D dimensions for the optimal split at each tree node, only a random subset of $d < D$ dimensions are considered. The trees are built to full size (i.e. until a leaf is reached containing samples from only one class) and are not pruned. During classification, unseen feature vectors are classified independently by each tree in the forest; each tree casts a unit class vote, and the most popular class can be assigned to the input vector. Alternatively, the proportion of votes assigned to each class can be used to provide a probabilistic labeling of the input vector.

3 Experimental Evaluation

We applied our approach to the DRIVE database, which contains 20 training retinograms and 20 test retinograms, each with expert annotated ground truth (see Figure 1). We applied a DT-CWT at 6 scales to all 40 images, giving a total of 72 (6 scales \times 6 sub-bands \times 2 complex components) features at each pixel. We found that expressing the complex values in (magnitude, phase) form gave the best results, so that is the approach we adopted in all the experiments reported here.

We considered several different approaches to using the information in the feature vectors, and trained a random forest classifier with 100 trees for each, using the 20 training images and the associated ground truth. In practice, we sampled around 3000 vessel pixels and 3000 background pixels randomly from each image in the training set – 120000 in total. We built classifiers using the following approaches (these results are representative, we tested other combinations that space does not allow us to report).

- Full feature vector: 72 dimensions
- Maximum sub-band – only the complex response with the largest magnitude across sub-bands at each scale: 12 dimensions.
- Reordered sub-bands – full feature vector, but cyclically reordered so that the maximum response is always first: 72 dimensions.
- 3x3 neighborhood – concatenation of all the feature vectors in a 3x3 neighborhood around the pixel: 648 dimensions.

Method	A_z	MAA	Kappa
Human Observer	n/a	0.9473(0.0048)	0.7589
Staal	0.9520	0.9442(0.0065)	0.7345
Current Method	0.9440	0.9336(0.0254)	0.6792
Niemeijer	0.9294	0.9416(0.0065)	0.7145
Zana	0.8984	0.9377(0.0077)	0.6971
Al-Diri	n/a	0.9258(0.0126)	0.6716
Jiang	0.9114	0.9212(0.0076)	0.6399
Martinez-Perez	n/a	0.9181(0.0240)	0.6389
Chaudhuri	0.7878	0.8773(0.0232)	0.3357
All Background	n/a	0.8727(0.0123)	0

Table 1: Comparison between methods applied to the DRIVE database [5].

We then applied these classifiers to the complete set of images (training and test), resulting in a vessel probability for each pixel. For each method we plotted a receiver operating characteristic (true positives vs false positives – ROC) for all the images in the test set, by thresholding at a series of levels and comparing the result to the ground truth. The ROC data was summarised by measuring A_z the area under the curve (an area of 1 indicates perfect classification). We also calculated the maximum average accuracy and kappa value for each method by establishing an optimal threshold using the training set, and applying that threshold to the test set to give a ‘best’ segmentation.

The best results were obtained using the full feature vector approach, which gave a maximum average accuracy of 0.934 and A_z of 0.944. These results are compared with others obtained for the DRIVE database in Table 1. Our results are second only to those obtained by Staal in terms of area under the ROC. Figure 3 shows typical vessel probability maps.




Figure 3: Classification result. (a) normal retinogram. (b) probability map of (a). (c) abnormal retinogram. (d) probability map of (c).

4 Discussion

Our results show that the DT-CWT coefficients capture a sufficiently rich representation of local structure to allow effective vessel/non-vessel classification. The performance of our method is comparable to the best method tested on the DRIVE database, even though the competing methods use far more contextual information. We expect to improve our results further by applying a contextual approach to our vessel probability images. As illustrated

in Figure 3, a significant proportion of our classification errors occur around the edge of the field of view, probably because ripples in the wavelet coefficients some distance from the very strong edge, produce responses similar to those from vessels. This requires further investigation, and may well be due to the fact that very few pixels will have been sampled from these regions during training. In practice, these errors could easily have been removed by shrinking the field of view by a few pixels, but we recognised that it was important to present results that were directly comparable with those in the literature.

In summary, the approach we have presented is computationally efficient (it takes a few minutes to train the classifier from scratch and a few seconds to segment each image), and produces encouraging results. It shows significant promise as a component of a complete system.

5 Acknowledgement

We are grateful to Nick Kingsbury for providing access to his DT-CWT code.

References

- [1] Leo Breiman. Random forests. In *Machine Learning*, pages 5–32, 2001.
- [2] Saaddine et al. Projection of diabetic retinopathy and other major eye diseases among people with diabetes mellitus. *Ophthalmology*, 126(12):1740–1747, 2008.
- [3] Neil H. Getz. A fast discrete periodic wavelet transform, 1992. EECS Department, University of California, Berkeley,UCB/ERL M92/138, <http://www.eecs.berkeley.edu/Pubs/TechRpts/1992/2238.html>.
- [4] N Kingsbury. Rotation-invariant local feature matching with complex wavelets. In *Proc. European Conf. Signal Processing (EUSIPCO)*, pages 4–8, 2006.
- [5] M. Niemeijer, J.J. Staal, B. van Ginneken, M. Loog, and M.D. Abramoff. Comparative study of retinal vessel segmentation methods on a new publicly available database. In *SPIE Medical Imaging*, pages 648–656, 2004. <http://www.isi.uu.nl/Research/Databases/DRIVE/>.
- [6] Ivan W. Selesnick, Richard G. Baraniuk, and Nick G. Kingsbury. The dual-tree complex wavelet transform, 2005. *IEEE SIGNAL PROCESSING MAGAZINE*,22(6),123-151.
- [7] Joes Staal, Michael D. Meindert Niemeijer, Max A. Viergever, and Bram van Ginneken. Ridge-based vessel segmentation in color images of the retina. *IEEE transactions on medical imaging*, 23(4):501–509, 2004.
- [8] T. Wong, R. Klein, A. Sharrett, B. Duncan, D. Couper, J. Tielsch, B.Klein, and L. Hubbard. Retinal arteriolar narrowing and risk of coronary heart disease in men and women the atherosclerosis risk in communities study. 287(9):1153–1159, 2002.