Application of Model Based Image Interpretation Methods to Diabetic Neuropathy

M.J. Byrne and J. Graham

Dept. of Medical Biophysics,
University of Manchester, Oxford Road,
Manchester, M13 9PT, UK.
Email: mjb@sv1.smb.man.ac.uk

Abstract

We present two applications of model based computer vision methods to measurement of image features significant in the diagnosis of diabetic neuropathy. The first involves the location of the boundaries of nerve fascicles in light microscope images. The second involves the segmentation of capillary cell regions using electron microscope images. In each case the boundaries required are of arbitrary shape and characterised by local texture or changes in textured regions. The fascicular boundary is located using an Active Contour Model responding to a texture measure based on edge directionality. A start position for the model is automatically generated. The capillary segmentation is performed using a region–based snake responding to a weighted combination of texture measures followed by a local boundary refinement using dynamic programming. These methods show that application of various types of Active Contour Model, accompanied by appropriate starting cues, or followed by local refinements, can locate robustly positioned and intuitively correct boundaries in these images. The aim of the work is the automation of diagnostic measurements currently performed manually. We discuss the implications of automated analysis for procedures in quantitative histology.

1. INTRODUCTION

There are various symptoms and side–effects of the disease diabetes. Important among these are the effects on the nervous system. The most frequent pattern of involvement of diabetes on the nervous system is a peripheral symmetric neuropathy (non–traumatic disorder of the peripheral nerves) of the lower extremities affecting both motor and sensory functions[14].

The major effect is degeneration of the insulating myelin sheath that surrounds each nerve fibre leading to deterioration in motor and sensory functions. Biopsies are taken from patients and images obtained using light and electron microscopy. Number densities and size distributions of myelineated nerve fibres are obtained from these images. Currently these fibre measurements are made manually[2] which is a time consuming process. We have developed automated methods for nerve fibre detection in both light and electron microscope images. We do not present these methods here, although we make use of the results in section 3.3.

Nerve fibres are grouped together to form fascicles (fig.1).

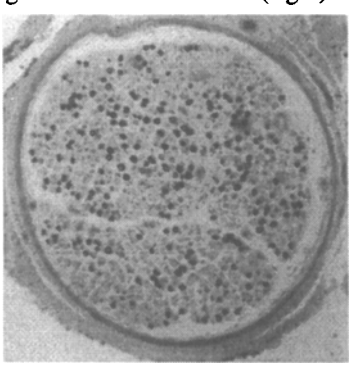

Fig.1. Example of nerve fibre fascicle

It is also necessary to locate the fascicle boundary so that an accurate measure of the fascicular area, and hence accurate fibre number densities can be obtained. In section 3 we describe a method for automatically locating the fascicular boundary using an Active Contour Model responding to an appropriate image measure with a starting cue generated as a result of the automated fibre detection.

Another symptom of interest is microangiopathy (disease of small blood vessels) with effects manifested within the nerve fibres themselves and in capillaries in the endoneurium (the interstitial connective tissue in peripheral nerves separating individual nerve fibres). The two main effects are;

i) A thickening of the basement membrane or an accumulation of basement membrane material in the capillaries (visible as an apparent contraction of the luminal area), and
ii) A proliferation of endothelial cell material together with basement membrane thickening. This manifests itself in arteries, arterioles and occasionally venules.

The structure of two endoenuerial capillaries as they appear in electron micrographs is shown in fig.2. Fig.2a is a healthy example whilst fig.2b displays a degree of diabetic neuropathy and shows these effects.

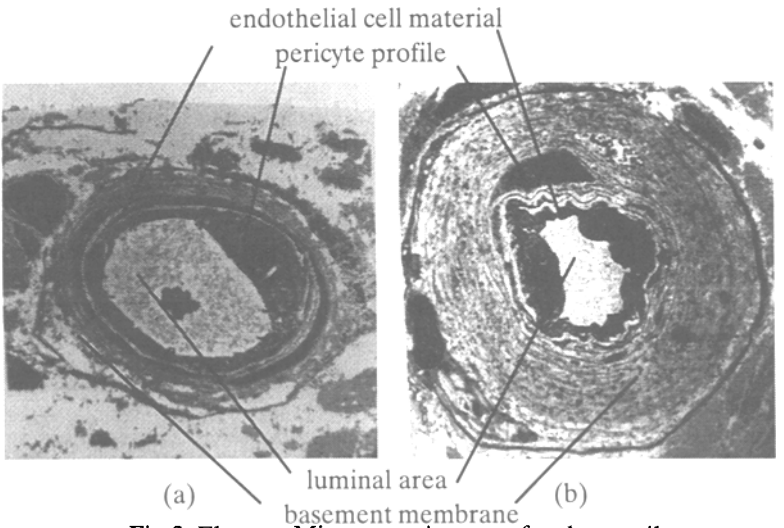

Fig.2. Electron Microscope images of endoneurail capillaries (a) normal and (b) showing neuropathy

Currently the various regions are delineated by hand[13,16]. In section 4 we will describe methods of automated segmentation of the three capillary regions using Active Contour Models followed by local boundary refinement using dynamic programming.

2. MODEL BASED METHODS

Both the fascicle and capillary images are complex. The image evidence defining the required regions involves several parameters which vary from image to image. Often the evidence is poor or missing requiring local interpolation of the data. Location and segmentation of the desired regions in both sets of images therefore require the use of some form of model based method. Statistically based methods such as Point Distribution Models (PDMs) have been sucessfully applied, as part of a constrained image search strategy[5] (Active Shape Models), to the location of poorly defined boundaries in noisy images[6]. PDMs rely on the ability to label a consistent set of landmark points representing the boundary shape in a set of training images in order to construct a statistical model of the expected boundary shape and its allowed degree of variation. In both problems described in this paper it is not possible to identify a set of landmark

points that are consistent from image to image. The lack of correspondence among training images introduces such a large degree of variability between different examples that statistics gathered on shape and appearance impose little constraint on the model. As a result, for the problems presented in this paper, Active Shape Models have not provided a useful approach to image segmentation.

2.1. SNAKES

Active Contour Models (snakes)[10] do not rely on a description of the expected region shape to constrain image search but instead impose internal constraints using a quasi–physical model to control the spacing and curvature of boundary elements. The snake combines these internal constraints with external "forces" derived from the image evidence and iteritively re–positions itself to achieve a minimum energy configuration.

In the absence of any external image forces a snake reaches equilibrium by collapsing either to a single point or line, as dictated by the internal constraints . Furthermore if the snake search is initiated too far away from the desired contour the snake will fail to converge to the correct solution. Cohen[4] added an extra inflationary term causing the snake to behave as a balloon. The inflationary term takes the form of an isotropic pressure potential resulting in an outward pressure force acting along the normal to each snake element. The balloon is inflated, expanding until trapped by strong image evidence (e.g. strong edges) but expanding through weaker evidence. The pressurised snake therefore has the advantage of being able to start its search a large distance away from the desired contour but has the disadvantage that an image feature must produce a strong response in order to overcome the snake's internal pressure force. Hence it is not always possible to segment an image adequately using a pressurised snake.

An adaptation proposed by Ivins and Porrill[9] is the statistical snake, an active region model linking the pressure term to image data within the region enclosed by the snake. An initial seed region is defined either through user interaction or through some form of cue generation. Within this region the means and variances for a suitable set of image measures are determined. These measures should be capable of distinguishing between the region of interest and those around it. The snake is allowed to expand from the boundary of the seed region until the boundary elements encounter pixels whose image measures differ significantly from those in the seed region.

An energy term for the region measure is obtained by multiplying the local change in area for each element by a goodness functional $\mathbf{G(I(x,y))}$ representing the goodness of fit of the region measure at an element of the snake positioned at $\mathbf{(x,y)}$ in an image $\mathbf{I}$. There are various choices for the goodness functional $\mathbf{G()}$. These are as follows for a region having a mean response μ and a range of allowed values within k standard deviations of μ.

Unary Pressure: The goodness functional $\mathbf{G()}$ is set to unity for pixels with image measures within the range specified by the seed region and zero for pixels outside this range.

Binary Pressure: The goodness functional $\mathbf{G()}$ is set to +1 for pixels with image measures within the range specified by the seed region and –1 for those with measures outside the range. When a snake element encounters pixels outside the seed region's range the direction of expansion at that element is reversed.

Linear Pressure. A normalised linear pressure term allows the model to reach equilibrium when its boundary elements encounter pixels at the statistical limits where the goodness functional and hence the pressure force evaluate to zero.

Several image features may be combined by use of a Mahalanobis pressure term.

3. FASCICULAR BOUNDARY LOCATION

To obtain images with sufficient resolution for fibre detection using light microscopy a magnification of 40 times is required. At this magnification several fields are required to represent an entire fascicle. As a result the fibre detection is carried out on a mosaic

of images connected using cross–correlation. The composite images produced typically have dimensions of 1000–2000 x 1000–2000 pixels.

Application of active contour models to boundary detection requires the choice of a suitable image force and a method of generating a suitable starting position for the model.

3.1 CHOICE OF IMAGE FORCE

The choice of image force is determined by the fascicular boundary structure which consists of a series of closely spaced, fairly parallel lines (fig.3a,b).

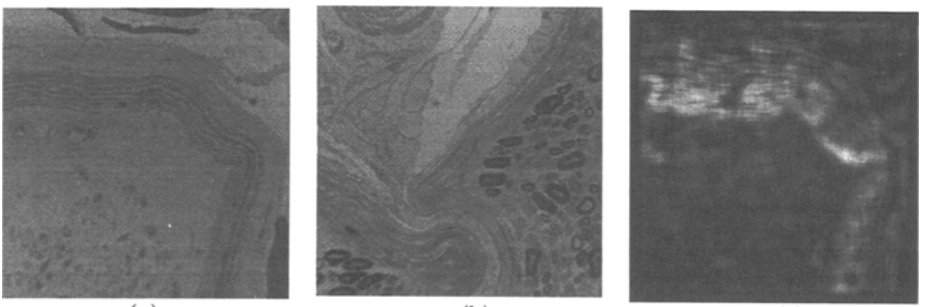

(a) (b) (c)

Fig. 3. Examples of fascicle boundary structure(a,b) and result of direction algorithm (c)

However the contrast between these lines and the background is often poor and simple measures based on edge magnitude or contrast fail. To achieve robust detection we have used a texture measure based on edge frequency and edge directionality.

Image Force Algorithm: The algorithm implemented makes uses of the response of a Canny edge detector[3].

The Canny response along the fascicular boundary to consists of a number of parallel edges. Generally the response over the remainder of the image is low, except around nerve fibres, where it shows very little local directionality.

The image feature used to generate the snake's image force is the modal value of the direction of the Canny output within a local neighbourhood. The following algorithm generates an intrinsic image based on this feature.

1. Apply Canny Operator
2. Threshold Canny output to retain only salient edges
3. Quantise edge directions to 16 values
4. Calculate modal direction value within local neighbourhood
5. Retain number of responses at modal direction as pixel output

A local neighbourhood half–width of 10 pixels was empirically found to give the most robust boundary response. An example of the algorithm's output is given in fig.3c.

3.2. GENERATION OF A START POSITION

A starting cue for the snake can be obtained by making use of myelineated nerve fibres detected by our automated method. The fascicle boundary lies in a region surrounding that containing the nerve fibres. In most cases the boundary is not a great distance from the fibre region. The limit of the fibre region is calculated using the distance transform[8] of the image containing the detected nerve fibres.

The outer boundary of the fibre region is determined by thresholding the distance transform at a range of increasing values until a single isolated contour is obtained. A single contour is typically obtained at a distance just greater than the maximum

distance between adjacent nerve fibres. This contour, after being smoothed using morphological closing, is used as the starting position of the snake. A degree of smoothing is required since the contour produced by the distance transform is extremely jagged.

3.3. RESULT OF FASCICLE SNAKE

A snake comprised of 61 elements, based on an algorithm by Williams and Shah[17] using the image force described in section.3.1, was applied to a series of sample images. In most examples the starting contour generated by the distance transform was fairly close to the actual boundary position. This allowed the snake to stabilise within 10−15 iterations achieving a close fit to the true boundary.

To assess the robustness of the boundary location with respect to the snake's starting point, an eroded version of the contour produced from the distance transform was obtained. This produced a starting point lying well within the fascicular boundary resulting in the snake having to cross regions of potentially confusing image evidence. Fig.4 shows an example of a snake using a cue eroded by 50 pixels. Fig.4a shows the detected fibres with the eroded position superimposed. Fig.4b shows the final position reached by the snake.

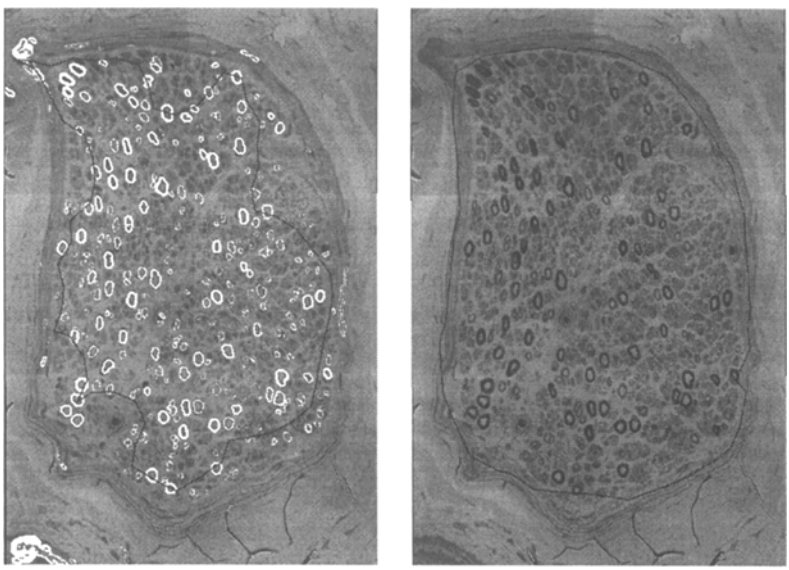

(a) (b)

Fig.4. Result of snake starting from eroded cue

While fig.4 demonstrates the robustness of the texture measure generating the image force, the greater search space and number of iterations required from the use of a distant starting point result in a substantial loss of efficiency. The advantage of having a starting point close to the true boundary position is in the generation of a rapid solution. In some cases exhibiting a high level of neuropathy the limit of the fibre region may be further away from the fascicle boundary requiring the use of an extended search space. The distance between the fascicle boundary and the limit of the fibre region may be a useful diagnostic measure in such cases.

4. SEGMENTATION OF CAPILLARY CELL REGIONS

The method used to segment the areas of interest is determined by the appearance of the three areas. The lumen is generally light in colour and is usually flat or shows only light texture. The endothelial cell material is dark in colour and shows a high degree

of structure. The basement membrane area is lighter in appearance and generally shows less structure with greater directionality than the endothelial cell area. At a coarse scale, across examples, the three regions can be characterised by specific textures. At a finer scale the region boundaries are characterised by a great deal of detailed structure.

4.1. APPROACH

A two stage strategy has been implemented. An initial approximation to the boundary is obtained by application of a region–based (statistical) snake starting from a user supplied seed region. The use of a region–based snake is appropriate since it makes use of the consistent texture within a region to locate the boundary. The only restriction on the snake's starting point is that it lies inside the region to be segmented. The approximate boundary obtained in this way is refined using a higher resolution search method based on dynamic programming which takes advantage of the approximation produced by the region–based snake. The dynamic programming method searches pixel by pixel near the approximate boundary using a measure based on the texture contrast between regions. This measure is capable of greater sensitivity to texture changes when applied close to the true boundary.

4.2. CHOICE OF IMAGE FORCE FOR REGION BASED SNAKE

An image measure is required that can distinguish between the various regions of interest in the capillary. A selection of texture measures was used with discriminant analysis applied to produce a weighted combination capable of producing good classification between either lumen and endothelial cell material or endothelial cell material and basement membrane area. The discriminant analysis is carried out separately for the two boundaries of interest producing a separate set of weights for each. The image measures used were:

- Local average luminance: the average grey level within a local neighbourhood.
- Gradient: [15] a measure of gradient as a function of distance between pixels using the distance dependent texture description function g(d) computed for a user specified distance d.
- Smoothness: A measure of the number of pixels within a local neighbourhood that lie within a specified grey level range of the central pixel value.
- Entropy: to distinguish between regions with little or no texture and regions with some degree of semi–random structure.
- Laws Texture Filters: Six combinations of Laws [11] texture filters were used. These are a well known set of 1D filters which can be combined to represent 2D texture primitives.

These measures were performed on the various regions of interest for a series of training images. A discriminant analysis was carried out on the measures obtained to produce a classification between the lumen and endothelial cell area and then the basement membrane and endothelial cell area. This produced a weighted combination of the region measures. A binary goodness functional based on the results of the discriminant analysis was used as the image force for the region based snake.

4.3. REFINEMENT USING DYNAMIC PROGRAMMING

Dynamic programming as a search technique has been applied to a variety of problems in machine vision[1]. An advantage of dynamic programming is that is always guaranteed to find the optimal path for a given objective function. It also compares well to other techniques such as heuristic search algorithms which depend critically upon the quality of the forward cost estimate. Its advantage as a refinement method is that the cost function is based on local measures, in contrast to the global energy function of the

snake methods. Lutkin[12] has shown that dynamic programming can be an effective method of assessing local image evidence based on an existing model boundary.

In order to refine the result of the region based snake, dynamic programming was applied to a "straightened" image constructed from single pixel spaced normals to the snake's boundary approximation. Each pixel on the straightened boundary corresponds to a node in the search graph[12]. The search then proceeds through the graph finding the best route as dictated by the cost function:

$$cost_i = \alpha\, bound_i + \beta\,(1 - (\theta_i - \theta_{i+1}))$$

where $bound_i$ is the normalised boundary response at node i and θ_i and θ_{i+1} are the angles at nodes i and i+1, and α and β are weighting constants. The cost function is designed to respond to the image evidence whilst maintaining a degree of compatibility with the initial boundary approximation. The compatibility constraint at a transition between nodes is a measure of the angle between the path from one node to the next and the direction of the approximate boundary. Since the dynamic programming is applied to a "straightened" version of the region based boundary approximation the second term in the cost function constrains the refined solution to remain close to this initial approximation.

Choice of Image Measure for Dynamic Programming: The image measure used is the difference in texture between two circular regions centred along the normal to the estimated boundary position. As in section 4.2, the texture measure is a weighted combination of the region measures, the weights being determined by discriminant analysis. At a true boundary point this difference between region responses should be maximised.

4.4. RESULTS

For the segmentation of the lumen from the endothelial cell area the starting point for the snake was within the lumen. For the boundary between the basement membrane and the endothelial cell area the region based snake was positioned in basement membrane and allowed to contract inwards towards the boundary. For these experiments the endothelial cell area was not used as a starting point since its structure is less consistent than that of the other two regions.

Results for the location of both boundaries are shown below. Fig.5 shows the region based snake applied to segmentation of the lumen endothelial cell area boundary. This boundary is very distinct and the region–based snake has produced a good approximation to the actual boundary position. Fig.6 shows an enlarged section of the capillary shown in fig.5 showing application of the boundary based refinement to the result of the region based snake.

Fig.7 shows the region based snake applied to locating the boundary between the basement membrane and the endothelial cell area. This boundary is less consistent than than the lumen endothelial boundary. In most places the boundary is distinct but in other places the image evidence is poor with the boundary appearing non–existent in some places. Fig.8 shows an enlarged version of fig.7 showing the detailed improvement achieved by the dynamic programming refinement.

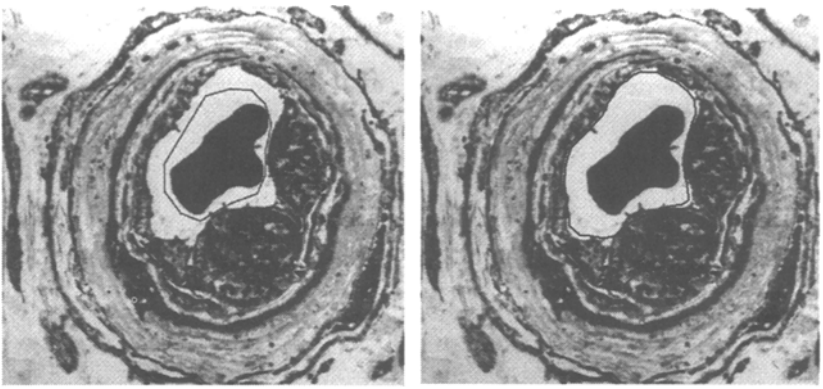

(a) (b)

Fig.5. Segmentation of lumen/endothelial cell area boundary
a)Starting position b)Result of region based snake

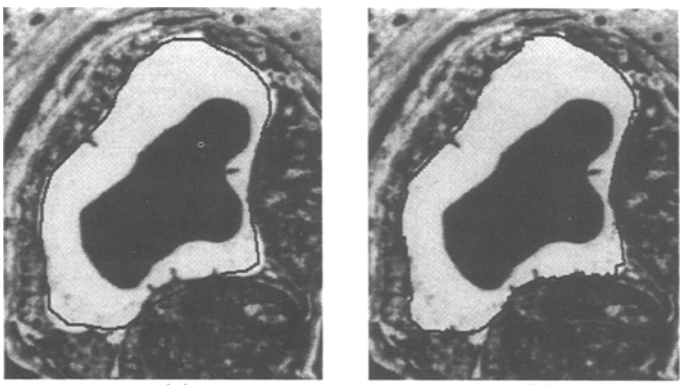

(a) (b)

Fig.6. Comparison of a)region–based snake result and
b)refinement due to dynamic programming

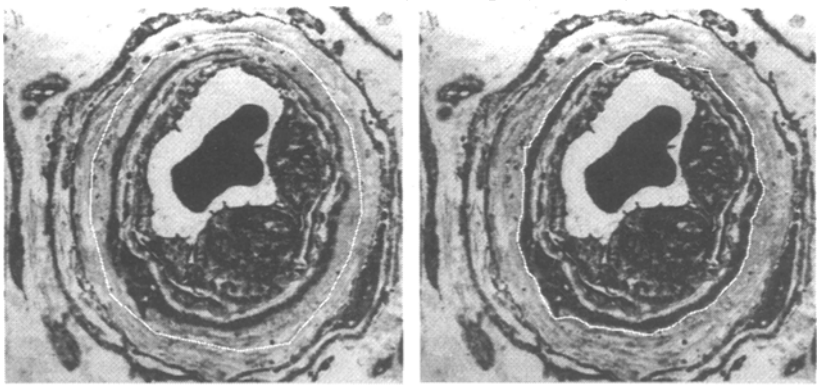

(a) (b)

Fig.7. Application of region based snake to location of basement membrane/endo-boundary. (a) starting point (b) result of region based snake.

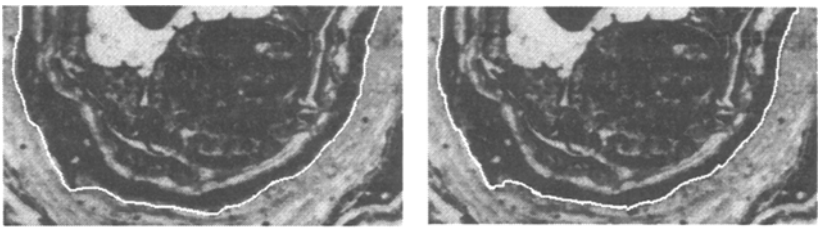

(a) (b)

Fig.8. Comparison of a)region–based snake and b)dynamic programming results.

4.5. APPLICATION OF PAIRED SNAKES

A problem encountered when segmenting the endothelial cell area from the basement membrane is the presence of small regions within the basement membrane showing similar texture to that of the endothelial cell area. Although they are small in comparison to the main body of the endothelial cell area, the region based snake can still be "trapped" by these regions (fig.9a). A snake placed within the endothelial cell area can encounter similar problems due to the inconsistent structure of this area (fig.9b).

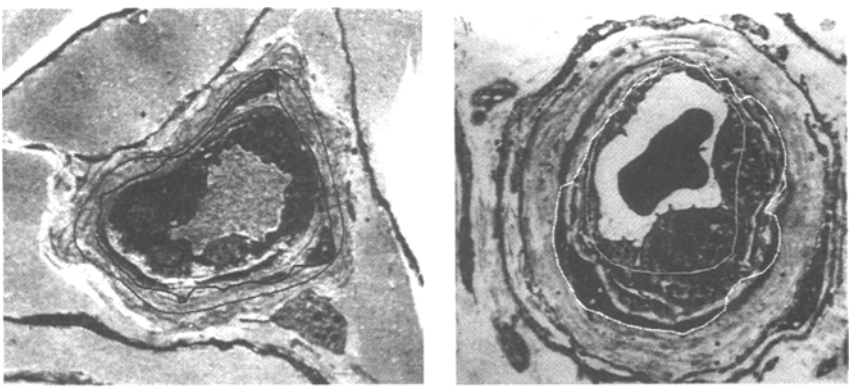

(a) (b)

Fig.9. Region based snake trapped by confusing local evidence
a)Starting from basement membrane b)Starting from endothelial cell area

This is an example of a general problem with snakes and arises from the fact that the only internal constraints on the snake are associated with smoothness. There is no way of preventing a smooth yet incorrect solution arising from confusing local evidence.

Problems with confusing evidence can be overcome by combining two or more independent assessments of the available evidence. We attempt to obtain two independent views of the evidence through the use of a pair of snakes running simultaneously from differing starting positions. In the absence of conflicting evidence both snakes would be expected to arrive at the same answer. Points where there is disagreement suggest the need for further analysis.

The snakes of section 4.4 were augmented by two further snakes initialised within the endothelial cell area, one contracting towards the lumen, the other expanding towards the basement membrane. Fig.10a shows the initial and final positions of a pair of snakes converging on the boundary between the basement membrane and the endothelial cell area. In many places the snakes arrive at an identical position. As in the example shown in fig.9a the outer snake has been trapped by the outlying regions around the endothelial area. However in these places the inner snake has generally arrived at a satisfactory solution

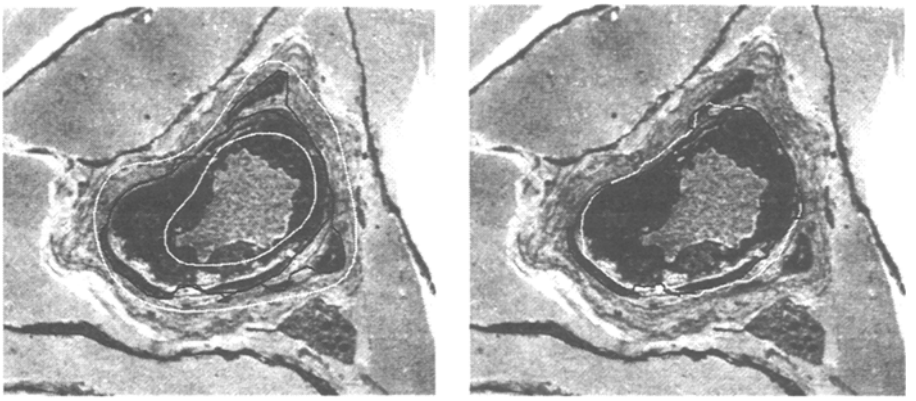

(a) (b)

Fig.10. a)Start(white) and end(black) positions of two region based snakes
b) Average snake position(white) and result(black) of application of dynamic programming

Our initial approach to resolving the conflicting evidence has been to use the average position of the two snakes as the model boundary for dynamic programming refinement. Fig.10b shows that the result of the refinement to be an acceptable representation of the region boundary.

5. CONCLUSIONS AND DISCUSSION

The overall aim of the work presented is the development on an automated system for measurement of diabetic neuropathy encompassing image capture and automated nerve fibre detection as well as the two applications discussed in this paper. The purpose of the fibre detection and fascicular boundary measurements is perform a study of the effects of diabetes on the number densities and size distributions of myelineated nerve fibres. This requires measurement to be made on samples from a large number of patients. Thus a suitably efficient and reliable automated system to replace the need for manual measurements is extremely desirable. Details of the diagnostic utility of the methods will be published elsewhere. The purpose of this paper is to describe the computer vision methods applied.

Fibre detection and fascicular boundary location is a fully automated process. Snake based methods have been shown to successfully locate the fascicular boundary using a starting cue generated from the results of automated fibre detection. The accuracy of boundary location is not sensitive to the effectiveness of this cue, but the fact that the starting point generated is generally close to the final position has a beneficial effect on the method's efficiency.

The capillary segmentation is intended to be as nearly automated as possible and is intended to replace manual delineation of region boundaries. The results shown in this paper have have been based on manually positioned starting points. It may be possible to generate image–based cues for this application as in the case of the fascicular boundary location. The generation of such cues has not been investigated as yet but a possible candidate might use the (usually) uniformly light luminal area.

Technically the achievement of this work has been the segmentation of structurally complex images. A principled approach to characterising texture boundaries has been taken based on trainable image features. These have been shown to be robust when used in conjunction with appropriate forms of Active Contour Model. Two problems which arise from from the use of snakes have proved particularly relevant to this work. Firstly the use of a global energy function means that they do not respond readily to detailed local boundary structure. Secondly the reliance on smoothness as the constraint on snake shape leads to a lack of shape specificity. In our case this means that confusing image evidence can lead to an incorrect solution. The first problem has been addressed through the use of an additional boundary refinement phase which takes locally detailed structure into account by using dynamic programming. The second problem has been addressed through the use of paired snakes to obtain two different views of the image evi-

dence. The use of this strategy in combination with the local refinement method produces adequate results. Further experiments will determine whether greater robustness is required. This could be achieved by more rigourous combination of the evidence. Gunn and Nixon[7], for example, have adopted an approach using paired snakes coupled together to encourage them to converge. Alternatively the dynamic programming refinement could weight the contribution made to its cost function by the approximate boundary according to the level of agreement between the paired snakes. At positions where the paired snakes disagree the local refinement could be allowed more freedom than at positions where the snakes are in agreement.

6. ACKNOWLEDGEMENTS

We would like to thank Dr. R.A. Malik of the Manchester Royal Infirmary for his assistance, and for providing the images used in this work.

7. REFERENCES

1. A.A.Amir, Using Dynamic Programming for Solving Variational Problems in Vision,*IEEE Trans. PAMI, 12(9),* 1990, pp.855–867.
2. S.T. Britland, R.J. Young, A.K. Sharma, B.F. Clarke, Acute and Remitting Painful Diabetic Neuropathy: A Comparison of Peripheral Nerve Fibre Neuropathy, *Pain, 48,* 1992, pp.316–370.
3. J. Canny, A Computational Approach to Edge Detection, *IEEE Trans. PAMI, 8(6),* 1986, pp679–698
4. L.D. Cohen, On Active Contours and Balloons, *CVGIP, Image Understanding, 53(2),* 1991, pp21–218
5. T.F. Cootes, C.J. Taylor, D.H. Cooper, J. Graham, Active Shape Models: Their Training and Application. *Computer Vision and Image Understanding 61,* 1995, pp.38–59.
6. T.F. Cootes, A. Hill, C.J. Taylor, J. Haslam. Use of Active Shape Models for Locating Structure in Medical Images, *Proc. IPMI (13),* 1993, pp33–47.
7. S.R. Gunn, M.S. Nixon, A Model–Based Dual Active Contour, *Proc. BMVC 94,* BMVC Press 1994, pp305–314.
8. RM. Haralick, L.G. Shapiro. Computer and Robot Vision Vol.1 *Addison–Wesley* 1992 pp221–223.
9. J. Ivins, J. Porrill, Statistical Snakes: Active Region Models, *Proc. BMVC 94,* BMVC Press 1994, pp377–386.
10. M. Kass, A. Witkin, D. Terzopoulos. Snakes: Active Contour Models. *Proc. 1st Intl. Conf. Computer Vision.* 1987, pp259–266.
11. K.I. Laws. Rapid Texture Identification. *SPIE Conf. on Image Processing for Missile Guidance.* vol.238 1980, pp376–380.
12. J.P. Lutkin. Interactive Segmentation of Medical Images. *Msc Thesis Manchester University,* 1994.
13. R.A. Malik, S. Tesfaye, S.D. Thompson, A. Veves, A.K. Sharma, A.J.M. Boulton, J.D. Ward. Microangiopathy in Human Diabetic Neuropathy: Relationship Between Capillary Abnormalities and the Severity of Neuropathy. *Diabetologia, 30,* 1989, pp.92–102.
14. R.A. Malik, P.G. Newrick, A.K. Sharma, A. Jennings, A.K. Ah–See, A.J.M. Boulton, J.D. Ward. Endoneurial Localisation of Microvascular Damage in Human Diabetic Neuropathy, *Diabetologia 36,* 1993 pp454–459.
15.R. Sutton, E. Hall. Texture Measures for Automatic Segmentation of Pulmonary Diseases, *IEEE Trans. Comp. c–21.* 1972, pp667–678.
16. R.G. Tilton, P.L. Hoffman, C. Kilo, J.R. Williamson. Pericyte Degeneration and Basement Membrane Thickening in Skeletal Capillaries of Human Diabetes, *Diabetes 30,* 1981, pp.326–334.
17. D.J. Williams, M. Shah. A Fast Algorithm for Active Contours and Curvature Estimation. *CVGIP: Image Undrstanding vol.55(1)* 1992, pp14–26.