

Optic Nerve Head Segmentation

James Lowell

Andrew Hunter

David Steel

Ansu Basu

Robert Ryder

Eric Fletcher

Abstract— **Reliable and efficient optic disk localization and segmentation are important tasks in automated retinal screening. General-purpose edge detection algorithms often fail to segment the optic disk due to fuzzy boundaries, inconsistent image contrast or missing edge features. This paper presents an algorithm for the localization and segmentation of the optic nerve head boundary in low-resolution images (about 20μ per pixel). Optic disk localization is achieved using specialized template matching, and segmentation by a deformable contour model. The latter uses a global elliptical model and a local deformable model with variable edge-strength dependent stiffness. The algorithm is evaluated against a randomly-selected database of 100 images from a diabetic screening programme. Ten images were classified as unusable; the others were of highly variable quality. The localization algorithm succeeded on all but one usable image; the contour estimation algorithm was qualitatively assessed by an ophthalmologist as having Excellent, Good or Fair performance in 83% of cases, and performs well even on blurred images.**

Keywords— **Optic Nerve Head, Diabetic Retinopathy, Active Contours, Deformable Models**

I. INTRODUCTION

THIS paper presents an algorithm for the automatic localization and segmentation of the optic nerve head in retinal images. No user intervention is required: the algorithm automatically selects the general location of the center of the optic nerve head, then fits a contour to the optic nerve head rim. Localization is achieved using a simple but effective specialized filter; segmentation by fitting an active contour to the optic nerve head rim using a specialized three phase global and local deformable model that exploits the specific characteristics of the optic nerve head's appearance. We evaluate the algorithm against alternative approaches using a set of 100 random images drawn from a diabetic screening programme, and present the results.

A. Motivation

The algorithm presented in this paper forms one component of our system for automated detection and grading of diabetic retinopathy, a common complication of diabetes mellitus and the leading cause of blindness in the working population of Western countries [1]. If appropriate treatment is given in the early stages blindness can be prevented in at least 60% of cases. There is a clear case for screening, and recently the UK National Screening Committee have suggested that annual screening by digital ocular fundal

photography is the preferred modality [2]. An expert human screener grades the images; this is a time-consuming, highly skilled task, susceptible to subjective variation and error. There is consequently a pressing need for reliable automated analysis of digitized fundal images.

Features suggestive of retinopathy include: haemorrhages and microaneurysms (dark spots); cotton wool spots and exudates (white spots); and vascular anomalies such as venous beading and neovascularisation [3] [4]. These must be detected against a highly variable background retinal image that includes: the optic nerve head, the vascular network, the macular, and distractors such as choroidal vessels, laser scars, light artifacts, and drusen. The variability of retinal appearance argues for a structured approach [5], where significant landmarks in the retina are identified, and the relationships between them exploited to increase confidence in the classification of each object.

Optic nerve head segmentation is a necessary step in this structured analysis for a number of reasons. First, the optic nerve head can itself act as a distractor: it is a large bright region that can be mistaken (by algorithms) for gross circinate exudation; the high-contrast rim also causes false responses to linear blood vessel filters [5]. Second, the vessels radiate from the optic nerve head, so vessel tracking algorithms [6] may start from there. Vessels are of direct importance in assessing vascular condition [7] [8] [9], and are useful in image registration [10]. Third, the optic nerve head is important in localization of the fovea, the central part of the retina that subserves fine vision. This lies at the center of a larger area, the macula. Retinopathy in this area, maculopathy, is associated with a high risk of visual loss. The macular is a dark approximately circular area, but the contrast is often quite small, and it may be obscured by exudates or blurring. Consequently a global correlational search often fails. The fovea is located approximately 2.5 disc diameters temporal to the temporal edge of the optic disc and between the major temporal retinal vascular arcades. These positional constraints can be used to identify a small search area for the macular, and to estimate the position if the search fails, although variation in the optic disk size compromises the reliability of this method. Littmann [11] has developed a technique to explicitly determine the distance from the disk center to the fovea, by correcting for the magnification factors of the fundus camera and the patient's eye. However, the ocular magnification factor depends mainly on anterior corneal curvature, refraction and axial length, which require measuring variables on the patient's eye – an important practical limitation in screening. The method can also give false values in the case of an abnormally high lens refractive power (e.g. by cataract formation) and so is not applicable in aphakic or pseudophakic eyes (10% of our diabetic

J. Lowell and A. Hunter, Dept. Computer Science, University of Durham, Science Labs, South Rd., Durham, DH1 3LE, UK. andrew.l.hunter@durham.ac.uk

D. Steel, Sunderland Eye Infirmary, Queen Alexandra Road, Sunderland, SR2 9HP

A. Basu and R. Ryder, Dept. Diabetes and Endocrinology, City Hospital, Dudley Rd., Birmingham, B18 7QH, UK

E. Fletcher, Dept. Computer Science, University of Sunderland, Tyne and Wear, UK.




Fig. 1. The Optic Nerve Head. a) Cross Section. b) A typical well-defined disk.

screening population). The “2.5 disk diameters” heuristic method is therefore more practical for our application, particularly as we do not require absolute accuracy.

We also note that segmentation of the optic nerve head is relevant to other diseases of the eye (e.g. glaucoma [12] [13]), but we do not consider these further in this paper.

B. Optic Nerve Head Appearance

Successful segmentation of the optic nerve head requires a careful analysis of its appearance (see figure 1). It is the extremity of the optic nerve in the interior of the eye, and also the entrance and exit site of the retinal vessels [14]. The shape is approximately elliptical, with a vertical principal axis (width $1.8 \pm 0.2mm$, height $1.9 \pm 0.2mm$) [15]. As the nerve fibres reach the optic nerve head they turn and exit through the optic nerve, leaving a small depression (the “cup”) in the center of the nerve head. There is often a brighter central region, the “pallor”, which if present usually includes the cup. The optic disc rim is judged to be the inner margin of the peripapillary scleral ring, seen as a thin white band encircling the optic disc.

In fundus images, the appearance varies quite substantially. The size and shape may vary significantly. The rim is usually visible as a bright boundary; the nasal side is usually less bright than the temporal side, and sometimes not visible at all. In some images the entire optic nerve head is brighter than the surrounding area, so that it appears as a disk; in others the appearance is of a hollow ring. In either case the pallor may appear as a smaller, brighter disk within the optic disk. There may also be bright areas just outside the rim caused by peripapillary atrophy, either distorting the shape or forming concentric elliptical arcs.

To complicate the issue further, departing vessels partially obscure the rim. The majority climb out on the nasal side and depart vertically; a smaller number depart nasally, with only a few fine vessels on the temporal side. Sometimes vessels turn at the nasal rim edge and run vertically, obscuring portions of the rim. A consequence of the nasal distribution of vessels is that the pallor, if visible, is mainly located to the temporal side.

The variability in appearance misleads obvious approaches. Large areas of circinate exudates, which have high contrast, act as strong distractors for correlation-based localization algorithms – algorithms that work well on images of healthy retina may fail on a screening pop-

ulation. Similar problems arise from reflection artifacts (young patients) and visible choroidal vessels.

Segmentation is complicated by the presence of strong distractors along the pallor and vessel edges, weakness of the rim and peripapillary atrophy [13]. However, these problems can be overcome by exploiting specific aspects of the appearance: the relative sharpness, reliability and lack of vascular intersections on the temporal side of the rim, and the approximately elliptical shape. The major contribution of this paper is in demonstrating how these specific characteristics of the optic nerve head may be exploited.

C. Screening Data

We have tested the algorithms using a random sample of 100 fundus images taken from 50 patients attending the diabetic retinal-screening programme at City Hospital, Birmingham. Some of the patients had been referred from family practitioners, and consequently demographic data was unavailable for two. The mean age of the remaining patients was 63.7 years (s.d. 14.8 years), with 65.5% ($n = 29$) male and 34.5% ($n = 19$) female. The patients were from various ethnic backgrounds (Asian 20%, Afro-Caribbean 16%, Caucasian 50%, Unknown 14%). 19 patients had type 2 diabetes mellitus, while the diabetes status was unavailable for the remaining 31. Given the characteristics of the regional diabetic population, type 2 diabetes mellitus is likely to predominate in this group too.

The images were acquired using a Canon CR6 45MNf fundus camera, with a field angle lens of 45 degrees, resolution 640×480 . Images were converted to grey-scale by extracting the Intensity component from the HSI representation. There is considerable variation in the images, with many characteristics that can affect the algorithms; these are summarized in Table I.

TABLE I
VISUAL CHARACTERISTICS. CHARACTERISTICS THAT PRESENT POTENTIAL PROBLEMS ARE ENUMERATED, IN THREE CATEGORIES: PROBLEMS RENDERING IMAGES POOR ENOUGH TO EXCLUDE; AFFECTING LOCALIZATION; AND AFFECTING SEGMENTATION. SOME IMAGES HAVE MULTIPLE CHARACTERISTICS.

Characteristic	No. images
No detectable optic nerve head	4
Severe Cataract	8
Moderate Cataract	2
Total potentially unusable	10
Exudates or laser scars	7
Light artifacts	7
Easily visible choroidal vessels	20
Total localization endangered	34
Some of rim blurred or missing	27
Severe peripapillary atrophy	6
Moderate peripapillary atrophy	19
Concentric peripapillary atrophy/artifacts	23
Strong pallor distractor	13
Total segmentation endangered	58

D. Contents of Paper

This paper introduces two new algorithms. In section II we discuss our optic nerve head localization algorithm, a simple but effective specialized correlation filter. The main advantage over previously suggested techniques is its simplicity combined with robust performance. In section III we discuss our segmentation algorithm, which exploits the optic nerve head’s structural characteristics to overcome distractors. This gives it good performance, particularly in the critical rim search phase, which is not considered by some previously published papers. In section IV we evaluate this algorithm against alternatives from the literature. Section V concludes the paper.

II. OPTIC DISK LOCALIZATION

A. Overview of Optic Disk localization Algorithms

Reliable optic nerve head localization is surprisingly difficult, due to the variable appearance; naïve approaches that work well for images of healthy retinæ often fail on screening images. Sinthanayothin *et al.* [15] use the rapid intensity variation between the dark vessels and the bright nerve fibres to locate the optic disk; we found that this algorithm often failed for fundus images with a large number of white lesions, light artifacts or strongly visible choroidal vessels. Akita *et al.* [14], trace the parent-child relationship between blood vessels segments, tracking back to the center of the optic disk; this presupposes robust detection of the blood vessels, which is difficult in images of diseased retinæ where even quite sophisticated algorithms detect false positives along the edges of white lesions and along the edge of the optic nerve head itself [5]. Lalonde *et al.* [16], uses pyramidal decomposition and Hausdorff-based template matching. The template-based matching technique is based on the edge map using a Hausdorff distance measure, guided by scale tracking of large objects using multiresolution image decomposition. This method is effective, but rather complex. We present a much simpler localization method with a detection performance of 99% on usable images from our data set.

B. A Correlational Filter

Our localization algorithm uses a specialized correlation filter, which matches key elements of the optic disc structure. The correlation peak gives the approximate center of the optic disc. The optic disc consists of a high intensity near-circular rim, with a significant vertically-oriented, roughly centrally located band of low intensity blood vessels; other parts of the disc (including the interior) are not reliable and are discounted. The template consists of a Laplacian of Gaussian with a vertical channel in the middle to correspond to the major vessel band; see figure 2.

The size of the optic nerve head varies significantly; in our data set had widths from 65–101 pixels (mean 78.5, standard deviation 7.6). This might suggest a need for filters at different scales. In fact the single filter used suffices; where the size mismatches the maximal correlation lies on an annulus around the optic nerve head center; the peak is




Fig. 2. Localization filter




Fig. 3. Localization Algorithm. a) Gross exudation. b) Strong pallor.

therefore still located within the optic nerve head, which is sufficient for our requirements.

The template is correlated with the intensity component of the fundus image. We use the full Pearson-R correlation, to account for variations in mean intensity and contrast, defined by:

$$C_{i,j} = \frac{\sum_{x,y} (f(x,y) - \bar{f}(x,y))(w(x-i, y-j) - \bar{w})}{\sqrt{\sum_{x,y} (f(x,y) - \bar{f}(x,y))^2 \sum_{x,y} (w(x-i, y-j) - \bar{w})^2}} \quad (1)$$

where \bar{w} is the mean value of the template, calculated once, and \bar{f} is the mean value of the area covered by w . The filter is prone to locate a point slightly to the temporal side, due to the characteristic asymmetry of the optic nerve head (see figure 1), which is convenient for the next stage of the algorithm.

C. Testing and Results

Performance is assessed against the 96 out of our 100 fundus images that have a discernable optic disc. An ophthalmologist labelled the center of each optic disc to create a gold standard. Images are graded by distance from the true center (see table II); the algorithm usually finds the center within 15 pixels, which is acceptable as a starting point for rim segmentation. The center point is within the optic disk area on all usable images bar one. Figure 3 shows some examples of its performance. This simple approach yields excellent results and is robust; its success lies in careful design of the filter.

TABLE II
LOCALIZATION PERFORMANCE

	0 – 5	6 – 15	16 – 25	26+
Fundus images	42	45	8	1
Percentage	43.8	46.9	8.3	1

III. OPTIC DISK BOUNDARY SEGMENTATION

A. Overview of Deformable Contour Algorithms

Active contours, or snakes [17], are deformable models that are fit to object edges under the control of two forces: *external forces* that pull the model towards image features such as edges; and *internal forces* that act as smoothing constraints or object model constraints. There are two main categories: freeform and parametric snakes. Freeform snakes provide many local degrees of freedom, whereas parametric deformable models encode a specific shape and its variations [18] [19]; the model can help to overcome problems in the images such as boundary gaps. The optic nerve head is well-suited to the parametric approach, as it has a simple global model and significant distractors.

Active contours seek points with a high gradient. There are two main problems: noise, which may trap the model in local minima; and the domain of attraction of the edges – if too small, the snake may not find the desired solution [20]. These issues are addressed by pre-filtering with a low-pass Gaussian filter [21]. Often the gradient magnitude image, $\|\nabla I\|$, is used; in gradient vector methods the direction of the vector is retained to help locate edges with an expected orientation.

B. Optic Nerve Head Contour Algorithms

There are several key problems in using active contours for optic nerve head segmentation. The blood vessels contribute powerful distractors along their edges, and obscure parts of the rim. The pallor edge may present a strong contrast boundary; it may also combine with the temporal edge of vessels on the nasal side to form a strong elliptical distractor. There may also be bright concentric arcs of peripapillary atrophy outside the rim. The rim – particularly on the nasal side – may be blurred, hidden by vessels or non-detectable. Two previous authors have reported the use of active contours to find the optic disk boundary.

Mendels [12] used a freeform snake, initialized as a circle centered on and inside the optic nerve head. The model tended to fit a convoluted boundary, following vessel edge distractors. Mendels addressed this problem using grey-scale morphological closure to remove most of the vessels. Mendel evaluated only nine images; the image illustrated in Mendel’s paper is of quite high quality, and it is to be expected that the algorithm would perform poorly on blurred images, or images with a low intensity rim or significant distractors at the pallor edge, due to the lack of a global model. Mendel used a gradient vector method, and noted that this helps to avoid distractors.

Lee [22] also applied an active contour model to high resolution images centered on the optic nerve head. Like

Mendel, he removed vessels morphologically. The control points and attractors were placed along radial spokes emanating from the center, thus partially imposing a global model. Attractor points were attached to edges detected using a Canny filter, which makes the system prone to complete failure where the rim is missing. Lee reported problems caused by the boundary of the pallor and by very faint or missing edges, and presented results on only four images.

These papers indicate the significant difficulties in segmenting the optic nerve head rim. However, neither author fully exploited the strongly consistent overall shape of the optic nerve head to constrain the contour, bypass distractors, and maintain shape where the rim is not visible. Both authors assumed that the model position was initialized quite accurately, as in their images the optic nerve head is always centered. In our less constrained problem domain initialization is much less reliable, and consequently our search algorithm must be much more robust.

Hu *et al.* [23] recently described a method for boundary extraction of cross-sectional blood vessels in 3D imaging. A circular global model is first fit to the boundary contour, then a local deformable model with variable stiffness fits closely to strong edge features, while maintaining a smooth contour close to the global model when edge features are weak or missing. There is a strong correspondence to our problem domain, and this paper extends the work of Hu *et al.* to optic nerve head segmentation. A number of alterations were necessary, including generic improvements (elliptical model, use of the gradient vector, and fast optimization), and domain specific improvements (exploitation of the optic nerve head topology).

In section III-C we describe Hu’s model, in section III-D we describe our generic alterations, and in section III-E we introduce the optic nerve head algorithm.

C. Hu’s Circular Deformable Model

In this section we describe Hu *et al.*’s deformable model [23], with some change of variables from the original paper to simplify the presentation. The global model is a circle with center \mathbf{c} , and radius r . The local model is defined by the center \mathbf{c} , and S evenly-spaced radial spokes, each with angle θ_i , and direction vector $\mathbf{s}_i = [\cos(\theta_i), \sin(\theta_i)]$. The model is defined by distances m_i from \mathbf{c} along each spoke, with corresponding (x, y) location vector $\mathbf{m}_i = \mathbf{c} + m_i \mathbf{s}_i$. The local model has a corresponding global model, with radius $r = \overline{m}_i$, the local model’s mean radial displacement.

The normalized negative gaussian-smoothed gradient magnitude image, $\Gamma = -\|\nabla I\| / \max(\|\nabla I\|)$, is sampled to produce radial gradient profiles; the radial displacement of the “peak gradient” within a limited search range about m_i , denoted g_i , provides an external attractor point. Since radial lines cross image pixels in an irregular manner, non-uniform sampling can occur. The gradient magnitude at distance ρ along spoke i is determined as follows. Let $\mathbf{p} = (p_x, p_y) = \mathbf{c} + \rho \mathbf{s}_i$. We calculate the radial gradient $\gamma_i(\rho)$ by bilinear interpolation of the image gradient magnitude at the point’s four neighboring pixels [24]; see




Fig. 4. Algorithm details. a) Bilinear interpolation of gradient. b) Stiffness factor function.

equation 2, figure 4. The “peak,” g_i , is chosen so that $\gamma_i = \gamma_i(g_i)$ is the minimum on the spoke.

$$\gamma_i(\rho) = - \sum_{x,y} w_{x,y} \Gamma(x,y) \quad (2)$$

$$x_d = |p_x - x| \quad y_d = |p_y - y| \quad (3)$$

$$w_{x,y} = \begin{cases} (1 - x_d)(1 - y_d) & (x_d < 1) \wedge (y_d < 1) \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

The contour is altered under the influence of a force, f_i , with external and internal components; see equation 5. All the forces work along the radial spokes. The external force drags the model towards the attractor points; see equation 6. The internal force limits model deformation using two components: the global force, which pulls the model towards the global shape; and the local force, which smoothes the model by penalizing differences in deformation between neighboring spokes. The internal force is modified by a stiffness factor, β_i , described further below.

$$f_i = f_i^{ext} - f_i^{int} \quad (5)$$

$$f_i^{ext} = g_i - m_i \quad (6)$$

$$f_i^{int} = \beta_i (d_i^{int} - \alpha \langle d_i^{int} \rangle) \quad (7)$$

where $d_i^{int} = m_i - r$ is deformation of the model from the global model on spoke i , and $\langle d_i^{int} \rangle$ is the mean deformation of neighboring spokes. The coefficient α balances the local versus global internal forces.

The stiffness parameter, β_i , controls the relative strength of the internal and external forces. On a strong edge a small stiffness value is used, the external force dominates, and localization is not unduly compromised. Where the edge is weak a larger stiffness value is assigned, and the contour is locally smoothed and attracted to the global model.

Hu *et al.* used a function shaped like that in figure 4b to determine the stiffness factor, but did not give the formula; our stiffness function is defined as follows. Let γ_i be the gradient magnitude at point g_i , (see equation 2), $\langle \gamma_i \rangle$ the mean gradient magnitude of the neighboring spokes. Define the squashing factor $\zeta = \chi + \langle \gamma_i \rangle$, with $\chi = 1.5$, which shifts the function to the left or right of the middle position. Let $\psi_i = 1 - \gamma_i$; then the stiffness factor β_i is given by:

$$\beta_i = \begin{cases} (2\gamma_i^3 - 3\gamma_i^2 + 1)^\zeta & \zeta \geq 0 \\ 1 - (2\psi_i^3 - 3\psi_i^2 + 1)^{(1-\zeta)} & \zeta < 0 \end{cases} \quad (8)$$




Fig. 5. Model fitting. a) The deformable model. b) Locking to the temporal edge

Hu *et al.* use the radial forces defined above to provide a center-shifting force, a global radius force, and local deformation forces. The shifting force is found by converting the radial forces to vector forces in the (x, y) plane, and averaging (equation 9), then the radial forces are then rectified by removing the shifting force projected in the radial direction (equation 10).

$$\mathbf{f} = \frac{1}{S} \sum_i f_i \mathbf{s}_i \quad (9)$$

$$f_i^* = f_i - \mathbf{s}_i \cdot \mathbf{f} \quad (10)$$

These forces are used to iteratively adjust the model, by adding \mathbf{f} to \mathbf{c} , and f_i^* to m_i (or, if a circular model is desired, f_i^* to r).

D. Alterations to the model

We introduce three modifications: an elliptical global model; the use of the vector gradient; and the use of energy functions to support fast non-linear optimization.

Our global model is an ellipse with a vertical principal axis, and a fixed aspect ratio, a ; see figure 5a. We retain the “radius” parameter, r , which is the ellipse width; the height is consequently ar . The elliptical model can be transformed into an equivalent circular model by scaling using the spoke ratios, a_i (equation 11), to “normalize” radial distances and forces. The local model has an associated global model, with radius defined by equation 12. We introduce variables for the model and attractor points normalized to the circular frame of reference, $m_i^n = m_i/a_i$, $g_i^n = g_i/a_i$.

$$a_i = \|\cos(\theta_i), a \sin(\theta_i)\| \quad (11)$$

$$r = \frac{1}{S} \sum_i \frac{m_i}{a_i} \quad (12)$$

We work with the smoothed normalized gradient vector image, $\Upsilon = \nabla \mathbf{I} / \max(\|\nabla \mathbf{I}\|)$, rather than its negative magnitude Γ . During profiling, the dot product of the gradient vector and the spoke direction vector determines the gradient coincident with the spoke – see equation 13; the weights are determined by equation 3. This yields a direction-sensitive gradient which ignores distractors at most vessel edges.

$$\gamma_i(\rho) = \left(\sum_{x,y} w_{x,y} \Upsilon(x,y) \right) \cdot \mathbf{s}_i \quad (13)$$

We use a fast non-linear optimization procedure, Quasi-Newton [25]. This requires an explicit energy function, in addition to a gradient function (the latter equivalent to Hu *et al.*'s forces). We define the energy functions as the sum-squared deformation of the model from the attractors:

$$E^{xy} = \frac{1}{2} \left\| \left(\frac{1}{S} \sum_i \mathbf{m}_i^n - \mathbf{c} \right) \right\|^2 \quad (14)$$

$$E = E^{ext} + E^{int} = E^{ext} + E^{glo} + E^{loc} \quad (15)$$

$$E^{ext} = \frac{1}{2} \sum_i (g_i^n - m_i^n)^2 \quad (16)$$

$$E^{glo} = \frac{1}{2} \sum_i \beta_i (m_i^n - r)^2 \quad (17)$$

$$E^{loc} = \frac{1}{2} \sum_i \beta_i \alpha (m_i^n - \langle m_i^n \rangle)^2 \quad (18)$$

Differentiating E^{xy} with respect to \mathbf{c} , and E with respect to m_i^n , and observing that r and $\langle m_i^n \rangle$ in equations 17 and 18 are both dependent on m_i^n , we obtain the gradients:

$$\frac{dE^{xy}}{d\mathbf{c}} = \frac{1}{S} \sum_i m_i^n \mathbf{s}_i \quad (19)$$

$$\frac{dE}{dm_i^n} = \frac{dE^{ext}}{dm_i^n} + \frac{dE^{glo}}{dm_i^n} + \frac{dE^{loc}}{dm_i^n} \quad (20)$$

$$\frac{dE^{ext}}{dm_i^n} = g_i^n - m_i^n \quad (21)$$

$$\frac{dE^{glo}}{dm_i^n} = \beta_i ((m_i^n - \overline{m_i^n}) - \overline{(m_i^n - \overline{m_i^n})}) \quad (22)$$

$$\frac{dE^{loc}}{dm_i^n} = \beta_i \alpha [(m_i^n - \langle m_i^n \rangle) - \langle [m_i^n - \langle m_i^n \rangle] \rangle] \quad (23)$$

The global and local gradients differ from the forces used by Hu *et al.* in the second term; they are identical at zero deformation, but the correct gradients are required for efficient optimization.

Let \mathbf{w} denote a vector of model parameters, $E(\mathbf{w})$ the error function for these parameters, and \mathbf{w}' the gradient of the error function with respect to the parameters; \mathbf{w} may be any desired combination of model parameters (e.g. m_i , \mathbf{c} , r); if we wish to optimize several parameters simultaneously, the appropriate parameter and gradient vectors may be concatenated and the error functions summed. We use standard non-linear optimization techniques to choose the parameters. We used two approaches: gradient descent with momentum (for the simpler global model phase) and Quasi-Newton BFGS (for the local model phase); both proved more efficient than the direct iterative technique.

In gradient descent with momentum, we choose a learning rate, η , and a momentum rate, μ , and update the model parameters \mathbf{w} at iteration $\tau + 1$ using:

$$\mathbf{w}_{\tau+1} = \eta \mathbf{w}'_{\tau} + \mu \mathbf{w}_{\tau} \quad (24)$$

Quasi-Newton optimization involves generating a sequence of matrices, \mathbf{G}_{τ} representing increasing accurate approximations to the inverse Hessian, \mathbf{H}^{-1} , using information from the first derivatives of the error function [26]. Using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) procedure, the approximation \mathbf{G}_{τ} of the inverse Hessian matrix, is updated using equation 29.

$$\mathbf{G}_0 = -\mathbf{I} \mathbf{w}'_0 \quad (25)$$

$$\mathbf{p} = \mathbf{w}_{\tau+1} - \mathbf{w}_{\tau} \quad (26)$$

$$\nu = \mathbf{w}'_{\tau+1} - \mathbf{w}'_{\tau} \quad (27)$$

$$\mathbf{u} = \frac{\mathbf{p}}{\mathbf{p}^T \nu} - \frac{\mathbf{G}_{\tau}^T \nu}{\nu^T \mathbf{G}_{\tau} \nu} \quad (28)$$

$$\mathbf{G}_{\tau+1} = \mathbf{G}_{\tau} + \frac{\mathbf{p} \mathbf{G}_{\tau}^T}{\mathbf{G}_{\tau}^T \nu} - \frac{(\mathbf{G}_{\tau} \nu) \nu^T \mathbf{G}_{\tau}}{\nu^T \mathbf{G}_{\tau} \nu} + (\nu^T \mathbf{G}_{\tau} \nu) \mathbf{u} \mathbf{u}^T \quad (29)$$

$$\mathbf{w}_{(\tau+1)} = \mathbf{w}_{(\tau)} + \xi_{\tau} \mathbf{G} \mathbf{w}' \quad (30)$$

The update procedure ensures that the approximation to the inverse Hessian is positive definite. The direction vector $-\mathbf{G}_{\tau} \mathbf{w}'$ is guaranteed to descend, and rapidly converges on the Newton direction. Optimization is by equation 30, where ξ_{τ} is found by a bracketing line search [25].

E. Optic Nerve Head Segmentation

Our localization and deformable model algorithms are very reliable, but the latter is sensitive to poor initialization. Hence, the most critical phase is the optimization of the global model; direct fitting of this proves unreliable. We therefore introduce a four phase algorithm, with the global fit carried out across the two middle phases. Phase two roughly locates the rim, by locking the global model onto the relatively reliable and distractor-free temporal rim; phase three fits it to the entire rim. Phase one (localization) was described in section II; phases two to four are described below.

E.1 Phase two – temporal lock

In this phase we define a global model using $S = 9$ spokes at angles -60° to $+60^\circ$ in steps of 15° , on the temporal edge only. The gradient image is smoothed with a gaussian filter ($\sigma = 5.0$), and the radial search range for γ_i is set to $m_i \pm 6$ pixels. The aspect ratio is fixed at $a = 1.03$ and the radius at $r = 40$; only the model center is adjusted. Optimization is by gradient descent (see equation 24), with $\eta = 1$, $\mu = 0.1$, $\mathbf{w} = \mathbf{c}$, $\mathbf{w}' = dE^{xy}/d\mathbf{c}$ (see equation 19).

The temporal bias of the localization algorithm ensures that the initial contour is usually just outside the temporal edge. Due to the lack of temporal vessels, their typically radial alignment, and the relative strength of the temporal edge, this phase locks on with high reliability; see figure 5b. The fixed radius ensures that the nasal edge of the model bypasses the distractors at the pallor and central blood vessels. A fixed radius is necessary since fitting an

elliptical model to 120° of arc on one side is prone to gross mis-estimation.

E.2 Phase three – global fit

This phase activates the entire global model, using $S = 24$ evenly-spaced spokes at 15° spacing; the radial search range is reduced to ± 4 pixels. Other parameters are identical to phase two. The radius of the model is now allowed to alter in addition to the center; the energy function becomes $E^{xy} + E^{ext}$, with the latter treated as a function of r ; $\mathbf{w} = [\mathbf{c}; r]$. The gradient of E^{ext} with respect to r is given by equation 31, derived from equation 21, bearing in mind that to enforce the global model, we set $m_i = a_i r$.

$$\frac{dE^{ext}}{dr} = \sum_i \frac{dE^{ext}}{dm_i^n} \frac{dm_i^n}{dr} = \sum_i (g_i - r) \quad (31)$$

Optimization is by a short burst of gradient descent, which “tweaks” the model onto the nasal edge. The use of the gradient vector (equation 13) ensures that the model is not pulled onto the nasal edges of the nasal vessels, which have opposite contrast direction to the rim, or the edges of near-orthogonal vertically-departing vessels. The only strong distractors are the pallor edge and peripapillary artifacts. The small search range usually avoids the former.

The fixed aspect ratio, $a = 1.03$ may lead to sub-optimal models (actual aspect ratios were in the range 0.94–1.15, mean 1.03, s.d. 0.038). To compensate for this, we iteratively recalculate the aspect ratio from the attractor points at the end of the phase, and repeat the optimization (which is extremely fast); three iterations suffice. The aspect ratio is calculated using equation 32, where \mathbf{u} and \mathbf{v} are the unit vectors along the x and y axes.

$$a = \sqrt{\frac{\sum_i m_i (\mathbf{v} \cdot \mathbf{s}_i)^2}{\sum_i m_i (\mathbf{u} \cdot \mathbf{s}_i)^2}} \quad (32)$$

E.3 Phase four – local deformation

When the global model reaches equilibrium the local model is activated. To improve contour localization the gradient image is recalculated with smoothing factor $\sigma = 1$. The model points, m_i , are allowed to vary in addition to the center, \mathbf{c} ; therefore, $\mathbf{w} = [\mathbf{c}; m_1, \dots, m_S]$. The full energy function from equation 15 is used, with \mathbf{w}' composed from the corresponding gradients. Local stage optimization is more challenging than the global stage; we therefore use the fast Quasi-Newton algorithm. The control parameter settings for the energy function were, $\alpha = 0.5$.

IV. TESTING AND RESULTS

We tested our algorithm against a number of alternative approaches, describe below. We used a subset of 90 images, excluding those with no discernable optic disk, or with severe enough cataract to prevent meaningful segmentation. To produce a “gold standard” segmentation four clinicians manually delimited the rim; we calculate their mean contours, and the radial standard deviations of these contours.

Let μ_i^j and σ_i^j summarize the clinician’s choice of rim location on spoke i of image j . We define the discrepancy, δ^j , on image j using equation 33. Division by σ compensates for uncertainty in rim position; $\varepsilon = 0.5$ is a small factor to prevent division by zero where the clinicians are in exact agreement.

$$\delta^j = \sum_i \frac{|m_i^j - \mu_i^j|}{\sigma_i^j + \varepsilon} \quad (33)$$

The obvious choice of overall summary statistic, the mean discrepancy, is uninformative as the distribution is heavily skewed (e.g. mean discrepancy 3.6, median 1.1 for the “temporal lock” algorithm). For illustrative purposes we instead plot the ogive of disparity, δ on a logarithmic scale (i.e. the number of images with disparity less than δ , $|\{j : \delta^j < \delta\}|$, versus δ). On these plots a superior model lies left of and above an inferior one; the number of images fit to any given level of accuracy can be read off the y-axis.

.4 Parameter Settings

We experimented with a wide range of parameter settings, including: initial model diameter and aspect ratio, gaussian smoothing factors, radial profile search sizes, and α in the deformable model. The values reported throughout the paper were heuristically selected for best performance. In most cases the algorithm’s overall performance is not very sensitive to changes in these parameters.

The initial diameter parameter has significantly performance implications for a small number of optic nerves heads of outlying size, as pallor and peripapillary distractors may induce errors. This might suggest running the algorithm with several initial diameters and selecting the lowest energy resulting contour. Unfortunately, strong distractors may actually have lower energy than the true contour, and a simple geometric rule to choose between alternative solutions is also elusive (pallor distractors are inside the true rim, but peripapillary distractors are outside), which makes the utility of even sophisticated minima-avoidance algorithms [27] questionable. We therefore use a fixed initial diameter which functions well in the vast majority of cases.

A. Performance of the algorithm

Figure 7a shows the performance graph of our final algorithm against a simple benchmark approach (“direct”) which proceeds directly from localization to local fitting without using the temporal lock or the vector gradient. Qualitatively, we define four categories (Excellent, Good, Fair and Poor) containing images with disparity up to one, two, five, or more, respectively. Figure 6 shows examples in each category. Table III summarizes performance on this subjective scale. These disparity ranges correspond reasonably well to a subjective assessment of quality.

B. Comparison with Alternative Algorithms

The temporal lock phase is a critical aspect of our algorithm that allows it to compensate for poor localization




Fig. 6. Sample Segmentations. a) Excellent. b) Good. c) Fair. d) Poor. Solid line: algorithm; dotted line: mean clinician boundary.

TABLE III
SUBJECTIVE CLASSIFICATION OF PERFORMANCE

(%)	<i>Excellent</i>	<i>Good</i>	<i>Fair</i>	<i>Poor</i>
<i>TemporalLock</i>	42	31	10	17
<i>Simple</i>	9	8	30	53
<i>DV – Hough</i>	39	22	20	19

while avoiding distractors. There are alternative methods: we can attempt to remove or mitigate distractors and use some other method to improve localization of the rim.

A method suggested by both Lee [22] and Mendels [12] to reduce distractors is to morphologically remove the vessels; we refer to this as “de-vascularization”, and implement it by grey-scale morphological closure with a disk-shaped structuring element, radius 7. Mendels also noted that the use of the vector gradient rather than magnitude improved his algorithm. Figure 7b shows the interaction between de-vascularization and the gradient vector, in the absence of a temporal lock phase. Either technique alone is beneficial, but there is no advantage in using both, as they compensate for the same distractors.

An alternative to the temporal lock phase is to attempt to fit a global model to the rim non-iteratively. We have experimented with the use of an elliptical Hough transform. The image is processed using a Canny edge-detector (upper and lower thresholds 0.2 and 0.08 respectively), and a Hough transform is used to locate the rim. Vessel edges provide a significant number of distractors, and we found the technique worked well only in conjunction with de-vascularization. We looked for ellipses of width 56–112, aspect ratio 1.03, and smoothed the Hough-space with a Gaussian kernel ($\sigma = 1.5$) before peak detection to account for deformations. We maintained our standard phase one localization algorithm, and used the Hough fil-




Fig. 7. a). Temporal Lock (TL) versus direct algorithm. b). Interaction between De-vascularization (DV) and vector (Vec)/magnitude (Mag) gradient versions of direct algorithm. c). Temporal Lock (TL), De-vascularized Hough (DV Hough), and De-vascularized direct (DV).

TABLE IV
ERROR TYPES

Source of error	Number
Localization failure	3
Temporal lock failure	5
Pallor distractor	3
Severe peripapillary atrophy	4

ter to tune the location within a search radius of 150 pixels, as we found the technique too unreliable to be used across the entire image; we also accepted the Hough center only if the peak in Hough space exceeded a threshold. Results are illustrated on figure 7c. In combination with de-vascularization and the vector gradient, we obtain a quite capable algorithm (“DV-Hough”), with performance close to that of the temporal lock algorithm. Using the one-sided sign test for paired medians, the difference in performance (temporal lock is better in 59/90 cases) is significant at the 0.5% level. The superiority of either algorithm over the simpler alternatives is equally clear-cut.

We have also tried integrating de-vascularization and the Hough search into the temporal lock algorithm, but these complications do not improve that algorithm.

Although the temporal lock algorithm performs very well in comparison to alternatives, there is room for improvement; in many cases the algorithm failed to detect the rim where it is unambiguously apparent to the human observer. Table IV attributes the cause of failure on the 15 “poor” performance images; two of the “temporal lock” failures are difficult to attribute cleanly, as localization was rather poor and probably contributed to the failure. Two of the “peripapillary atrophy” failures subjectively look “fair.”

V. CONCLUSIONS AND DISCUSSION

We have presented algorithms for localization and segmentation of the optic nerve head, an important stage in structured analysis of the retina, which can be used in diagnosis of eye diseases such as diabetic retinopathy.

Although a number of methods have been published for optic nerve head localization, many are unreliable when confronted with images of diseased retinæ including strong distractors, and the reliable methods tend to be quite computationally complex. We have presented a simple but effective algorithm for localization.

Optic nerve head segmentation by active contours has not been extensively examined in the past. There are significant problems in dealing with distractors along blood vessels edges and the pallor, and with the very variable appearance of the optic nerve head. Previously published techniques require careful initialization of the model position, pre-processing of the image using morphological operations, and perform badly where the rim is faint or undetectable. In contrast, the algorithm presented in this paper exploits specific features of the optic nerve head anatomy to achieve good localization while avoiding distractors. The temporal lock algorithm exploits the natural shape of the neuroretinal rim to bypass blood vessels and avoid the pallor, and the global and local deformable model deals effectively with weak areas of rim and vessel crossings. Defining energy functions and using a Quasi-Newton optimization strategy makes the algorithm reasonably fast.

We have conducted an experimental comparison with a range of alternative approaches, using a fairly large randomly-selected experimental set, and have demonstrated the superiority of the proposed approach. We note that the algorithm still does not have perfect performance, owing to the variable nature of the images and the presence of distractor boundaries concentric with the desired rim, which may be located either inside or outside the rim. Future work to improve the algorithm to deal with these distractors will require reasoning at a higher level.

Although our current approach focuses on detecting the optic nerve head boundary as a stage in the structured analysis of images for diabetic retinopathy fundus images, the approach is also of relevance in diagnosis of other diseases, and particularly in screening for glaucoma; in the future we will evaluate the method for use in detection of the rim in glaucomatous images.

ACKNOWLEDGMENTS

This project was supported by Diabetes UK, Project Grant No. BDA:RD00/0002033.

The authors would like to thank the referees for their detailed and useful suggestions, which have helped to greatly improve this paper.

REFERENCES

- [1] A. Manivannan, P.F. Sharp, R.P. Phillips, and J.V. Forrester, "Digital fundus imaging using a scanning laser ophthalmoscope," *Physiol Meas*, vol. 14, pp. 43–56, 1993.
- [2] D.E. Singer, D.M. Nathan, H.A. Fogel, and A.P. Schachat, "Screening for diabetic retinopathy.," *Ann Intern Med*, vol. 116, pp. 660–71, 1992.
- [3] J.J. Kanski, *Clinical Ophthalmology 3rd Edition*, Butterworth-Heinemann, 1994.
- [4] E.M. Kohner, *Diabetes and the Eye, Oxford Textbook of Ophthalmology Vol. 2*, Oxford University Press, 1999.
- [5] A. Hoover, V. Kouznetsova, and M. Goldbaum, "Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response," *IEEE Transactions on Medical Imaging*, vol. 19, pp. 203–210, 2000.
- [6] T. Teng, M. Lefley, and D. Claremont, "Progress towards automated diabetic ocular screening: a review of image analysis and intelligent systems for diabetic retinopathy," *Medical and Biological Engineering and Computing*, vol. 40, pp. 2–13, 2002.
- [7] A. Daxer, "Characterization of the neovascularization process in diabetic retinopathy by means of fractal geometry: diagnostic implications," *Graefe's Arch. Clin. Exp. Ophthalmol.*, vol. 231, pp. 681–686, 1993.
- [8] W. Lotmar, A. Freiburghaus, and D. Bracher, "Measurement of vessel tortuosity on fundus photographs," *Graefe's Arch. Clin. Exp. Ophthalmol.*, vol. 211, pp. 49–57, 1979.
- [9] P.H. Gregson, Z. Shen, R.C. Scott, and V. Kozousek, "Automated grading of venous beading," *Computers and Biomedical Research*, vol. 28, pp. 291–304, 1995.
- [10] M.H. Goldbaum, V. Kouznetsova, B.L. Coté, W.E. Hart, and M. Nelson, "Automatic registration of digital ocular fundus images for comparison of lesions," *SPIE 1877, Ophthalmic Technologies III*, pp. 291–304, 1995.
- [11] H. Littmann, "Zur bestimmung der wahren grösse eines objektes auf dem hintergrund des lebenden auges," *Klin Monstbl Augenheilkd*, p. 180:286, 1982.
- [12] F. Mendels, C. Heneghan, P.D. Harper, R.B. Reilly, and J-Ph. Thiran, "Extraction of the optic disc boundary in digital fundus images," in *Proc. 1st Joint BMES/EMBS Conf.*, 1999, p. 1139.
- [13] K. Yogesan, C.J. Barry, et al., "Software for 3-d visualization/analysis of optic-disc images," *IEEE Engineering in Medicine and Biology*, vol. 18, pp. 43–49, 1999.
- [14] K. Akita and H. Kuga, "A computer method of understanding ocular fundus images," *Pattern Recognition.*, vol. 15, pp. 431–443, 1982.
- [15] C. Sinthanayothin, J.A. Boyce, H.L. Cook, and T.H. Williamson, "Automated localisation of the optic disc, fovea, and retinal blood vessels from digital colour fundus images," *Br J. Ophthalmol.*, vol. 83, pp. 902–910, 1999.
- [16] M. Lalonde, M. Beaulieu, and L. Gagnon, "Fast and robust optic disc detection using pyramidal decomposition and hausdorff-based template matching," *IEEE Trans. Medical Imaging*, vol. 20, pp. 1193–1200, 2001.
- [17] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes. active contour models," *Int. J. Comput. Vision*, vol. 1, pp. 321–331, 1987.
- [18] D. Terzopoulos and D. Metaxas, "Dynamic 3d models with local and global deformations: deformable superquadrics," *IEEE Trans. PAMI*, vol. 13, pp. 703–714, 1991.
- [19] A.K. Jain, Y. Zhong, and M. Dubuisson-Jolly, "Deformable template models: A review," *Signal Processing*, vol. 71, pp. 109–129, 1998.
- [20] F.A. Velasco and J.L. Marroquin, "Robust parametric active contours: the sandwich snakes," *Machine Vision and Applications*, vol. 12, pp. 238–242, 2001.
- [21] H. Eviatar and R.L. Somorjai, "A fast, simple active contour algorithm for biomedical images," *Pattern Recognition Letters*, pp. 969–974, 1996.
- [22] S. Lee, *Visual Monitoring of Glaucoma*, Ph.d., Robotics Research Group Department of Engineering Science, University of Oxford, 1991, Available on micro-fiche.
- [23] Y-L. Hu, W.J. Rogers, D.A. Coast, C.M. Kramer, and N. Reichek, "Vessel boundary extraction based on a global and local deformable physical model with variable stiffness," *Magnetic Resonance Imaging*, vol. 16, pp. 943–951, 1998.
- [24] A. Blake and M. Isard, *Active Contours*, Springer, 1998.
- [25] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, *Numerical Recipes in C*, 1992.
- [26] C.M. Bishop, *Neural Networks for Pattern Recognition*, Clarendon Press, Oxford, 1995.
- [27] M. Lalonde and L. Gagnon, "Variable neighbourhood search for geometrically deformable templates," in *Proc. Int. Conf. Patt. Rec.*, Québec, Canada, 2002, pp. 689–692.