

Accepted Manuscript

A Recursive Bayesian Approach To Describe Retinal Vasculature Geometry

Fatmatülzehra Uslu, Anil Anthony Bharath

PII: S0031-3203(18)30362-5
DOI: <https://doi.org/10.1016/j.patcog.2018.10.017>
Reference: PR 6685



To appear in: *Pattern Recognition*

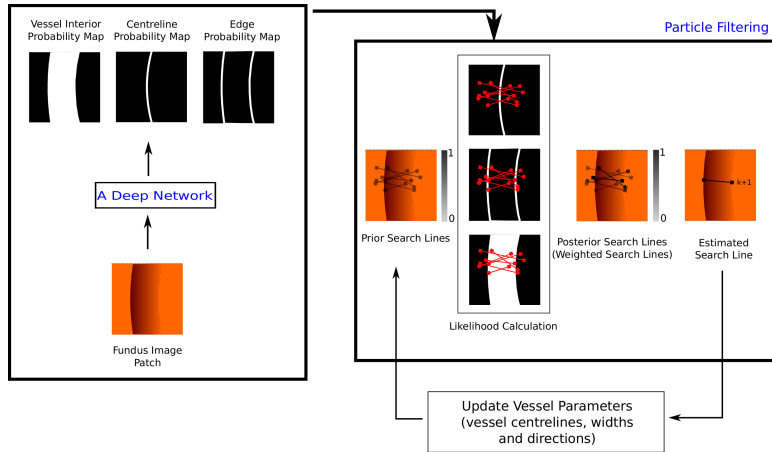
Received date: 28 November 2017
Revised date: 5 October 2018
Accepted date: 13 October 2018

Please cite this article as: Fatmatülzehra Uslu, Anil Anthony Bharath, A Recursive Bayesian Approach To Describe Retinal Vasculature Geometry, *Pattern Recognition* (2018), doi: <https://doi.org/10.1016/j.patcog.2018.10.017>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Highlights

- A deep belief net (DBN) is trained to detect vessel interior, centreline and edges.
- Particle filtering is used to quantify vasculature by using the output of the DBN.
- A vessel profile is represented with probability profiles of centreline and edges.
- The appearance of vessels in fundus images is considered in a vessel geometry model.
- The lack of labelled data for segmentation was tackled using a probabilistic approach.



A Recursive Bayesian Approach To Describe Retinal Vasculature Geometry

Fatmatülzehra Uslu^{a,*}, Anil Anthony Bharath^a

^a*Bioengineering Department, Imperial College London, South Kensington Campus, SW7 2AZ, UK*

Abstract

Deep networks have recently seen significant application to the analysis of medical image data, particularly for segmentation and disease classification. However, there are many situations in which the purpose of analysing a medical image is to perform parameter estimation, assess connectivity or determine geometric relationships. Some of these tasks are well served by probabilistic trackers, including Kalman and particle filters. In this work, we explore how the probabilistic outputs of a single-architecture deep network may be coupled to a probabilistic tracker, taking the form of a particle filter. The tracker provides information not easily available with current deep networks, such as a unique ordering of points along vessel centrelines and edges, whilst the construction of observation models for the tracker is simplified by the use of a deep network. We use the analysis of retinal images in several datasets as the problem domain, and compare estimates of vessel width in a standard dataset (REVIEW) with manually determined measurements.

Keywords: Particle filtering, deep neural network, Deep Belief Net, fundus image, width estimation, tracking

1. Introduction

*Corresponding author

Email addresses: fzehrauslu@gmail.com (Fatmatülzehra Uslu),
a.bharath@imperial.ac.uk (Anil Anthony Bharath)

With the increased adoption of deep learning for semantic labelling, it is easy to forget that the labelling task may not be an end goal in itself. For example, the segmentation of a medical image is often merely an input to some further process, such as rendering or measurement. In this paper, we investigate the use of deep networks as inputs to tackle the problem of *tracking*. A key insight is that probabilistic outputs of networks can be useful inputs to probabilistic tracking algorithms, and methods of parameter estimation. We focus this paper on a problem domain in which measurement is important, and for which meaningful comparisons might be made about the performance of tracking and parameter inference. This domain – retinal image analysis [1, 2] – is one that contains well-established and open datasets for which results of manual and automated measurement processes are available.

~~The retina provides a convenient way to image fine vasculature optically. Both organ-specific diseases, such as diabetic retinopathy or glaucoma and also systemic disorders such as diabetes, hypertension and cardiovascular diseases induce early changes in the retina. Moreover, the brain, which is closely located to the eye, shares similar characteristics to retinal vasculature.~~

Potential associations between the changes on the retinal vasculature geometry and the presence of some diseases have been the subjects of demographic studies [3, 4, 5, 6]. The changes on the vasculature, which may be subtle, can be on vessel width, curvature and branching angles. In order to detect and measure these changes, quantitative analysis is helpful. This involves steps beyond mere segmentation, such as estimating the locations of vessel centerlines [7] and edges to sub-pixel precision, at a basic level, and vessel widths and curvatures, at a more sophisticated level. However, the quality of fundus images captured in large screening programmes varies. The variable image quality [8] can present a challenge even to the best segmentation algorithms, let alone any attempt to obtain automatic estimates of vessel width. This may be sometimes complicated by the complex topology and global geometry of the vasculature, where vessels can be very close or overlap; moreover, vessel appearance may be affected by pathologies, the central light reflex or uneven illumination.

In the literature, the majority of methods for the quantitative analysis of vasculature use binary vessel maps obtained by segmentation as *a priori* information to locate vasculature. Techniques for obtaining quantifiable measures include model fitting [9, 10, 11, 12, 13], graphs [14] or active contours [15]. The location of pixels inside vessels can be obtained from vasculature skeletons generated by thinning binary vessel maps [10, 11]. Because of reliance on segmentation, these approaches may not take into account missing or false vessels inherited from the segmentation.

On the other hand, a few methods, not included in the former group, require prior knowledge of vessel parameters only to start a largely autonomous estimation process. For example, tracking methods ([see Salih *et al.* \[16\] for a general review of visual tracking, and \[17\] for techniques that integrate deep networks for video-tracking](#)) sequentially estimate vessel trajectory and geometry parameters, given a prior estimate of a point on a vessel and relevant parameters [18, 19, 20, 21]. These approaches are well-aligned with the principles of Bayesian estimation. As an example, Chutatape *et al.* [18] used an extended Kalman filter, Yin *et al.* and Zhang *et al.* Maximum a Posteriori (MAP) and, Wu *et al.* generalized particle filters [21]. In all of these approaches, a key factor influencing the performance of tracking seems to relate to how well the likelihood function reflects the actual vessel geometry and appearance in the original image data [19, 20].

To date, approaches to Bayesian vessel tracking have used likelihood functions that describe the appearance of vessel cross-section in the original, or at most band-pass filtered, image data. For example, the cross-sectional vessel intensity profile has been approximated with one-dimensional Gaussian functions, [18, 19, 20]. However, the intensity profile can be easily affected by many factors, such as the presence of pathologies, the central light reflex, uneven illumination of the retina, noise, the contrast of vessels, the focus of the camera; the latter may lead choroidal vessels being superimposed on retinal vasculature. These factors may make the Gaussian approximation too optimistic for vessel parameter estimation. As reported by Zhang *et al.* [20], adding “vesselness” – a

measure of how likely a particular pixel is to belong to a blood vessel – information in a longitudinal direction to the likelihood function led to an improvement
 65 in the tracking performance obtained by Yin *et al.* [19], who used only vessel profiles in the likelihood function.

In contrast, this study introduces a new way to represent vessel cross-sections: using probability maps for vessel interior, centerline and edge locations.
 70 The probability maps are produced by a single deep network, thus maintaining the relations between the vessel parts, significantly simplifying the construction of a likelihood function. Then, the likelihood function is used in a sequential Bayesian method, particle filtering, which extracts a comprehensive and precise representation of vasculature, through the estimation of parameters that
 75 describe vessel geometry. We introduce a new model for particle filtering to improve the flexibility in the search of best fitting parameters to actual vessel geometry.

Finally, we suggest a new way to evaluate the performance of tracking in vasculature, which considers the dependency of vessel profiles in a vessel segment.
 80 This approach to performance evaluation is suggested to better identify vessels where there is large disagreement between reference and estimated values. The experiments indicate that the method can cope with various characteristics of a dataset without re-training the network on unseen datasets. Moreover, where vessels are very small and of low contrast, the method appears to be more reliable than human observers at detecting vessel boundaries. Figure 1 presents a
 85 general overview of the approach.

2. A Probabilistic Tracking Method For Retinal Vasculature

2.1. Problem Definition

Tracking of retinal vasculature can be described as a recursive estimate of
 90 vessel geometry parameters, considering the smooth variations on vessel thickness and curvature over a vessel branch. The recursion can be initiated with an initial estimate of the geometry parameters given at iteration $k = 0$. For

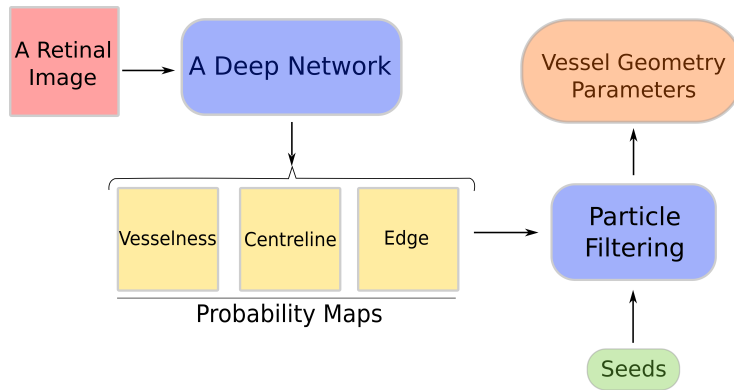


Figure 1: A general overview of the proposed method.

example, at $k = 2$, the geometry parameters are initially estimated given the parameter estimates at $k = 1$, and then corrected by evaluating the fitness
 95 of the initial predictions to measurements obtained from the vessel part under tracking. The evolution of the geometric parameters and the relation between the geometry parameters and the measurements can be modeled with (1-2).

$$\mathbf{z}_{k+1} = G(\mathbf{z}_k, \mathbf{v}_k) \quad (1)$$

$$\mathbf{y}_{k+1} = O(\mathbf{z}_{k+1}, \mathbf{u}_{k+1}) \quad (2)$$

where $G(\cdot)$ is a model of vessel geometry, which captures the evolution during tracking of the parameters describing vessel geometry in a recursive way. $O(\cdot)$
 100 is an observation model, which relates a set of the geometry parameters to a set of measurements. The state vector \mathbf{z} , contains geometric parameters describing the vessel being tracked, and \mathbf{v} and \mathbf{u} , respectively, represent the uncertainty in the geometry and observation models.

2.2. Bayesian Approach to Solution

105 The estimation of the geometry parameters given the observations can be defined with the prediction of the posterior probability distribution of the geometry parameters $P(\mathbf{z}_{k+1}|y_{k+1})$ from the Bayesian point of the view. The

posterior probability distribution can be calculated from Bayes' rule in (3).

$$P(\mathbf{z}_{k+1}|\mathbf{y}_{0:k+1}) \propto P(\mathbf{z}_{k+1}|y_{0:k}) \cdot P(\mathbf{y}_{k+1}|\mathbf{z}_{k+1}) \quad (3)$$

where $P(\mathbf{z}_{k+1}|y_{0:k})$ and $P(\mathbf{y}_{k+1}|\mathbf{z}_{k+1})$ respectively [denote show](#) the prior probability distribution of the geometry parameters and their likelihood at iteration $k + 1$.

The prior probability distribution reflects our belief about vessel parameters, which is represented with the updated initial estimate of the geometry parameters at the start of each iteration, k . When $k = 0$, this prior could be initialised by manual input, or a method that detects vessel tracks as they leave the optic disc [20]. As iterations proceed, the prior probability distribution could be evolved according to the geometry model $G(\cdot)$ in (1), where the posterior probability distribution of the geometry parameters at iteration k is used as the prior probability distribution of the geometry parameters for iteration $k + 1$. The role of the noise, \mathbf{v} , in (1) is to explain how much change in vessel geometry is foreseen over an iteration.

The prior probability distribution over geometric parameters could then be updated with (3), according to the fitness of this distribution to measurements obtained from the image, described by the likelihood. The likelihood, $P(\mathbf{y}_{k+1}|\mathbf{z}_{k+1})$ incorporates the observation model, $O(\cdot)$, in (2).

2.2.1. Particle Filtering

Particle filtering [22], a technique based on recursive Bayesian estimation, can capture an arbitrary posterior probability distribution of geometry parameters at iteration $k + 1$, with a set of particles $\{\mathbf{z}_{k+1}^n\}_{n=1}^N$ and their weights $\{W_{k+1}^n\}_{n=1}^N$. Each particle, \mathbf{z}_{k+1}^n , hypothesizes a set of geometric parameters (see Section 2.3) at iteration $k + 1$. According to the Law of Large Numbers [23], a sufficient number of particles can approximate the distribution; $P(\mathbf{z}_{k+1}|\mathbf{y}_{k+1}) \approx \sum_{n=1}^N W_{k+1}^n \delta(\mathbf{z} - \mathbf{z}_{k+1}^n)$.

Particle filtering often relies on importance sampling [24] to estimate the posterior probability distribution. Importance sampling initially samples parti-

cles from a proposal distribution, then the weights of the particles are updated with the importance weights [25]. In the case of the proposal distribution being the prior probability distribution, and using resampling after each weight update, the recursive update to the importance weights is simplified to (4) [24]:

$$W_{k+1}^n \propto P(y_{k+1} | \mathbf{z}_{k+1}^n) \quad (4)$$

140 where W_k^n and W_{k+1}^n are the weights of n^{th} particle at iteration k and $k + 1$ consecutively. $P(y_{k+1} | \mathbf{z}_{k+1}^n)$ denotes the likelihood of the n -th particle at iteration $k + 1$.

The expectation of the posterior probability distribution can be calculated with (5), returned as the geometry parameters estimated at iteration $k + 1$:

$$\bar{\mathbf{z}} \approx \sum_{n=1}^N W_{k+1}^n \mathbf{z}_{k+1}^n \quad (5)$$

2.3. A Geometry Model

The three-dimensional shape of a vessel can be assumed to be a tube with
 145 a width (diameter) w and a centerline location C , and oriented in a direction D . The change on its geometry over a small distance s (related to the step size during tracking) can be assumed to be smooth, even though sudden changes can occasionally occur due to pathologies such as vessel beading. Regarding the geometry model in (1), the change on the geometry can be modelled by using
 150 a normal distribution, which accounts for the noise \mathbf{v} . Although we can model the vessel geometry with the diameter w , centerline location C , direction D and step size s with the tube model, it is difficult to infer these parameters from the appearance of vessels in fundus images without considering factors influencing its appearance, such as optical characteristics, imaging parameters, noise and
 155 pathologies.

To date, many Bayesian tracking methods [18, 19, 20] have modelled the appearance of the cross-section of a vessel segment conditional on centerline location, diameter and orientation, by Gaussian functions. Though analytically convenient, it is not realistic when the shape of the intensity profile changes due

160 to uneven illumination, pathologies or other noise components. The differences
between the shape of the intensity profile of a vessel without the central light
reflex, and a Gaussian function describing it may appear as (i) the intensity
profile is skewed to one of vessel edges [10] and (ii) the intensities at edges of
the same vessel may be different [11]. These characteristics of the intensity
165 profile have been addressed by Lupacscu *et al.* and Araujoa *et al.*[10, 11] in
their parametric models to improve the estimates of vessel widths. However,
these characteristics have not been applied in Bayesian tracking methods.

In this study, we factorise the probability profile of vessel cross-section to
vessel centerline and edge probability profiles. This factorisation allows us to
170 relax the symmetry in appearance models implied in using Gaussians to model
vessel profile; in other words, the peaks of Gaussian functions being in the middle
of the intensity profile. Therefore, we can model a skewed vessel appearance
over a profile by using two parameters for vessel width: the distance between
an arbitrary location, A , inside the vessel and left edge w^L and that between
175 the location and right edge w^R . The state vector of the vessel parameters at
iteration k can be written $\mathbf{z}_k = [D_k, A_k, w_k^L, w_k^R]$. The evolution of the state
vector over iterations is given in (6)-(9):

$$D_{k+1} = D_k + \epsilon_D \quad (6)$$

$$A_{k+1} = A_k + D_{k+1} \cdot s + \epsilon_A \quad (7)$$

$$w_{k+1}^R = w_k^R + \epsilon_{w^R} \quad (8)$$

$$w_{k+1}^L = w_k^L + \epsilon_{w^L} \quad (9)$$

180 where ϵ represents a normal distributed noise variable (ϵ_D for direction vector,
 ϵ_{w^R} and ϵ_{w^L} for width and ϵ_A for the arbitrary interior location). s is a con-
stant and denotes the step size for tracking. $(\mathbf{x})_\perp$ denotes a direction vector
perpendicular to arbitrary vector, \mathbf{x} .

2.4. An Observation Model

In this study, we introduce three sources of information, namely: from vessel centerline, edge, and interior probability maps. How these maps are generated will be explained in Section 2.4.3. We utilise profiles through the edge and centerline probability maps, rather than from the intensity images, to perform tracking. Figure 2 compares the probability profiles for a large and thin vessel, obtained along the red lines from the centerline and edge probability maps shown in Figure 9, with those taken from reference vessel maps. The former figure exemplifies the typical characteristics of the probability profiles: (i) these curves are virtually free of either noise components, or intensity variations that are present within the fundus images in Figure 9. (ii) these curves have maxima at locations which almost overlap with significant points (e.g. centerline locations) of the reference vessel profiles. The small disagreement between reference and estimated profiles is acceptable: locating precise boundaries is difficult even for human observers, and subject to significant inter-observer variability [26, 27, 15].

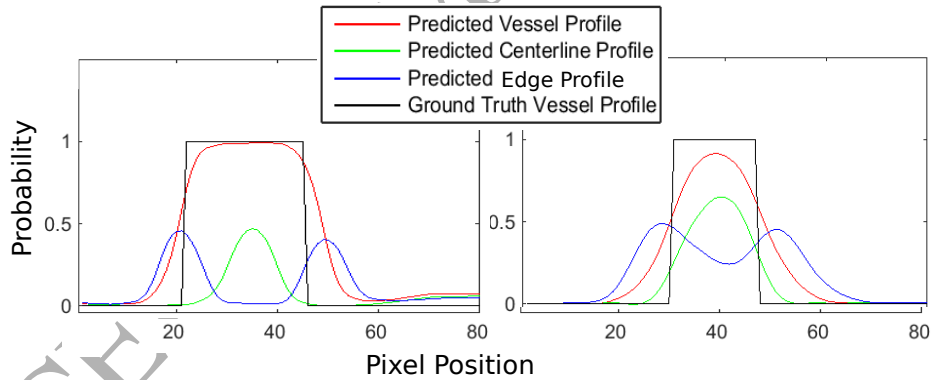


Figure 2: Probability profiles in the left side belong to the larger vessel and those in the right side to the thinner one in Figure 9. (Best viewed in color.)

The probability profiles are sampled from an arbitrary search region. Apart from Yin *et al.* [19] who used an adaptable semi-elliptical search curve for each iteration, our search region contains hundreds of adaptable lines (see Figure 3(a)). The spatial distribution of these lines is driven by the normally dis-

tributed noise in (6)-(9). The ends of a line correspond to hypothesised edge locations, calculated with (11) with respect to a parameter set sampled from the prior probability distribution. A search line is divided into 4 segments of equal length by sampling at 3 predetermined locations. The locations were selected in a way that they can capture the overall shape of the probability profiles illustrated in Figure 2.




Figure 3: (a) Some hypotheses over a vessel part sampled from the prior probability distribution according to (6)-(9). Each search line accounts for a set of hypothesized vessel parameters (e.g. centerline location, width and direction) and is named with H . In order to show hypothetical vessel edges represented by each hypothesis, the lines are accompanied with perpendicular short lines at hypothetical edges. (b) Edges $E_{k+1}^{L,R}$ and the sampling locations χ_j ; $j = 1 : 3$ in (10) are shown on the line H_n . The distance between neighbouring sample locations given a particular hypothesis is the same and equal to one fourth of the length of the line $|L|$. (Figure best viewed in color.)

The observation model, providing the likelihood of each hypothesis, is given in (10):

$$\begin{aligned}
P(y_{k+1}|\mathbf{z}_{k+1}) &= P'_e \cdot P'_c \cdot P_c(A_{k+1}) \cdot P_s & (10) \\
P'_e &= \prod_{i=\{L,R\}} P_e(E_{k+1}^i) \cdot (1 - P_c(E_{k+1}^i)) \\
P'_c &= \prod_{j=\{1,2,3\}} P_c(\chi = \chi_j) \cdot (1 - P_e(\chi = \chi_j)) \\
P_s &= \left| z_{k+1}^D \cdot \vec{E}(A_{k+1}) \right|
\end{aligned}$$

where P_e and P_c are profiles respectively obtained from edge and centerline
210 probability maps between hypothesized edge locations E_{k+1}^L and E_{k+1}^R . P_s
denotes the similarity of hypothesized vessel direction D_{k+1} to eigenvector
 $\vec{E}(A_{k+1})$ at hypothesized arbitrary interior location A_{k+1} . Sampling locations
from the probability profiles are $\chi_1 = \frac{(3E_{k+1}^L + E_{k+1}^R)}{4}$, $\chi_2 = \frac{(E_{k+1}^L + E_{k+1}^R)}{2}$ and
 $\chi_3 = \frac{(E_{k+1}^L + 3E_{k+1}^R)}{4}$ (see Figure 3(b)).

215 Equation (10) has three main components to be maximized: P'_e , P'_c and P_s .
The first one calculates the probability of edge locations and increases when the
edge estimates of the tracker become closer to vessel edge locations, where edge
probabilities and the complements of centerline probabilities are the maximum.
The second one aims to maximize centerline probability in order to make sure
220 of the centerline estimates of the tracker to be inside the vessel under tracking.
This component is important in terms of avoiding the tracker tracing boundaries
of two different vessels that are in close proximity. When the centerline esti-
mates of the tracker are inside the vessel, P'_c is far larger than when it is between
vessels. This is due to the centerline probabilities and the complement of edge
225 probabilities being large inside the vessel and almost zero outside the vessel.
The third component measures the similarity between an eigenvector indicating
the direction of the vessel and the direction hypothesized by the tracker. Both
vectors are obtained from the vessel interior probability map at hypothesized
centerline locations. Even if the hypothesized edge locations fall on the edges
230 of a vessel, this does not guarantee that the tracker will estimate accurate ge-
ometry parameters unless the orientation of the probability profiles consistently

aligns with the orientation of the vessel cross-section. Therefore, the term P_s in (10) contributes to the consistency of width estimations by assigning larger likelihoods to hypothesized edge locations that have similar alignment to the vessel cross-section.

2.4.1. Estimating Vessel Edges

After calculating the expectation of the posterior probability distribution with (5), the locations $E_{k+1}^{L,R}$ for the left and right edges of a vessel may then be calculated using the state parameters according to

$$E_{k+1}^{L,R} = A_{k+1} \pm (w^L, w^R)_{k+1} \cdot (z_{k+1}^D)_\perp \quad (11)$$

2.4.2. Strong and Weak Hypotheses

Figure 4 exemplifies the likelihoods of one strong and one weak hypothesis. The strong hypothesis estimates edge locations closer to the peaks of the edge probability profile, while the weak one predicts one of the edges outside the vessel and the other inside the vessel. In order to simplify the comparison of their likelihoods, two assumptions are made for both hypotheses: (i) the orientation of hypothesized vessel cross-sections is the same, (ii) hypothesized centerline locations are in the middle of the hypothesized edge locations. Considering the heights of the arrows denoting probabilities, it is obvious that the tracker assigns significantly far lower likelihood to the weak hypothesis (Figure 4 (b)) than the strong hypothesis (Figure 4 (a)) by reducing its contribution to the expectation of the posterior probability distribution; it is this weighting that ultimately determines the eventual estimate of the geometry parameters.

Figure 5 illustrates three cases, where (i) the search line in blue finds relatively better edge locations, (ii) the search line in red is located between edges of different vessels and (iii) the search line in green is oriented parallel to the vessel. Among these search lines, the blue one has significantly larger likelihood. Also, the likelihood of the red line is much lower than that of the green line, which indicates that the proposed observation model can convincingly remarkably discriminate edges of the same vessel from ~~than~~ those of different vessels.

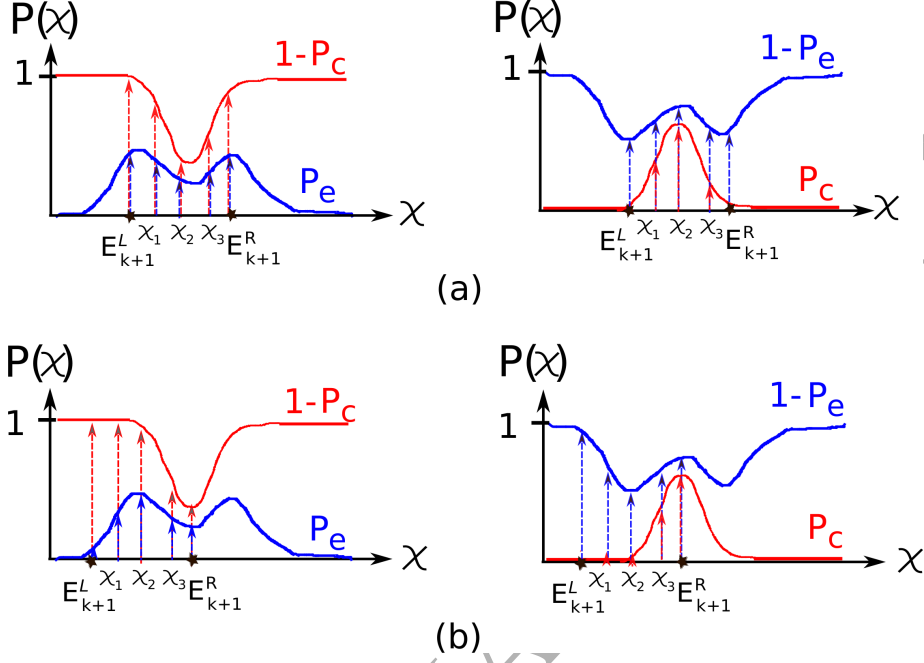


Figure 4: The calculation of likelihoods of two hypotheses, on (a) for a strong and (b) for a weak hypothesis, which are shown with H_4 and H_8 in Figure 3(a) consecutively. E_{k+1}^L and E_{k+1}^R denote hypothesized edge locations and χ_1, χ_2 and χ_3 sampling locations in (10). centerline probability profile $P(c)$ and its complement $1 - P(c)$ are shown with red solid lines, and samples from these profiles with red arrows. Similarly, an edge profile $P(e)$ and its complement $1 - P(e)$ are demonstrated with blue solid lines, and samples with blue arrows. The right and left plots respectively represent the calculation of P'_e and P'_c in (10). Ignoring the effect of $P_c(A_{k+1}) \cdot P_s$ in (10), the likelihood of the strong hypothesis is a million times larger than that of the weak hypothesis. (Figure best viewed in color.)

2.4.3. Generating Probability Maps For the Observation Model

In this study, we aim to generate probability maps of vessel interior, centerline and edge pixels with a single network. The latter two probability maps are directly used in the likelihood calculation in (10) while the former one is utilised to calculate eigenvectors to estimate vessel directions. The network selected was a specific version of Deep Belief Nets (DBNs) [28], trained to transform one image to another. With this approach, for example, fundus images can be converted to their vessel probability maps, a way of describing the segmenta-




Figure 5: Search lines and their likelihoods:(i) the blue search line is between edges of the vessel under tracking, (ii) the red search line is between edges of two vessels and (iii) the green search line is inside the vessel but oriented parallel to it. (b) Corresponding edge P_e and centerline P_c probability profiles are indicated by dashed and solid lines consecutively, and with colors matching those of (a). The likelihood is $1.00e^{-5}$ for the blue one, $9.15e^{-12}$ for the red one and $9.56e^{-8}$ for the green one. Note that P_s in (10) was not included in the likelihood calculation of these lines. (Figure best viewed in color.)

265 tion task. This network was observed to require considerably less training time but producing comparable segmentation performance on the detection of vessel interior pixels in pilot experiments (a performance comparison can be found in Section 4.2), when compared to alternative approaches to the same task (e.g. U-Net [29] and other CNN based segmentation methods [30, 31]).

270 *Training the network.* DBNs are often initially trained as a deep feature extractor in a generative manner, then fine-tuned with stochastic gradient descent to accomplish a target task, usually classification. In this study, a DBN is initially trained in a generative way to learn the relation between fundus images and their label maps (e.g. vessel interior, centerline and edges), which may be expressed as $f : \{I_f, I_i, I_c, I_e\} \rightarrow \{\overline{I_f, I_i, I_c, I_e}\}$, where $f(\cdot)$ is a function to learn the representation of its input. $\{I_f, I_i, I_c, I_e\}$ represents the concatenation of fundus images, I_f , vessel interior, I_i , centerline, I_c and edge, I_e label maps, respectively, and $\{\overline{\cdot}\}$ denotes the joint representation learned by the DBN.

The concatenation of these images is demonstrated in Figure 6, where the concatenated images correspond to a sample in the training dataset. Then, the
 280 DBN is fine-tuned to transform fundus images to their label maps by learning a function $g : I_f \rightarrow \{I_i, I_c, I_e\}$, where $g(\cdot)$ converts a fundus image to a trio of probability maps representing vessel interior, I_i , centerline I_c and edge I_e .

Denoising. The interaction between pixels at the same locations, but belong-
 285 ing to either the fundus image or one of its label maps, can be more efficiently learned by combining the generative training of DBNs with the denoising proposed by Vincent *et al.* [32]. The spirit of the denoising is to hide some information in the training data on which a network is trained, encouraging the network to predict the missing information in the training data. Sticking to this
 290 spirit, we randomly replace either a fundus patch or its label maps with zeros in a training sample, but motivate the network to estimate unaltered pixel values in the sample; $f(\{0, I_i, I_c, I_e\} \text{ or } \{I_f, 0\}) = \{I_f, I_i, I_c, I_e\}$. Because this type of denoising is applied at pixel level, we call it image-wise denoising. Figure 7(a) shows how to combine the image-wise denoising with the training of the first
 295 hidden layer of a DBN. The second and upper hidden layers of the DBN can be trained by applying denoising in a unit-wise manner as originally proposed [32], as demonstrated in Figure 7(b). Denoising is only introduced during generative training.

After completing the training, the DBN layer-wise (see Figure 8(a)), its
 300 weights connecting layers are “unfolded” [33], resulting in a deep autoencoder (see Figure 8(b)). Finally, the unfolded DBN is modified by removing weights not contributing to the image transformation task (see Figure 8(c)) and, the modified network is fine-tuned with a simple stochastic gradient descent algorithm with L_2 loss.




Figure 6: Generating noisy training samples: A training sample consists of a fundus image and its vessel interior, centerline and edge label maps. In order to generate noisy training samples, a sample can be multiplied with $\{0, 1, 1, 1\}$ or $\{1, 0, 0, 0\}$, where 1 is represented with white squares while 0 with black squares.

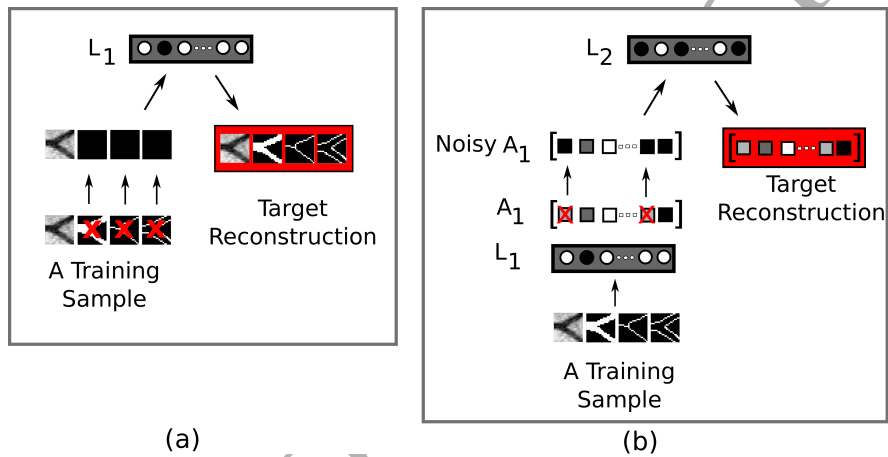


Figure 7: Integrating denoising with the training of a DBN: (a) The training of the first hidden layer L_1 of the DBN, where denoising is applied image-wise. Red crosses show images to replace with zero for denoising; (b) The training of the second hidden layer L_2 of the DBN, where denoising is applied unit-wise. Squares illustrate the activations of units in the first hidden layer when they are fed with a training sample. The vector of activations is denoted by A_1 . The units whose activations are suppressed for denoising are shown with red crosses. This type of denoising can be applied to following layers of the DBN (e.g. the third hidden layer). (Figure best viewed in color.)

305 3. Experimental Setup and Results

3.1. Dataset

The REVIEW dataset [34] contains images collected in the diabetic retinopathy clinic at Sunderland Eye Infirmary during clinical routine. This dataset has 4 sub-datasets: the high resolution image set (HRIS), the vascular disease image




Figure 8: (a) The DBN trained in a generative way as demonstrated in Figure 7, (b) Unfolding the DBN in (a), (c) Modifying the unfolded DBN in (b) for fine-tuning.

310 set (VDIS), the central light reflex image set (CLRIS) and the kick point image set (KPIS). Some properties of the sub-datasets are summarized in Table 1.

Table 1: The properties of the sub-datasets in the REVIEW dataset.

Name	Camera	FOV	Resolution (pixels)	No. Of Images		No. Of Ves. Segment	No. Of Profiles
				Healthy	Diseased		
HRIS	Cannon 60 UV	60°	3584 x 2438	-	4	90	2368
VDIS	Zeiss & JVC 3CCD	50°	1360 x 1024	2	6	79	2249
CLRIS	Zeiss FF 450 & JVC 3CCD	50°	2160 x 1440	-	2	21	285
KPIS	Canon 60UV	60°	3300 x 2600	2	-	3	164

HRIS and KPIS provide a performance evaluation at a sub-pixel accuracy because images in these datasets were down-sampled by a factor of 4 after receiving estimates of widths by observers; therefore, the accuracy of width
 315 estimation was limited to an error of ± 0.25 pixels. In HRIS, two of images were graded with severe and one with moderate and the other one with minimal

non-proliferative retinopathy. Apart from other sub-datasets, vessel profiles in KPIS were marked through detecting kick points on thicker and non-tortuous vessel segments between bifurcation locations.

320 VDIS contains images with pathologies and higher noise. The images were observed to have a larger variance of vessel profiles provided by observers [34]. 6 of the images were captured from patients with various types of Diabetic Retinopathy. CLRIS consists of images showing signs of atherosclerosis, which are the exaggeration of the central light reflex and changes in vessel walls.

325 *Marking Vessel Edges.* Vessel profiles were detected by three experts, two of them with an experience in retinal vessel analysis and the other one trained to locate vessel edges. The experts independently located vessel edges at the same vessel segments. Then, the edge locations were edited by an algorithm to ensure even spaces between neighbor profiles [34]. The average of edge locations marked
330 by the experts is used as reference data to reduce inter-subject variability on detected vessel boundaries.

3.2. Experimental Settings

3.2.1. Overview

Although the performance of the tracker was evaluated on the REVIEW
335 dataset, this dataset does not include the vessel maps required for segmentation. In order to deal with the problem, we used knowledge transfer, where we trained the network with a well-known fundus image dataset, the DRIVE [26] then used the trained network to generate label maps for the REVIEW dataset. The main challenge with this approach is that the resolution of the REVIEW dataset is
340 much larger than that of the DRIVE. Therefore, the resolution of the REVIEW dataset was reduced by down-sampling before being fed into the network and its resolution was brought the original level prior to being used for tracking. This solution was acceptable for the present research, because the probability maps
345 of the posterior probability distribution with relative fitness of the hypothesized

geometry parameters. The sub-sampling factors were 2 for VDIS, 3 for CLRIS and 4 for HRIS. No sub-sampling was applied to KPIS.

The centerline and edge images required for training were generated by applying a standard thinning algorithm [35] and a Prewitt edge detection algorithm [36] respectively to the reference vessel maps. The training of the network was realized patch-wise, where we randomly selected patches from each fundus image map and its corresponding vessel interior, centerline and edge label maps, at the same locations. The size of an image patch was 16 by 16 pixels. Because the DRIVE has two vessel maps for each image in its training set, we used the maps produced by the first expert as reference, complying to the general practice [26].

In order to increase the representation of vessel pixels in the training dataset, we performed denser patch sampling inside the Field-of-View (FOV) masks as follows: initially, we multiplied FOV masks with the green channels of the fundus images, so pixels outside FOV regions became zero. After randomly and densely sampling image patches, we removed patches completely outside the FOV masks. The number of the patches in the final stage was roughly 1,800,000. The fundus image patches were normalized patch-wise in the range of $[0, 1]$, which was visually observed to better reveal vessels on patches with lower contrast.

3.2.2. Network Parameters

Because the network goes through different types of training, the configuration of the network is altered accordingly; this is typical of DBN methods [37, 28]. Initially, the network has the input layer of $256 \times 4 = 1024$ units and 3 hidden layers, with each having 400 units during the generative training. After 'unfolding' and reshaping it for fine-tuning, the network contains the input layer of 256 units, the output layer of $256 \times 3 = 768$, units and 5 hidden layers of 400 units.

The network was trained with mini-batches of 100 sets of image patches for both the generative training and fine-tuning. In the generative training,

the network weights were initialized by sampling from a normal distribution $\mathcal{N}(0, 0.001)$. Then, the network was trained for 50 epochs with a learning rate of 0.005. A momentum of 0.5 was initially used in the first 5 epochs. Later, this number was increased to 0.9. The learning rate for fine-tuning was 0.08 for 120
 380 epochs. The squeezing function for all layers was sigmoid.

3.2.3. Tracking Parameters

The number of particles for tracking was 700 and step size was 2 pixel. The standard deviations of the normal distributed noise variables, with zero mean, ϵ_D , ϵ_A , ϵ_w^R and ϵ_w^L in (6)-(9) were experimentally determined. The standard
 385 deviations of the noise distributions were respectively 0.1 for ϵ_D , 0.5 for ϵ_w^R and ϵ_w^L , and 0.5 for ϵ_A . When $k = 0$, vessel centerline and width parameters of each vessel segment were initialized with the reference data in the REVIEW dataset. In the same way, vessel direction is assigned the direction from the first reference centerline location to second one, similar to [20]. Therefore, the prior
 390 probability distribution for each vessel parameter at $k = 0$ becomes a normal distribution with a fixed standard deviation given above and with a mean value depending on the reference data. In the case of ground truth data not to be available, an expert can pick a location inside a vessel and can roughly estimate vessel width at the location to initialise tracking. The direction of the vessel
 395 can be assumed to be the eigenvector with larger eigenvalue [38], obtained as a result of eigen-analysis of Hessian matrix of related vessel interior probability map. We prefer vessel interior probability maps over original fundus images because the former ones are usually far less noisier than the latter ones.

3.3. Evaluation Criteria

We used the standard deviation, σ_w , of error in estimating vessel width, (12), and mean absolute error, Δ_w , in (13), which are commonly used to assess estimates of vessel widths to reference widths [11, 12, 14, 15].

$$\sigma_w = \text{std}(w_r - w_e) \quad (12)$$

$$\Delta_w = \text{mean}(|w_r - w_e|) \quad (13)$$

400 where w_r denotes reference width while w_e represents the estimate.

In vessel width estimation, relative widths, rather than actual widths, over a fundus image are usually used in assessment [9]. One good argument for this is to have relative assessments of vessel diameter changes (within subject) that are approximately independent of optical magnification. Al-Diri *et al.* [34] 405 also points out that consistent biases or scale factors in measurement can be removed by simple linear transforms of width estimates. In line with other studies, we use σ_w to assess the success of a method, rather than Δ_w because of the possible effect of constant bias on the Δ_w measure. Therefore, even though Δ_w is reported in following experiments, it is primarily for completeness.

410 In addition to σ_w and Δ_w , the percentage of vessel profiles whose width estimates are meaningful was also reported in previous studies [9, 10, 12, 14, 15, 39]. This measure indicates the ability of a method to deal with different challenges (e.g. noise, pathology) without compromising the performance of width estimation [15].

415 4. Results

4.1. Generated Probability Maps by the Network

The trained network is used to produce probability maps for vessel interior, centerline and boundaries as shown in Figure 9, where two image patches, each containing a thick or a thin vessel, are shown, along with their aforementioned 420 probability maps. Because the patches are normalized in the range of $[0, 1]$, the contrast of vessels seems very similar despite the significant difference between their thickness. Thus, despite varying vessel thickness and noise levels, the network-generated smooth probability maps are consistent with the ground truth vessel masks.

425 4.2. Segmentation

Because there is no direct way to assess how well probability maps represent vessel parts in an image, we consider the performance of the network on vessel

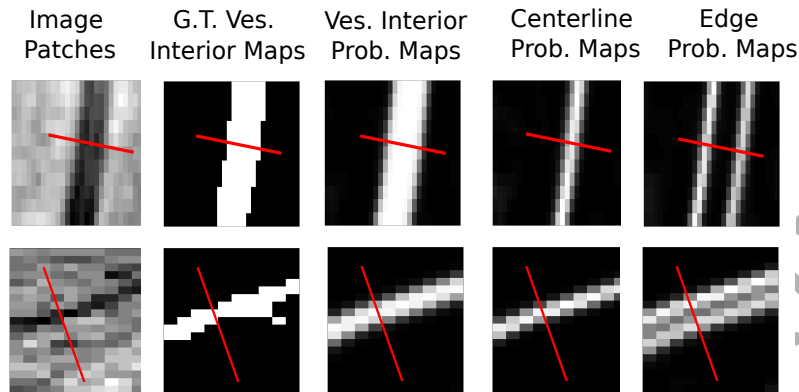


Figure 9: Fundus image patches and the typical probability maps generated by the deep network: the top row belongs to a large vessel and bottom row to a fine vessel. Columns from the left to right side show a normalized fundus image patch, its manually labeled vessel interior map, generated vessel interior, centerline and edge probability maps. Red lines show the locations that profiles in Figure 2 are obtained. (Best viewed in color.)

interior segmentation as an overall indicator of its performance on probability maps and measured its performance with standard performance metrics of accuracy, sensitivity, specificity and Area-Under-Curve (AUC). Apart from the former three metrics, AUC is calculated *independently* of a subjective threshold, and shows the certainty of the method on discriminating vessel pixels from non-vessel pixels; in other words, it is an indication of the robustness and quality of the vessel interior probability maps. Table 2 compares the performance of the network and that of recent state of the art methods using supervised methods, regarding AUC, accuracy, sensitivity and specificity, whose definitions can be found in [40]. Referring to the table, the present study is among the best performing studies regarding AUC.

4.3. Width Estimation

We now assess the performance of the tracker on the estimation of vessel width, in the REVIEW dataset. The conventional way to evaluate the reliability of vessel width estimation is to compare estimated widths to the reference ones over predetermined profiles where the estimated and reference profiles either

Table 2: The comparison of segmentation performance of the proposed method with the performances of state of the art methods on the DRIVE dataset.

Year	Method	AUC	Accuracy	Sensitivity	Specificity
2017	The proposed method	0.9761	0.9542	0.7752	0.9800
2017	U-net [29] *	0.9790	-	-	-
2016	Liskowski and Krawiec [30]	0.9710	0.9515	0.7520	0.9806
2015	Wang <i>et al.</i> [41]	0.9475	0.9767	0.8173	0.9733
2015	Li <i>et al.</i> [40]	0.9738	0.9527	0.7569	0.9816
2014	Cheng <i>et al.</i> [31]	0.9648	0.9474	0.7252	0.9798
2012	Fraz <i>et al.</i> [42]	0.9747	0.9480	0.7406	0.9807

* The result was taken from <https://github.com/orobix/retina-unet> not from [29].

share only the same centerline locations or both centerlines and orientations
 445 [10, 11, 12, 14, 15]. However, such a one to one comparison was not possible
 for the current method because both centerline locations and the orientations
 of the vessel cross-sections were autonomously estimated by the tracker.

In order to reduce a potential discrepancy between the locations of widths
 estimated by the method and those given in the reference data as much as
 450 possible, we used bi-cubic spline interpolation to sample 100 locations, from both
 reference and estimated widths. Then, the evaluation criteria were calculated
 over interpolated values. However, it should be noted that because the locations
 of interpolated profiles from both reference and estimations were not guaranteed
 to overlap, this evaluation also has potential limitations. Also, vessel segments
 455 with less than 2 reference profiles could not be traced because these profiles are
 either used to start tracking or to stop it.

4.3.1. Evaluation Over A Dataset

Traditionally, the performance of width estimation is compared over all pro-
 files in the sub-dataset, regardless of vessel identification [11, 12, 14, 15]. Table
 460 3 compares the performance of the proposed method to that of previous studies,
 reporting Δ_w , σ_w , and the percentage of meaningful width estimations. The

proposed method is able to obtain meaningful estimates of width for the entire set of vessel profiles, whereas the majority of methods fails to predict plausible widths for some profiles in CLRIS, HRIS and VDIS; this includes the best performing methods of Zhang *et al.* [20] and Yin *et al.* [19].

The best σ_w results were obtained by Araújo *et al.* [11], who used a supervised model-fitting method. However, because they did not present the percentage of meaningful width estimations, it is not clear if the results reflect the performance – in terms of σ_w – of overall sub-datasets, or only those corresponding to the successfully estimated profiles. The performance of Araújo *et al.*'s method [11] is followed by that of tracking methods [19, 20], whose performance mostly surpasses that of other supervised and unsupervised methods [9, 10, 14, 15, 39].

Amongst tracking approaches (see Table 3), the performance of the proposed method closely follows that of Zhang *et al.*'s method in the measure of σ_w . This dataset has been generally found challenging by many methods [20], due to containing the central light reflex along vessel segments. For the HRIS dataset, the σ_w values of the proposed method are slightly worse than those of Zhang *et al.* [20] and of Yin *et al.* [19].

For VDIS, the proposed method has larger σ_w values than that of Zhang *et al.* [20] and that of Yin *et al.* [19], but the proposed method predicts all widths for this dataset while Zhang *et al.* [20] and Yin *et al.* successfully estimated only 94.2% and 92.7% of widths, respectively. For KPIS, the proposed method gives the lowest σ_w values among tracking methods, which is in the same range as that among supervised and unsupervised methods.

4.3.2. Evaluation For Each Vessel Segment

In contrast to the traditional performance evaluation, which implicitly accepts that vessel profiles are independently sampled from a dataset and summarizes the performance with a single number, we also assessed the performance of the proposed approach for each vessel segment in REVIEW dataset, based on the fact that profiles selected from the same vessel segment are highly prob-

Table 3: The performance comparison of the proposed method with those of previous studies on the REVIEW dataset, where % represents the percentage of meaningful width estimations

Method	Year	CLRIS			HRIS			VDIS			KPIS		
		Δ_w	σ_w	%	Δ_w	σ_w	%	Δ_w	σ_w	%	Δ_w	σ_w	%
Observer 1		0.61	0.57		0.23	0.29		0.35	0.54		0.34	0.42	
Observer 2		0.11	0.70		0	0.26		0.06	0.62		0.11	0.32	
Observer 3		0.72	0.57		0.23	0.29		0.3	0.67		0.23	0.33	
The proposed method	2017	0.92	1.15	100	0.32	0.40	100	1.01	0.85	100	1.51	0.34	100
Zhang <i>et al.</i> [20]	2014	0.37	1.13	98.3	0.08	0.30	100	1.37	0.59	94.2	0.74	0.37	100
Tracking M. Yin <i>et al.</i> * [19]	2012	0.77	1.41	93.1	0.01	0.39	96.2	1.41	0.56	92.7	0.69	0.43	100
Zhou <i>et al.</i> ** [39]	1994	7.5	4.14	98.6	0.54	0.90	99.6	3.07	2.11	99.9	2.57	0.4	100
Araújo <i>et al.</i> [11]	2017	0.01	0.56	-	0.00	0.22	-	0.00	0.69	-	0.00	0.30	-
Aliahmad and Kumar [12]	2016	0.33	1.56	98	0.24	0.65	99.4	0.45	1.14	97.8	0.72	0.45	100
Supervised M. Lupasçu <i>et al.</i> [10]	2013	0.00	1.15	100	0.00	0.44	100	0.02	1.07	100	0.02	0.32	100
Xu <i>et al.</i> [14]	2011	0.08	1.78	94.3	0.21	0.567	100	0.53	1.43	96	1.14	0.67	99.4
Al-Diri <i>et al.</i> [15]	2009	1.9	1.47	93	0.28	0.42	99.7	0.05	0.77	99.6	0.96	0.33	100
Unsupervised M. Lowell <i>et al.</i> ** [9]	2004	6.8	6.02	26.7	0.17	0.70	98.9	2.26	1.33	77.2	1.65	0.34	100

* These results were taken from [20].

** These results were taken from [12].

able to have similar widths, and also similar types of problems, such as the presence of the central light reflex. Therefore, we calculated σ_w for each vessel segment and obtained σ_w distributions for each sub-dataset to observe if poor or good performance of the method on the sub-dataset may be related to the performance on specific vessel segments.

Figure 10 shows the distributions of the σ_w values of vessel widths produced for CLRIS, HRIS and VDIS with box-plots. Because KPIS does not have a sufficient number of vessel segments for this demonstration, its results are, instead, summarized in the text. The figure shows three outliers for both CLRIS and HRIS, and four outliers for VDIS, with σ_w of over 1 pixel, which indicates abnormal disagreements between reference and estimated widths on the vessel segments responsible for the outliers. It should be noted that this information was not revealed in Table 3. The values of σ_w for CLRIS, HRIS and VDIS in the table are, respectively, 1.15, 0.40 and 0.85 pixels, which are far larger than the medians of the σ_w distributions illustrated in the figure. Also, according to the figure, the median σ_w for CLRIS and that for VDIS are almost the same; however, the σ_w values of outliers in CLRIS are much larger than those

in VDIS. Obviously, these extreme outliers can be taken to be responsible for
 510 CLRIS having larger σ_w values than that of VDIS in Table 3.

In addition to results presented in the figure, KPIS was observed to have
 σ_w values of 0.35, 0.31 for the first and second vessel segment in the first image
 respectively and, that of 0.37 for the single segment in the second image. These
 results are mainly consistent with those for σ_w given in Table 3.

515 We argue that the proposed way of calculating the evaluation criteria is more
 appropriate than the traditional way [11, 12, 14, 15], because the former can
 identify vessel segments for which a given method yields significantly different
 widths from reference data. The identification of the challenging vessel segments
 in this manner is useful when developing new approaches to tracking or width
 520 estimation. We will closely examine vessel segments that have been found to
 be challenging in order to appreciate the sources of disagreement between our
 estimations and the reference data.

Outliers in the Box-Plot. The vessel segments producing largest σ_w values for
 HRIS, VDIS and KPIS in Figure 10 are demonstrated with estimated widths
 525 and reference ones on both fundus images and edge probability maps in Figure
 11.

Figure 11(a) shows a vessel segment from CLRIS, with a σ_w value of 1.3. The
 subtle change on vessel width seems not to be captured by human observers, in
 contrast to the proposed method. This situation can also be observed in Figure
 530 11(b), which demonstrates an image from VDIS, with a σ_w value of 0.98. This
 vessel segment can be characterized with abnormal width changes. Similar to
 former image pair, vessel widths are estimated reasonably consistently by the
 proposed method, following the changes on actual vessel width in pathologies.

The superiority of the current method to human observers becomes more
 535 obvious in Figure 11(c), where an image pair from HRIS, with the σ_w value of
 0.98, is demonstrated. The vessel edges in this figure are mis-detected by the
 observers.

Apart from imperfection of human observers to accurately estimate vessel

widths, the discrepancy between estimates and reference data can also be due to
 540 predicting widths along slightly different profiles as appearing in Figure 11(d).
 This figure illustrates a vessel segment from VDIS, with a σ_w value of 0.97.




Figure 10: The distributions of the σ_w values for vessel width estimates calculated for individual vessel segments in CLRIS, HRIS and VDIS. The maximum length of each whisker is 1.5 times of the interquartile range of related distribution. (Best viewed in color.)

4.3.3. Challenging Vessels

In the literature, some vessels have been found more challenging than others; such as (i) those with the central light reflex [9, 15], (ii) close vessels [15, 43],
 545 (iii) those in junction regions [39, 43] and (iv) those with high curvature or, (v) low contrast or high noise. We now discuss each of these in turn.

(i) The central light reflex is a bright strip around the vessel centerline, which may be confused with vessel edges by some methods [44]. In the presence of this reflection, the intensity profile across vessels deviates from a Gaussian-like appearance, which may be compensated for by increasing the complexity of models for the intensity profile, such by using multiple Gaussian functions
 550 [10, 11, 13]. Another way may be to combine a method using the intensity profile as the main information source with additional sources, for example line detector responses [20]. On the other hand, the proposed method does not need to take
 555 extra measures to deal with the reflection, which is naturally suppressed during the network training. Figure 12 shows a vessel with the central light reflex from CLRIS and the probability maps generated by the network, where no sign for




Figure 11: Estimated vessel profile locations (green lines) and interpolated locations (red lines). The interpolated locations provide a means of comparing our width estimates with prior work (see text for details), and are shown on the vessel segments with the worst performance of the proposed method: (a) segment 2 – 8 in CLRIS, (b) segment 8 – 3 in VDIS, (c) segment 1 – 26 in HRIS, (d) segment 8 – 10 in VDIS. The left and right images for each image couple respectively show a fundus image and its edge probability map. The images in the top row have the size of 60×60 pixels and 122×122 pixels respectively while those in the bottom row have the size of 66×66 pixels and 100 by 100 pixels consecutively. (Figure best viewed in color.)

the reflection appears. Figure 13(a)-(c) show estimated and reference profiles for vessels with the reflection. As may be seen, the presence of the reflection
 560 does not degrade the consistency of width estimations.

(ii) In some cases, vessels can be close to each other, which makes it difficult to estimate vessel geometry: it may not possible to identify vessel edges due to the presence of the other nearby vessels [15, 43]. However, the proposed method can be observed to successfully trace a vessel with the central light reflex and low
 565 contrast, without being distracted by a closely passing one, as shown in Figure 13(c). When the figure is examined closely, it appears that the edge probability map of the traced vessel is affected by the nearby vessel to a large extent, which




Figure 12: The intensity profile $I(x,y)$ of a vessel with the central light reflex, a part of the vessel in Figure 13(a), and its probability profiles $P_i(x,y)$, $P_c(x,y)$ and $P_e(x,y)$, which respectively denote vessel interior, centerline and edge locations.

is manifested with far lower and diffused probabilities for the left edge of the traced vessel. Despite the large uncertainty along this edge, the tracker manages
 570 to identify both vessel boundaries correctly and tracks the vessel without any disruption. This may be attributed to the existence of the prior probability distribution keeping the memory of the previously traced path.

(iii) Vessels in junction regions are difficult to analyze because vessel boundaries may become completely indistinguishable [39, 43]. The edge pixels in
 575 junctions can have lower edge probabilities, which is visible in Figure 13(e). Because of the prior information implicit in the tracking process, the region of lower edge probabilities can be traced confidently.

(iv) Curvy vessels may pose a big challenge for tracking methods due to their fast changing directions. However, vessels with high curvature are observed to
 580 not pose significant problems for the proposed method, because of the use of estimated vessel direction, and incorporated in the P_s term of 10). Figure 13(f) shows estimated widths for a curvy vessel.

(v) Human observers can also fail at estimating the right locations for vessel edges, particularly, if the contrast of vessels is poor or their calibers are small.

585 Figure 13(d),(g)-(h) illustrate profiles located by the observers and the proposed
method. On fundus images, both reference and estimated locations seem ac-
ceptable to the naked eye. This may align with the inter-observer variability
in locating vessel boundaries on ground truth images, a problem acknowledged
in [26, 27]. However, considering the disagreement between the values in the
590 edge probability maps at the edge locations estimated by the observers (see
Figure 13(g)-(h)) and those of the proposed tracker, we argue that our method
estimates better edge locations in these images than human observers.

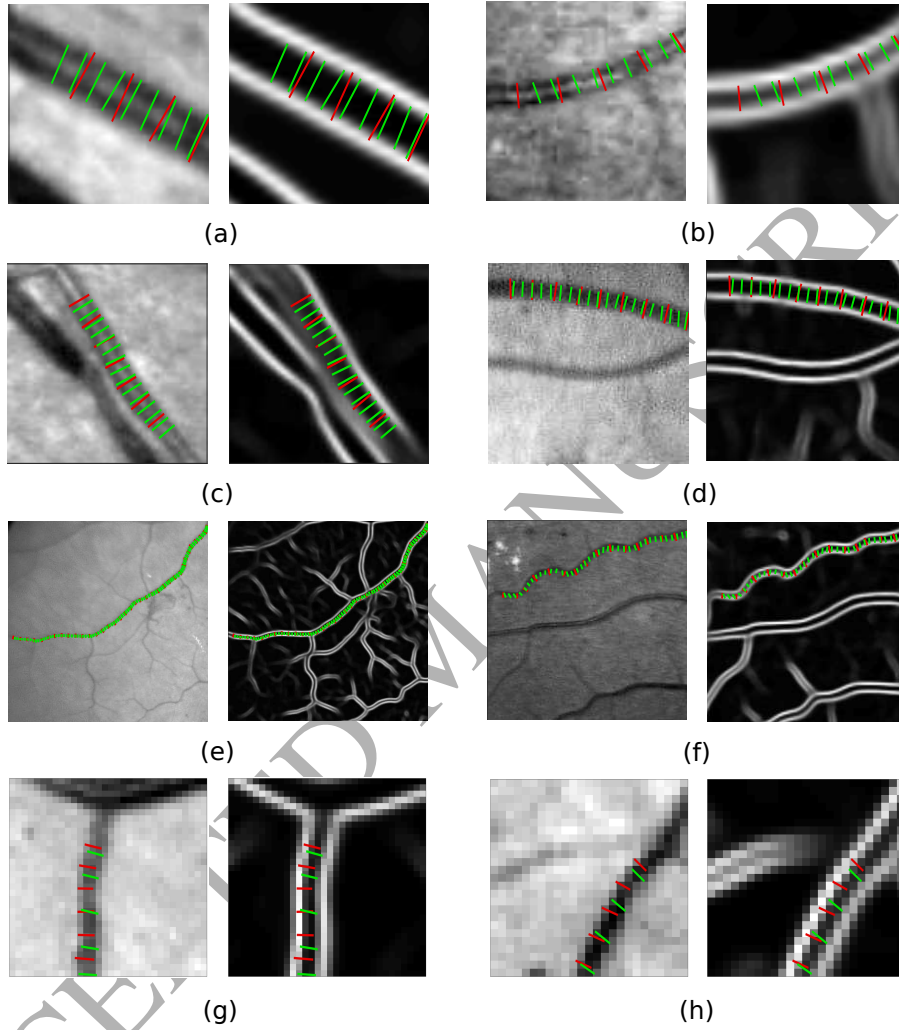


Figure 13: Estimated vessel profile locations (green lines) and interpolated locations (red lines, see text for details) for challenging vessels: (a)-(c) from CLRIS, (d)-(f) from VDIS and (g)-(h) from HRIS. The left and right images for each image couple respectively show a fundus image and its edge probability map. The sizes of images are respectively 63×63 pixels in (a), 64×64 pixels in (b), 98×98 pixels in (c), 130×130 pixels in (d), 424×424 pixels in (e), 207×207 pixels in (f), 31×31 pixels in (g) and 25×25 pixels in (h). (Best viewed in color.)

5. Conclusion

In this study, we propose a Bayesian method ¹ to estimate vessel geometry parameters, and evaluated the performance of the method on width estimation on the REVIEW dataset. In contrast to previous methods, which have used the intensity profiles across vessel edges for parameter estimation, we utilized probability profiles sampled from vessel interior, centerline and edge location probability maps generated by a single deep network. As far as we are aware, it is the first method for retinal vessel analysis – or any vascular data – that uses probability maps as inputs to a tracking and vessel width estimation process. The method addressed four challenges encountered in width estimation.

The first is that probability maps for vessel parameter estimation can better explain the uncertainty and subjectivity at detecting vessels and, particularly, edge locations, appearing in the ground truth data [26, 34]. Due to these maps serving a Bayesian method, particle filtering, the uncertainty in these maps is efficiently utilized for vessel parameter estimation. For instance, the proposed approach could make reasonable estimates of vessel width: even when there is not sufficient information available for vessel parameter estimation (e.g. in junction regions), or when the information is vague (e.g. for thin and low contrast vessels, and vessels in a close proximity to another vessel). Estimating all sets of vessel profiles in the REVIEW dataset reinforces the effectiveness of this approach. Moreover, having edge probability maps facilitated an evaluation of the consistency of reference vessel profiles.

Secondly, despite the lack of training data for vessel segmentation specific to the REVIEW dataset, the proposed method was able to generate useful probability maps for vessel geometry. Specifically, training the network with low resolution and almost healthy fundus images, provided in the DRIVE, was observed to produce sufficient quality of probability maps for the REVIEW

¹The code producing results reported in this paper can be found at <https://bitbucket.org/fzehra/a-recursive-bayesian-approach-to-describe-retinal-vasculature>.

620 dataset, which has high resolution and mostly pathological images. This success may be attributed to two factors. Firstly, the generalization capability of the network was adequate for the purpose of the presented method. Secondly, using particle filtering for vessel parameter estimation might compensate any imperfections in the probability maps.

625 Thirdly, to date, the performance of a method on width estimation has been assessed with the independent evaluation of the profiles. However, this approach ignores the spatial dependence of vessel profiles for particular vessel segments. In this study, we also assessed the performance of the method on individual vessel segments, which allowed us to immediately spot disagreements
630 between reference measurements in the REVIEW data and estimates from the proposed method at the level of vessel segments. Moreover, we could evaluate the reliability of the reference data to some extent: our analysis seems to have uncovered some errors in the reference data, revealed by our way of performance evaluation.

635 Finally, the proposed method can be viewed as being superior to methods based on supervised model fitting [10, 11, 12], whose performance strongly depends on the characteristics of the training datasets. If the datasets have any errors or bias in their reference data, the estimates made by the supervised methods [10, 11] are highly probable to have the same issues, despite having
640 close agreement with reference data. See [45] for another interesting approach that implicitly addresses the difficulty in constructing likelihood functions.

Currently, this work only considers tracing vessel segments for parameter estimation. However, we are working on a method to detect junction locations in fundus images, which will be integrated with the proposed method to trace
645 complete vessel trees. [More generally, the principle of using the probabilistic outputs of deep networks as inputs to probabilistic trackers seems a promising approach to address the challenge of constructing *ad hoc* likelihood functions for observation models.](#)

6. Acknowledgement

650 The authors thank to Republic of Turkey Ministry Of National Education
for their financial support.

References

References

- [1] J. Zhang, Y. Chen, E. Bekkers, M. Wang, B. Dashtbozorg, B. M. t. H.
655 Romeny, Retinal vessel delineation using a brain-inspired wavelet transform
and random forest, *Pattern Recognition* 69 (2017) 107–123.
- [2] Y. Q. Zhao, X. H. Wang, X. F. Wang, F. Y. Shih, Retinal vessels seg-
mentation based on level set and region growing, *Pattern Recognition* 47
(2014) 2437–2446.
- 660 [3] M. K. Ikram, C. Y. Cheung, M. Lorenzi, R. Klein, T. L. Jones, T. Y.
Wong, Retinal vascular caliber as a biomarker for diabetes microvascular
complications, *Diabetes Care* 36 (2013) 750–759.
- [4] J. Ding, K. L. Wai, K. McGeechan, Retinal vascular caliber and the de-
velopment of hypertension: a meta-analysis of individual participant data,
665 *Journal of Hypertension* 32 (2014) 207.
- [5] B. Gopinath, J. Chiha, A. J. Plant, A. Thiagalingam, G. Burlutsky,
P. Kover, G. Liew, P. Mitchell, Associations between retinal microvas-
cular structure and the severity and extent of coronary artery disease,
Atherosclerosis 236 (2014) 25–30.
- 670 [6] C. Y.-l. Cheung, Y. T. Ong, M. K. Ikram, S. Y. Ong, X. Li, S. Hilal, J.-A. S.
Catindig, N. Venketasubramanian, P. Yap, D. Seow, Microvascular network
alterations in the retina of patients with Alzheimer’s disease, *Alzheimer’s
& Dementia* 10 (2014) 135–142.

- [7] A. Sironi, E. Türetken, V. Lepetit, P. Fua, Multiscale centerline detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38 (2016) 1327–1341.
- [8] K. Daniel Chaves Viquez, O. Arandjelovic, A. Blaikie, I. Ae Hwang, Synthesising wider field images from narrow-field retinal video acquired using a low-cost direct ophthalmoscope (arlight) attached to a smartphone, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 90–98.
- [9] J. Lowell, A. Hunter, D. Steel, A. Basu, R. Ryder, R. L. Kennedy, Measurement of retinal vessel widths from fundus images based on 2-D modeling, *IEEE Transactions on Medical Imaging* 23 (2004) 1196–1204.
- [10] C. A. Lupaşcu, D. Tegolo, E. Trucco, Accurate estimation of retinal vessel width using bagged decision trees and an extended multiresolution hermite model, *Medical Image Analysis* 17 (2013) 1164–1180.
- [11] T. Araújo, A. M. Mendonça, A. Campilho, Estimation of retinal vessel caliber using model fitting and random forests, in: *Proc. SPIE Medical Imaging*, volume 10134, pp. K1–8.
- [12] B. Aliahmad, D. K. Kumar, Adaptive Higuchi’s dimension-based retinal vessel diameter measurement, in: *Proc. the IEEE Engineering in Medicine and Biology Society (EMBC)*, IEEE, pp. 1308–1311.
- [13] X. Gao, A. Bharath, A. Stanton, A. Hughes, N. Chapman, S. Thom, A method of vessel tracking for vessel diameter measurement on retinal images, in: *Proc. the International Conference on Image Processing (ICIP)*, volume 2, IEEE, pp. 881–884.
- [14] X. Xu, M. Niemeijer, Q. Song, M. Sonka, M. K. Garvin, J. M. Reinhardt, M. D. Abràmoff, Vessel boundary delineation on fundus images using graph-based approach, *IEEE Transactions on Medical Imaging* 30 (2011) 1184–1191.

- [15] B. Al-Diri, A. Hunter, D. Steel, An active contour model for segmenting and measuring retinal vessels, *IEEE Transactions on Medical imaging* 28 (2009) 1488–1497.
- 705 [16] Y. Salih, A. S. Malik, Comparison of stochastic filtering methods for 3d tracking, *Pattern Recognition* 44 (2011) 2711–2737.
- [17] P. Li, D. Wang, L. Wang, H. Lu, Deep visual tracking: Review and experimental comparison, *Pattern Recognition* 76 (2018) 323–338.
- [18] O. Chutatape, L. Zheng, S. M. Krishnan, Retinal blood vessel detection and tracking by matched Gaussian and Kalman filters, in: *Proc. the IEEE Engineering in Medicine and Biology Society (EMBC)*, volume 6, IEEE, pp. 3144–3149.
- 710 [19] Y. Yin, M. Adel, S. Bourenmane, Retinal vessel segmentation using a probabilistic tracking method, *Pattern Recognition* 45 (2012) 1235–1244.
- [20] J. Zhang, H. Li, Q. Nie, L. Cheng, A retinal vessel boundary tracking method based on Bayesian theory and multi-scale line detection, *Computerized Medical Imaging and Graphics* 38 (2014) 517–525.
- 715 [21] A. Wu, Z. Xu, M. Gao, M. Buty, D. J. Mollura, Deep vessel tracking: A generalized probabilistic approach via deep learning, in: *Proc. International Symposium on Biomedical Imaging (ISBI)*, IEEE, pp. 1363–1367.
- 720 [22] P. M. Djuric, J. H. Kotecha, J. Zhang, Y. Huang, T. Ghirmai, M. F. Bugallo, J. Miguez, Particle filtering, *IEEE signal processing magazine* 20 (2003) 19–38.
- [23] C. M. Grinstead, J. L. Snell, *Introduction to probability*, American Mathematical Soc., 2012.
- 725 [24] M. S. Arulampalam, S. Maskell, N. Gordon, T. Clapp, A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking, *IEEE Transactions on Signal Processing* 50 (2002) 174–188.

- 730 [25] A. Doucet, N. de Freitas, N. Gordon, An introduction to Sequential Monte Carlo methods, in: A. Doucet, N. de Freitas, N. Gordon (Eds.), *Sequential Monte Carlo Methods in Practice*, Springer New York, 2001, pp. 3–14.
- [26] J. Staal, M. D. Abràmoff, M. Niemeijer, M. A. Viergever, B. Van Ginneken, Ridge-based vessel segmentation in color images of the retina, *IEEE Transactions on Medical Imaging* 23 (2004) 501–509.
- 735 [27] A. Hoover, V. Kouznetsova, M. Goldbaum, Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response, *IEEE Transactions on Medical imaging* 19 (2000) 203–210.
- [28] I. Fasel, J. Berry, Deep belief networks for real-time extraction of tongue contours from ultrasound during speech, in: *Proc. International Conference on Pattern Recognition (ICPR)*, IEEE, pp. 1493–1496.
- 740 [29] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, pp. 234–241.
- 745 [30] P. Liskowski, K. Krawiec, Segmenting retinal blood vessels with deep neural networks, *IEEE Transactions on Medical Imaging* 35 (2016) 2369–2380.
- [31] E. Cheng, L. Du, Y. Wu, Y. J. Zhu, V. Megalooikonomou, H. Ling, Discriminative vessel segmentation in retinal images by fusing context-aware hybrid features, *Machine Vision and Applications* 25 (2014) 1779–1792.
- 750 [32] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, P.-A. Manzagol, Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion, *Journal of Machine Learning Research* 11 (2010) 3371–3408.
- 755 [33] G. Hinton, S. Osindero, Y.-W. Teh, A fast learning algorithm for deep belief nets, *Neural Computation* 18 (2006) 1527–1554.

- [34] B. Al-Diri, A. Hunter, D. Steel, M. Habib, T. Hudaib, S. Berry, Review-a reference data set for retinal vessel profiles, in: Proc. the IEEE Engineering in Medicine and Biology Society (EMBC), IEEE, pp. 2262–2265.
- [35] L. Lam, S.-W. Lee, C. Y. Suen, Thinning methodologies-a comprehensive survey, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14 (1992) 869–885.
- [36] G. Shrivakshan, C. Chandrasekar, A comparison of various edge detection techniques used in image processing, *IJCSI International Journal of Computer Science Issues* 9 (2012) 272–276.
- [37] G. E. Hinton, R. R. Salakhutdinov, Reducing the dimensionality of data with neural networks, *Science* 313 (2006) 504–507.
- [38] A. F. Frangi, W. J. Niessen, K. L. Vincken, M. A. Viergever, Multiscale vessel enhancement filtering, in: Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI), Springer, pp. 130–137.
- [39] L. Zhou, M. S. Rzeszotarski, L. J. Singerman, J. M. Chokreff, The detection and quantification of retinopathy using digital angiograms, *IEEE Transactions on Medical Imaging* 13 (1994) 619–626.
- [40] Q. Li, B. Feng, L. Xie, P. Liang, H. Zhang, T. Wang, A cross-modality learning approach for vessel segmentation in retinal images, *IEEE Transactions on Medical Imaging* 35 (2015) 109–118.
- [41] S. Wang, Y. Yin, G. Cao, B. Wei, Y. Zheng, G. Yang, Hierarchical retinal blood vessel segmentation based on feature and ensemble learning, *Neurocomputing* 149 (2015) 708–717.
- [42] M. M. Fraz, P. Remagnino, A. Hoppe, B. Uyyanonvara, A. R. Rudnicka, C. G. Owen, S. Barman, An ensemble classification-based approach applied to retinal blood vessel segmentation, *IEEE Transactions on Biomedical Engineering* 59 (2012) 2538–2548.

- [43] E. Bekkers, R. Duits, T. Berendschot, B. ter Haar Romeny, A multi-orientation analysis approach to retinal vessel tracking, *Journal of Mathematical Imaging and Vision* 49 (2014) 583–610.
- [44] N. Chapman, N. Witt, X. Gao, A. Bharath, A. Stanton, S. Thom, A. Hughes, Computer algorithms for the automated measurement of retinal arteriolar diameters, *British Journal of Ophthalmology* 85 (2001) 74–79.
- [45] Y. Sui, X. Zhao, S. Zhang, X. Yu, S. Zhao, L. Zhang, Self-expressive tracking, *Pattern Recognition* 48 (2015) 2872–2884.

Fatmatulzehra Uslu completed her PhD in bioengineering in Imperial College London in 2017, under the supervision of Dr. Anil Bharath. She received her master degree in biomedical engineering in 2012 from Erciyes University, Turkey. She is interested in applying deep learning and tracking methods to medical images.

Anil Anthony Bharath (FIET) is a Reader at Imperial College London. He is a Fellow of Imperial College Data Science Institute, and a co-founder of Cortexica Vision Systems. His research interests are in deep architectures for visual inference.