



PERGAMON

Pattern Recognition 34 (2001) 795–809

PATTERN RECOGNITION

THE JOURNAL OF THE PATTERN RECOGNITION SOCIETY

www.elsevier.com/locate/patcog

Segmentation of macular fluorescein angiographies. A statistical approach

A. Simó^{a,*}, E. de Ves^b

^a*Departamento de Matemáticas, Univ. Jaume I, Campus Penya Roja, 12071-Castellón, Spain*

^b*Instituto de Robótica, Univ. de Valencia, Spain*

Received 7 October 1989; received in revised form 3 December 1999; accepted 19 January 2000

Abstract

This paper is concerned with the use of Bayesian methods in the segmentation of macular fluorescein angiographies. Fluorescein angiography is used in ophthalmic practice to evaluate vascular retinopathies and choroidopathies: Sodium fluorescein is injected in the arm's cubital vein of the patient and its distribution is observed along retinal vessels at certain times. In this task a previous and essential step is the segmentation of the image into its relevant components. In order to obtain this segmentation Bayesian methods can be used because a previous knowledge about the spatial structure of the scene to be segmented is available in this kind of images. The stochastic model assumed for the observed intensities is a simple model with a Gaussian noise process which is statistically independent between pixels. The process of labels x is modelled as a Markov random field with a space-dependent external field expressing the anatomy of the ocular fundus and higher order interactions encouraging blood vessels to be thin and large. This procedure is applied to different cases of diabetic retinopathy and vein occlusions. Two algorithms have been used to estimate x , simulated annealing and iterated conditional modes. In order to evaluate the accuracy of the estimation several error measures have been calculated. © 2001 Pattern Recognition Society. Published by Elsevier Science Ltd. All rights reserved.

Keywords: Random field models; Bayesian image analysis; MCMC methods; Ocular fundus angiography

1. Introduction

Fluorescein angiography is a common procedure in ophthalmic practice, mainly to evaluate vascular retinopathies (diabetic retinopathy, vein occlusions), and choroidopathies (choroidal neo-vascularisation, central serious choroidopathy). The technique consists of an injection of sodium fluorescein in the arm's cubital vein followed by observation of its distribution along retinal vessels at certain times. After 8–10 s, the fluorescent colorant reaches the ocular vasculature and its distribution can be observed and photographed through adequate exciter and barrier filters to show the fluorescence phenomenon. Only abnormal retinal vessels leak sodium

fluorescein, causing progressive increment in fluorescence in certain areas. This exploration allows the clinician to plan and evaluate laser treatment. When important diagnostic and therapeutic decisions are to be made based on these images, direct visual inspection might not be reliable enough and more accurate, qualitative analysis might be required.

A previous and essential step in the analysis of macular angiographies is the segmentation of the image into relevant components. This relevant components are blood vessels (arteries and veins), fovea (avascular foveal region) and the rest of the ocular fundus, that we will call background. The segmentation is important for many reasons. Firstly, the level of fluorescence is not the same at different times in these parts, for example, the grey-level increment due to the fluoresceinic process appears later in veins vessels than in arteries. Secondly, the type and seriousness of lesions depends on its location, in our particular case the nearer the fovea the graver the lesions

* Corresponding author. Tel.: +34-64-72-83-92; fax: +34-64-72-84-29.

E-mail address: simo@mat.uji.es (A. Simó).

are. And finally, some manifestations of retinal diseases are detected by observation of abnormal blood vessels (too thin, too thick or tortuous vessels).

In order to obtain the segmentation we have used photographs taken when the fluorescein appears in the choroidal circulation and the leakage has not yet appeared (from 8 to 20 s after injection). Fig. 1 shows four fluorescein angiographies they were taken from diabetic retinopathy and vein occlusion patients.

Although the segmentation problem has been broadly treated at the computer vision literature its applications to the ophthalmic context are not very abundant.

Goldbaum and coworkers [1] proposed the design of specific filters matched for the segmentation of ocular fundus images. These filters are different for every relevant element in the retina. An inverted Gaussian-shaped matched filter rotated at discrete angles is used to segment and identify piecewise-linear segments of blood vessels. In the same way, a circularly symmetric Gaussian-shaped convolution kernel is used to detect the optic nerve. The fovea is also localised with matching filters, just looking for it in some appropriate positions. There is an obvious disadvantage in this technique: the segmentation process is divided in several pre-processing steps,




Fig. 1. Four fluorescein angiographies.

one for each single element instead of using just one algorithm to segment the image as a whole.

Other approaches have been previously reported to segment blood vessels in medical images. Liu and Sun [2] suggested an algorithm which provides as output pixels belonging to arteries. This algorithm is based on a tracking procedure under a detection–deletion scheme. Kutka and Stier [3] used for the extraction of blood vessels a tracking algorithm which takes into account line properties. As we are interested in separating arteries from veins, these latest techniques are inappropriate, since they cannot provide such a distinction.

The detection of the fovea is very important, as we said above, because of the seriousness of lesions near it. Zana et al. [4] use mathematical morphology to the segmentation of the fovea and they propose an algorithm based on region merging. Ibañez et al. [5] apply Bayesian methodologies to detect the fovea contour.

This paper shows the use of Bayesian image analysis for the segmentation of macular angiographies.

Bayesian methods are useful in medical image analysis because in this kind of data a previous knowledge about the spatial structure of the scene to be segmented is always available. Human anatomy and the use of well-defined protocols for data acquisition guarantee a common image structure for each particular type of exam. Bayesian methods provide a way to use this previous knowledge in order to design the segmentation process.

In this framework the segmentation task is reformulated as a problem of statistical parameter estimation by considering the observed image y as a noisy version of the segmented image x , i.e. the segmented image is regarded as a parameter of y distribution and our goal is the estimation of x . The previous knowledge about x is incorporated by means of a prior distribution. The natural prior probability models for x are Markov random fields since they can model contextual information. Contextual information plays an important role in segmentation problems because the true label of a pixel is normally compatible with labels of neighbours pixels in some sense. Context represents our a priori assumptions about the physical world, we expect images to be continuous and smooth at almost every pixel.

Bayesian methods have been extensively used in some fields of medical image analysis such as the reconstruction from single-photon emission computed tomography data (Refs. [6–10], are just a few examples), the detection of ventricular contours [11], the segmentation of 3D X-ray computerised tomography scenes and detection of microcalcifications in mammograms [12].

This paper is organised as follows. In the next section, we establish some notation and describe the general methodology of Bayesian image analysis. In Section 3 we introduce the stochastic model for the observed data. The prior distribution is introduced in Section 4. In Section 5 we describe the way to carry out the estimation of the

segmented image. Section 6 deals with parameter estimation. Section 7 shows the results of applying our methodology to some images. And finally, in Section 8 we state conclusions of the methodology described in this paper.

2. Preliminaries

This section presents a brief overview of Bayesian approach to image segmentation. A further study in the subject can be found in Cressie [13], Guyon [14] and Winkler [15] books.

Let $y = (y(s))_{s \in S}$ be the observed image and let $x = x(s)_{s \in S}$ the labelled image where $x(s)$ denotes the true (unknown) segment or region which pixel s belongs to. The observed $y(s)$ (usually grey level) is interpreted as a realisation of a random variable whose parameters are a function of $x(s)$. The probability distribution of the observed image y has a density $f(y/x)$. This density describes the “forward problem” of image formation, including both the stochastic noise inherent in observing y and the deterministic influence of x on y .

Given a prior distribution $\pi(x)$ of x , the posterior distribution of x given y is

$$P(x/y) \propto f(y/x)\pi(x)$$

and the segmentation x is usually estimated by choosing a \hat{x} that maximise the posterior distribution.

This estimator is called the maximum a posteriori estimator, MAP estimator. MAP estimators are the Bayes estimators for the 0–1 loss function. Note that the image x is estimated as a whole, we look for a global maximum. In particular, contextual requirements incorporated in the prior (such as homogeneity of regions) are inherited by posterior distribution and, thus, they will influence in the estimator \hat{x} . In this way, some pixels may be misclassified for the sake of homogeneity. We have used this estimator in our experiments.

If we were not interested in regular structures but only in small error rate another reasonable estimator could be the marginal posterior mode estimator (MPME), that corresponds to the error rate loss function [15]. Finally, in other applications such a tomographic reconstruction the mean value or expectation of the posterior distribution is a more convenient estimator (MMSE). This is the Bayes estimator for the mean-square loss function [15].

As we said above, we use the prior model of the “true” image x to express prior knowledge about the physical world like spatial interaction and human anatomy. Markov random fields models (see Refs. [13,16]) have been shown to be appropriate tools for modelling spatial context and so they represent natural models for a priori distribution of x .

To define a MRF, it is necessary to introduce a neighbourhood relation in S , the set of pixels. A subset c of S is

called a clique if any pair of elements of c are neighbours. Thus, the Hammersley–Clifford theorem (see Ref. [13]) shows that the probability distribution corresponding to a MRF can be expressed as

$$\pi(x) = \frac{1}{Z} \exp(Q(x)) = \frac{1}{Z} \exp\left(-\sum_{c \in C} V_c(x_c)\right), \quad (1)$$

where $Q(x)$ is called the negpotential function, $x_c = (x(s) : s \in c)$, c is a clique, and the summation is over the set of all cliques C . The V_c 's are called “potential functions”, lower potentials correspond to higher probabilities and vice versa.

As we have stated, Bayesian estimators of x are based on the posterior distribution. Apart from very particular cases, this distribution cannot be determined completely. For this reason, in order to obtain such estimators the methodology of Markov chain Monte Carlo has been adopted [17]. In this methodology we can compute the posterior distribution empirically by the simulation of computationally feasible Markov chains with limit distribution $P(x/y)$. Widely accepted methods used for this purpose are the Metropolis dynamic and the Gibbs sampler. If we are interested in the MAP estimator we can use simulated annealing algorithm [18] with the Gibbs or Metropolis updates.

Another choice to estimate x is the use of the iterated conditional modes estimator algorithm [19]. The ICM estimator is a local maximum of $P(x/y)$. Although the convergence of this algorithm is faster than the SA, the ICM might fall into a local maximum without possibilities to escape from it whereas the SA is behaved better at these situations by jumping to a randomly chosen solution which avoids falling in a local extreme.

3. A stochastic model

To apply statistical methods as described in Section 2, a stochastic model for the observed intensities has to be assumed. We have considered a simple model with a Gaussian noise process statistically independent between pixels.

Let S be the grid and $n = |S|$, the number of pixels. At each pixel or location $s \in S$ we denote

- $y(s)$ the observed intensities at pixel $s \in S$,
- $x(s) \in \{0, 1, 2, 3\}$ the labels of kind of region, conventionally: 0 for background, 1 for artery, 2 for vein and 3 for fovea.

We will assume that the observed intensity given x follows the model

$$y(s) = \alpha(s) + \varepsilon(s), \quad (2)$$

where $\{\alpha(s)\}_{s \in S}$ are constant within each type of region, i.e., $\alpha(s) = \alpha_{x(s)}$ and $\varepsilon(s)$ are (conditionally) independent

normal variables with null mean and constant variance into each region, $\sigma_{x(s)}^2$ (Gaussian noise).

Then the likelihood is

$$\begin{aligned} f(y/\alpha) &= \prod_{s \in S} f(y/\alpha(s)) = \prod_{s \in S} f(y(s)/x(s)) \\ &= \prod_{s \in S} \frac{1}{\sqrt{2\pi\sigma_{x(s)}^2}} \exp\left\{-\frac{1}{\sigma_{x(s)}^2} (y(s) - \alpha_{x(s)})^2\right\}. \end{aligned}$$

In this likelihood function the only information handled is based on pixel intensity which we have proved to be not sufficient for the segmentation of our images since brightness is very sensitive to the illumination conditions in the capture process. Mainly, blood vessels are the most affected because of their small widths. To solve this problem we can incorporate further knowledge about the types of objects that are being sought in fluorescein angiographies. Thus, it is known that blood vessels usually appear as curvilinear structures that can be locally modelled as lines.

The most obvious approach to modelling this shape difference would be through the use of a more complex prior distribution than that considered in the next section. That would lead to consider a higher order neighbourhood, which, from the point of view of computation, would be rather inefficient.

Instead of that, our approach consists in extracting a linearity parameter from the image data and incorporate this information into the likelihood.

Several algorithms can be found in the literature to achieve a good measure of linearity in each single pixel, understood as the degree of belief that a straight line crosses over a given pixel. Most of them are based on the gradient operation which indicate whether a pixel belongs to an edge (line, high value of the result) or not and, if so, the angular direction of the line associated with it. In a recent work, Díaz et al. [20] proposed a method to detect corners based on statistical analysis of the gradient orientations in a circular neighbourhood of the point considered as a possible corner. This method, even originally developed for detecting corners in an image also indicates whether a pixel belongs to an edge. According to this method, an edge point s in the image can be defined as the point of the image whose gradient is not null and the observed orientations of the gradient for other pixels in its neighbourhood are similar enough. In fact, they are compacted around a privileged orientation μ . Noise causes the observed angles to vary around this unknown μ . It is assumed that the orientations θ of the gradient follow a von Mises distribution [21] that plays the same role in circular statistics as the Gauss distribution in standard statistics.

We have taken advantage of these ideas in order to get the required linearity parameter. Hence, we need to compute the gradient in every pixel and threshold it by a given threshold E . All pixels which have a value for the

modulus of the gradient higher than E are studied by computing their orientation gradient histogram. Thus, for every pixel s in the analysed image, the parameters (κ, μ) of a von Mises distribution for the orientations in its neighbourhood are estimated and the null hypothesis is tested: “The distribution of the orientations in a neighbourhood of s is a von Mises distribution with parameters (κ, μ) ” by using a test of Watson–Stephens [22]. The obtained p -value, $l(s)$, will be used as the linearity parameter since it denotes the similarity of our data to the assumed model, a von Mises distribution.

We will admit conditional independence of the variables y , the intensity, and l , the linearity parameter. This assumption has to be validated because, as we mentioned, the calculus of the linearity parameter is based on the intensity gradient. Using training data, the hypothesis of independence (in fact not lineal dependence) has been tested, finding that at level of significance $\alpha = 0.01$ it cannot be rejected. The resulting correlation coefficients and their significance are shown in Table 1. In the same way, we have tested the hypothesis of Gaussian noise using training data. The obtained p -values are shown in the Table 1 which indicates that our assumption could be not rejected.

With all these ingredients, the likelihood is expressed as follows:

$$f(y, l/x) = \prod_{s \in S} \frac{1}{\sqrt{2\pi\sigma_{x(s)}^2}} \exp \left\{ -\frac{1}{\sigma_{x(s)}^2} (y(s) - \alpha_{x(s)})^2 \right\} \times f(l(s)/x(s)), \quad (3)$$

where f denotes the probability distribution of the linearity parameter given the label. This distribution is unknown and it will be estimated in Section 6.2 using density estimation.

4. A prior model for the labels

The process followed by labels x is modelled as a Markov random field (1). We know in advance the types of objects that appear in the ocular fundus; therefore the features of these objects can be incorporated into the model.

Table 1
 p -values corresponding to the normality and uncorrelation tests

Label	Normality test p -value	Correlation coefficient	p -value
Background	0.0827	0.0661	0.2978
Fovea	0.0319	0.0697	0.0347
Veins	0.0907	0.1474	0.0293
Arteries	0.0339	−0.0673	0.2891

The blood vessels usually appear as curvilinear and thin structures that branch or cross. The intensity of the arteries is lighter than the retina because of the fluorescein. The intensity of the veins is almost always darker than adjacent retina. The fovea (about 500 μm diameter) is an avascular zone and generally appears as a dark spot near the centre of the image.

More specifically we want that the prior model expresses three conditions:

- (i) Neighbouring pixels are likely to belong to the same region, i.e. our image is smooth.
- (ii) Pixels in the middle of the image are likely to belong to the fovea, and pixels in the periphery are likely to belong to blood vessels, i.e. the distribution does not exhibit spatial homogeneity.
- (iii) Pixels that belong to the blood vessels are arranged in large and thin structures, meanwhile fovea is a convex and compact region.

In order to model pixel interaction to express above conditions, it is apparent that simple pairwise interactions are inappropriate and higher order cliques have to be introduced. We have defined the neighbourhood system as the union of two kinds of neighbourhood:

- *Second-order neighbourhood*: The neighbours of a pixel s are those pixels at a Euclidean distance less than 2 from s .
- *Fourth-order neighbourhood*: The neighbours of a pixel s are those pixels whose coordinates are at a distance less than or equal to 4 from the corresponding coordinate of s .

Thus we will use the first type of neighbourhood to express condition (i) and the second one to express condition (iii) (taking into account that veins and arteries usually have a thickness lesser than eight pixels), meanwhile the non-homogeneity will be expressed by means of a space-dependent external field.

For the sake of simplicity, only four different types of cliques will contribute to the negpotential function, namely: cliques with a single pixel (that we will call type-1-cliques), cliques with two second-order neighbour pixels (type-2-cliques), cliques with five fourth-order neighbour pixels like that in Fig. 2(a) and its rotations (type-3-cliques) and finally cliques with 25 fourth-order neighbour pixels like that in Fig. 2(b) (type-4-cliques).

As it was previously commented, we intend to model the particular spatial structure of fundus images by means of a space-dependent external field of the MRF model. In order to simplify we have considered again a very simple space-dependent external field:

We have defined four partitions of S_i a partition for each type of region: (S_k, S'_k) , $k = 0, 1, 2, 3$, $S_0 = S$, because the background is spread to the whole image, $S_1 = S_2$ are those pixels which are within a distance r_1 of




Fig. 2. Some cliques which contribute to the negpotential function.

the centre of the image, because veins and arteries are both in the periphery, and S_3 are those pixels within a distance r_2 of the centre of the image, because the fovea is in the center.

The external field that we have assumed, for a k fixed, is constant within each member of the partition.

All above considerations lead us to assume the next prior distribution for the segmented image x :

$$\pi(x) \propto \exp \left\{ \sum_{s \in S} \sum_{k=0}^3 a_k^s 1_k(x(s)) + \sum_{0 \leq k \leq l \leq 3} b_{kl} n_{kl} + \sum_l d_l n_l + e_l m_l \right\}, \quad (4)$$

where n_{kl} is the number of second-order neighbour pairs of pixels labelled as k , and l , respectively, n_l is the number of type-3-cliques labelled as l , m_l is the number of type-4-cliques labelled as l , and $a_k^s = a_k$ if $s \in S_k$ and a_k^s otherwise.

The parameters b_{ij} are thought to encourage type-2-cliques to have the same label in order to have a smooth segmented image. The parameters d_l intend to encourage all the pixels in a type-3-cliques to have the same label and, on the contrary, the parameters e_l discourage all the pixels in a type-4-cliques to have label vein or artery. By means of these parameters we will achieve that veins and arteries are large and thin and the fovea be compact.

We shall have to impose restrictions on the parameters to make the model identifiable. The usual restriction is to consider $V_c(x_c) = 0$ if $x(s) = 0$ for some $s \in c$. So, we have fixed $a_0 = d_0 = e_0 = 0$ and $b_{0l} = 0 \forall l$.

We have fixed also $a_l = 0.1a'_l$ for $l = 1, 2$ and $a_3 = 0.01a'_3$.

This conditional distribution result in the joint posterior:

$$P(x/y, l) \propto \frac{1}{(\sqrt{2\pi\sigma^2})^n} \exp \left\{ \sum_{s \in S} \frac{1}{2\sigma^2} (y(s) - \alpha_{l(s)})^2 \right.$$

$$\left. + \sum_{s \in S} \sum_{k=0}^3 a_k^s 1_k(x(s)) + \sum_{0 \leq k \leq l \leq 3} b_{kl} n_{kl} + \sum_l d_l n_l + e_l m_l \right\} \prod_{s \in S} f(l(s)/x(s)). \quad (5)$$

The posterior distribution of x is also a Markov random field with respect to the same neighbourhood system.

5. Estimation of x

Our goal is to find the estimate of the vector of labels x . In order to make inferences about these parameters we should work with the posterior distribution $P(x/y, l)$. This distribution has a number of additional unknown parameters (hyperparameters) required to fully define the model. These parameters will be now assumed fixed and known. We will speak about them in next section.

In our experiments we have used two methods in order to estimate x , simulated annealing and ICM. These methods require only access to the full conditional posterior distributions, given by

$$\begin{aligned} P(x(s)/x(t) \ t \neq s, y, l) \\ \propto \exp \left\{ a_{x(s)}^s + \sum_{x(s) < k} b_{x(s)k} n_{s,k} + \sum_{k < x(s)} b_{kx(s)} n_{s,k} + d_{x(s)} t_{s,x(s)} \right. \\ \left. + e_{x(s)} m_{s,x(s)} \frac{1}{2\sigma^2} (y(s) - \alpha_{x(s)})^2 \right\} f(l(s)/x(s)), \end{aligned} \quad (6)$$

where $n_{s,k}$ is the number of second-order neighbours of s with label k , $t_{s,k}$ is the number of type-3-cliques in which s is included having all other pixels labeled as k , $m_{s,k}$ is the number of type-4-cliques in which s is included having all other pixels labeled as k .

The simulated annealing algorithm we have used is the stochastic relaxation algorithm proposed by Geman and Geman [18]. It is based on simulated annealing and the Gibbs sampler. It is known [14] that this algorithm is well adapted when the optimisation is done in a product space and the function to minimise is Markovian.

The Gibbs sampler's visiting scheme used is a raster scan and the cooling schedule $T(k)$ is geometric.

The scheme is as follows:

Step 1: Initialisation. $k = 1$. An initial estimate of x is introduced. The temperature $T(1)$ is initialised.

Step 2: Full update. Run a full sweep over all the image x in such a way that the new $x(s)$ is randomly chosen with probability $\exp\{(1/T(k))\ln(P(x(s)/x(t) \neq s, y, l))\}$.

Step 3: Update the control parameters. $k = k + 1$, $T(k)$ is updated.

Step 4: Return to 2 for a fixed number of cycles or until approximate convergence of x .

In practice, we have obtained good results by using $1/T(k) = 1.01^{-k}$, where k is the number of full sweeps.

Simulated annealing has a slow convergence. As an alternative we have also tried with the deterministic relaxation algorithm known as iterated conditional modes proposed by Besag [19]. The scheme is now

Step 1: Initialisation. $k = 1$. An initial estimation of x is introduced.

Step 2: Full update. Run a full sweep over all the image x in such a way that the new $x(s)$ is chosen to maximise $P(x(s)/x(t) \neq s, y, l)$. $k = k + 1$.

Step 3: Return to 2 for a fixed number of cycles or until approximate convergence of x .

5.1. Initial estimates

In order to use the above iterative procedures some initial estimates are needed. The initial estimate of x has been obtained using a simple segmentation algorithm based on local thresholds.

The initial estimates of images in Fig. 1 are shown in Fig. 3.

5.2. Monitoring convergence

We are going to monitor convergence by plotting the number of pixels with each type of label against the iteration number. Some of these plots are shown in Fig. 6.

6. Parameters and hyperparameters

As we said above the model has a great number of unknown parameters. If we denote θ the vector of parameters, in the Bayesian paradigm in which we are operating, there are several ways to estimate them:

- *Empirical Bayesian approximation:* In an empirical Bayesian way we consider that the parameters are unknown but fixed.
- *Fully Bayesian approximation:* In the fully Bayesian approximation the parameters result as random samples drawn from some hyperprior distribution. This procedure is used in Refs. [10,23–25] advises the use of this Bayesian approach when the models (likelihood




Fig. 3. Initial estimates for the segmentation of images (a) and (b) of Fig. 1.

and prior) are not a very correct description of the processes.

In this way the desired inference about x should be based on the posterior distribution which arises from integrating out all the variables in the full joint posterior. This marginalisation is a very difficult task. In addition to that, if this marginalisation were possible we would still be left with a high dimensional and, may be, more complicated marginal.

In this work an empirical Bayesian approximation has been adopted and we will consider that the parameters are fixed and unknown and so they have to be estimated.

Again, there are two alternative ways: using training data if they are available, or not using them because either they are not available or because their use is not suitable for our particular type of problem. As a result of the different types of parameters used by our model we have to distinguish three cases.

6.1. Parameters in the prior

Training data can be used in order to estimate the parameters in the prior distribution because all our images are fundus images and so they share common structure. Maximum likelihood method could be used but the normalisation constant ($\pi(x)$ distribution) neither is available or can be computed. Several algorithms have been proposed to overcome this difficulty. We can use, for example, those due to Younes [26], Geyer and Thompson [27], Geyer [28] and the Monte Carlo Newton–Raphson approach [13,14].

Nevertheless, the pseudolikelihood approach has been used here because of its lower amount of computation. The pseudolikelihood function of the x prior distribution has the next expression

$$PL(a_k, b_{kl}, d_l, e_l) = \frac{\exp\{\sum_{s \in S} (\sum_{k=0}^3 a_k^s 1_k(x(s))) + \sum_{0 \leq k \leq l \leq 3} 2b_{kl} n_{kl} + \sum_l 5d_l n_l + 25e_l m_l\}}{\prod_{s \in S} \sum_{x=0}^3 \exp\{a_x^s + \sum_{x < k} b_{xk} n_{s,k} + \sum_{k < x} b_{kx} n_{s,k} + d_x t_{s,x} + e_x m_{s,x}\}}. \quad (7)$$

In order to maximise this pseudolikelihood the conjugate gradient method of Fletcher–Reeves–Polak–Ribiere (see Ref. [29]) has been used.

6.2. Distribution of the linearity parameter

The density function of the linearity distribution is completely unknown, unlike the other cases, we have to estimate the function and not only some parameters. By using again training data we can use a kernel estimator to obtain $\hat{f}(l(s)/x(s))$.

6.3. Parameters of the distribution of the intensity

Training data is not useful to estimate the parameters of the distribution of the intensity because there is a great variability in grey levels among images.

We choose to follow an iterative procedure (see Refs. [19,30]) that given a current estimate of x , \hat{x} a new value for the parameters of y distribution is obtained by maximising the conditional density of y :

$$\hat{\alpha}_k = \frac{1}{n_k} \sum_{s \in S: \hat{x}(s)=k} y(s),$$

$$\hat{\sigma}^2 = \frac{1}{n} \sum (y(s) - \hat{\alpha}_{\hat{x}(s)})^2,$$

where $n_k = \#\{s \in S: \hat{x}(s) = k\}$.

With these new values of the parameters, new estimators for x are obtained. This process is repeated until convergence.

7. Experimental results

The proposed methodology has been applied to several fundus ocular images. All these images have been taken by a Canon CF-60U and Nikon NFC 50 fundus cameras, transformed into video signal with a photo video camera PHV-A7E, Sony and digitalised with a Matrox MVP-AT card with 512×380 pixels. They belong to patients suffering from different lesions such as diabetic retinopathy, vascular occlusion and macular edema.

Four images from this database are shown in Fig. 1.

As we have explained earlier, prior distribution parameters have been estimated using training data. We have considered $d_1 = d_2$ and $e_1 = e_2$ because of the similar behaviour of the arteries and veins. The rest of parameters values are shown in Table 2. Two different optimisation algorithms have been applied: the

simulated annealing and the ICM which have been coded in VISTA, a set of image processing libraries written in C.

The segmented images obtained by using these schemes are shown in Figs. 5 and 4, which corresponds,

Table 2
Parameter estimates using training data

$a_1 = -0.02$	$a_2 = 0.07$	$a_3 = 0.31$
$b_{11} = 1.28$	$b_{12} = -1.4$	$b_{13} = -0.25$
	$b_{22} = 1.7$	$b_{23} = -0.15$
		$b_{33} = 1.8$
$d_1 = 1.1$	$d_3 = 0.21$	
$e_1 = -0.98$	$e_3 = 0.1$	




Fig. 4. Results using 50 iterations of ICM algorithm.

respectively, to the simulated annealing after 100 iterations and ICM after 50 iterations.

With regards to convergence, in Fig. 6 we can see the plot of the number of fovea pixels against the number of iterations (full sweeps of the image) corresponding to the second image in Fig. 1. We see that ICM reaches the convergence in about 40 iterations meanwhile SA in about 90 iterations. The same behaviour is observed in all the plots. The CPU times involved with ICM is about 25 s by iteration and 30 s with SA.

In order to test our methodology these automated segmented images are compared with the images segmented by expert clinicians manually. These pattern images are shown in Fig. 7. From now on, we refer to images segmented by the expert clinicians as the true images.

In order to evaluate the accuracy of the estimation several error measures have been used. These error measures have been separately calculated for each labelled segment of the image (fovea, vein and artery), for each




Fig. 5. Results using 100 iterations of SA algorithm.

image and for each optimisation algorithm. Let A be the true image and B the estimated image and $n(S)$ = the number of pixels in S , we calculate:

- The false positive rate defined by

$$\alpha(A, B) = \frac{n(B \setminus A)}{n(S_k \setminus A)}, \quad (8)$$


where $k = 1$ for artery, $k = 2$ for vein and $k = 3$ for fovea and S_k $k = 1, 2, 3$ is defined in Section 5.

- The false negative rate defined by


$$\beta(A, B) = \frac{n(A \setminus B)}{n(A)}. \quad (9)$$

- The Pratt's [31] "figure of merit" defined by

$$FOM(A, B) = \frac{1}{\max\{n(A), n(B)\}} \sum_{x \in B} \frac{1}{1 + \kappa d(x, A)}, \quad (10)$$



(a)



(b)

Fig. 6. Convergence plots for the image (b) of Fig. 1. (a) SA algorithm. (b) ICM algorithm.

where $d(x, S)$ denotes the function distance of S and κ is a scaling constant (usually set to $\frac{1}{9}$).

False positive and false negative rates are detection performance measures, i.e. discrepancies between A and B are measured by the number of disagreements, regardless of the pattern. FOM is a localisation performance measure.

Results are shown in Table 3. Examination of this table leads to the following points:

In all cases the measure error values are acceptable enough, the worst values are reached with Fig. 1(d) because of their poor quality.

The FOM values are always greater than 0.5 even in the image of Fig. 1(d).

Table 3

Error measures for each labelled segment of the images of Fig. 1

Image	Label	ICM			SA		
		α (%)	β (%)	FOM	α (%)	β (%)	FOM
a	Fovea	22.4	10.3	0.71	27.7	6.5	0.67
	Arteries	16.7	26.5	0.79	18.7	19.9	0.76
	Veins	10.5	41.3	0.55	10.6	41.0	0.6
b	Fovea	24.0	3.0	0.8	43.4	0.0	0.67
	Arteries	1.4	28.2	0.8	2.6	24.5	0.85
	Veins	1.6	20.3	0.56	1.7	21.7	0.58
c	Fovea	23.9	7.48	0.65	29.4	6.6	0.6
	Arteries	4.7	43.11	0.78	9.8	29.01	0.84
	Veins						
d	Fovea	82.2	0.0	0.6	90.1	0.0	0.53
	Arteries	61	42.5	0.77	8.7	38.7	0.73
	Veins	0.5	15.8	0.78	0.73	6.51	0.78

Table 4

Error measures for noisy images of Fig. 8

Noise level	Label	ICM			SA		
		α (%)	β (%)	FOM	α (%)	β (%)	FOM
Low	Fovea	29.4	0.2	0.78	32.4	3.4	0.759
	Arteries	1.7	36.4	0.89	2.4	40	0.85
	Veins	1.7	25.3	0.56	1.9	26.7	0.59
Medium	Fovea	24.2	8.0	0.81	28.7	2.1	0.78
	Arteries	1.8	41.7	0.84	2.5	39.0	0.85
	Veins	1.8	23.4	0.57	1.9	28.9	0.59
High	Fovea	23.4	8.7	0.81	29.6	0.96	0.78
	Arteries	2.8	41	0.85	2.4	40.3	0.85
	Veins	1.8	23.2	0.57	1.9	30.6	0.57

The lowest values for the false positive rate usually correspond to the fovea. That means that the estimated fovea is always a little bigger than the true fovea. Because of the seriousness of the lesions near the fovea, this is not a serious error in the opinion of the experts clinicians. For clinical practice, it is better to obtain a bigger fovea than a smaller one.

In relation to the false negative rate, usually, it is bigger in vessels than fovea because some little branches are lost. These branches are not very important to the clinical practice.

On the other hand, if we compare SA and ICM algorithm we can see that both of them provide good results with slight differences. These slight differences can be due




Fig. 7. Images segmented by expert clinicians manually.

to the fact that ICM stops when the algorithm reaches a local maximum meanwhile SA reach the global one.

We have found that in the most of the cases the images obtained by the SA algorithm have a similar estimate for veins and arteries: false positive rate and false negative rate and FOM are similar. However, the estimates for the fovea use to give a false positive error rate bigger than this obtained by the ICM algorithm. We could conclude that there is not much difference between results ob-

tained by the SA algorithm and those obtained using the ICM algorithm. However, the convergence time with the former is 3 times the time needed with the later.

7.1. Robustness against noise

In order to prove the robustness against noise of these algorithms, salt and pepper noise has been added to one of the images under study. Three different levels of noise

have been introduced that we will call low, medium and high. Low noise means that 5% of randomly chosen points in the image have been modified, medium means 8% and high means 10%. The noisy images are shown in Fig. 8. The segmented images obtained from these noisy version after applying the ICM and SA algorithms can be seen in Figs. 9 and 10. Table 4 shows false positive, false negative and FOM computed from these angiographies. The performance of these algorithms are very similar to the one obtained from the original image (with no noise). This points out that the proposed procedures are robust against noise.

8. Conclusions

We have proposed a method to obtain a segmentation of the eye fundus angiographies based on Bayesian statistical methodologies. In contrast to other previous approaches, it provides a global segmentation, i.e. veins, arteries and fovea are obtained by using only one algorithm.

In order to assess the performance of the method it has been applied to several eye fundus angiographies and the results have been compared with those obtained manually by expert clinicians. These experimental results




Fig. 8. Image 1 (b) with salt and pepper noise. (a) Low, (b) medium, (c) high.




Fig. 9. Segmented images obtained from noisy images of Fig. 8 using ICM algorithm.




Fig. 10. Segmented images obtained from noisy images of Fig. 8 using SA algorithm.

showed that the method achieves a good segmentation of the main vessels but little branches are lost. The estimate of the fovea is good but is a little bigger than those obtained manually. These results are considered good enough for clinical practice by the expert clinicians. We have proved also that the method exhibits immunity against noise.

We could conclude that although we can find in the literature other methods which obtain a better segmentation of one of the different elements in the eye fundus, we have showed that our method provides a segmentation that is good to the clinical practice and that it segments the image as a whole using just one algorithm.

References

- [1] M.H. Goldbaum, N.P. Katz, S. Chaudhuri, M. Nelson, Image understanding for automated retinal diagnosis, in: L. Kingsland (Ed.), *Proceedings of the IEEE for the 13th Annual Symposium for Computer Application in Clinical Medicine*, Computer Soc. Press, 1989, pp. 756–760.
- [2] I. Liu, Y. Su, Recursive tracking of vascular networks in angiograms based on the detection-deletion scheme, *IEEE Trans. Med. Imaging* 12 (2) (1993) 334–341.
- [3] R. Kutka, S. Stier, Extraction of line properties based on direction fields, *IEEE Trans. Med. Imaging* 15 (1) (1996) 51–58.
- [4] F. Zana, I. Meunier, J.C. Klein, A region merging algorithm using mathematical morphology: application to macula detection, *ISMN98*, 1998, pp. 423–430.
- [5] M.V. Ibañez, A. Simó, Bayesian detection of the fovea in eye fundus angiographies, *Pattern Recognition Lett.* 20 (1999) 229–240.
- [6] S. Geman, D.E. McClure, Statistical methods for tomographic image reconstruction, *Proceedings of 46th Session Inst. Stat. Inst. Bulletin ISI*, Vol. 52, 1987.
- [7] R.G. Aykroyd, P.J. Green, Global and local priors, and the location of lesions using gamma-camera imagery, *Philos. Trans. Roy. Soc. London Ser. A* 337 (1991) 323–342.
- [8] P.J. Green, Bayesian reconstruction from emission tomography data using a modified em algorithm, *IEEE Trans. Med. Imaging* 9 (1990) 84–93.
- [9] I.S. Weir, P.J. Green, Modeling data from single-photon emission computed tomography, in: K.V. Mardia (Ed.), *Statistics and Images*, Carfax, Abingdon, 1994, pp. 313–338.
- [10] I.S. Weir, Fully Bayesian reconstruction from single photon emission computed tomography data, *J. Amer. Statist. Assoc.* 92 (437) (1997) 49–60.
- [11] M. Teles, Bayesian estimation of ventricular contours in angiographic images, *IEEE Trans. Med. Imaging* 11 (3) (1992) 416–429.
- [12] N. Karssemeijer, Application of Bayesian methods to segmentation in medical images, in: P. Frigessi Barone, M. Piccioni (Eds.), *Lecture Notes in Statistics, Stochastic Models, Statistical Methods, and Algorithms in Image Analysis*, 74, Springer, Berlin, 1992.
- [13] N.A.C. Cressie, *Statistics for Spatial Data*, Wiley Series in Probability and Mathematical Statistics, Wiley, New York, 1993.
- [14] X. Guyon, *Random Fields on a Network. Modeling Statistics and Applications*, Springer, Berlin, 1995.
- [15] G. Winkler, *Image Analysis, Random fields Models and Dynamic Monte Carlo Methods*, Springer, Berlin, 1996.
- [16] J.E. Besag, Statistical analysis of non lattice data, *Statistician* 24 (1975) 179–195.
- [17] P.J. Green, Mcmc in image analysis, in: W.R. Gilks et al. (Eds.), *Markov chain Monte Carlo in Practice*, Chapman & Hall, London, 1996, pp. 381–400.
- [18] S. Geman, D. Geman, Stochastic relaxation, gibbs distributions, and the Bayesian restoration of images, *IEEE Trans. Pattern Anal. Mach. Intell.* 6 (1984) 721–741.
- [19] J.E. Besag, On the statistical analysis of dirty pictures, *J. Roy. Statist. Soc. B* 48 (1986) 259–302.

- [20] M.E. Diaz, J. Domingo, G. Ayala, A gray-level 2-d feature detector using circular statistics, *Pattern Recognition* 18 (1997) 1083–1089.
- [21] K. Mardia, *Statistics of Directional Data*, Academic Press, New York, 1972.
- [22] G. Watson, Goodness of fit tests on a circle, *Biometrika* 49 (1962) 57–63.
- [23] V.E. Johnson, A model for segmentation and analysis of noisy images, *J. Am. Statist. Assoc.* 89 (1994) 230–241.
- [24] A. Mollie, Bayesian mapping of disease, in: W.R. Gilks, S. Richardson, S. Spiegelhalter (Eds.), *Markov chain Monte Carlo in Practice*, Chapman & Hall, London, 1996 (Chapter 20).
- [25] R. Molina, On the hierarchical Bayesian approach to image restoration: applications to astronomical images, *IEEE Trans. Pattern Anal. Mach. Intell.* 16 (1994) 1122–1128.
- [26] L. Younes, Estimation and annealing for Gibbsian fields, *Ann. Ins. Henri Poincaré. Prob. Stat.* 24 (1988) 269–294.
- [27] C.J. Geyer, E.A. Thompson, Constrained Monte Carlo maximum likelihood for dependent data (with discussion), *J. Amer. Statist. Assoc.* 56 (1994) 261–274.
- [28] C.J. Geyer, On the convergence of Monte Carlo maximum likelihood calculations, *J. Roy. Statist. Soc. Ser. B* 56 (1994) 261–274.
- [29] W.H. Press, S.A. Teulkolsky, W.T. Vetterling, B.P. Flannery, *Numerical Recipes in C*, Cambridge University Press, Cambridge, 1992.
- [30] W. Qian, D.M. Titterington, Bayesian image restoration: an application to edge-preserving surface recovery, *IEEE Trans. Pattern Anal. Mach. Intell.* 15 (1993) 748–752.
- [31] W.K. Pratt, *Digital Image Processing*, Wiley, New York, 1977.

About the Author—AMELIA SIMÓ is an Associate Professor of Statistics in the Department of Mathematics, University Jaume I, Castellón, Spain. She received her M.S. degree and her Ph.D. degree in Mathematics from the University of Valencia, Spain in 1990 and 1995, respectively. Her current research interest include Statistical methods in Pattern Recognition, Markov Models and Markov Chain Monte Carlo Methods.

About the Author—ESTHER DE VES was born in Almansa (Spain) in 1971. She received the M.S. degree in Physics and the Ph.D. in Computer Science from the University of Valencia in 1993 and 1999, respectively. Since 1994 she has been with the Department of Computer Science from the University of Valencia where she is now an Assistant Professor. Her current interests are Image Registration, Medical Imaging and Pattern Recognition.