

A BAYESIAN APPROACH TO INFERRING VASCULAR TREE STRUCTURE FROM 2D IMAGERY

Elke Thönnies², Abhir Bhalerao¹, Wilfrid Kendall², and Roland Wilson¹

Departments of ¹Computer Science and ²Statistics
University of Warwick, UK

{elke|wsk}@stats.warwick.ac.uk

{abhir|rgw}@dcs.warwick.ac.uk

ABSTRACT

We describe a method for inferring tree-like vascular structures from 2D imagery. A Markov Chain Monte Carlo (MCMC) algorithm is employed to produce approximate samples from the posterior distribution given local feature estimates, derived from likelihood maximisation for a Gaussian intensity profile. A multiresolution scheme, in which coarse scale estimates are used to initialise the algorithm for finer scales, has been implemented and used to model retinal images. Results are presented to show the effectiveness of the method.

1. INTRODUCTION

The problem of inferring vascular structure from image data is an important one, especially in the area of surgical planning, which requires both efficient computation and effective use of prior knowledge. Previous work in the area has tended to focus on the modelling of specific vascular features [1] or to use approaches such as adaptive thresholding [4].

The aim of the work described here is to formulate a general method for the inference, which can be applied in two or three dimensions and makes effective use of prior knowledge, yet which is sufficiently general to be applied to a wide range of problems. The common statistical methods for such medical image analysis have typically used likelihood techniques, such as Expectation-Maximisation (EM) [6, 5]. Although EM methods can be efficient computationally, they have only limited scope for incorporating prior knowledge. A more powerful way of including prior information is to use a Bayesian method, such as *maximum a posteriori* (MAP) estimation. The principal difficulty with Bayesian techniques is a computational one: they normally require the use of Markov chain Monte Carlo (MCMC) algorithms, which may run for hundreds of thousands of iterations to yield reliable results [3]. This has restricted their

use in applications involving large data sets, such as medical images.

The method we have adopted combines local likelihood maximisation using a Gaussian model of the spatial intensity profile, on a multiresolution grid and global structure determination using a Bayesian technique derived from a general model of vasculature as a collection of tree structures. By approaching the problem in this way, we combine the computational efficiency of likelihood techniques with the power and generality of a Bayesian approach. After a brief description of the estimation algorithms, we present results of two dimensional structure inference on real data. The paper is concluded with some observations on the technique.

2. A STATISTICAL MODEL OF VASCULAR STRUCTURE

To draw inferences about global structure, we employ a Bayesian formalism: the data are modelled as a random tree-like structure and we then use an MCMC algorithm [2] to sample from the posterior distribution, which is conditioned on the data. The sampling distribution is an approximate equilibrium of a random process whose configuration space is the space of tree-like structures and whose equilibrium is designed to be the target conditional distribution. As well as gaining information about the global structure, variation in the posterior samples enables us to quantify uncertainties about the image interpretation.

The prior distribution defines the global structure as a *forest* of a random number of trees. Each such tree is binary: branches divide only into two sub-branches at any given position. A physical realisation of such a tree needs each vertex to be located in space. Unfortunately, the simplistic approach of displacing each vertex from its parent by a Gaussian displacement of zero mean (a “random-walk” tree) leads to a tangled local structure (left hand figure 1). We therefore introduce a correlation by allowing the mean

This work was supported by the UK EPSRC.

displacement to be a small linear multiple of the displacement of the parent vertex from the grandparent vertex (an ‘‘AR(1)’’ tree) (right hand of figure 1). To model intensity, with each vertex in the tree we associate a Gaussian kernel that represents the spatial grey level profile of the corresponding vessel segment.

The posterior distribution for a random number N of trees τ_1, \dots, τ_N is given by

$$\begin{aligned} \pi(\tau_1, \dots, \tau_n | f) &\propto \\ &\mathbb{P}(N = n) \prod_{i=1}^n \prod_{\nu \in \tau_i} p_{v(\nu)} \\ &\phi(\theta_\nu | \theta_\eta, \eta \in \mathcal{A}(\nu)) \\ &\times L(\tau_1, \dots, \tau_n; f), \end{aligned} \quad (1)$$

where f is the image. The distribution of N , $\mathbb{P}(N = n)$, penalises the number of trees; in the examples, a Poisson distribution was used. This ensures that a ‘minimal’ explanation of the data is found. p_0, p_1, p_2 are the probabilities defining the degree of the branching process and $v(\nu)$ is the degree of the vertex ν . $\phi(\theta_\nu | \theta_\eta, \eta \in \mathcal{A}(\nu))$ is the distribution of the parameters of the vertex, θ_ν . In the examples, this is an autoregressive (AR(1)) process. The parameters θ_ν represent the positions of vertices, x_ν , and the amplitude and width parameters of the edges, while $L(\cdot; f)$ is the likelihood function. The observation model is based on the approximation of linear structures by a sum of Gaussian kernels:

$$\tau_i(x) = \sum_{\nu \in \tau_i} a_\nu \mathcal{N}\left(x - \frac{x_{\text{parent}(\nu)} + x_\nu}{2}, \Sigma_\nu\right) \quad (2)$$

where a_ν is a magnitude scale factor and x_ν, Σ_ν the mean and covariance parameters of the kernel. Rather than sample directly on the covariance parameters, they are sampled in a *normalised* form, in terms of the orientation and the standard deviations of their principal components. The orientation is defined by the line joining a vertex with its parent. The likelihood is then computed based on a normal model of the observations within a region

$$f(x) = \sum_i \tau_i(x) + v(x) \quad (3)$$

where $v(x)$ represents the noise in the observations, which again is assumed to be normally distributed, with zero mean and independent at each pixel. The simulation, whose configuration at any one time is a collection of random trees, is designed to have *moves* which take it from one configuration to the next. In order to arrive at a Maximum a Posteriori (MAP) estimate from a given image, it is necessary to sample from the above posterior distribution. This is done by a simulation involving a number of different *moves*, affecting the state of the forest at each iteration. As long as each

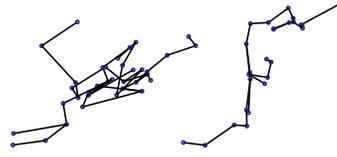


Fig. 1. Illustration of a random-walk tree and an AR(1) tree

move has an ‘opposite’ (e.g. add versus delete, split versus graft) and the chances of each move are balanced against its opposite, it is straightforward to compute the required equilibrium distribution and to design the move probabilities to give the required posterior as equilibrium, using the Metropolis-Hastings technique (MH) [3]. The MH method iterates in two steps: the first proposes a move from the collection of moves and the second accepts or rejects the move so as to ensure that the equilibrium coincides with the desired posterior. The randomised decision whether to accept or reject a move depends on whether the resulting new tree will be a more adequate representation from the posterior than the current tree.

The efficiency of the algorithm depends crucially on the proposed moves. As all moves influence the likelihood only locally, the likelihood evaluation can be implemented efficiently. Global structure that can be inferred with high certainty locally will lead to a tree structure that is stable over time, while low local certainty results in a volatile tree-structure that alternates between different explanations for the global structure. A summary of the moves currently implemented for the algorithm is shown in figure 2.

Each iteration of the sampler consists of the random selection of one such move and its random acceptance or rejection, according to posterior probability.

3. EXPERIMENTS

Figure 3 illustrates results of the ML model estimation. We have used part of a 2D retinal angiographic image size 404×404 pixels (fig. 3(a)) for our 2D experiments. To improve the efficiency of the estimation, we start from a set of maximum likelihood local feature estimates, based on the Gaussian kernel of (3), within a multiresolution framework, combined with a piecewise linear model of the background intensity to account for the obvious large scale variation in intensity across the images such as fig. 3(a). In other words, each region in a quadtree tessellation of the image is modelled as the sum of a linear background term and a Gaussian intensity profile, whose parameters are estimated using an

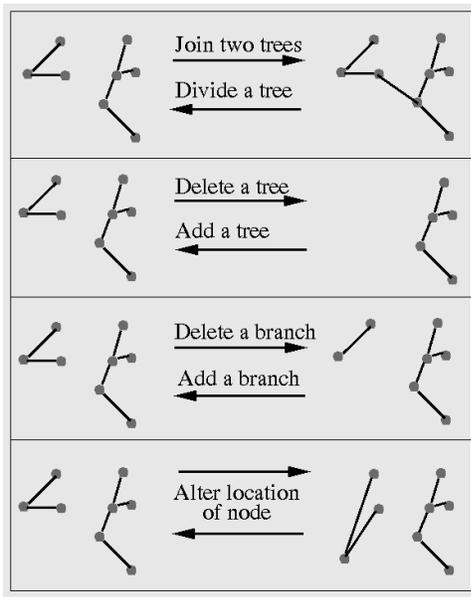


Fig. 2. Moves used in MCMC simulation on trees.

iterative, EM-type algorithm to maximise likelihood

$$a_i^{t+1} = \sum_x f_w(x) \mathcal{N}(x - x_i^t, \Sigma_i^t) / \sum_x \mathcal{N}^2(x - x_i^t, \Sigma_i^t) \quad (4)$$

$$x_i^{t+1} = \frac{\sum_x x f_w(x) \mathcal{N}(x - x_i^t)}{\sum_x f_w(x) \mathcal{N}(x - x_i^t, \Sigma_i^t)} \quad (5)$$

$$\Sigma_i^{t+1} = 2 \frac{\sum_x (x - x_i^t)(x - x_i^t)^T f_w(x) \mathcal{N}(x - x_i^t, \Sigma_i^t)}{\sum_x f_w(x) \mathcal{N}(x - x_i^t, \Sigma_i^t)} \quad (6)$$

where the data $f(x)$ are windowed with a cosine window, whose size is twice the block width at a given scale, to give the data $f_w(x)$ used in the estimator. The index t denotes iteration number; typically 4-5 iterations are sufficient to give accurate estimates. Figures 3(b)-(c) show reconstructions using the 2D Gaussians in each block (at corresponding block sizes) based on the ML feature estimates at block sizes of 64×64 and 16×16 respectively. Clearly, at lower spatial resolutions, the model cannot easily describe the presence of multiple vessels within the window, such as occur at bifurcations, and the resulting low-amplitude, isotropic Gaussians are locally the ‘best’ description of these regions. However, these blocks can be modelled accurately at higher spatial resolutions. The second set of images shows how the local estimates from different window scales is used in a *coarse-fine* stochastic simulation, in order to get a Bayesian estimate of the forest

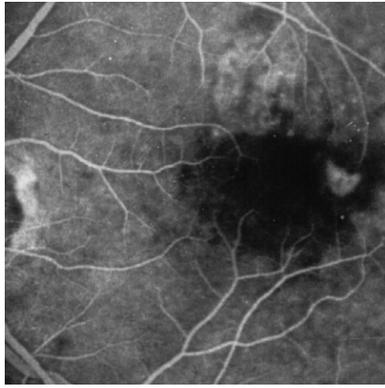
structure. After the first 200 iterations, the scale is halved and the appropriate local feature estimates are used to guide the sampler; after 3000, the scale is halved again, as it is after 6000 iterations, at which point, the highest spatial resolution is reached. This approach has been found to speed convergence to the equilibrium distribution, while avoiding becoming trapped in local modes, in a similar manner to many coarse-fine algorithms. Note that one iteration consists of the generation and acceptance of a single proposal (for ‘editing’ the tree). In other words, 100000 iterations is comparable, in terms of computation, to a single scan through the image. It has been noted from experiments that equilibrium is reached in approximately 50000 iterations, a comparatively low burden computationally.

4. CONCLUSIONS

Some encouraging preliminary results have been achieved using the approach described in section 2, demonstrating its potential for modelling vascular structure globally in a computationally efficient way. Fine-tuning the algorithm will lead to significant improvements. These will include, for example, the use of the local estimates to produce initial configurations for the MCMC algorithm. Such improvements are currently being implemented. The work is also being tested on other types of data and extended to three dimensions.

5. REFERENCES

- [1] P. Datlinger A. Pinz, S. Bernogger and A. Kruger. Mapping the human retina. *IEEE Trans. Medical Imaging*, 17(1):606–619, 1998.
- [2] S.P. Brooks. The Markov Chain Monte Carlo Method and its Application. *The Statistician*, 47:69 – 100, 1998.
- [3] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter. *Markov Chain Monte Carlo in Practice*. Chapman & Hall, 1996.
- [4] M. E. Martinez-Perez, A. D. Hughes, A. V. Stanton, A. S. Thom, A. A. Bharath, and K. H. Parker. Segmentation of retinal blood vessels based on the second directional derivative and region growing. In *Proc. of IEEE ICIP-99*, pages 173–176, Kobe, Japan, 1999.
- [5] W. M. Wells, R. Kikinis, W. E. L. Grimson, and R. Jolesz. Adaptive segmentation of MRI data. *IEEE Trans. Medical Imaging*, 15:429–442, 1996.
- [6] D. L. Wilson and J. A. Noble. An adaptive segmentation algorithm for extracting arteries and aneurysms from time-of-flight MRA data. *IEEE Trans. Medical Imaging*, 18(10):938–945, 1999.



(a)

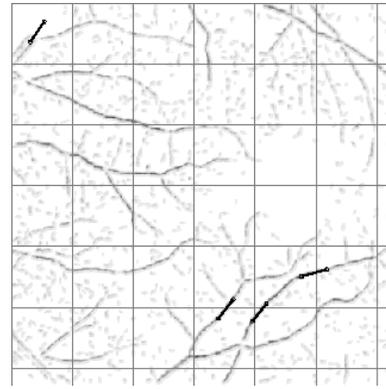


(b)

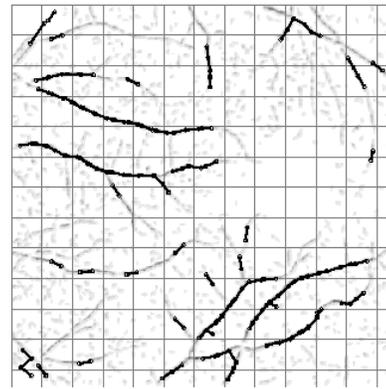


(c)

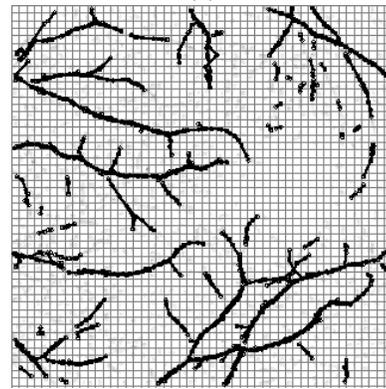
Fig. 3. (a) 2D retinal angiogram size 404×404 pixels. (b) Reconstruction of data from model parameters estimates Θ_{ML} for block sizes of 64 and (c) 16.



(a)



(b)



(c)

Fig. 4. Estimates from the tree-based sampler, at (a) 200, (b) 1000 and (c) 50000 iterations, showing how use is made of the local multiresolution feature estimates, in a coarse-fine approach to the MCMC algorithm. After 50000 iterations, few changes occur.