ISSN 1398-599X MPS-RR 2000-17

Bayesian contour detection in a time series of ultrasound images through dynamic deformable template models

Martin B. Hansen and Jesper Møller

Aalborg University, Denmark

Frede Aa. Tøgersen†

Danish Institute of Agricultural Sciences, Denmark

Summary. Deformable template models and Markov chain Monte Carlo methods are used for analyzing a space-time process of intracoronary ultrasound images in order to detect the artery contour and various other characteristics as a function of time.

Keywords: Bayesian image analysis; Deformable templates; Intracoronary ultrasound images; Markov chain Monte Carlo; Space-time modeling

1. Introduction

In medical imaging, in vivo video films of dynamic behavior such as blood oxygenation in the brain under stimuli, beating hearts, and pulsating artery walls, object detection in both space and time is required, see e.g. Aizenman (2000). In this paper we develop an automatic method for contour detection in a sequence of high-frequency intra coronary ultrasound (ICUS) images, using a fully Bayesian approach. To the best of our knowledge, the paper provides the first published attempt of a fully Bayesian analysis of a space-time process based on deformable template models and Markov chain Monte Carlo simulations.

The ICUS images show the dynamic behavior of the coronary artery; the data set is further described in Section 2. Extracting information from the ICUS images is important e.g. in assessing the extent of arteriosclerosis and flow response to various stimuli, and one feature of interest is the dynamic behavior of the cross sectional area of the artery at a given place. Analyzing the images for anomalies by manually following the contour of the artery in a digitized version of each video sequence is both time consuming and observer dependent. A quick and accurate method for automatic contour detection and feature extraction would be an appreciated aid.

Many papers are dealing with automatic detection of the contour of an object in a single image. In the engineering literature the active contour approach has been popular (Kass *et al.*, 1987), where the outline is normally modeled by a connected and sufficiently smooth curve. The best fit is found as the curve that minimizes a functional which balance the degree of misfit, some smoothness criterion, and a quantity describing the closeness to some important image features. Yet other methods have been developed for medical

†*Address for correspondence:* Frede Aa. Tøgersen, Department of Agricultural Systems, Danish Institute of Agricultural Sciences, P.O. Box 50, DK-8830 Tjele, Denmark.
E-mail: FredeA.Togersen@agrsci.dk

image analysis, see Aizenman (2000) and the references therein. Many of the methods produce a good "point estimate" of the contour, but say nothing about the uncertainty of the estimate, a crucial requirement when it comes to diagnosis. This has led to a number of more statistical formulations of the contour detection problem, usually using a Bayesian setting.

Possibly the most popular class of prior models for contours are provided by the deformable polygonal template approach of Grenander and others, see, for example, Grenander *et al.* (1991), Grenander and Miller (1994), and Kent *et al.* (1996). Another approach is based on marked point processes for object detection (Baddeley and van Lieshout, 1993; Pievatolo and Green, 1998). The two approaches of deformable templates and marked point processes are combined in Rue and Syversveen (1998), Rue and Husby (1998), and Hurn *et al.* (1999). Likelihood based methods for template matching have recently been studied in Hobolth and Jensen (1999) and Hobolth and Pedersen (1999).

There is an extensive engineering literature dealing with space-time images, including Bayesian approaches based on deformable templates, see Kervrann (1998, 1999) and the references therein. Aizenman (2000) extends ideas from deterministic spatial feature detection to the temporal domain, and applies this to boundary detection in echocardiography. Kervrann (1999), in particular, combines a spatial segmentation procedure based on deformed templates with a fixed number of vertices, with edge-based temporal tracking. The tracking procedure over long image sequences was supplemented with Kalman filtering.

We also use a deformable template as prior distribution, but do not find it necessary to include global deformations. In order to account for uncertainty of the parameters, we impose prior distributions on the parameters of the deformation model in contrast to Kervrann (1999) who estimates them by maximum marginal likelihood. Opposite to Kervrann's approach based on maximum a posteriori estimates, we use the posterior means as estimates of the contour and to predict the templates. Furthermore, we supply the posterior mean calculations of the cross sectional area of the artery and other parameters with calculations of credibility intervals.

The paper is organized as follows. Section 2 considers the data and the observation model corresponding to the likelihood term of our posterior. Deformable template priors related to our application are studied in Section 3. Section 4 is concerned with Bayesian inference and Markov chain Monte Carlo algorithms. The empirical results are discussed in Section 5.

2. Data and observation model

The data originates from a study of patients with angina pectoris (pain in the chest) without angiographic evidence of arteriosclerosis. The study was conducted by the Department of Cardiology, Aarhus University Hospitals, Skejby, with the purpose to diagnose whether or not chest pain is caused by inappropriate responses of the coronary artery to various stimuli. Notice that the present paper does not deal with the diagnosis of the patients but merely presents a way of doing the tracing and feature extraction automatically.

The data we consider arises from a 2 seconds long video film recorded by ultrasound scanning at a prespecified place in a coronary artery, approximately perpendicular to the blood flow. The frame rate is 25 images per second, yielding 49 ultrasound images which have been digitized into grey-level pictures. Figure 1 shows two of these grey-level images. Additional figures, which show the development over time, are given in Figures 5–8.

Consider the right image in Figure 1. The intracoronary scanning is performed by inserting into the artery (the black part in the middle of the image) a catheter (indicated by 3) with a rotating ultrasound transducer at the top. The brightest grey-levels represents the greatest reflection of the ultrasound. The drop outs at 1 and 2 (compare with the left image in Figure 1) may be true image artifacts or offsprings of vessels.

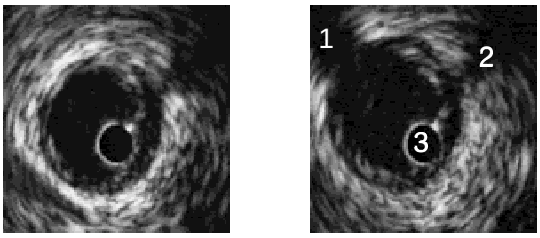

Fig. 1. Two examples of digitized ICUS images of a coronary artery. Right image: 1 and 2: drop outs; 3: the ultrasound transducer.

Each grey-level image is made up of $N^2 = 128^2$ square pixels, which we identify with unit squares. We denote the union of these pixels by $W = [0, N] \times [0, N]$. The grey-levels of a given ultrasound image is denoted by $y = (y_s)_{s \in S}$, where $S = \{0.5, 1.5, \ldots, N - 0.5\} \times \{0.5, 1.5, \ldots, N - 0.5\}$ is the set of centers of pixels. Although the range of grey-level is $\{0, 1, \ldots, 255\}$ (where 0 is black and 255 is white), y_s is conceptually viewed as being real-valued.

For each ultrasound image, we use a simple observation model (likelihood) given the true coronary artery $V \subset W$: the y_s are assumed to be independent and normally distributed with mean $\mu_s(V)$ and variance $\sigma_s^2(V)$, where $(\mu_s(V), \sigma_s^2(V)) = (\mu_i, \sigma_i^2)$ if $s \in S_i$, where $S_1 = S \cap V$ (the inside of the coronary artery) and $S_2 = S \setminus V$ (the outside of the coronary artery). So the likelihood for the data y given $\theta = (\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)$ and V is specified by

$$\begin{aligned} L(y|V,\theta) &= (\sigma_1^2)^{-n_1/2}(\sigma_2^2)^{-n_2/2} \exp\left(-\mathrm{SS}_1/(2\sigma_1^2) - \mathrm{SS}_2/(2\sigma_2^2)\right) \\ &\times \exp\left(-n_1(\bar{y}_1 - \mu_1)^2/(2\sigma_1^2) - n_2(\bar{y}_2 - \mu_2)^2/(2\sigma_2^2)\right) \end{aligned} \tag{1}$$

where

$$n_i = \#(S_i \cap V), \quad \bar{y}_i = \sum_{s \in S_i} y_s/n_i, \quad \mathrm{SS}_i = \sum_{s \in S_i} (y_s - \bar{y}_i)^2, \quad i = 1, 2. \tag{2}$$

As the grey-levels inside the coronary artery in general are supposed to be smaller than those outside, the domain of variation for the parameter θ is assumed to be

$$0 < \mu_1 < \mu_2 < 255, \quad \sigma_1^2 > 0, \quad \sigma_2^2 > 0. \tag{3}$$

Note that owing to the discretization into pixels, there is for each fixed value of (y, θ) an upper bound and a strictly positive lower bound on the likelihood,

$$0 < c_1(y, \theta) \leq L(y|V, \theta) \leq c_2(y, \theta) < \infty. \tag{4}$$

In the sequel, we assume that V is a polygon specified by n vertices $v_1, \ldots, v_n$. Its boundary

$$\partial V = \bigcup_{j=1}^{n} [v_{j-1}, v_j]$$

is a closed self-avoiding planar contour, where $v_0 = v_n$ and $[v_{j-1}, v_j]$ is the closed line segment with endpoints v_{j-1} and v_j ("self-avoiding" means that for all $1 \leq j < k \leq n, [v_{j-1}, v_j] \cap [v_{k-1}, v_k]$ equals $\{v_k\}$ for $k = j + 1$ and $\emptyset$ otherwise). In some applications of object restoration it is relevant to let n be random (see, for example, Pievatolo and Green (1998) and Rue and Hurn (1999)). In our case, where ∂V is expected to be close to a circular n-gon, we found it sufficient and much simpler to fix $n = 16$.

3. Deformable template priors

For each single ultrasound image, we assume that the contour ∂V and the parameters σ_1^2, σ_2^2, and (μ_1, μ_2) are mutually independent, where σ_j^2 follows a flat IG(α_j, β_j)-distribution, $j = 1, 2$ (the inverse gamma distribution with shape parameter α_j and scale parameter β_j), while (μ_1, μ_2) is uniformly distributed on the set where $0 < \mu_1 < \mu_2 < 255$, cf. (3). The hyper parameters α_j and β_j are specified in Subsection 5.1. In Subsections 3.1 and 3.2 we specify two kinds of deformable template priors for ∂V. As discussed in Section 4, this can be combined with the likelihood (1) to obtain a joint distribution for the entire sequence of ultrasound images. Note that by (4) and the independence between ∂V and θ, posterior propriety is equivalent to prior propriety for ∂V.

3.1. A simple deformable template prior

The contour in the ith image is modeled as a deformation of a template

$$\partial V^0 = \bigcup_{j=1}^{n} [v_{j-1}^0, v_j^0]$$

which in our case is expected to be close to a circular n-gon. If $i = 1$, the vertices of the template are found by manual tracing of the artery wall. This is done by means of a computer screen and mouse where the operator clicks in the vertices on the displayed image such that the produced n-gon in a best way capture what is believed to be the artery wall, cf. the figures in Subsection 5.3. For $i > 1$, the template is given by the (estimated) posterior mean of the vertices of the deformed template in the $(i-1)$th image, cf. Subsection 4.2.

Prior models for deformable templates are usually obtained by specifying a conditional cyclic auto-regression (CAR) model on the vertices or edges before conditioning on closeness of the contour. Grenander *et al.* (1991), Grenander (1993), and Grenander and Miller (1994) propose a multiplicative model for the edges, while additive models for vertices and edges are considered in Kent *et al.* (1996). The CAR model we suggest is for the edges, and it is obtained as a certain limit of proper first-order cyclic CAR models so that closeness is automatically satisfied. The advantage of using this template model compared to those considered in the abovementioned references is explained below.

Specifically, let

$$e_{*l} = (e_{1l}, \ldots, e_{nl})^{\mathsf{T}}, \quad e_{*l}^0 = (e_{1l}^0, \ldots, e_{nl}^0)^{\mathsf{T}}, \quad l = 1, 2,$$

be the vectors of first ($l = 1$) and second ($l = 2$) coordinates of the edges in ∂V and ∂V^0 as defined by

$$e_j = v_j - v_{j-1}, \quad e_j^0 = v_j^0 - v_{j-1}^0, \quad j = 1, \ldots, n.$$

By closeness of the contours,

$$\sum_{j=1}^{n} e_{jl} = \sum_{j=1}^{n} e_{jl}^0 = 0, \quad l = 1, 2. \tag{5}$$

Before conditioning on that ∂V is closed and self-avoiding, we let e_{*1} and e_{*2} be independent and normally distributed with means e_{*1}^0 and e_{*2}^0, and a common precision matrix $\kappa^{-1}R(\gamma)$, where $\kappa > 0$ is the variance. Setting $r_{1,0} = r_{1,n}$ and $r_{n,n+1} = r_{n,1}$, the non-zero entries of $R(\gamma)$ are

$$r_{jj} = 1, \quad r_{j,j+1} = r_{j,j-1} = -\gamma/2, \quad j = 1, \ldots, n.$$

As $R(\gamma)$ is circulant (see Davis (1979)),

$$R(\gamma) = \frac{1}{n} F \mathtt{diag}\{d_1(\gamma), \ldots, d_n(\gamma)\} \overline{F}, \tag{6}$$

where $d_j(\gamma) = 1 - \gamma \cos(2(j-1)\pi/n)$, $j = 1, 2, \ldots, n$, are the eigenvalues of $R(\gamma)$, $F = \{\exp(2\pi i(j-1)(k-1)/n)\}_{j,k=1,\ldots,n}$ is the discrete Fourier transform matrix, and '$^-$' denotes complex conjugate. Hence $R(\gamma)$ is positive definite if $|\gamma| < 1$, but only positive semi-definite if $\gamma = 1$. Below we argue why we set $\gamma = 1$ in this model and we specify the corresponding covariance matrix.

Consider for the moment the case where $|\gamma| < 1$. Then the full conditional distribution of e_{jl} given the rest is normal with mean $e_{jl}^0 + \gamma(e_{j-1,l} - e_{j-1,l}^0 + e_{j+1,l} - e_{j+1,l}^0)/2$ and variance $1/\kappa$. However, when we condition on the linear constraint (5), the CAR model loses its simple dependence on nearest neighbors. In fact the conditional distribution of e_{*l} becomes normal with mean e_{*l}^0 and covariance matrix $\Sigma(\kappa, \gamma) = \kappa(I - E)R(\gamma)^{-1}(I - E)$, where I is the identity matrix and E has entries equal to $1/n$. By (6), $\Sigma(\kappa, \gamma) \to \kappa R(1)^-$ as $\gamma \to 1$, where $R(1)^-$ is the generalized inverse of $R(1)$ of rank $n - 1$, which we decompose as

$$R(1)^- = (n^{1/2} F \Lambda^{1/2})(n^{1/2} F \Lambda^{1/2})^\star \tag{7}$$

where $\Lambda^{1/2} = \mathtt{diag}(0, 1/d_2(1)^{1/2}, \ldots, 1/d_n(1)^{1/2})$ and '*' denotes conjugate transpose.

For our application, simulations from the posterior distribution (using a uniform prior on $]-1, 1[$ for γ and remaining priors as specified below) showed γ to be concentrated on values very close to 1. For these simulations one needs to find a square root of the precision matrix for each new value of γ, which is found in a way similar to (8) below.

In the sequel we assume that $\gamma = 1$, or more precisely that e_{*l} has covariance matrix $\kappa R(1)^-$. Then (5) is "automatically" satisfied. By (7), we can set

$$(e - e^0) = (n\kappa)^{1/2} F \Lambda^{1/2} \begin{bmatrix} 0 \\ z \end{bmatrix}, \tag{8}$$

where $e = e_{*1} + ie_{*2}$, $e^0 = e_{*1}^0 + ie_{*2}^0$, and $z = z_{*1} + iz_{*2}$ with z_{*1} and z_{*2} independent $(n-1)$-dimensional standard normally distributed. For the calculation of the right hand side in (8), we use a radix-2 FFT-algorithm (van Loan, 1992). Incidentally, we remark that

in many Bayesian applications (including many others than deformable template modeling) an improper prior density is used instead, provided the posterior becomes proper, see, for example, Besag and Kooperberg (1995), Besag and Higdon (1999), and Kent *et al.* (1996). The improper prior density for (e_{*1}, e_{*2}) is proportional to

$$\exp\left(-\sum_{l=1}^{2}\sum_{j=1}^{n}\left(e_{jl}-e_{jl}^0-(e_{j-1,l}-e_{j-1,l}^0)\right)^2/(4\kappa)\right)$$

(when κ is fixed). As noticed, in our case, this causes impropriety of the posterior; see also the discussion in Besag and Kooperberg (1995).

Finally, we condition on self-avoidingness of ∂V. Clearly this is only a condition on x and not on any other parameters like v_1 and θ. Note that ∂V is determined by (x, v_1, e^0) where $x = \sqrt{\kappa}z$. we assume that z, v_1, κ, θ are mutually independent, v_1 is uniformly distributed on W, and κ follows a flat $\mathrm{IG}(\delta, \epsilon)$ distribution. Note that we do not require any knowledge of the location v_1^0 of the template. The hyper parameters δ, ϵ, ω are specified in Subsection 5.1.

3.2. A refined deformable template prior

As mentioned in Section 2, in some of the images parts of the contour of the artery wall is missing. It turns out that if the simple prior of Subsection 3.1 is used, deformations of the template with sharp interior angles are likely because of artifacts, in particular lost parts of contours such as the drop outs shown in Figure 1. In order better to incorporate the prior knowledge that the coronary artery is close to a circular n-gon, we consider (before conditioning on ∂V being self-avoiding) a refined prior density for x defined as follows.

Given (κ, ω), where ω is a hyper parameter, the density of the refined prior is

$$\pi(x|\kappa,\omega) \propto \exp\left(-(||x_{*1}||^2+||x_{*2}||^2)/(2\kappa)-\omega\sum_{j=1}^{n}(\phi_j-(n-2)\pi/n)^2\right), \qquad (9)$$

where ϕ_j is the interior angle of V at the vertex v_j (measured in radians) and $(n-2)\pi/n$ is the interior angle in a regular n-gon. Setting $\omega = 0$ we obtain the prior in Subsection 3.1, but in accordance with our prior belief we shall assume that $\omega > 0$.

We compare in Subsection 5.3 the results of the Bayesian analysis using either the simple or refined prior.

Pievatolo and Green (1998) consider a somewhat similar model, but without incorporating a deformable template and the circular structure. Another approach was suggested by Grenander (1993, Remark 16.2.3). There the lost part of the boundary is modeled by a stochastic destructive deformation field acting on the template. Rue and Husby (1998) show that such a field is well suited to model partly destroyed objects. However, this approach complicates the statistical modeling and is heavily computer intensive. We therefore prefer the more simple approach described above.

4. Bayesian inference and simulation

4.1. Analyzing a single ultrasound image

For each image we have a hierarchical model given the edges e^0 of the template and the hyper parameters $\alpha_1, \alpha_2, \beta_1, \beta_2, \delta, \epsilon, \omega$. At the first level, the conditional distribution of y

given (x, v_1, κ, θ) depends only on $(\partial V, \theta)$ as described in Section 2. At the next level the conditional distribution of (x, v_1, θ) given κ is specified and at the third level the distribution of κ is specified, where $x/\sqrt{\kappa}, v_1, (\mu_1, \mu_2), \sigma_1^2, \sigma_2^2, \kappa$ are mutually independent with marginal distributions as described in Section 3. Thereby, using either the simple or refined prior, we obtain the posterior distribution of $r = (x, v_1, \kappa, \theta)$ given (y, e^0). Its density is

$$
\begin{aligned}
&\pi(r|y, e^0) \qquad (10)\\
&\propto \mathbf{1}\,[\partial V \text{ is self-avoiding}]\,\pi(x|\kappa,\omega)\pi(\kappa|\delta,\epsilon)\pi(v_1)\pi(\mu_1,\mu_2)\pi(\sigma_1^2|\alpha_1,\beta_1)\pi(\sigma_2^2|\alpha_2,\beta_2)L(y|V,\theta)\\
&= \mathbf{1}\,[\partial V \text{ is self-avoiding}]\,\kappa^{-\delta-n}(\sigma_1^2)^{-\alpha_1-n_1/2-1}(\sigma_2^2)^{-\alpha_2-n_2/2-1}\\
&\quad\times \exp\left(-(||x_{*1}||^2+||x_{*2}||^2)/(2\kappa)-\epsilon/\kappa-\beta_1/\sigma_1^2-\beta_2/\sigma_2^2-\mathrm{SS}_1/(2\sigma_1^2)-\mathrm{SS}_2/(2\sigma_2^2)\right)\\
&\quad\times \exp\Big(-n_1(\bar{y}_1-\mu_1)^2/(2\sigma_1^2)-n_2(\bar{y}_2-\mu_2)^2/(2\sigma_2^2)-\omega\sum_{j=1}^{n}\big(\phi_j-(n-2)\pi/n\big)^2\Big)
\end{aligned}
$$

with $x \in \mathbb{C}^{n-1}$, $v_1 \in W$, $\kappa > 0$, $0 < \mu_1 < \mu_2 < 255$, $\sigma_j^2 > 0$, $j = 1, 2$. Here it is only V and the sufficient statistics in (2) which depend on e^0. The posterior (10) is clearly analytical intractable, but Markov chain Monte Carlo (MCMC) methods provide a feasible way for performing Bayesian inference (see for example Gilks *et al.* (1996) for background material on MCMC).

We use a Metropolis-Hastings algorithm with proposals corresponding to updates of either

1 the means (μ_1, μ_2),

2 the variances (σ_1^2, σ_2^2),

3 the position v_1 of the deformed template,

4 the x associated to the edges of the deformed template, or

5 the hyper parameter κ.

Each of these five types of updates are chosen with equal probability $1/5$. For types **1**, **3**, and **4**, we use Metropolis random walk algorithms with the following kind of proposals: in **1**, a uniform distribution defined on a square of side length $a > 0$ centered around the current value of (μ_1, μ_2); similarly in **3**, where $b > 0$ denotes the side length of the square centered at the current value of v_1; and in **4**, a multivariate Gaussian distribution with mean equal to the current x value, independent coordinates (for the real and imaginary parts), and a common variance $c > 0$. The proposals are then accepted or rejected in accordance with the usual Metropolis-Hastings criterion. The user-specified parameters a, b, c are adjusted in order to obtain reasonable acceptance probabilities, see Subsection 5.2. Finally, for types **2** and **5**, we simply use Gibbs sampling updates corresponding to the following full conditionals obtained from (10): in **2**, $\sigma_i^2|\cdots \sim \mathrm{IG}(\alpha + n_i/2, \beta + n_i(\bar{y}_i - \mu_i)^2/2 + \mathrm{SS}_i/2)$ for $i = 1, 2$, and independent; in **5**, $\kappa|\cdots \sim \mathrm{IG}(\delta + n - 1, (||x_{*1}||^2 + ||x_{*2}||^2)/2)$.

The proposals are accepted or rejected in accordance to the usual Metropolis-Hastings criterion based on the Metropolis-Hastings ratio, and the algorithm can be shown to be ergodic with equilibrium distribution $r|(y, e^0)$. Note that many of the terms in (10) cancel when the Metropolis-Hastings ratio is calculated. The updates in **1** and **2** are particularly simple as $L(y|V, \theta')/L(y|V, \theta)$ depends on V only through the sufficient statistics in (2) which are all known from the previous update. The updates in **5** are likewise rather simple. The most time consuming steps are **3** and **4** because of the need for calculating a new

likelihood term, although this calculation can be restricted to pixels where the suggested update yields a change in the grey-level.

Preliminary runs of the algorithm showed that our 'parameterization' leads to better mixing properties than if we use the more immediate parameterization where x is replaced by z. In the latter case, we are modifying type **4** updates in the algorithm by updating z instead of x, using a proposal $z' \sim \mathcal{N}(z, cI)$. Our experiments showed that the proposal z' is likely to be rejected for modest and large values of κ, since $x' = \sqrt{\kappa}z' \sim \mathcal{N}(z, c\kappa I)$ is frequently not sufficiently small.

4.2. Analyzing the time series of ultrasound images

It is convenient to analyze the images of the sequence of ultrasound images $y(i)$, $i = 1, \ldots, 49$, in turn using the Metropolis-Hastings algorithm in Subsection 4.1 with equilibrium distribution $r(i)|(y(i), e^0(i))$ given by (10), where $r(i) = (x(i), v_1(i), \kappa(i), \theta(i))$ are the 'parameters' and $e^0(i)$ is the template of edges associated to the ith image. Under the model assumptions imposed below, we show in (11) that this is equivalent to analyzing the posterior distribution based on all images.

Assume that the sequence $u(i) = (y(i), r(i), e^0(i))$, $i = 1, \ldots, 49$, forms a Markov chain, which can be viewed as a directed graphical model (Lauritzen, 1996) with a conditional dependence structure as shown in Figure 2. Note that for each $i > 1$, $(y(i), r(i))$ is conditionally independent of $(u(1), \ldots, u(i-1))$ given $e^0(i)$. Further, assume that $e^0(1)$ is given and define recursively

$$e^0(i) = \mathbf{E}[e(i-1)|y(i-1), e^0(i-1)], \quad i = 2, \ldots, 49,$$

i.e. the edges of the template for each image $i > 1$ are given by the posterior mean of deformed edges from the previous image. Since $(e^0(2), \ldots, e^0(49))$ is determined by $(y(1), \ldots, y(49), e^0(1))$, it follows from Figure 2 that $r(1), \ldots, r(49)$ are conditionally independent given $(y(1), \ldots, y(49), e^0(1))$. Finally, assume that for each $i \geq 1$, $y(i)|(r(i), e^0(i))$ is distributed as $y(i)|(V(i), \theta(i))$ (see Section 2) and $r(i)$ follows the same distribution as $r = (x, v_1, \kappa, \theta)$ (see Section 3). Then the posterior distribution given the data $(y(1), \ldots, y(49))$ and the first template of edges $e^0(1)$ has density

$$\pi(r(1), \ldots, r(49)|y(1), \ldots, y(49), e^0(1)) \propto \prod_{i=1}^{49} \pi(r(i)|y(i), e^0(i)), \tag{11}$$

that is the product of posterior densities (10) associated to each image $i = 1, \ldots, 49$. Below we consider estimates of posterior means like $\mathbf{E}\left[g(r(i))|y(i), e^0(i)\right]$ based on simulated samples from $r(1)|(y(1), e^0(1)), \ldots, r(49)|(y(49), e^0(49))$, respectively, using in each case the Metropolis-Hastings algorithm in Subsection 4.1. Ignoring the fact that in practice we need succesively to estimate the templates as explained below, this is equivalent to simulate samples from the posterior distribution (11) for all images.

One may ask for an extension of the model for $(u(1), \ldots, u(49))$ so that posterior information from the $(i-1)$th image is taken more into account when modeling the ith image. For example, in Figure 2 the hyper parameters $\alpha_1, \alpha_2, \beta_1, \beta_2, \delta, \epsilon, \omega$ are left unchanged for all images, but as reported in Subsection 5.1 the effect of incorporating posterior information about e.g. (β_1, β_2) was minor in our application. In fact, as explained below we do incorporate posterior information when starting our MCMC runs.

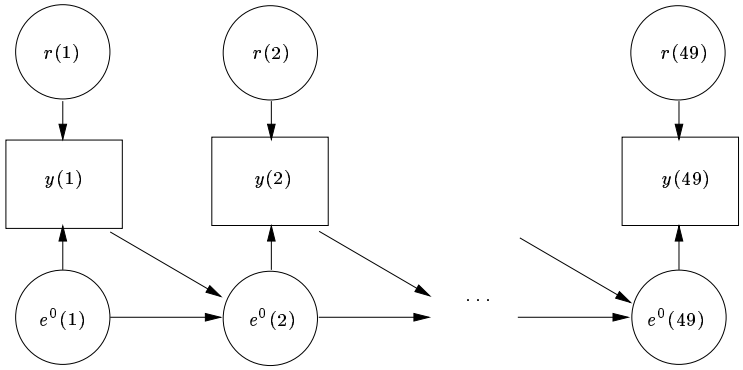

Fig. 2. Graphical representation of a joint model for all images and associated parameters.

We describe now in more detail how simulations from the posterior (11) are performed.

For the first image the template is initially determined manually using a computer mouse as described in the beginning of Subsection 3.1. As this estimate of the template is observer dependent, we generate first a chain $r_1^{(0)}(1), \ldots, r_J^{(0)}(1)$ of length $J > 0$ by the Metropolis-Hastings algorithm with equilibrium distribution $r|(y(1), e^0(1))$ and initial values specified as follows: θ is the moment estimate $\theta_0 = (\mu_{01}, \mu_{02}, \sigma_{01}^2, \sigma_{02}^2)$ obtained from the likelihood (1) with $(y, V) = (y(1), V^0(1))$; $x = 0$ and v_1 is the first vertex 'clicked in'; κ is set equal to its prior mean. Then we update the template by the Monte Carlo estimate based on $r_L^{(0)}(1), \ldots, r_J^{(0)}(1)$ using an appropriate burn-in $L < J$. This estimate is used throughout when generating further chains for the first image.

Next we are running through the images $i = 1, \ldots, 49$, generating a chain of length $K > L$ for each image i. Let $r_1(i), \ldots, r_K(i)$ denote the chain for the ith image with equilibrium distribution $r|(y(i), e^0(i))$. As the initial value of x we use again $x = 0$ for all images. If $i = 1$, as the initial value of v_1 we use the Monte Carlo estimate based on $r_L^{(0)}(1), \ldots, r_J^{(0)}(1)$, while the initial values of θ and κ are specified as before. If $i > 1$ then we estimate $e^0(i)$ and use as the initial values of θ, v_1, κ the estimated Monte Carlo posterior means obtained from $r_L(i-1), \ldots, r_K(i-1)$.

In order to obtain more accurate results we simulate M independent realizations

$$\left(r_1^{(u)}(1), \ldots, r_K^{(u)}(1), \cdots, r_1^{(u)}(49), \ldots, r_K^{(u)}(49)\right), \; u = 1, \ldots, M,$$

in the same way as we generate $(r_1(1), \ldots, r_K(1), \cdots, r_1(49), \ldots, r_K(49))$. Then, for any statistic $g(r(i))$ associated to the ith image, we obtain a Monte Carlo estimate of its posterior mean,

$$\mathbf{E}\left[g(r(i))|y(i), e^0(i)\right] \approx \sum_{u=1}^{M} \sum_{t=L}^{K} g(r_t^{(u)}(i))/(M(K-L+1)). \tag{12}$$

The values of J, K, L, M are specified in Section 5.2, and (12) is used in Figures 5–8 when estimating posterior means and percentiles.

5. Empirical results

In this section we discuss some empirical results based on the Bayesian analysis described in Section 4. Unless otherwise stated the refined prior of Subsection 3.2 has been used.

5.1. Specification of hyper parameters

For our application we found it appropriate to use the same value of the hyper parameters $\alpha_1, \beta_1, \alpha_2, \beta_2, \epsilon, \delta, \omega$ for all images. For the $\mathrm{IG}(\delta, \epsilon)$ prior of κ, preliminary runs showed that the posterior (10) is robust against different choices of $\delta \leq 2$ and $\epsilon > 0$ corresponding to a flat prior. Note that $\delta \leq 2$ is equivalent to assume an infinite variance of $\kappa \sim \mathrm{IG}(\delta, \epsilon)$. We have chosen $\delta = 2$ and $\epsilon = 10$. Similarly, for $\sigma_j^2 \sim \mathrm{IG}(\alpha_j, \beta_j)$, we set $\alpha_j = 2,\ j = 1, 2$. Then β_j is the prior mean of σ_j^2, so we set $\beta_j = \sigma_{0j}^2,\ j = 1, 2$, where σ_{0j}^2 is the moment estimate based on the first image as described in Subsection 4.2. Finally, we found $\omega = 1000$ to be an appropriate value for the refined prior. The results presented in the following seemed rather robust for choices of $500 \leq \omega \leq 3000$.

We also investigated the effect of changing the values of the hyper parameters for each image $i = 1, \ldots, 49$. For example, if $i > 1$, each $\sigma_{0j}^2(i)$ may be estimated by the MCMC estimate of the posterior mean $\mathbf{E}[\sigma_{0j}^2(i-1)|y(i-1)]$ from the $(i-1)$th image. However, the results were not seriously effected by such a change.

5.2. Output analysis of MCMC runs

The parameters a, b, c in the random walk Metropolis algorithms are chosen to obtain reasonable acceptance rates. This is exemplified in Figure 3 for the special case of c when the chain $r_1^{(0)}(1), \ldots, r_K^{(0)}(1)$ with $K = 50000$ is generated.

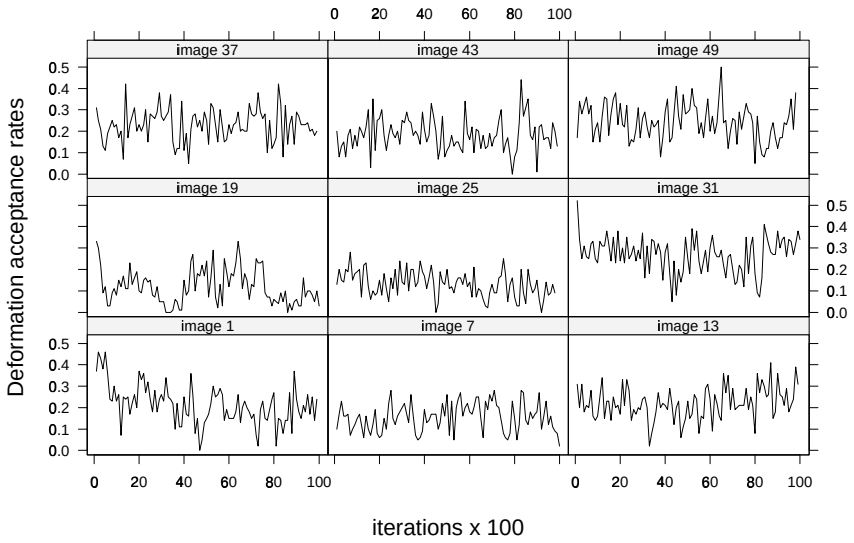

Fig. 3. Acceptance rates of updates of template deformations for images $1, 7, \ldots, 49$. The horizontal axis refers to each one hundred of type **4** updates in the Metropolis-Hastings algorithm when the total number of type **1**–**5** updates is 50000.

The performance of the Metropolis-Hastings algorithm was monitored with respect to various statistics for the different images. Time series plots (not shown here) indicate fast convergence with respect to θ and κ. For these parameters except the mean (μ_1, μ_2), estimated autocorrelations (not shown here) are effectively 0, while for the mean they are nearly 0 at lags larger than 20. Times series for the area of the deformed template (not shown here) are heavily correlated and show a much slower convergence, but reliably results are obtained by repeated runs (i.e. when $M \gg 1$). Since the template scales and moves from image to image, a value of $K \gg 1$ is needed. In conclusion we found it appropriate to use $J = 50000$, $K = 10000$, $L = 1000$, and $M = 20$ (these values and (12) are used in Figures 5–8 when estimating posterior means and percentiles). For example, Figure 4 shows time series for the area of the deformed template when $L = 1000$ and $M = 20$, but the chain length is either $K = 5000$, 10000, 15000, or 20000. The posterior means and 95% credibility intervals shown in Figure 4 are rather stable for all the time series with values of $K \geq 10000$, while the result for $K = 5000$ is clearly different from the other cases.

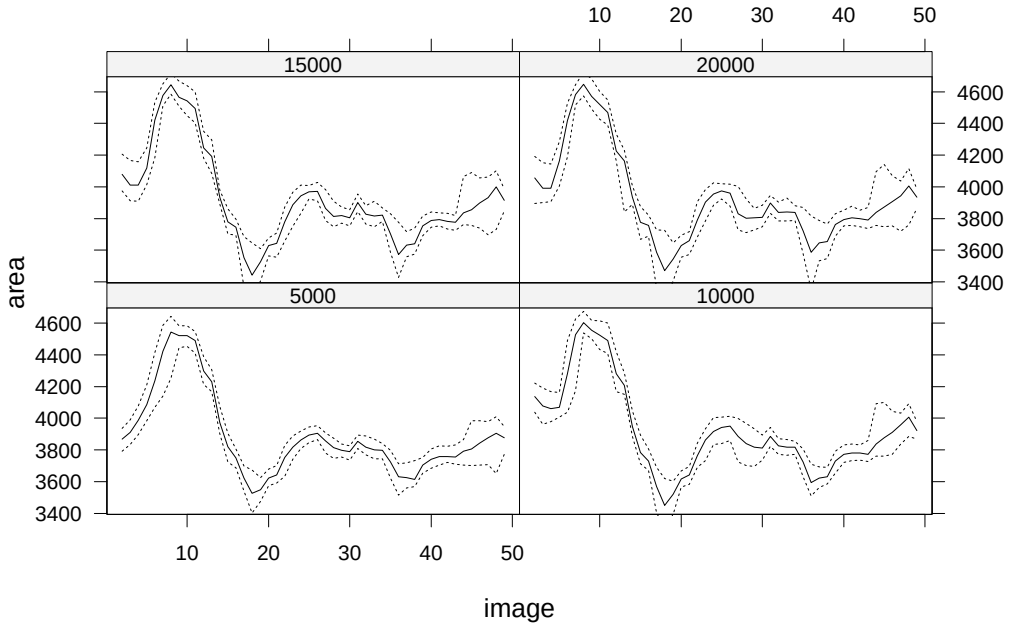

Fig. 4. Posterior means and 2.5% and 97.5% percentiles for the area of the deformed template for each image when different chain lengths are used. From bottom left to top right: $K = 5000$, $K = 10000$, $K = 15000$, $K = 20000$.

5.3. Comparison of results for the simple and refined priors

We have compared runs using the simple prior ($\omega = 0$) with those for the refined prior ($\omega = 1000$). For the simple prior, Figure 5 shows every sixth image with the initial template (lines with circles) and posterior means of the deformed templates (lines with crosses) superimposed. The deformed templates show clearly a departure from a circular n-gon,

and especially where the contour is broken, the deformed templates seem to include areas outside the missing contour. Using the refined prior, much better results are obtained, see Figure 6. Now the deformed templates resembles circular n-gons and they do not tend to include areas outside the missing contour.

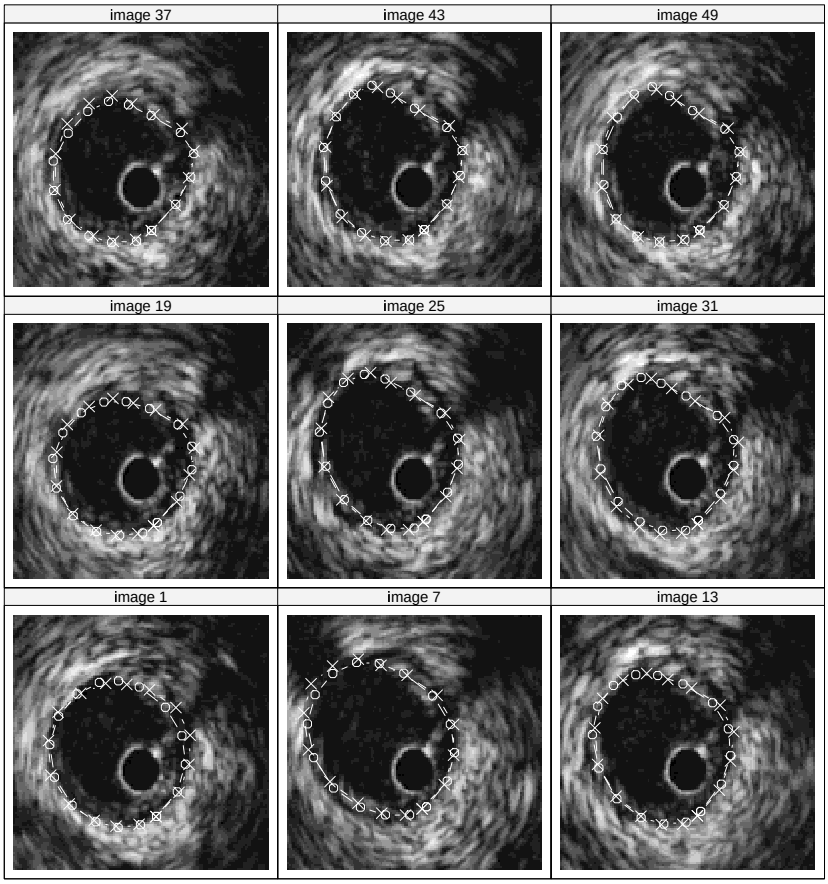

Fig. 5. Images $1, 7, \ldots, 49$ with superimposed template (lines with circles) and posterior means of deformed templates (lines with crosses). The simple prior model was used.

5.4. Conclusions

We have outlined how a fully Bayesian analysis can be performed for the time sequence of ultrasound images. Owing to the results reported in Subsection 5.3 we shall base our inference on the refined prior.

The dynamic behavior of the artery wall may be summarized by plots of the area of V and κ over time, while plots of θ may show the dynamics of the equipment dependent circumstances under which the images were taken. Figure 7 shows time series of the posterior means and 95% credibility intervals for the area of V and the parameters $\mu_1, \mu_2, \sigma_1^2, \sigma_2^2, \kappa$.

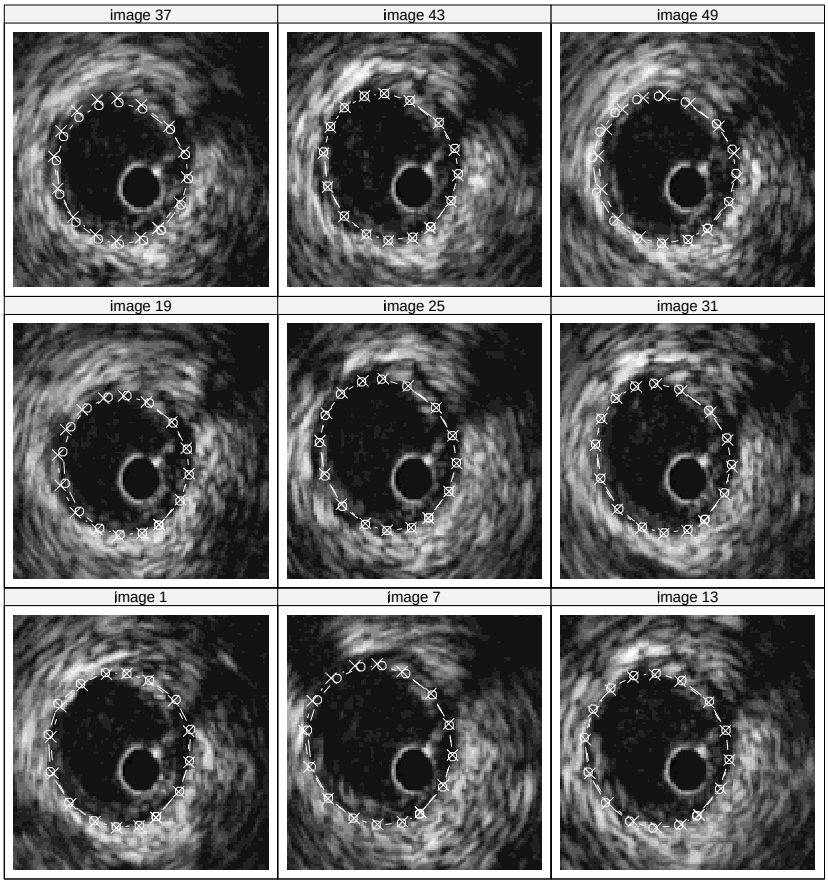

Fig. 6. Images $1, 7, \ldots, 49$ with superimposed template (lines with circles) and posterior means of deformed templates (lines with crosses). The refined prior model was used.

The size of the credibility interval is rather stable in each of the six plots. The posterior distribution of κ is effectively the same for all images. The curves for $\mu_1, \mu_2, \sigma_1^2, \sigma_2^2$ show a somewhat parallel behavior, where σ_2^2 is at least twice as large as σ_1^2, and μ_2 is clearly larger than μ_1 also. There is no clear relationship between the plot of the area and the other plots.

The area of the coronary artery on images 6–9 seems to be overestimated in Figure 7. Figure 8 shows a plausible reason for this: owing to the drop outs, the estimated posterior means of ∂V for images 6–9 are less circular than those for images 5 and 10, and the estimated posterior means of ∂V for images 6–9 seem to include area belonging to the region outside the contour of the artery. Using a larger value of ω would partly correct for this, but would also necessitate larger runs, when the random walk scale parameter c has to be adjusted to a smaller value in order to obtain acceptance probabilities of the same value as before. A more realistic model that takes into account broken contours may be needed

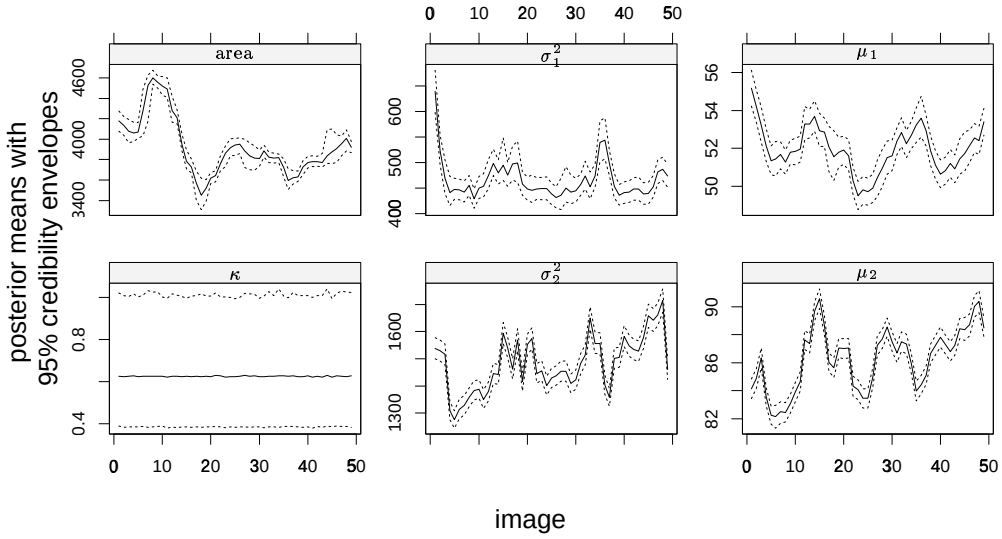

Fig. 7. Posterior means and 2.5% and 97.5% percentiles for each image and different statistics. From top left to bottom right: the area of V and the parameters $\sigma_1^2, \mu_1, \kappa, \mu_2, \sigma_2^2$.

to fully remedy the problem.

In conclusion plots like in Figures 6–8 provide useful posterior information about the dynamic behavior of the artery and the quality of the results.

6. Acknowledgments

We would like to thank Bent Østergaard Kristensen, Department of Cardiology, Aarhus University Hospitals, Skejby, for providing the ICUS image sequence and for helpful comments. Our colleagues Claus Dethlefsen and Søren Lundbye-Christensen are acknowledged for introducing us to the problem and for helpful discussions. The research has been supported by the European Union's research network "Statistical and Computational Methods for the Analysis of Spatial Data. ERB-FMRX-CT96-0095". MBH and JM were supported by MaPhySto, Centre for Mathematical Physics and Stochastics, funded by a grant from The Danish National Research Foundation. JM and FAaT were supported by the Danish Informatics Network in the Agricultural Sciences.

References

Aizenman, M. (2000). 2D+T acoustic boundary detection in echocardiography. *Medical Image Analysis* To appear.

Baddeley, A. J. and van Lieshout, M. N. M. (1993). Stochastic geometry models in high-level vision. In: *Statistics and Images, Advances in Applied Statistics, a supplement to the Journal of Applied Statistics* (eds. K. Mardia and G. Kanji), volume 20, chapter 11, Carfax publishing, Abingdon, 235–256.

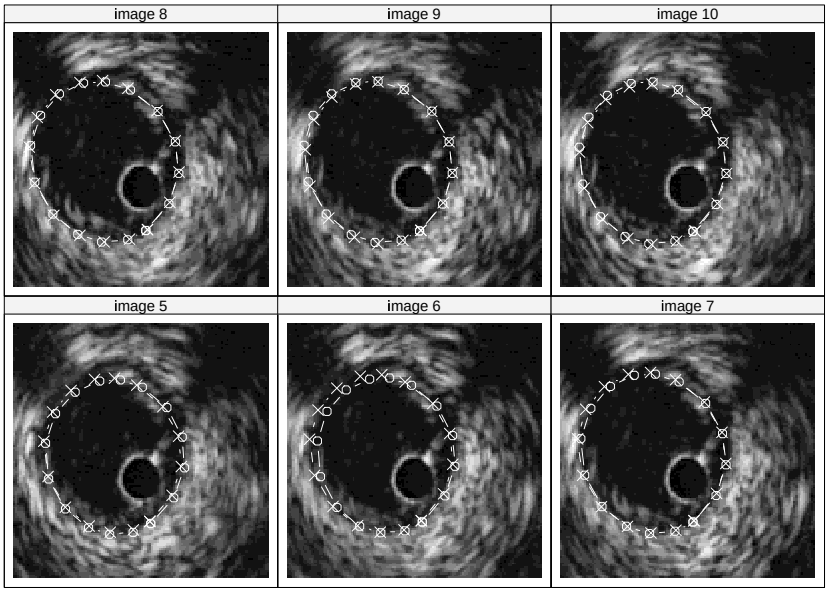

Fig. 8. Images 5–10 with superimposed template (lines with circles) and posterior means of deformed templates (lines with crosses).

Besag, J. and Higdon, D. M. (1999). Bayesian analysis of agricultural field experiments (with discussion). *J. R. Statist. Soc.,* B **61**, 691–746.

Besag, J. and Kooperberg, C. (1995). On conditional and intrinsic autoregressions. *Biometrika,* **82**, 733–746.

Davis, P. J. (1979). *Circulant Matrices*. Wiley, New York.

Gilks, W., Richardson, S. and Spiegelhalter, D., eds. (1996). *Markov Chain Monte Carlo in Practice*. Chapman and Hall, London.

Grenander, U. (1993). *General Pattern Theory*. Oxford University Press, New York.

Grenander, U. and Miller, M. I. (1994). Representations of knowledge in complex systems (with discussion). *J. R. Statist. Soc.,* B **56**, 549–603.

Grenander, U., Chow, Y. and Keenan, D. M. (1991). *Hands. A Pattern-Theoretic Study of Biological Shapes*. Research Notes in Neural Computing. Springer, Berlin.

Hobolth, A. and Jensen, E. B. V. (1999). Modelling stochastic changes in curve shape, with an application to cancer diagnostics. Research Report no. 5, Laboratory for Computational Stochastics, Department of Mathematical Sciences, University of Aarhus.

Hobolth, A. and Pedersen, J. (1999). Parameter estimation in a class of periodic, gaussian residual processes. Research Report no. 8, Laboratory for Computational Stochastics, Department of Mathematical Sciences, University of Aarhus.

Hurn, M. A., Steinsland, I. and Rue, H. (1999). Parameter estimation for a deformable template model. Research Report no. 5, Norwegian University of Science and Technology, Trondheim, Norway.

Kass, M., Witkin, A. and Terzopoulos, D. (1987). Snakes: active contour models. *Internat. J. Comp. Vis.* **1**, 321–331.

Kent, J. T., Mardia, K. V. and Walder, N. A. (1996). Conditional cyclic Markov random fields. *Adv. Appl. Prob.,* **28**, 1–12.

Kervrann, C. (1998). A hierarchical Markov modeling approach for the segmentation and tracking of deformable shapes. *Graphical Models and Image Processing,* **60**, 173–195.

Kervrann, C. (1999). Statistical deformable model-based segmentation of image motion. *IEEE Trans. Image Process* **8**, 583–588.

Lauritzen, S. L. (1996). *Graphical Models*. Clarendon Press, Oxford, UK.

Pievatolo, A. and Green, P. J. (1998). Boundary detection through dynamic polygons. *J. R. Statist. Soc.,* B **60**, 609–626.

Rue, H. and Hurn, M. A. (1999). Bayesian object identification. *Biometrika,* **86**, 649–660.

Rue, H. and Husby, O. K. (1998). Identification of partly destroyed objects using deformable templates. *Stat. Comp.,* **8**, 221–228.

Rue, H. and Syversveen, A. R. (1998). Bayesian object recognition with Baddeley's delta loss. *Adv. Appl. Prob.,* **30**, 64–84.

van Loan, C. (1992). *Computational Frameworks for the Fast Fourier Transform*. SIAM, Philadelphia.