NORGES TEKNISK-NATURVITENSKAPELIGE
UNIVERSITET

Estimating Blood Vessel Areas in Ultrasound Images Using a Deformable Template Model

by

Oddvar Husby and Håvard Rue

PREPRINT
STATISTICS NO. 17/2001
ISSN: 0804-9173

NORWEGIAN UNIVERSITY OF SCIENCE AND
TECHNOLOGY
TRONDHEIM, NORWAY

This report has URL http://www.math.ntnu.no/preprint/statistics/2001/S17-2001.ps
Håvard Rue has homepage: http://www.math.ntnu.no/~hrue
E-mail: hrue@math.ntnu.no
Address: Department of Mathematical Sciences, Norwegian University of Science and Technology,
N-7491 Trondheim, Norway.

Estimating blood vessel areas in ultrasound images using a deformable template model

Oddvar HUSBY and Håvard RUE
Department of Mathematical Sciences
Norwegian University of Science and Technology
Norway

October 9, 2001

Abstract

We consider the problem of obtaining interval estimates of vessel areas from ultrasound images of cross sections through the carotid artery. Robust and automatic estimates of the cross sectional area is of medical interest and of help in diagnosing atherosclerosis, which is caused by plaque deposits in the carotid artery. We approach this problem by using a deformable template to model the blood vessel outline, and use recent developments in ultrasound science to model the likelihood. We demonstrate that by using an explicit model for the outline, we can easily adjust for an important feature in the data: strong edge reflections called specular reflection. The posterior is challenging to explore, and naive standard MCMC algorithms simply converge to slowly. To obtain an efficient MCMC algorithm we make extensive use of computational efficient Gaussian Markov Random Fields, and use various block-sampling constructions that jointly update large parts of the model.

1 INTRODUCTION

Ultrasound is widely used in medical imaging, mainly because of its real-time imaging capability, low cost, relative safety, and the availability of portable units. Increasingly, medical ultrasound images are used in complicated applications such as surgery (Langø, 2000) and evaluation of cardiac diseases (Mulet-Parada and Noble, 2000). Common to these and other applications is the need for detecting anatomical boundaries, eg. for detecting the position of a surgical instrument, measuring response to some stimuli, or estimating quantities such as ejection fraction and wall motion. To focus the discussion, we consider the ultrasound images in Figure 1a and b, showing cross sections through the carotid artery of a single individual. The carotid artery is susceptible to plaque deposits, a condition called atherosclerosis. The condition is usually diagnosed by angiography, but this method does not work for all patients. An alternative is to use ultrasound, since the vessel walls appear more clearly in that particular image modality. It is known that healthy arteries will dilate in response to

Oddvar Husby is PhD, and Håvard Rue is Professor at Department of Mathematical Sciences, Norwegian University of Science and Technology, N-7491 Trondheim, Norway. O. Husby was supported by a grant from the Research Council of Norway. The authors thank Dr. Thomas Langø at Sintef Unimed for providing the ultrasound data, and Sveinung Erland for helpful comments and suggestions.

infusion of acetylcholine, and thus comparison of the cross sectional vessel areas before and after infusion may help in diagnosing atherosclerosis. For this method to be of any use, not only point estimates, but also uncertainty estimates must be provided, and we therefore aim at producing interval estimates of the quantity of interest.

We construct a Bayesian model which is analyzed using MCMC. In order to obtain reliable estimates we need both a good model for the data formation process, ie. the likelihood, and a well-designed MCMC sampling scheme. In particular, we need to update jointly all or nearly all parameters in our model, an approach which is computational feasible due to recent advances in sampling and design of Gaussian Markov Random Fields (Rue, 2001; Rue and Tjelmeland, 2001).

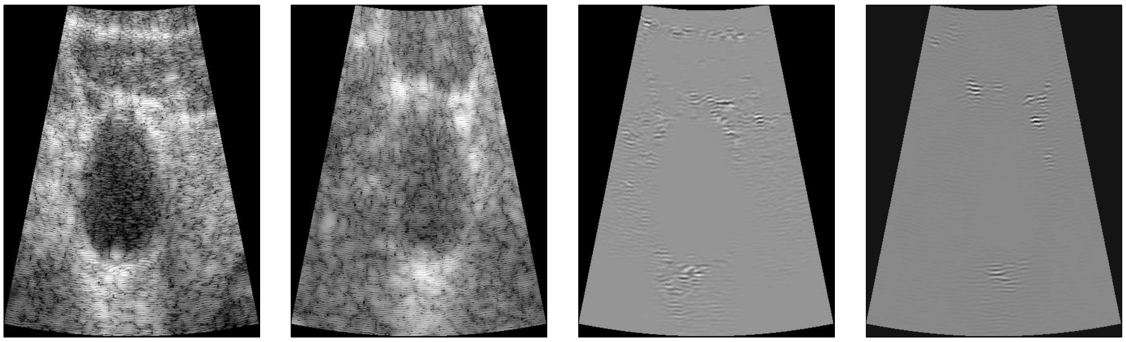

(a) (b) (c) (d)

Figure 1: Ultrasound images of cross sections through the carotid artery. Panels (a)-(b) show the log-compressed intensity images as they appear on the screen, while panels (c)-(d) show the raw data – or radio frequency image – collected by the ultrasound scanner.

There exists a vast literature on contour detection in medical images, see eg. Duncan and Ayache (2000) and Pham, Cu and Prince (2000) for recent reviews. Most of these methods, including the popular active contour models (Kass, Witkin and Terzeopoulos, 1988), are deterministic, thus only providing point estimates of the contour. In a more statistical setting, a popular approach is to use the deformable template models introduced by Ulf Grenander, see Grenander (1993) for a thorough review. Deformable template models have been used for such diverse tasks as recognizing hands (Grenander, Chow and Keenan, 1991), galaxies (Ripley and Sutherland, 1990), potatoes (Grenander and Manbeck, 1993), cells (Grenander and Miller, 1994; Rue and Syversveen, 1998; Rue and Hurn, 1999), magnetic domains (Qian, Titterington and Chapman, 1996), mushrooms (Mardia, Qian, Shah and de Souza, 1997), fish (de Souza, Kent and Mardia, 1999), and roads (Stoica, Descombes and Zerubia, 2000). In ultrasound images there exists a few Bayesian approaches to contour detection, using eg. deformable templates (de Figueiredo and Leitao, 1992; Hansen, Møller and Tøgersen, 2001), point distribution models (Glasbey, 1998), and line processes (Kao, Pan, Hiller and Chen, 1998).

In ultrasound imaging a pulse of ultra high frequency sound is sent into the body, and the backscattered signal is measured after some time delay corresponding to the desired depth. Compared to most other image modalities, contour detection in ultrasound images is particularly difficult because of the low resolution caused by noise, blur, edge reflections, and image artifacts called speckle. Speckle is caused by the coherent detection of the acoustic echoes, and give ultrasound images their characteristic granular appearance. Another important feature is the strong reflections at tissue interfaces being approximately perpendicular to the incoming pulse. This effect is called specular reflection, and is clearly seen in Figure 1a and b; note in particular the strong signal on the upper and lower part of the vessel. Speckle and specular reflection seriously affects the quality, accuracy and robustness of those contour detection methods in which the likelihood is defined through intensity gradients in the image. This is because intensity gradients in ultrasound images often have large variations, are missing or are artificial. It is tempting to resolve the data modeling issue based on the intensity data available on the screen, but this approach is hard for several reasons and is still not solved satisfactory.

In our approach we use a model that takes into account how speckle patterns and specular reflections are formed. Contrary to what is common, we do not use the processed ultrasound image as it appears on the screen, but instead collect the raw data received by the scanner, before any pre-processing is done. The corresponding data for Figure 1a and b are provided in Figure 1c and d. The raw data – or radio frequency image – contains more information than the commonly used intensity image, and moreover there exists a tractable, physical model for how these data are formed (Goodman, 1975; Wagner, Insana and Brown, 1987). One drawback of using the raw data is that they are not easily accessible but require some engineering work to tap the signal in the correct place. Hokland and Kelly (1996) and Husby, Lie, Langø, Hokland and Rue (2001) have investigated image restoration using the raw data with binary and gray-level pixel prior models, but neither are able to provide us with robust interval estimates of the vessel area.

The purpose of this work is to demonstrate how a deformable template model for the vessel contour naturally fits into the data model for the radio frequency image using the Bayesian paradigm. Further, a direct prior model for the contour makes it easy to correct for specular reflections. However, the posterior density is challenging to explore using standard MCMC methodology, and the design of a fast and robust sampling scheme was a challenge; even quite involved schemes gave severely biased interval estimates for the vessel area. We have made use of three important ingredients to construct a robust and fast MCMC sampling scheme: the ability of Gaussian Markov Random Fields (GMRF) to approximate stationary Gaussian fields on a lattice (Rue and Tjelmeland, 2001), fast sampling of GMRFs based on numerical algorithms for sparse matrices (Rue, 2001), and knowledge on how to construct joint proposals for use in MCMC algorithms for situations with a few hyperparameters controlling one or many Gaussian or near Gaussian fields (Knorr-Held and Rue, 2002).

The article is organized as follows. In section 2 we present details of the model, including the deformable template model for the vessel wall, and the likelihood model for ultrasound image formation. Inference and MCMC sampling are discussed in section 3, and results are presented in section 4.

2 MODEL FORMULATION

2.1 Data Formation

2.1.1 A Model for Radio Frequency Ultrasound Images A statistical model for radio frequency signals was first described in Goodman (1975), and the model has later been shown to be reasonable for body liquids and most soft tissues (Insana, Wagner, Garra, Brown and

Shawker, 1986). Under the model a body liquid or tissue is seen as a collection of point scatterers lying in a uniform non-scattering medium. The incoming ultrasound beam is reflected at each scatterer, and the sum of these reflections is the received signal. The dominating part of the scattering is called diffuse scattering, which occurs when there is a large number of randomly located scatterers of roughly equal size. It is assumed that the spatial variation in density is small relative to the resolution of the image. Then the diffuse scattering signal $\boldsymbol{x}$ consists of n independent Gaussian random variables x_i having zero mean and a variance dependent on the acoustical properties of the tissue or body liquid. The variance will only in ideal situations be constant within each tissue type or body liquid, but will in practice be spatially smooth.

In our case the true scene consists of two natural regions, one corresponding to the interior of the blood vessel and the other to the surrounding muscle tissue. Let $\boldsymbol{t}$ define the boundary of the blood vessel, and $\boldsymbol{\nu}_0$ and $\boldsymbol{\nu}_1$ be the log variance of $\boldsymbol{x}$ in the exterior and interior region, respectively. The assumption of independence between the x_i's is not realistic in all situations, but reasonable in our case.

2.1.2 Modeling Edge Reflections This signal $\boldsymbol{x}$ is the main component of the raw data. However, at interfaces between different tissue types or at tissue/blood interfaces, there is a strong reflection component due to the abrupt change in acoustical impedance, which is the mass density of the tissue times the speed of sound. The effect is called specular reflection and is strongest when the tissue interface is approximately perpendicular to the incoming pulse. The specular reflection is determined by the reflection coefficient (Christensen, 1988)

$$R = \frac{\frac{Z_2}{\cos\theta_2} - \frac{Z_1}{\cos\theta_1}}{\frac{Z_2}{\cos\theta_2} + \frac{Z_1}{\cos\theta_1}}, \tag{1}$$

where Z_1 and Z_2 are the acoustic impedances on the incident and transmitted side of the interface, respectively; and θ_1 and θ_2 are the angles between the interface and the incoming and transmitted beams, see Figure 2a. The incident and transmitted beam angles θ_1 and θ_2 are related by Snell's law $\sin\theta_1/\sin\theta_2 = c_1/c_2$, where c_1 and c_2 are the beam velocities. The coefficient in Eq. (1) measures the total reflection, but the amount of the reflection returning to the ultrasound transducer will depend on the angle of incidence. We follow Hokland and Kelly (1996) and model the specular reflection returned to the transducer as

$$r_i = \frac{\frac{Z_2}{\cos\theta_2} - \frac{Z_1}{\cos\theta_1}}{\frac{Z_2}{\cos\theta_2} + \frac{Z_1}{\cos\theta_1}} \cos^{\gamma}\theta_1, \tag{2}$$

for a site i on the tissue interface. Figure 2b shows the reflection coefficient R measured counter-clockwise along a circle, starting on the middle right of the circle. We use $Z_1 = 162.0\ \mathrm{kg/(s \cdot cm^2)}$ and $Z_2 = 166.6\ \mathrm{kg/(s \cdot cm^2)}$ corresponding to muscle and blood, respectively (Christensen, 1988). Note that when the beam velocity is higher in the transmitted medium than in the incident medium, there exists critical angles for which no energy is transmitted.

The raw data is, after correcting for specular reflection, $\boldsymbol{x} + \rho\boldsymbol{r}$, where the constant ρ is the relative magnitude of the specular and diffuse scattering component. ρ can be estimated from the generalized spectrum of the observed raw data (Varghese, Donohue and Chatterjee, 1995), but we find it easier to treat ρ as an unknown parameter. Note that the specular reflection is clearly visible in Figure 1c and d, as parts of the data has non-zero mean.

It is difficult to model the specular reflection without an explicit model for the contour of the vessel, as we need to know both its position and tangent to properly correct for the specular reflection. This is one motivation for the later use of an explicit model for the contour of the vessel using a deformable (polygon) template as a prior. Note however that the model for specular reflection does not correct for all degradation effects, such as constructive and destructive interference due to the spacing of the specular reflectors, but these effects are in general more difficult to model.

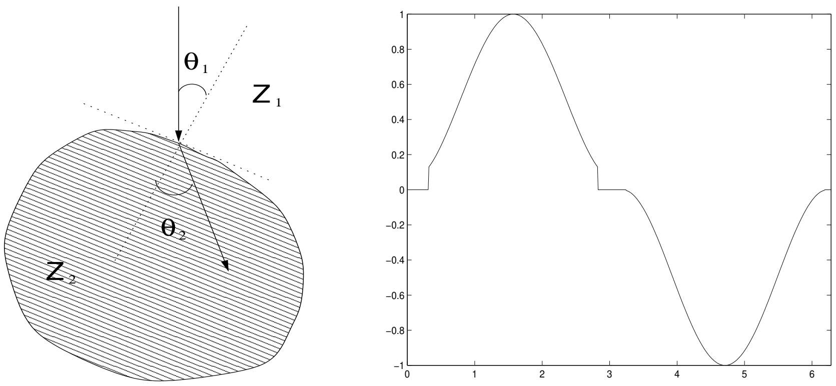

(a) (b)

Figure 2: Specular reflection. (a) Diagram showing the path of the ultrasound pulse through a boundary between tissues having acoustic impedances Z_1 and Z_2. (b) The specular reflection coefficient R measured counter-clockwise from the middle right of a circle. The impedances inside and outside the circle are $Z_2 = 166.6$ kg/(s $\cdot$ cm^2) and $Z_1 = 162.0$ kg/(s $\cdot$ cm^2), respectively.

2.1.3 Observation Model The observed radio frequency image $\boldsymbol{y}$ will at site i be the superposition of signals from a small neighborhood around i. We model $\boldsymbol{y}$ as a convolution of $\boldsymbol{x} + \rho\boldsymbol{r}$ with a point spread function $\boldsymbol{h}$, with additive independent zero mean Gaussian noise with unknown variance $1/\lambda$. We use the point spread function

$$h_{kl} \propto \exp\left(-\frac{k^2}{2\sigma_1^2} - \frac{l^2}{2\sigma_2^2}\right) \cos\frac{2\pi k}{\omega}, \tag{3}$$

where k and l indexes the radial and lateral dimensions of the image, respectively. Experiments have verified this to be a good approximation (Ødegård, 1995) for images similar to ours. To avoid over-parameterising the model we fix ω, σ_1 and σ_2 from inspection of the observed image and its derived frequency spectrum. Experiments indicate that similar methods are insensitive to slight misspecifications of these parameters (Langø, Lie, Husby and Hokland, 2001).

2.2 Prior model

2.2.1 A prior model for the acoustical parameters In section 2.1.1 we defined $\boldsymbol{\nu}^{(0)}$ and $\boldsymbol{\nu}^{(1)}$ as the log variances of $\boldsymbol{x}$ in the exterior and interior, respectively. To model their smooth behavior, we let each field be Gaussian with exponential correlation functions, ie. $\boldsymbol{\nu}^{(k)} \sim \mathrm{N}(\boldsymbol{\mu}^{(k)}, \boldsymbol{\Sigma}^{(k)})$ for $k = 0, 1$, where

$$\Sigma_{ij}^{(k)} = \sigma_k^2 \exp\left(-3\|i - j\|/r_k\right), \qquad k = 0, 1. \tag{4}$$

Although Gaussian fields are natural candidates for smooth log variance fields, they are not computationally convenient due to full matrices giving complexity of order $\mathcal{O}(n^3)$ for vital operations. We make use of the results in Rue and Tjelmeland (2001) which demonstrate how we can fit a GMRF with covariance function very close to eg. the exponential using only a 5×5 neighborhood around each site. The computational complexity for the fitted GMRF is only $\mathcal{O}(n^2)$ (Rue, 2001), giving a speed-up of $\mathcal{O}(n)$. The density for $\boldsymbol{\nu}^{(k)}$ is

$$\pi(\boldsymbol{\nu}^{(k)} \mid \mu_k, \tau_k) \propto \exp\left(-\frac{1}{2}\tau_k \left(\boldsymbol{\nu}^{(k)} - \mu_k \mathbf{1}_n\right)^T \boldsymbol{Q}^{(k)} \left(\boldsymbol{\nu}^{(k)} - \mu_k \mathbf{1}_n\right)\right), \qquad k = 0, 1, \tag{5}$$

where the $\boldsymbol{Q}^{(k)}$'s are the normalized precision matrices obtained by the approximation method in Rue and Tjelmeland (2001), and $\mathbf{1}_n$ is a vector with n ones. The parameters τ_0 and τ_1 are the precisions and are given Gamma priors. The levels μ_0 and μ_1 are given uninformative constant prior densities.

Note that $\boldsymbol{\nu}^{(0)}$ and $\boldsymbol{\nu}^{(1)}$ are defined on the whole lattice, but only observed exterior and interior to the vessel contour, respectively. The next step is to define the explicit model for the vessel contour.

2.2.2 A Deformable Template Model for the Vessel Wall We assume that the shape of the vessel contour is well approximated by an m-sided simple polygon. In the polygonal deformable template approach (see eg. Grenander et al. (1991)) the template for an object is a polygon representing a typical instance of its shape, and variability in appearance is achieved by applying transformations to the template polygon. In our application it is natural to let the prototype shape be a circle; thus the template is represented by a set of vectors $\boldsymbol{g}_0, \boldsymbol{g}_1, \cdots, \boldsymbol{g}_{m-1}$ defining the edges of a circular polygon. With this representation, the template does not have any location information, and so the first vertex is located at a point $\boldsymbol{c}$. The positions of the remaining vertices are $\boldsymbol{c} + \boldsymbol{g}_0$, $\boldsymbol{c} + \boldsymbol{g}_0 + \boldsymbol{g}_1$, and so on. Finally, to ensure that the polygon is closed we must impose the constraint $\sum_{i=0}^{m-1} \boldsymbol{g}_i = \mathbf{0}$. To model shape variability, the template is subject to local deformations constituting a scaling and rotation of each edge, and we write

$$\boldsymbol{s}_i \boldsymbol{g}_i = \boldsymbol{g}_i + \zeta_i \begin{pmatrix} \cos\theta_i & \sin\theta_i \\ -\sin\theta_i & \cos\theta_i \end{pmatrix} \boldsymbol{g}_i = \begin{pmatrix} 1 + t_i^{(0)} & t_i^{(1)} \\ -t_i^{(1)} & 1 + t_i^{(0)} \end{pmatrix} \boldsymbol{g}_i, \tag{6}$$

where the variables ζ_i and θ_i represents the scale and angle of rotation for the change $\boldsymbol{s}_i \boldsymbol{g}_i - \boldsymbol{g}_i$. Letting ζ_i be proportional to a Rayleigh random variable, and θ_i be uniformly distributed on the interval $[0, 2\pi)$, $t_i^{(0)}$ and $t_i^{(1)}$ can be seen to be independent and identically distributed zero mean Gaussian random variables. To impose smoothness in the deformations we let the *unconstrained* deformation vector $\boldsymbol{t} = (t_0^{(0}, \dots, t_{m-1}^{(0)}, t_0^{(1)}, \dots, t_{m-1}^{(1)})$ be Gaussian with zero mean and precision matrix $\boldsymbol{I}_2 \otimes \boldsymbol{P}$, where $\boldsymbol{P}$ has entries

$$P_{ij} = \begin{cases} \frac{\kappa}{m} + 6\eta m^3 & i = j \\ -4\eta m^3 & i = j-1, j+1 \bmod m \\ \eta m^3 & i = j-2, j+2 \bmod m. \end{cases} \tag{7}$$

This parameterization was suggested in Grenander (1993, Chap. 11), and its behavior is approximately independent of the number m of edges. The hyperparameters κ and η are given Gamma priors. Hobolth and Jensen (2000) derived the limiting Gaussian process as $m \to \infty$, and based on their results we can choose hyperparameters such that the priors are reasonably informative. Imposing the closure constraint $\sum_i \boldsymbol{s}_i \boldsymbol{g}_i = \boldsymbol{0}$ reduces the parameter dimension to $2m-2$, and destroys the Markov structure of the model. However, for the purpose of MCMC simulation the unconstrained density is sufficient (when κ and η are kept fixed) since we only need to evaluate the ratio of the density at two positions $\boldsymbol{t}$ and $\boldsymbol{t}'$. As the constraint is linear in $\boldsymbol{t}$, the ratio of two constrained densities equals the ratio of the corresponding unconstrained densities.

2.3 Summary of the posterior

The posterior is proportional to the product of the prior and the likelihood, and is summarized in Figure 3. The first part of the model is the deformable template $\boldsymbol{t}$ defining the

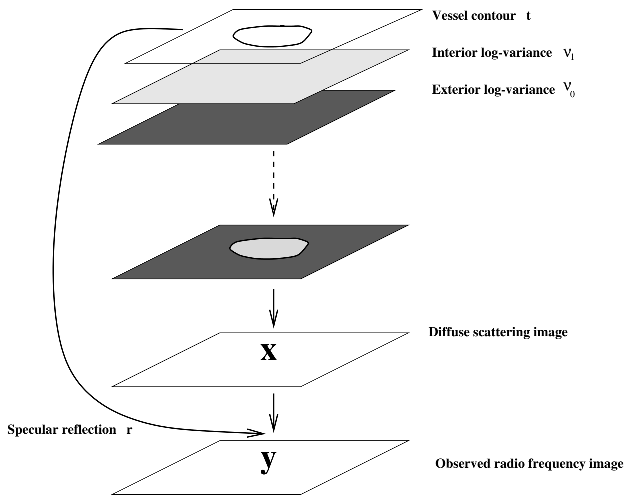

Figure 3: Graphical representation of the model for contour estimation in ultrasound images.

contour of the vessel. The contour defines which parts of the log variance fields $\boldsymbol{\nu}^{(\cdot)}$ are visible, and hence we have defined the log variance of $\boldsymbol{x}$ for each site in the image. The contour and its tangent function defines the specular reflection ρr_i at each site. Sampling the components of $\boldsymbol{x}$ independently with the corresponding variance and adding $\rho \boldsymbol{r}$, then blurring with the point spread function $\boldsymbol{h}$ and adding independent zero-mean Gaussian noise, gives us the observed raw data.

3 CONSTRUCTION OF MCMC ALGORITHMS

Our next task is to construct an MCMC algorithm that is able to explore the posterior defined in section 2.3, and return samples with the correct limiting distribution. For each sample from the posterior we can compute the corresponding cross sectional area (CSA) of the blood vessel using Green's theorem, thus producing a sample from the posterior distribution for the CSA. Using these samples we estimate the marginal density of the CSA, and compute its posterior mean estimate and estimated credibility interval. Although the route ahead might seem standard, there are severe problems constructing a reasonable MCMC algorithm.

The model consists of several fields ($\boldsymbol{x}$, $\boldsymbol{\nu}^{(0)}$ and $\boldsymbol{\nu}^{(1)}$), each with their respective hyperparameters; and the deformable template $\boldsymbol{t}$ with its hyperparameters. The sampling difficulties are due both to the strong interaction between each field and its hyperparameters, and to the interdependency between the fields including $\boldsymbol{t}$. Moreover, there are computational difficulties due to the high dimension of the fields, so a naive implementation will not be computationally feasible.

The Markov chain is easiest and most commonly constructed using *single site* updating, which consists of sampling parameters sequentially from their full conditional distributions. For models with strong spatial interactions single site samplers are known to have poor mixing and convergence properties, and significant computational improvements can be gained by block updating the components of the random fields (Liu, Wong and Kong, 1994; Knorr-Held and Rue, 2002). Knorr-Held and Rue (2002) demonstrated that the best effect occurs when the random fields are updated jointly with their hyperparameters by constructing joint updates: Use a simple random walk, say, for the hyperparameters, update the field (or more than one field) by sampling from its conditional density (or an approximation) and then accept or reject jointly. The computational costs of such blocking schemes have until now been prohibitive, but efficient computations are available for GMRFs (Rue, 2001). This includes both efficient sampling and computations of the log-posterior gradient for use in Langevin proposals; the speedup is in both cases $\mathcal{O}(n)$. The full conditionals for $\boldsymbol{x}$, $\boldsymbol{\nu}^{(0)}$ and $\boldsymbol{\nu}^{(1)}$ all take advantage of the GMRF representation. The full conditional of the deformation parameter $\boldsymbol{t}$ of the template is somewhat more complicated due to inclusion and exclusion of data which are either outside or inside of the deformed template, and due to the correction for specular reflection. For $\boldsymbol{\nu}^{(0)}$, $\boldsymbol{\nu}^{(1)}$ and $\boldsymbol{t}$ we construct block-updates based on the Langevin proposal, ie.

$$\boldsymbol{\nu}^{(0)\prime} \sim \mathrm{N}\left(\boldsymbol{\nu}^{(0)} + \frac{1}{2} h \nabla \ln \pi\left(\boldsymbol{\nu}^{(0)} \mid \ldots\right), h\boldsymbol{I}\right) \tag{8}$$

where h is a scaling parameter. For $\boldsymbol{x}$ we use that the full conditional is a GMRF.

We have implemented various sampling schemes to investigate whether the choice of sampler affect the estimated CSA of the blood vessel, and the robustness of each scheme over various datasets. All schemes follow the setup in Knorr-Held and Rue (2002), as previously described. The hyperparameters are all updated in the same manner; if they are positive, we propose a new value by scaling with an s having density proportional to $1 + 1/s$ on $[1/f, f]$, where $f > 1$; otherwise we sample a new proposal uniformly in a neighborhood around the old value. We tune the proposals so the acceptance rate is about 25%.

Scheme 1 This scheme updates each field ($\boldsymbol{x}$, $\boldsymbol{\nu}^{(0)}$, $\boldsymbol{\nu}^{(1)}$ or $\boldsymbol{t}$) jointly with each one of its hyperparameters. The fields are either updated using a GMRF ($\boldsymbol{x}$) or a Langevin proposal

($\boldsymbol{\nu}^{(0)}$, $\boldsymbol{\nu}^{(1)}$ and $\boldsymbol{t}$).

Scheme 2 This scheme is similar to scheme 1, but we update the fields $\boldsymbol{x}$ and $\boldsymbol{t}$ jointly together with each one of their hyperparameters in the following sense: A new value θ' for one of the hyperparameters θ is proposed. Conditioned on θ' and $\boldsymbol{x}$ we propose a new value $\boldsymbol{t}'$ for $\boldsymbol{t}$, and then a new value $\boldsymbol{x}'$ for $\boldsymbol{x}$ is drawn from its full conditional. Finally $(\boldsymbol{x}', \boldsymbol{t}', \theta')$ are accepted or rejected jointly.

Scheme 3 This scheme is similar to scheme 2 except that we update all fields in a single block together with each hyperparameter in turn.

The motivation behind scheme 2 is the belief that $\boldsymbol{x}$ and $\boldsymbol{t}$ are the fields having the strongest interaction, and we therefore update them jointly. If there are strong between-field dependencies, scheme 2 and 3 may possibly reduce their effect by improving the mixing properties of the Markov chain.

As the hyperparameters change in each proposal, the acceptance proposal will in most cases involve the evaluation of the normalizing constant of the GMRF fields. However, the normalizing constant can be efficiently computed, see Rue (2001) for details.

For a 64×128 image, a 500MHz PC used approximately 0.18 seconds per iteration for scheme 1, and 0.16 seconds per iteration for schemes 2 and 3. One iteration consists of updating all hyperparameters in the model, each jointly with one or more of the fields.

4 RESULTS

In this section we present the estimated CSAs in the images shown in Figure 1, and summarize our experience regarding specification of hyperparameters, performance of the different sampling schemes and the effect of correcting for specular reflection.

4.1 Specification and effect of priors

We use quite informative priors for the unknown parameters in the deformable template due to our specific knowledge of the object in the image. Based on analysis of the limiting correlation structure (Hobolth and Jensen, 2000), we found that κ about 100 and η about 10^{-3} gave a reasonable behavior of realizations from the prior. Hyperpriors were then selected having these values as means. All other parameters in the model were given near un-informative Γ-priors and uniform priors, for parameters defined on $\mathbb{R}^{+}$ and $\mathbb{R}$, respectively. We found the interval estimate for the vessel area to be rather insensitive to different but sensible choices of the hyperparameters.

4.2 Performance of the sampling schemes

Performance of the three different sampling schemes was monitored by inspecting trace plots and autocorrelations for the parameters and other statistics of the Markov chain. Especially the trace of the CSA, which is our main target, where studied carefully. It was our experience (but not a surprise), that the CSA was the slowest converging statistics of those we studied.

Quite interestingly, we found the interval estimate for a fixed run-length to depend on the choice of sampling scheme. Scheme 1 performed poorly, exhibiting slow convergence. Repeated runs revealed that the Markov chain is prone to get stuck in local minima, or even drifting off; in this sense the sampling scheme is not robust. The behavior of scheme 1 is somewhat surprising considering the complexity of the sampling scheme. The reason is the strong interdependency between the radio frequency field $\boldsymbol{x}$ and the template $\boldsymbol{t}$. Thus, updating one field conditionally on the other may cause the chain to move very slowly. Single-site schemes behave even worse.

Scheme 2 and 3 perform better in that their behavior is consistent over repeated runs, several datasets and different initializations, and that the estimated contour and CSA were reasonable compared to "reference" estimates obtained with very long runs. Inspection of autocorrelation plots and comparison of the asymptotic variances for the CSA, reveal that scheme 2 is better than scheme 3. In the last scheme the number of parameters to update is more than doubled, and the dependence between the $\boldsymbol{\nu}$-fields and the other fields does not seem to be strong enough to counterbalance the increase in complexity. The estimated autocorrelation function for scheme 2 is shown in Figure 4, indicating reasonably good mixing of the chain.

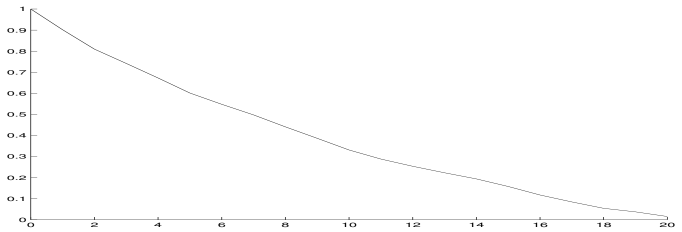

Figure 4: Estimated autocorrelation function for the CSA using scheme 2. The function is calculated using every 100th sample after a burn-in of 25 000 iterations.

4.3 Effect of the specular reflection

The fact that we could adjust for specular reflection by having an explicit model for the outline of the blood vessel, was one main motivation for using a deformable template model. Modeling specular reflection is important in ultrasound imaging, as the large variations in signal magnitude around the contour may cause problems. The importance of the correction is most prominent when data is of low quality, as is the case here. Figure 5 illustrates this effect, and shows 50 samples taken with a separation of 100 iterations in MCMC runs without correction for specular reflection (left image), and with correction for specular reflection (right image). Without the correction, the deformable template can take strange shapes trying to over-fit the data.

4.4 Results

Figure 7 shows the histogram of the CSA for the images in Figure 1, using the model for specular reflection and the prior settings described above. The corresponding 95% credibility intervals are $[2690, 3086]$ for Figure 1a and $[1062, 1530]$ for Figure 1b. We have also computed point-estimates of the vessel walls by using the distance average for random closed sets (Baddeley and Molchanov, 1998). Visually the point-estimates (Figure 6) seem reasonable, without carrying too much spurious detail. The results have been evaluated by cardiologists,

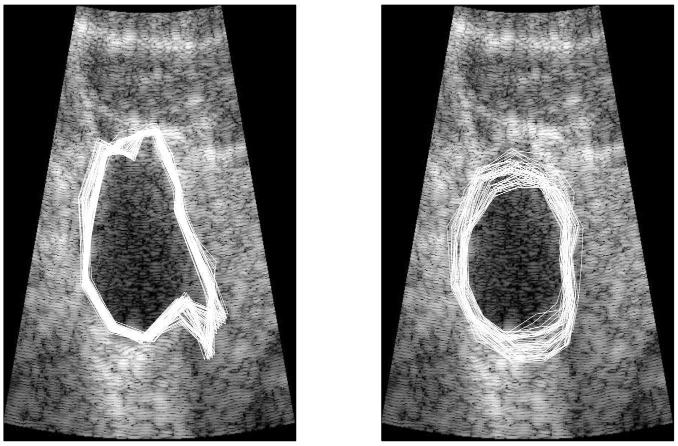

(a) (b)

Figure 5: Samples from the posterior distribution for the vessel wall in Figure 1a. We have collected every 100th sample after a suitable burn-in period.

who have found them to be in good agreement with their knowledge of vessel shapes.

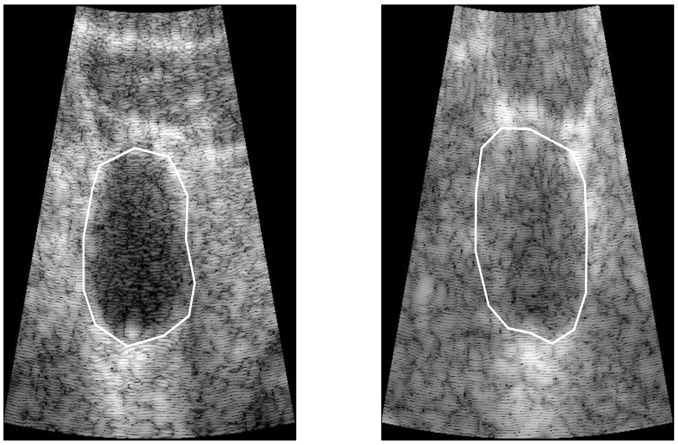

(a) (b)

Figure 6: Point-estimates of the vessel wall for the images in Figure 1a and Figure 1b, respectively.

References

Baddeley, A. J. and Molchanov, I. (1998). Averaging of random sets based on their distance functions, *Journal of Mathematical Imaging and Vision* **8**(1): 79–92.

Christensen, D. A. (1988). *Ultrasonic Bioinstrumentation*, Wiley.

de Figueiredo, M. T. and Leitao, J. M. N. (1992). Bayesian estimation of ventricular contours in angiographic images, *IEEE Transactions on Medical Imaging* **11**(3): 416–429.

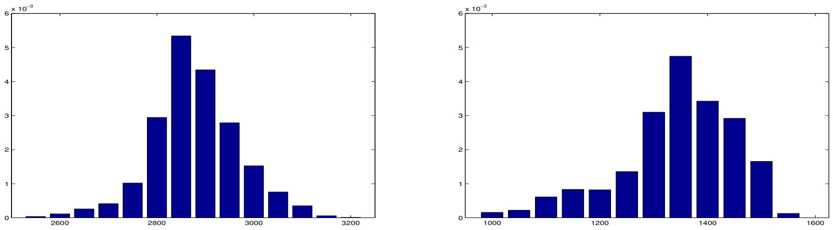

(a) (b)

Figure 7: Histograms of the CSA for the images in Figure 1a and Figure 1b, respectively.

de Souza, K. M. A., Kent, J. T. and Mardia, K. V. (1999). Stochastic templates for aquaculture images, *Journal of the Royal Statistical Society, Series C* **48**(2): 211–227.

Duncan, J. S. and Ayache, N. (2000). Medical image analysis: Progress over two decades and the challenges ahead, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**(1): 85–106.

Glasbey, C. A. (1998). Ultrasound image segmentation using stochastic templates, *Journal of Computing and Information Technology* **6**: 107–116.

Goodman, J. (1975). Statistical properties of laser speckle patterns, *in* J. Dainty (ed.), *Laser Speckle and Related Phenomena*, Springer Verlag, Berlin.

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

Grenander, U., Chow, Y. and Keenan, D. M. (1991). *Hands: a Pattern Theoretic Study of Biological Shapes*, Research Notes on Neural Computing, Springer Verlag, Berlin.

Grenander, U. and Manbeck, K. M. (1993). A stochastic model for defect detection in potatoes, *Journal of Computer Graphics and Statistics* **2**: 131–151.

Grenander, U. and Miller, M. I. (1994). Representations of knowledge in complex systems (with discussion), *Journal of the Royal Statistical Society, Series B* **56**(4): 549–603.

Hansen, M. B., Møller, J. and Tøgersen, F. A. (2001). Bayesian contour detection in a time series of ultrasound images through dynamic deformable template models, *Biostatistics* . (to appear).

Hobolth, A. and Jensen, E. B. V. (2000). Modelling stochastic changes in curve shape, with application to cancer diagnostics, *Advances in Applied Probability (SGSA)* **32**(2): 344–362.

Hokland, J. and Kelly, P. (1996). Markov models of specular and diffuse scattering in restoration of medical ultrasound images, *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **43**(4): 660–669.

Husby, O., Lie, T., Langø, T., Hokland, J. and Rue, H. (2001). Bayesian 2d deconvolution: A model for diffuse ultrasound sc attering, *IEEE Trans. on Ultrason. Ferroelec. Frec. Contr.* **48**(1): 121–130.

Insana, M. F., Wagner, R. F., Garra, B. S., Brown, D. G. and Shawker, T. H. (1986). Analysis of ultrasound image texture via generalized Rician statistics, *Optical Engineering* **25**: 743–748.

Kao, C.-M., Pan, X., Hiller, E. and Chen, C.-T. (1998). A Bayesian approach for edge detection in medical ultrasound images, *IEEE Transaction on Nuclear Science* **45**(6): 3089–3096.

Kass, M., Witkin, A. and Terzeopoulos, D. (1988). Snakes: Active contour models, *International Journal of Computer Vision* **1**(4): 321–331.

Knorr-Held, L. and Rue, H. (2002). On block updating in Markov random field models for disease mapping, *Scandinavian Journal of Statistics* . (to appear).

Langø, T. (2000). *Ultrasound guided surgery: Image processing and navigation*, PhD thesis, Norwegian University of Science and Technology, Trondheim, Norway.

Langø, T., Lie, T., Husby, O. and Hokland, J. (2001). Bayesian 2d deconvolution: Effect of using spatially invaria nt point spread functions, *IEEE Trans. on Ultrason. Ferroelec. Frec. Contr.* **48**(1): 131–141.

Liu, J. S., Wong, W. H. and Kong, A. (1994). Covariance of the Gibbs sampler with applications to the comparison of estimators and augmentation schemes, *Biometrika* **81**(1): 27–40.

Mardia, K. V., Qian, W., Shah, D. and de Souza, K. M. A. (1997). Deformable template recognition of multiple occluded objects, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **19**(9): 1035–1042.

Mulet-Parada, M. and Noble, J. A. (2000). 2D+T acoustic boundary detection in echocardiography, *Medical Image Analysis* **4**(1): 21–30.

Ødegård, L. (1995). *Phase aberration correction in medical ultrasound imaging*, PhD thesis, The Norwegian Institute of Technology, Trondheim, Norway.

Pham, D. L., Cu, C. Y. and Prince, J. L. (2000). Current methods in medical image segmentation, *Annual Review of Biomedical Engineering* **2**: 315–.

Qian, W., Titterington, D. M. and Chapman, J. N. (1996). An image analysis problem in electron microscopy, *Journal of the American Statistical Association* **91**(435): 944–952.

Ripley, B. D. and Sutherland, A. I. (1990). Finding spiral structures in images of galaxies, *Philosophical Transactions of the Royal Society of London A* **332**: 477–485.

Rue, H. (2001). Fast sampling of Gaussian Markov random fields with applications, *Journal of the Royal Statistical Society, Series B* **63**(2): 325–338.

Rue, H. and Hurn, M. A. (1999). Bayesian object identification, *Biometrika* **86**(3): 649–660.

Rue, H. and Syversveen, A. R. (1998). Bayesian object recognition with Baddeley's delta loss, *Advances in Applied Probability (SGSA)* **30**(1): 64–84.

Rue, H. and Tjelmeland, H. (2001). Fitting Gaussian Markov random fields to Gaussian fields, *Scandinavian Journal of Statistics* . (to appear).

Stoica, R., Descombes, X. and Zerubia, J. (2000). A Markov point process for road extraction in remote sensed images, *Technical Report 3923*, INRIA.

Varghese, T., Donohue, K. D. and Chatterjee, J. P. (1995). Specular echo imaging with spectral correlation, *Proc. IEEE Ultrason. Symp.*, pp. 1315–1318.

Wagner, R., Insana, M. and Brown, D. (1987). Statistical properties of radio-frequency and envelope-detected signals with application to medical ultrasound, *J. Opt. Soc. Am, A* **4**(5): 910–922.