Pattern Recognition 47 (2014) 634–642

Contents lists available at ScienceDirect

Pattern Recognition

journal homepage: www.elsevier.com/locate/pr

Boundary reconstruction in binary images using splines

Larissa Stanberry [a,*], Julian Besag [b,1]

[a] University of Bristol, Bristol, UK
[b] University of Bath, Bath, UK

ARTICLE INFO

Article history:
Received 29 June 2012
Received in revised form
1 March 2013
Accepted 16 July 2013
Available online 26 July 2013

Keywords:
Shape estimation
Boundary reconstruction
B-splines
Mammography
Oriented distance functions
Set inference

ABSTRACT

In image analysis, it is often required to reconstruct the boundary of an object in a noisy image. This paper presents a new method, which relies on flexibility and computational simplicity of B-spline curves, to reconstruct a smooth connected boundary in a noisy binary image. Boundary inference is based on oriented distance functions yielding the estimator which is interpreted as a posterior expected boundary of the underlying random set. The performance of the method and its dependence on the image quality and model specification are studied on simulated data. The method is applied to reconstruct the skin-air boundary in digitised analogue mammogram images.

1. Introduction

Boundary reconstruction in a noisy image is a frequently encountered problem in image analysis. Examples are abound and include applications in computer vision, robotics, medical imaging, and remote sensing. Current work was motivated by boundary reconstruction in mammography, where the tissue outline is required to estimate the radiographic density and to calculate the asymmetry measure of a breast [7,59,50,9,30].

In digital processing, planar images are stored as two- or three-dimensional arrays with values corresponding to the intensity or the colour of the pixels, respectively. The goal of boundary reconstruction is to recover the boundary of an object in a noisy image. Note that the information loss due to discretisation implies that an object and its boundary cannot be uniquely identified from the image.

Standard methods for boundary reconstruction include intensity thresholding and edge detectors [10,23]. While readily available, the methods are sensitive to the contrast-to-noise ratio of an image, identifying discontinuous boundaries and false edges. The quality of the reconstruction can be improved by careful preprocessing of the image and tuning of the parameters.

In the pattern algebra, local boundary patterns are described by generators and the density of the configuration is specified by interactions between pairs of bonds attached to each site [26]. This model is inconsistent on a rectangular lattice with a second-order neighbourhood structure and rigid interactions, because certain bond configurations are deterministically excluded, violating the positivity condition. The limitation, however, can be resolved on a hexagonal lattice of pixels [56].

In general, boundary models based on local interactions are sensitive to noise and often result in nonsmooth estimators and spurious edge patterns. To improve the quality of the reconstruction, one can consider deformable template models with generators given by geometrical elements (line segments, arcs, etc.) [27]. Boundary templates can be further extended to describe the global geometry/shape of an object [1,28,33]. For a comprehensive review of the methods and their extensions, we refer to [29,32,11].

In contrast to deformable template models, active contour models do not model the boundary per se, but rather rely on parametric or free-form curves for estimation [6,8,18,20,37,40,49,53,54]. The modelling curve is attracted to edges by local image forces, while its smoothness is controlled by user-imposed regularity constraints. The locality of the gradient forces driving the curve leads to a strong dependence of the estimator on initialisation, a CNR of the image and convexity of the shape. Various extensions were proposed to alleviate these problems and improve performance [14,18,58,13,60].

Developed by [21], a stochastic model for boundary detection uses the coarsened image lattice and assigns the boundary and partitioning labels to blocks of pixels, rather than individual lattice sites. The distribution of label configurations is based on a disparity measure between identified pixel blocks and the reconstruction is

* Corresponding author. Present address: Seattle Children's Research Institute, 1900 9th Ave, Seattle, WA 98101, USA. Tel.: +1 425 418 9186.
E-mail addresses: larissa.stanberry@seattlechildrens.org, larissa.stanberry@gmail.com (L. Stanberry).

[1] Sadly, Julian Besag passed away after the completion of the manuscript.

0031-3203/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.patcog.2013.07.007

given by the solution of a constrained optimisation problem, where the constraints specify forbidden label configurations. To preserve the continuity and avoid spurious patterns, [22] proposed to recover the boundary by means of a biased random walk. In turn, [44] models the boundary of a magnetic domain by polygons.

The polygonal model was further extended to polygons with a varying number of vertices and side lengths [43]. The optimal number of vertices is determined using a reversible jump Markov Chain Monte Carlo [24], thus balancing the complexity and the flexibility of the model. The final estimator of an object is given by a greyscale image, with intensities corresponding to the posterior mean state of pixels.

In this paper, we introduce a new method to reconstruct a smooth connected boundary of an object in a noisy image. As an example, we consider binary images which arise in molecular imaging, geology, etc. when, for example, the grey level information is limited or the image reflects a presence/absence status of an object [46]. Note that in this paper, we chose the binary noise model mainly for its simplicity. The method however generalises to different types of images. For example, substituting Gaussian noise for binary channel is straightforward and allows estimation of the boundary, the noise variance, and intensity values of the foreground and background without much added complexity.

In this paper, we model the boundary using B-spline curves, without any shape constraints. The flexibility and computational simplicity of B-spline curves make them well-suited for contour modelling [42]. Further, we address the question of posterior inference for sets and their boundaries and describe appropriate loss functions. To obtain the boundary estimator, we use an approach based on oriented distance functions (ODFs), which allows us to make inference not only about an object, but also about its boundary. The resulting estimator is interpreted as the posterior expected boundary. Notably, the loss functions for set inference are not restricted to planar domains and can be applied to high-dimensional problems. Similarly, the ODF inference can be implemented to accommodate hypersurfaces.

This paper is structured as follows. In Section 2, we describe the statistical model and give a brief overview of B-spline curves. Section 3 discusses the loss functions and estimators based on them. The algorithm for sampling from the posterior is given in Section 4. In Section 5, we study the performance of the method on a simulated image and explore the sensitivity of the results to image quality and model parameters. Section 6 gives the results of boundary reconstruction in a mammogram image. A discussion in Section 7 concludes the paper.

2. The model

2.1. Notation

Let I be the true (noise-free) image and I^{D} be its observed degraded version. We assume that both I and I^{D} are binary. Without loss of generality, we can assume that the foreground is white and the background is black, describing them as $I_{\mathrm{F}} = \{x : I_x = 1\}$ and $I_{\mathrm{B}} = \{x : I_x = 0\}$, respectively; here, subscript x indexes image pixels. We assume that the boundary of an object extends from the left edge to the right edge of the image and define the foreground (background) as the area located below (above) the boundary. Note that under these assumptions, reconstructing the boundary in I^{D} is equivalent to reconstructing the true image I.

2.2. Likelihood

We assume that the degradation mechanism of the observed image I^D is described by a binary noise channel; i.e. the colour of a foreground pixel is reversed with probability α_{F} and preserved with probability $1-\alpha_{\mathrm{F}}$, independent of other pixels. Similarly, a background pixel becomes white with probability α_{B} and remains black with probability $1-\alpha_{\mathrm{B}}$, independent of other pixels, so that α_{F} and α_{B} are the chances of pixel-wise error in the foreground and background, respectively.

Assume that the true planar boundary is described by some curve $\mathcal{C}$. The likelihood of the acquired image I^{D}, given the boundary $\mathcal{C}$ and error probabilities α_{F} and α_{B}, is written as

$$\mathcal{L}(I^{\mathrm{D}}|\mathcal{C}, \alpha_{\mathrm{F}}, \alpha_{\mathrm{B}}) = \alpha_{\mathrm{F}}^{N_{\mathrm{FB}}}(1-\alpha_{\mathrm{F}})^{N_{\mathrm{FF}}}\alpha_{\mathrm{B}}^{N_{\mathrm{BF}}}(1-\alpha_{\mathrm{B}})^{N_{\mathrm{BB}}}, \tag{1}$$

where $N_{\mathrm{FF}}(N_{\mathrm{FB}})$ is the number of pixels in I_{F} that remain white (become black) in I^{D}, and $N_{\mathrm{BB}}(N_{\mathrm{BF}})$ is the number of pixels in I_{B} that remain black (become white) in I^{D}. Here, we assume that the boundary is thin and classify a pixel as a foreground (background), if its centre is below (above) the boundary, so that $N_{\mathrm{FF}} + N_{\mathrm{FB}} + N_{\mathrm{BB}} + N_{\mathrm{BF}}$ is the total number of pixels in the image.

Assuming mutual independence of the boundary and error probabilities, the Bayes theorem implies that the posterior distribution of the boundary and noise parameters, given the observed image, is

$$\pi(\mathcal{C}, \alpha_{\mathrm{F}}, \alpha_{\mathrm{B}}|I^{\mathrm{D}}) \propto \mathcal{L}(I^{\mathrm{D}}|\mathcal{C}, \alpha_{\mathrm{F}}, \alpha_{\mathrm{B}})\rho(\mathcal{C})\rho(\alpha_{\mathrm{F}})\rho(\alpha_{\mathrm{B}}),$$

where $\rho(\mathcal{C}), \rho(\alpha_{\mathrm{F}}), \rho(\alpha_{\mathrm{B}})$ are the prior distributions for the boundary $\mathcal{C}$, and the error probabilities α_{F} and α_{B}, respectively. Merely for later clarity, here we denote the priors by ρ and posteriors by π.

2.3. Prior distributions

We assume that the boundary of an object is a twice differentiable curve and denote by $\mathcal{C}$ a planar parametric curve, given by a mapping $C : [0, 1] \mapsto \mathbb{R}^2$. In general parametrisation, vector $C(t) = (C_1(t), C_2(t))$ is the position vector along the curve at time t and $(C_1(t), C_2(t))$ are Cartesian coordinates of point $C(t)$ in $\mathbb{R}^2$. For a natural parametrisation $C(s)$ with arc length s as a parameter, let $T(s) = C'(s)$ be the unit tangent vector of $\mathcal{C}$ at point $C(s)$ [38]. The direction of $T(s)$ depends on the orientation of the curve and points in the direction of increasing parameter values s. For the curvature vector $K(s) = T'(s) = C''(s)$, the curvature of $\mathcal{C}$ at point $C(s)$ is given by $\kappa(s) = |K(s)|$. Thus, $\kappa(s)$ measures the rate of change of the tangent vector, or the deviation of $\mathcal{C}$ from the tangent line in the neighbourhood of $C(s)$. Correspondingly, the global behaviour of the curve of length L is captured by its integrated squared curvature $\int_0^L \kappa^2(s)\,\mathrm{d}s$. In the mathematical theory of elasticity, the integrated squared curvature represents the potential energy stored in the elastic beam [39].

To control the behaviour of the curve, we define the prior

$$\rho(\mathcal{C}) \propto \exp\left\{-\beta \int_0^L \kappa^2(s)\,\mathrm{d}s\right\},$$

where the deformation constant $\beta > 0$ determines the elasticity of the curve. The curve prior is analogous to the roughness penalty in nonparametric smoothing and functional regression, where the penalty has a form of the integrated squared second derivative with respect to the argument [25,45]. The specified prior assigns higher probabilities to slowly varying curves, thus favouring them over erratic contours. The prior is improper because it is invariant with respect to translations within the image window. The integrated curvature is also used in active contour models as a regularisation term that controls the smoothness of the resulting curve [37].

For computational reasons, we prefer an arbitrary parametrisation $C(t)$ with parameter t, in which case the squared curvature is given by

$$\kappa^2(t) = \{|C'(t)|^2|C''(t)|^2 - \langle C'(t), C''(t)\rangle^2\}/|C'(t)|^6, \tag{2}$$