

# Maximum *a posteriori* estimation of image boundaries by dynamic programming

C.A. Glasbey

Biomathematics and Statistics Scotland

JCMB, King's Buildings, Edinburgh, EH9 3JZ, Scotland

and M.J. Young

Animal Biology Division, SAC

West Mains Road, Edinburgh EH9 3JG, Scotland

September 25, 2001

## Abstract

We seek a computationally-fast method for solving a difficult image segmentation problem: the positioning of boundaries on medical-scanner images to delineate tissues of interest. We formulate a Bayesian model for image boundaries such that the maximum *a posteriori* (MAP) estimator is obtainable very efficiently by dynamic programming. The prior model for the boundary is a biased random walk and the likelihood is based on a border appearance model, with parameter values obtained from training images. The method is applied successfully to the segmentation of ultrasound images and X-ray computed tomographs of sheep, for application in sheep breeding programmes.

**Key words:** Bayesian image analysis, Border appearance model, Image segmentation, Random walk, Ultrasound image, X-ray computed tomography.

## 1 Introduction

The primary application of non-invasive medical imaging techniques such as ultrasound scanning, X-ray computed tomography (CT) and magnetic resonance imaging (MRI), is diagnosis, by detection of physiological abnormalities. However, they also have the potential to be used quantitatively, to estimate organ volumes and tissue proportions, for example. This necessitates that the images be segmented, to delineate areas to be measured. To illustrate, consider images of sheep, obtained as part of a sheep breeding programme (Simm, 1992).




Figure 1: *Non-invasive images of sheep obtained using two imaging techniques, together with tissue labels and, in white, manually-positioned boundaries which identify the tissues of interest: (a) ultrasound scan ( $301 \times 341$  pixels in size,  $1 \text{ pixel} \equiv 0.16\text{mm}$ ), (b) X-ray CT lumbar image (cropped to  $176 \times 176$  pixels in size,  $1 \text{ pixel} \equiv 2.0\text{mm}$ ).*

- Fig 1(a) shows an ultrasonic scan of the back of a live sheep, with tissue areas labelled. From this image, sheep breeders want to measure the average depth of the subcutaneous layers of fat. Therefore, they manually positioned boundaries on the image along the top and bottom of the fat layers. These are shown as thin white lines superimposed on the image.
- Fig 1(b) shows an X-ray CT image of the lumbar region of a second sheep, after some low-level processing to remove from the image the U-shaped plastic cradle in which the sheep was lying, and stretching of the pixel grey-scale to emphasise differences between soft tissues. Again, tissue areas are labelled and a manually-positioned boundary has been superimposed in white. In this case, the boundary separates the sheep's internal organs from the fat and muscle areas that the breeders want to measure.

Our aim is to automate the positioning of these boundaries, as the current, manual method is both slow and tedious. Also, as there are many images to segment, we seek a computationally-fast method. Although we focus on these two applications, the methodology is generic and could be used with other subjects and imaging modalities.

Segmentation is often the most difficult step in image analysis, and the positioning of boundaries in images such as Fig 1 is a challenging problem because many features of these images have to be taken into account. There are a large number of low-level algorithms for image segmentation (see, for example, Glasbey and Horgan, 1995, chapter 4). In practice, some human intervention is frequently needed to achieve an accurate segmentation, unless high-level knowledge can be built into the approach. Bayesian methodology offers one formal framework for incorporating such domain-specific knowledge.

The Bayesian approach to image restoration and image segmentation received its initial impetus from the seminal papers of Geman and Geman (1984) and Besag (1986). In essence, prior information on an appropriate image model is combined with image data, in order to derive the posterior distribution of realisations of the model. The most commonly used priors are Markov random fields (Geman and Geman, 1984; Besag, 1986), and deformable templates which specify the boundaries of objects (Grenander and Miller, 1994; Phillips and Smith, 1994; Cootes et al., 1995). An estimator is obtained from the posterior distribution by minimising a particular loss function. For example, under a 0-1 loss function the optimal estimator maximises the posterior distribution; this is the maximum *a posteriori* (MAP) estimator. Unfortunately, whatever the loss function, the optimal estimator is usually difficult to compute, and simulated annealing (Geman and Geman, 1984), gradient descent (Besag, 1986) or approximate maximisation methods (Ferrari et al., 1995) have been used. One notable exception is that, for binary Markov random fields, Greig et al. (1989) showed that the MAP estimator can be obtained efficiently using a maximum-flow algorithm. Loss functions are often chosen pragmatically for computational convenience, though Rue (1997) illustrates the effectiveness of alternative choices, albeit at far greater computational cost.

In this paper, we formulate another category of Bayesian model for which the MAP can be computed efficiently, in this case using a dynamic programming algorithm. In §2 we construct a Bayesian model appropriate for a particular class of image boundary, which includes the images in Fig 1. Then, in §3 we show how the MAP estimator can be obtained efficiently by dynamic programming. In §4 we apply the methodology to the sets of images. Finally, in §5 we discuss generalisations.

## 2 Bayesian model

The class of boundary problems we consider is where a boundary intersects each column of an image precisely once and the distribution of pixel values in that column is affected by the boundary only to the extent that pixel positions need to be defined relative to this point of intersection. In the case of Fig 1(a) a ‘column’ is literally a column in the image. However, more generally, ‘columns’ can refer to any set of pre-specified lines or simple curves. These may have been obtained from a preliminary, approximate segmentation (Brej1 and Sonka, 2000), or by a geometric transformation of the image space, as we illustrate below with Fig 1(b). We are not aware of anyone else having specified such a class.

We have available 143 ultrasound images additional to Fig 1(a), a further 197 lumbar X-ray CT images, and similar numbers of X-ray CT images obtained at two other anatomical positions. For all these images we also know where the boundaries would be positioned manually. We use half these images for training, to identify a Bayesian model, and retain half for validation, in §4.

Boundaries are one pixel wide, and we use integer  $b_j$  to denote the row where a boundary intersects the  $j$ th column, where columns are indexed from 0 to  $J$ . Rows are similarly indexed from 0 at the top of an image to  $I$  at the bottom. Our novel prior model for  $b$  is a one-

dimensional, biased, random walk on the integers, a type of deformable template. Let  $\beta_j$  be an integer that denotes a measure of the location of  $b_j$ . The change in row position between columns  $(j-1)$  and  $j$  is constrained to lie within  $\pm 1$  of the change in  $\beta$  between columns  $(j-1)$  and  $j$ , although in applications where it were appropriate it would be straightforward to permit changes in row position in excess of 1. We specify the probability density to be

$$f(b) \propto \exp \left[ -\frac{\lambda}{2} \sum_{j=1}^J \phi(\{b_j - b_{j-1}\} - \{\beta_j - \beta_{j-1}\}) \right], \quad b \in \mathcal{B}, \quad (1)$$

where  $\lambda$  is a non-negative constant,  $\phi$  is the function defined by

$$\phi(x) = \begin{cases} |x| & \text{if } x = 0, \pm 1, \\ \infty & \text{otherwise,} \end{cases}$$

and  $\mathcal{B}$  specifies the set of permissible boundaries. As we have a finite state-space, this is a proper distribution. We assign values to  $\beta$  by averaging the manually-identified boundaries in the training set:

$$\beta_j = \text{int} \left\{ \frac{1}{N} \sum_{n=1}^N b_j^{(n)} \right\}, \quad j = 0, \dots, J,$$

where  $\text{int}\{x\}$  denotes the nearest integer to  $x$ ,  $b^{(n)}$  denotes the manually-drawn boundary in the  $n$ th training image in a set of size  $N$ . Parameter  $\lambda$  controls the roughness of the boundary, ranging from a value of 0 which gives equal probability to  $(\{b_j - b_{j-1}\} - \{\beta_j - \beta_{j-1}\}) = 0, \pm 1$ , up to a value of  $\infty$  which constrains the boundary to lie parallel to  $\beta$ . We could obtain a suitable value for  $\lambda$  from the boundaries in the training images, but, the partial, quasi-likelihood we use for the data (see below) is likely to be inappropriately scaled relative to such a prior density. Therefore, we instead select  $\lambda$  in the pragmatic way explained at the end of §3. We note that Phillips and Smith (1994) reduced the pixellation in their images of faces from  $256 \times 256$  to  $64 \times 64$ , presumably to prevent the likelihood, based on the unrealistic assumption of independently distributed pixel values, from totally swamping the prior in the posterior distribution.

We choose  $\mathcal{B}$  from the distribution of manual boundaries in the training images. For the ultrasound images we found that distributions of  $b_j^{(n)}$  were skew, as often occurs with variables that are constrained to be positive. So, we chose to measure variability on a log-scale, with standard deviation ( $\tau_j$ ) in column  $j$  given by:

$$\tau_j = \sqrt{\frac{1}{N-1} \sum_{n=1}^N (\log b_j^{(n)} - \log \beta_j)^2}, \quad j = 0, \dots, J.$$

We specify bounds for  $b_j$  to be  $\beta_j e^{\pm m \tau_j}$ , with  $m = 2, 3$  or  $4$ . The way we choose  $m$  is again explained at the end of §3. By specifying hard limits for  $b$ , we have implicitly defined an expectation for each  $b_j$ , although explicitly we have only given expectations for differences  $(b_j - b_{j-1})$ . Alternatively, we could have specified the full distribution of  $b$ , but then the algorithm in §3 would have run more slowly as we would have had to consider a greater number of possible values for each  $b_j$ .

We construct a separate model for each boundary in an image, so, for the ultrasound images we have two models. Fig 2 shows, for both ultrasound boundaries,  $\beta$  obtained for the 72 ultrasound




Figure 2: *Illustration of prior model for upper and lower boundaries in ultrasound images: (—) denotes  $\beta$ ; (- - -) possible ranges for  $\mathcal{B}$ , namely  $\beta e^{\pm 2\tau}$ ; (···) manually-drawn boundaries in Fig 1(a).*

training images, together with bounds for  $\mathcal{B}$  given by  $\beta e^{\pm 2\tau}$ . Also shown are the boundaries for the image in Fig 1(a), from which we can see that the biased random walk specified by (1) looks to be a reasonable prior model. We defer further discussion on the precise choice of  $\mathcal{B}$  to the end of §3.

For the X-ray CT images, we note that the boundary in Fig 1(b) is approximately star-shaped, meaning that there is an interior point from which all radiating spokes intersect the boundary precisely once. Therefore we can adopt the same approach as for the ultrasound images, but with  $j$  denoting a polar angle and  $i$  a radius, as shown in Fig 3. The polar centre has been located automatically, by extrapolating a line from the centre of the backbone through the centre of the sheep to an equal distance beyond. This choice adjusts for variations in orientation of animals, because  $j = 180$  will always coincide with the backbone. Also, by positioning the polar centre away from the backbone, we ensure that  $b$  is single valued, especially where the boundary traces around the muscle above the backbone. For simplicity, in what follows we will continue to refer to  $j$  as a ‘column’ index and  $i$  as a ‘row’ index.

We use a border appearance model (Brejl and Sonka, 2000) for the data, conditional on  $b$ . We assume that the pixel which is  $k$  units below the boundary in column  $j$ , denoted  $Y_{b_j+k,j}$ , is Gaussian distributed with mean  $\mu_{k,j}$  standard deviation  $\sigma_{k,j}$ . We form a partial, quasi-likelihood by considering only those pixels that lie within  $K$  units of the boundary and by forming a product of marginal likelihoods, thereby ignoring the dependence between pixel values. Therefore, the probability of the data, given  $b$ , is

$$f(Y|b) \propto \exp \left[ -\frac{1}{2} \sum_{k=-K}^K \sum_{j=0}^J \left( \frac{Y_{b_j+k,j} - \mu_{k,j}}{\sigma_{k,j}} \right)^2 \right]. \quad (2)$$




Figure 3: *Polar coordinates used with X-ray CT lumbar images.*

In cases where  $b_j + k$  falls outside the range 0 to  $I$ ,  $Y_{b_j+k,j}$  is set to its mean value. Again, we obtain values for the means and standard deviations from the training data:

$$\mu_{k,j} = \frac{1}{N} \sum_{n=1}^N Y_{b_j+k,j}^{(n)}, \quad \sigma_{k,j} = \sqrt{\frac{1}{N-1} \sum_{n=1}^N (Y_{b_j+k,j}^{(n)} - \mu_{k,j})^2}, \quad k = -K, \dots, K; \quad j = 0, \dots, J.$$

Fig 4 shows these arrays for the ultrasound and X-ray CT images, with separate blocks for negative and positive  $k$  to aid interpretation. These may be thought of as a grey-level templates for the boundaries. On either side of the upper ultrasound boundary, we see approximately horizontal bands of high and low values for both  $\mu$  (Fig 4(a)) and  $\sigma$  (Fig 4(b)), which will serve as a signature for the automatic positioning of this boundary. For the lower ultrasound boundary the pattern is simpler, with pixels below the boundary having smaller means and standard deviations than those above (Figs 4(c) and (d)). For the polar-transformed X-ray CT boundary, because of correction for sheep rotation, all the backbones coincide to produce a clear pattern for  $\mu$  in Fig 4(e). However, because the size of the backbone varies between sheep,  $\sigma$  (Fig 4(f)) exhibits high values around the edge of the backbone and at other tissue boundaries, and the lowest values of  $\sigma$  occur in the centres of tissues and in the air surrounding the sheep. We defer discussion on the choice of  $K$  to the end of §3.

By Bayes theorem, the posterior density for  $b$  conditional on the data, is proportional to the product of the prior probability and the likelihood:

$$f(b|Y) \propto f(Y|b)f(b).$$

Therefore, combining (2) and (1), to within an additive constant, twice the negative log-posterior density is given by

$$\begin{aligned} F(b) &= -2 \log f(b|Y) \\ &= \sum_{k=-K}^K \sum_{j=0}^J \left( \frac{Y_{b_j+k,j} - \mu_{k,j}}{\sigma_{k,j}} \right)^2 + \lambda \sum_{j=1}^J \phi(\{b_j - b_{j-1}\} - \{\beta_j - \beta_{j-1}\}), \quad b \in \mathcal{B}. \end{aligned} \quad (3)$$




Figure 4: Arrays of means ( $\mu$ ) and standard deviations ( $\sigma$ ) of pixel values on either side of a boundary: (a)  $\mu$  for upper ultrasound boundary, (b)  $\sigma$  for upper ultrasound boundary, (c)  $\mu$  for lower ultrasound boundary, (d)  $\sigma$  for lower ultrasound boundary, (e)  $\mu$  for polar-transformed X-ray CT lumbar images, (f)  $\sigma$  for polar-transformed X-ray CT. (Each  $\sigma$  display is individually scaled, with the smallest value displayed as black and the largest value as white).

This is the function from which we wish to make inference about  $b$ . How we do this depends on our choice of loss function. The most appropriate choice from the perspective of the application is probably the mean-square error of  $b$ , the criterion we use in §4 in assessing the method. However, we would have to resort to very computationally-intensive methods such as MCMC to obtain the optimal estimator. Therefore, in common with many others, we make a choice that simplifies the computation. In our case, we choose the MAP estimator because we can minimise (3) with respect to  $b$  very efficiently using dynamic programming.

### 3 Dynamic programming algorithm

Equation (3) can be re-expressed as

$$F(b) = Y_{b_0,0}^* + \sum_{j=1}^J \left[ Y_{b_j,j}^* + \lambda \phi(\{b_j - b_{j-1}\} - \{\beta_j - \beta_{j-1}\}) \right], \quad b \in \mathcal{B}, \quad (4)$$

where the array  $Y^*$  is obtained from  $Y$  by the application of a non-homogeneous filter:

$$Y_{i,j}^* = \sum_{k=-K}^K \left( \frac{Y_{i+k,j} - \mu_{k,j}}{\sigma_{k,j}} \right)^2 \quad i = 0, \dots, I; \quad j = 0, \dots, J. \quad (5)$$

Note that it is non-homogeneous in the sense that a different operation is applied in each column. Fig 5 shows the effects of applying the filters appropriate for the upper and lower ultrasound boundaries to the data in Fig 1(a). The boundaries that minimise (4) correspond to the darkest paths between the left and right sides of Figs 5(a) and (b), subject to a smoothness constraint imposed by the term involving  $\lambda$  in (4). We can see that these darkest paths are similar to the manual boundaries in Fig 1(a).

Because (4) is a sum of separate costs, one for each column, it satisfies the Principle of Optimality (Bellman, 1957) and its global minimum can be found by dynamic programming. Such algorithms have already been used to find boundaries in images by Amini et al. (1990), Buckley and Yang (1997), and Mortensen and Barrett (1998). Recursively, for  $j = 0, \dots, J$ , we find the minimum cost of a path from column 0 to pixel  $(i, j)$ , which we denote by  $S_{i,j}$ , for  $i \in \mathcal{B}_j$ , where  $\mathcal{B}_j$  is the set of permissible values for  $b_j$ . In column 0 we have simply that

$$S_{i,0} = Y_{i,0}^*, \quad i \in \mathcal{B}_0.$$

Costs to column  $j$  are computed from those to column  $(j - 1)$  by

$$S_{i,j} = \min_{l \in \mathcal{B}_{j-1}} \left[ S_{l,j-1} + Y_{i,j}^* + \lambda \phi(\{i - l\} - \{\beta_j - \beta_{j-1}\}) \right] \quad i \in \mathcal{B}_j. \quad (6)$$

However, as  $\phi(x) = \infty$  if  $|x| > 1$ , we only need consider the three or fewer elements in  $\mathcal{B}_{j-1} \cap \{i - \{\beta_j - \beta_{j-1}\}, i - \{\beta_j - \beta_{j-1}\} \pm 1\}$ . The minimum-cost path reaches  $(i, j)$  via pixel  $(l, j - 1)$  with  $l$  chosen as the minimising argument in (6), which we denote by  $P_{i,j}$ . Upon completion of the recursion, the end-point of the minimum-cost path can be found as

$$b_J = \arg \min_{i \in \mathcal{B}_J} S_{i,J}.$$




Figure 5: Results of applying non-homogeneous filter defined by (5) to ultrasound image Fig 1(a): **(a)** filter to detect upper boundary, **(b)** filter to detect lower boundary. (For each array, the smallest value is displayed as black and the largest value as white).

Then, we recursively track the boundary back to column 0, by

$$b_j = P_{b_{j+1}, j+1} \quad j = (J - 1), \dots, 0.$$

For the X-ray CT images, we would also like the boundary to form a closed contour (i.e.  $b_J = b_0$ ). This can be achieved by modifying the algorithm so that the boundary is constrained to start and finish at row  $i$ , and repeating the optimisation separately for all  $i \in \mathcal{B}_0$ , or alternatively by applying a maximum-flow algorithm (Boykov et al., 2001) from the centre of the sheep to its exterior. In practice, we found that the solution to the unconstrained problem usually formed a closed contour, so we did not implement these modifications.

The Bayesian formulation does not give us a way to choose  $\mathcal{B}$  or  $K$ , and, as discussed in §2, we prefer not to choose a value for  $\lambda$  directly from the training data. Instead, we adopt a pragmatic approach to these choices: we tune the algorithm to get the best agreement between its output and the manually-identified boundaries. We are not aware of anyone else having used this approach. We concede that it is more akin to a penalised likelihood approach (Green and Silverman, 1994) than to a Bayesian one. It is also computationally-intensive, though the intention is that the algorithm is calibrated on a relatively small set of training images and then applied routinely to a large number of images from the same population. Table 1 shows the results of applying the algorithm for a range of settings for the upper and lower boundaries in the ultrasound images. These yield values for  $(\mathcal{B}, K, \lambda)$  of  $(\beta e^{\pm 2\tau}, 10, 2)$  and  $(\beta e^{\pm 3\tau}, 10, 3)$ , respectively, as the parameter settings for which the MAP estimated boundaries are closest to the manual ones.

$\mathcal{B}$ limits	$\beta e^{\pm 2\tau}$			$\beta e^{\pm 3\tau}$			$\beta e^{\pm 4\tau}$		
$K$	10	30	50	10	30	50	10	30	50
upper boundary									
$\lambda = 0$	1.69	1.99	2.04	1.81	1.99	2.04	3.03	2.73	2.04
1	1.44	1.93	1.98	1.47	1.93	1.98	2.49	2.14	1.98
2	<b>1.41</b>	1.96	1.99	1.44	1.96	1.99	2.15	1.96	1.99
3	1.43	1.98	2.06	1.44	1.98	2.06	2.15	1.98	2.06
4	1.43	2.06	2.07	1.44	2.06	2.07	1.44	2.06	2.07
5	1.44	2.08	2.08	1.46	2.08	2.08	1.46	2.08	2.08
6	1.45	2.09	2.12	1.45	2.09	2.12	1.45	2.09	2.12
7	1.48	2.10	2.16	1.48	2.10	2.16	1.48	2.10	2.16
8	1.52	2.13	2.15	1.53	2.13	2.15	1.53	2.13	2.15
9	1.57	2.14	2.17	1.57	2.14	2.17	1.57	2.14	2.17
10	1.62	2.17	2.19	1.61	2.17	2.19	1.61	2.17	2.19
lower boundary									
$\lambda = 0$	6.77	8.21	8.25	6.82	10.00	10.63	9.43	11.74	13.27
1	5.95	7.03	6.65	5.65	8.32	8.97	6.02	9.95	11.67
2	5.70	6.62	6.18	5.07	6.72	6.68	5.56	7.73	9.41
3	5.27	6.27	6.07	<b>4.23</b>	6.62	6.84	4.27	6.72	9.00
4	5.29	6.14	5.96	4.27	6.45	6.93	4.31	6.69	10.36
5	5.28	5.92	5.84	4.28	6.34	6.90	4.33	6.39	10.39
6	5.39	5.99	5.92	4.49	6.28	7.30	4.54	6.55	10.63
7	5.02	6.03	6.42	4.28	7.38	7.41	4.33	7.73	10.70
8	5.06	6.09	6.89	4.32	7.48	8.69	4.37	7.82	11.37
9	5.08	5.94	6.93	5.47	7.52	9.00	5.50	7.88	11.59
10	5.12	6.02	7.23	5.50	7.56	9.04	5.53	7.94	11.62

Table 1: *Root-mean-square differences (measured in pixel units) between manually-drawn and automatically-chosen boundaries, averaged over 72 ultrasound training images, for a range of values of algorithm settings ( $\mathcal{B}, K, \lambda$ ) with the smallest value highlighted.*

## 4 Results

Fig 6 shows the results of applying the algorithm to the images in Fig 1. These results are typical of those obtained with the other training images. The algorithm took 2.5 seconds CPU to find each boundary in an ultrasound image, running Fortran77 on a SunUltra5. For the X-ray CT images, optimal choices of  $K$  and  $\lambda$  were 20 and 0, respectively, and the set  $\mathcal{B}$  was simply defined to be all boundaries  $b$  that lay inside the sheep.

The algorithm was applied to validation sets of a further 72 ultrasound images and a further 99 X-ray CT images. It was also applied to training and validation data of sheep X-ray CT images at two other anatomical positions, the ischium and thorax. Fig 7 illustrates the polar transformations used with the ischium and thorax images and shows typical results. For the ischium, the centres of the two femurs were located and then the polar centre was taken to be




Figure 6: *Results for images in Fig 1: (—) denotes automatically-chosen boundaries, (●) some points on the manually-drawn boundaries; (a) ultrasound scan, (b) X-ray CT lumbar image.*

	Ultrasound boundary		X-ray CT anatomical location		
	upper	lower	lumbar	ischium	thorax
training	1.41	4.23	1.56	0.75	0.69
validation	1.42	†4.94	1.88	0.80	0.77

Table 2: *Root-mean-square differences between manually-drawn and automatically-chosen boundaries in training and validation data sets. (†One image, shown in Fig 8, has been omitted from this summary: a sheep fatter than any in the training data for which the algorithm performed particularly poorly on the ultrasound image, detecting the lower boundary with a root-mean-square error of 64 pixel units.)*

the third vertex to form an equilateral triangle, because for this image our aim is to remove the sex organs. For the thorax, the centre of the sheep was located, and zero angle was determined as the axis of symmetry of the image.

Table 2 summarises all the results, both for the training and validation data sets. As we would expect, the fit to the validation data is not quite so good as for the training data, but is still very good, and that for the ultrasound images is an improvement over that achieved using simpler methods (Glasbey et al., 1996; Glasbey, 1998), as can be seen from figures in these papers. However, there is no guarantee that the algorithm will not occasionally fail spectacularly, as it did for the lower boundary of an ultrasound image of one sheep, shown in Fig 8. This sheep was fatter than any in the training data, and the algorithm confused a boundary within the fat layer with the bottom layer of fat.




Figure 7: *Polar parameterisations used with X-ray CT images at other anatomical positions on sheep, together with typical results of algorithm: (—) denotes automatically-chosen boundaries, (●) some points on the manually-drawn boundaries; (a) ischium, (b) thorax.*




Figure 8: *Results for an ultrasound scan for which the algorithm performed badly: (—) denotes automatically-chosen boundaries, (●) some points on the manually-drawn boundaries.*

## 5 Discussion

The proposed method is computationally efficient for finding a particular class of image boundary, where a boundary intersects each ‘column’ of an image precisely once and the distribution of pixel values in that column is affected by the boundary only to the extent that pixel positions need to be defined relative to this point of intersection. It is also easily calibrated to new data sets. For example, if we are faced with X-ray CT images of a different breed of sheep, we simply have to manually segment  $N$  images and use these to modify the model parameters.

The method is also open to many generalisations. The prior model could include a constraint for the boundary to be symmetric, monotonic or unimodal, and could allow the boundary to be more or less variable in different columns, or we could relax the hard-limits on  $b$ , as discussed in §2. We could also introduce a second-order term involving  $\{b_j - 2b_{j-1} + b_{j-2}\}$  in (1), such as that used in ‘snakes’ (Kass et al., 1988), for which the minimisation problem could be solved using a time-delayed shortest path algorithm (Amini et al., 1990), though Buckley and Yang (1997) also show the benefits of working at sub-pixel resolution. For the border appearance model, a mixture distribution could be used to represent the diversity of physical shapes in a population of sheep, either by including transformations in the single model or by partitioning the training images into homogeneous groups and modelling each separately. The likelihood could be based on transformed pixel values, such as the output from an edge filter or histogram-based classification into different tissue types, or could take into account cross-correlations between pixel values in the same column. However, the model formulation does not permit us to introduce correlations between pixel values in different columns. The border appearance model could include a magnification factor, which could also be estimated by dynamic programming, by finding a path in a higher-dimensional space. Finally, we note that the dynamic programming algorithm can be modified so that the boundary is constrained to pass through a fixed point. By varying the point along a column and seeing by how much  $F(b)$  is increased, we gain an indication of how certain our positioning of the boundary is in this column.

## Acknowledgements

This work was funded by the Ministry of Agriculture, Fisheries and Food, the Meat and Livestock Commission and the Scottish Executive Rural Affairs Department as part of the LINK Sustainable Livestock Production Programme.

## References

Amini, A. A., Weymouth, T. E., and Jain, R. C. (1990). Using dynamic programming for solving variational problems in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:855–867.

- Bellman, R. (1957). *Dynamic Programming*. Princeton University Press, Princeton.
- Besag, J. (1986). Statistical analysis of dirty pictures (with discussion). *Journal of the Royal Statistical Society, Series B*, 48:259–302.
- Boykov, Y., Veksler, O., and Zabih, R. (2001). Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. (in press).
- Brejl, M. and Sonka, M. (2000). Object localization and border detection criteria design in edge-based image segmentation: Automated learning from examples. *IEEE Transactions on Medical Imaging*, 19:973–985.
- Buckley, M. and Yang, J. (1997). Regularised shortest-path extraction. *Pattern Recognition Letters*, 18:621–629.
- Cootes, T. F., Taylor, C. J., Cooper, D. H., and Graham, J. (1995). Active shape models – their training and application. *Computer Vision and Image Understanding*, 61:38–59.
- Ferrari, P. A., Frigessi, A., and Gonzaga de Sa, P. (1995). Fast approximate maximum *a posteriori* restoration of multicolour images. *Journal of the Royal Statistical Society, Series B*, 57:485–500.
- Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–735.
- Glasbey, C. A. (1998). Ultrasound image segmentation using stochastic templates. *Journal of Computing and Information Technology*, 6:107–116.
- Glasbey, C. A., Abdalla, I., and Simm, G. (1996). Towards automatic interpretation of sheep ultrasound scans. *Animal Science*, 62:309–315.
- Glasbey, C. A. and Horgan, G. W. (1995). *Image Analysis for the Biological Sciences*. Wiley, Chichester.
- Green, P. J. and Silverman, B. W. (1994). *Nonparametric Regression and Generalized Linear Models*. Chapman and Hall, London.
- Greig, D. M., Porteous, B. T., and Seheult, A. H. (1989). Exact maximum *a posteriori* estimation for binary images. *Journal of the Royal Statistical Society, Series B*, 51:271–279.
- Grenander, U. and Miller, M. I. (1994). Representations of knowledge in complex systems (with discussion). *Journal of the Royal Statistical Society, Series B*, 56:549–603.
- Kass, M., Witkin, A., and Terzopoulos, D. (1988). Snakes: Active contour models. *International Journal of Computer Vision*, 1:321–331.
- Mortensen, E. N. and Barrett, W. A. (1998). Interactive segmentation with Intelligent Scissors. *Graphical Models and Image Processing*, 60:349–384.
- Phillips, D. B. and Smith, A. F. M. (1994). Bayesian faces via hierarchical template modelling. *Journal of the American Statistical Association*, 89:1151–1163.

Rue, H. (1997). A loss function model for the restoration of grey level images. *Scandinavian Journal of Statistics*, 24:103–114.

Simm, G. (1992). Selection for lean meat production in sheep. In Speedy, A. W., editor, *Recent Advances in Sheep and Goat Research*, pages 193–215. CAB International.