HIGH–LEVEL IMAGE PRIORS IN CONFOCAL MICROSCOPY APPLICATIONS

MERRILEE HURN*
School of Mathematical Sciences
University of Bath, U.K.

HÅVARD RUE
Department of Mathematical Sciences
NTNU, Norway

March 1997

ABSTRACT

The incorporation of high–level prior information in Bayesian imaging has been proposed via the theoretical frameworks both of Grenander and Miller (1994) and Baddeley and Van Lieshout (1993). Recent developments in Markov chain Monte Carlo algorithms for varying dimension problems (Green, 1995) have made the application of these high–level models more widely accessible, although this remains more complicated than using the usual low–level Markov random fields models. One application area in which incorporating information on the object level in a scene would be of particular benefit is in confocal microscopy. In this paper we outline two such applications, the first addresses uncertainty in segmenting confocal microscope images, and the second is an object recognition problem.

1 INTRODUCTION

High–level models specify an image in global terms, describing it as a list of its constituent objects rather than as the list of pixels on which the data are recorded or digitised. The complexity of such models is far greater than that of the commonly used Markov random field low–level models, in part because the dimension of the parameter space may not be fixed. Two main approaches have been proposed; Grenander and Miller (1994) construct a varying dimension model as the union of models containing fixed but different numbers of deformed templates, each template representing one object. Moves between scenes containing different numbers of objects are handled by a jump–diffusion algorithm. Baddeley and Van Lieshout (1993) suggest using a marked point process model as an object prior, with points representing the locations of objects and the marks being the variables needed to describe the objects themselves (for example a deformable template). Here spatial birth and death algorithms are used to vary the number of objects.

Microscopy is one area in which the scientific questions of interest in biological applications may mean that a high–level model is appropriate; for example, how many cells of a certain type are present

*ADDRESS FOR CORRESPONDENCE: School of Mathematical Sciences, University of Bath, Bath BA2 7AY.
E–MAIL: m.a.hurn@maths.bath.ac.uk and havard.rue@math.ntnu.no

in a specimen? what are their typical shapes? how is a particular cell's volume changing over time? Several authors have now considered applications in this area (Grenander and Miller, 1994; Qian and Mardia, 1995; Rue and Syversveen, 1995; Hurn, 1996; Qian et al., 1996; Rue and Hurn, 1997). In this paper, we will concentrate on one particular form of microscopy, namely confocal fluorescence microscopy. This technique images optical rather than physical sections through the specimen, allowing detailed study of living cells in 3–D and over time. As such, it is an important tool in biological studies. The basic principle is that the object under study is stained with a fluorescent dye, then using a system of lenses and pinholes, a laser is focused at a point in the specimen. Any dye at this focal point will respond, and a record is made of the fluorescence coming from that point alone. The focal point of the laser is then scanned through the specimen, building up a point–by–point image of the object as a function its dye concentration. Out–of–focus information should be removed by this optical set–up, thereby providing the desired optical sectioning property.

2 SEGMENTATION UNCERTAINTY IN MEASURING CELLS

Our first example concerns the estimation of various size and shape attributes of leaf cells in a plant science application (Hurn, 1996). The full aim of this study is to follow how pairs of stomatal guard cells change shape and volume during an opening and closing sequence, assessed by taking images at 20 minute intervals. By staining and thus imaging different cells in the leaf, it may be possible to learn about the mechanism of the opening response (which is thought to be driven by the movement of water around the leaf).

The first step in the measurement process is to segment the cells from the background. Unfortunately, the nature of confocal microscopy complicates this segmentation as the images it produces tend to be both blurred and noisy to some extent (see Figure 1(a)). An important point to note is that the image degradation is due partially to the specimen under study because it forms part of the optical system (through diffraction and attenuation effects, amongst others). When the segmentation is not clear–cut, there will be some uncertainty in the size and shape estimates derived from it. The aim of the Bayesian approach applied here is to assess some of the variability associated with the estimates due to this segmentation uncertainty.

In this application, the image is known to contain a pair of cells; for this reason, a high–level prior would seem an appropriate way to include this information. The marked point process framework of Baddeley and Van Lieshout (1993) is used, although its full generality is not required since the dimensions of the problem are fixed. We shall represent the scene $\boldsymbol{X}$ as a list of its constituent cells $\boldsymbol{X} = \{X_1, X_2\}$. Each cell X_p will be represented by an $2N$–sided simple polygon. These polygons can be parameterised in a number of ways; the first one we shall use, fitting obviously into the marked point process framework, is to parameterise the polygon by its centroid $\boldsymbol{o}$ together with a full set of Fourier descriptors $\{a_k, b_k, c_k, d_k\}$, so that the i and j coordinates of its $2N$ vertices may be written

$$i_n = \boldsymbol{o}_i + \sum_{k=1}^{N} a_k \cos(\pi kn/N) + \sum_{k=1}^{N-1} b_k \sin(\pi kn/N)$$

$$j_n = \boldsymbol{o}_j + \sum_{k=1}^{N} c_k \cos(\pi k n/N) + \sum_{k=1}^{N-1} d_k \sin(\pi k n/N)$$

for $n = 1, ..., 2N$. We will use the Fourier descriptors as marks since it is possible to interpret a prior distribution on the descriptors in terms of smoothness of the resulting polygon (see Hurn, 1996, for more details). We can then define an object prior for $\boldsymbol{X}$ by specifying that the two centroids are uniformly distributed in the window and that the objects have a hard-core interaction, that is the polygons are not permitted to overlap. (This distribution is with respect to a Poisson process conditioned to have exactly two points.) Notice that it is possible to simulate from this distribution using Markov chain Monte Carlo methods (Hastings, 1970), by proposing changes either to the locations s, or to global properties such as size and shape (via some combination of the Fourier descriptors), or to the individual vertex locations (since the required reparameterisation is linear, the associated Jacobian will cancel in the Metropolis-Hastings step).

The likelihood model for how the data $\boldsymbol{y}$ arise is based on approximating the following: The microscope is actually measuring low numbers of fluorescence photon counts which in theory come only from the focal point, but in practice they are affected by specimen dependent 3–D spatially varying scatter and attenuation effects (the net result of which is to blur the images). On top of this, the counts are adjusted by an operator–determined additive constant known as the black–level, and are then recorded subject to instrument white noise. We begin by assuming that each cell has a typical fluorescence level μ_p, but that the pixel lattice of expected Poison counts are formed after blurring these values by a Gaussian point spread function. We then model the record at a pixel (i, j) as a Gaussian random variable with mean equal to the black–level plus the expected Poisson count for (i, j), and with variance equal to the variance of the instrument noise again plus the expected Poisson count. This introduces additional model parameters, namely the typical fluorescence levels, the parameter of the blurring kernel, the black level and the variance of the instrument noise. To estimate these parameters well requires a good initial estimate of the segmentation; since the aim is to address some of the segmentation uncertainty, we will treat the problem in a Bayesian way, assigning vague priors to all these parameters (the $\{\mu_p\}$ may be regarded as additional marks). Each parameter can then be sampled in turn by drawing from its distribution conditional on the remaining parameters and the current realisation of $\boldsymbol{X}$. See Hurn (1996) for more details of a similar approach applied to a simpler likelihood approximation.

Figure 1 shows the results for a single section of stomatal guard cell data. The initial segmentation, Figure 1(b), has been obtained by thresholding the data then followed by some morphological processing; it is used to initialise $\boldsymbol{X}$. Figures 1(c) and (d) show a realisation of $\boldsymbol{X}$ overlaid on the data displayed in reverse–video, and the corresponding pixel classification into cell and background. The sequence of realised segmentations can be used to generate interval estimates for the size and shape attributes of interest, for example Figure 2 shows the areas and perimeters (in pixel units) for the two cells observed in a long run of MCMC sampling, other shape attributes such as compactness could also be monitored.

The high–level prior outlined here can be extended to deal with the case of an unknown number of objects. Details of a number of useful MCMC proposals, including births, deaths and split and merge type moves, can be obtained from the authors.

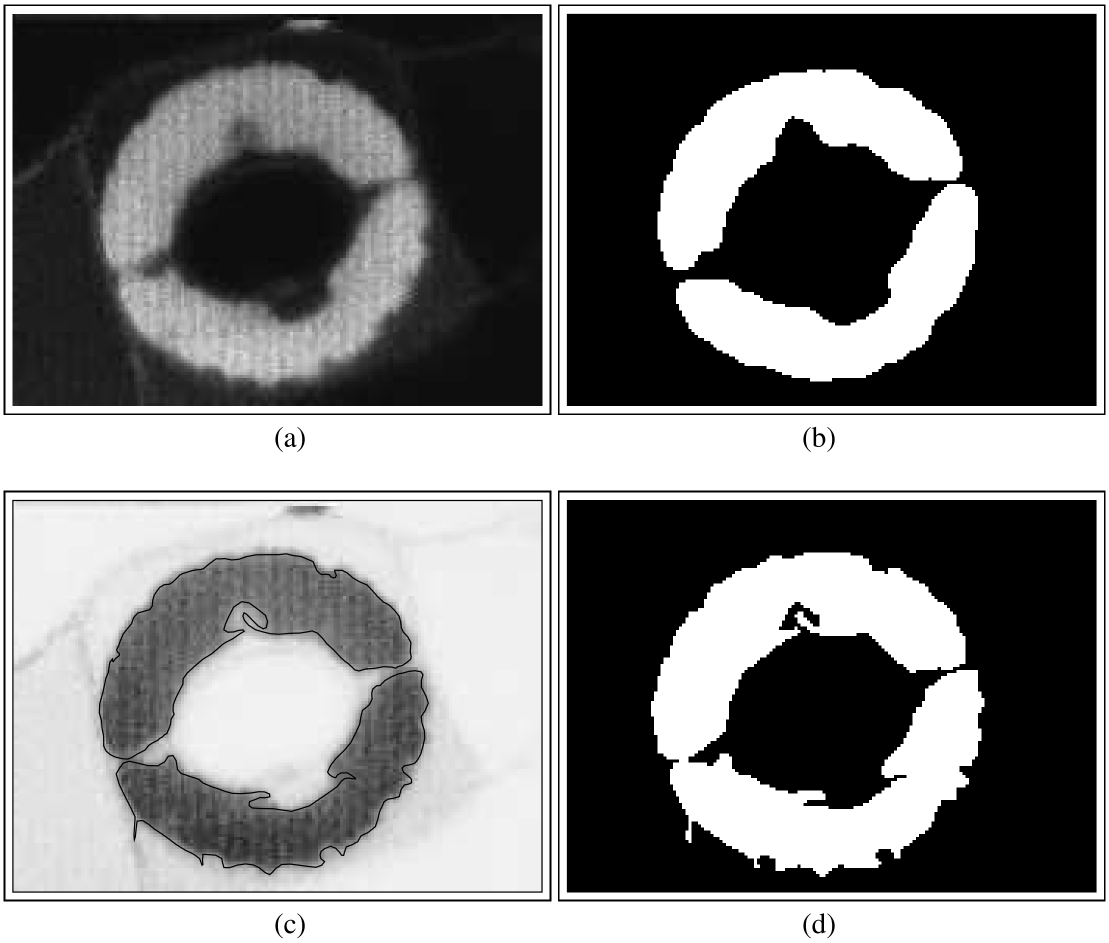

Figure 1: (a) An optical section through a pair of stomatal guard cells; (b) an initial segmentation into cell and background; (c) a realisation of $\boldsymbol{X}$ overlaid on the data (displayed in reverse video); (d) the segmentation corresponding to (c).

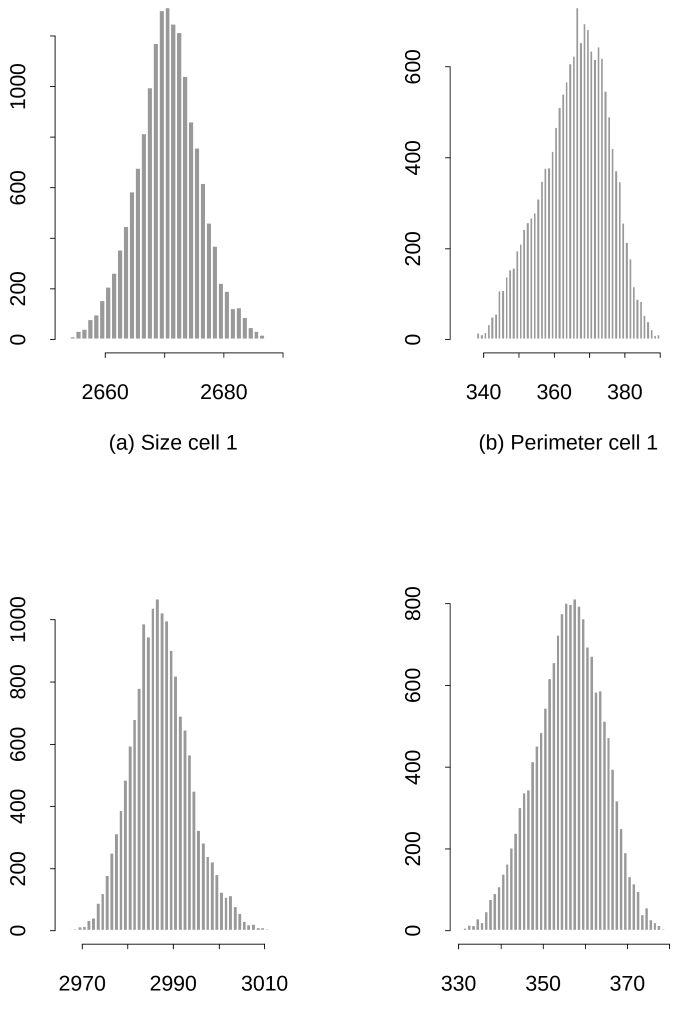

(a) Size cell 1

(b) Perimeter cell 1

(c) Size cell 2

(d) Perimeter cell 2

Figure 2: Histograms of the area and perimeter values observed in the sequence of segmentations derived from MCMC sampling the posterior distribution $\pi(\boldsymbol{x}|\boldsymbol{y})$.

3 CELL IDENTIFICATION

The second application considered is an object recognition problem where the aim is to locate and identify an unknown number of cells in the data. This differs from the first application in a number of ways which might affect the way we choose to approach it. Firstly, it may well be the case that we are less interested in detailed shape estimates for each identified cell, and more interested in the location and classification of possible cells in the data. Secondly, we now have a set of possible cell types into one of which each cell will be classified on the basis of its size, shape, fluorescence level, or some other discriminating factor. In the previous application, very little prior information about the cell shape was assumed; here shape may be one of the main discriminatory factors, it may be known in advance what a typical cell of each class looks like, and learning about the shape of any individual cell is not of primary concern. These differences suggest modelling the set of cell types by a sequence of templates which encode the main discriminating features, and then to model individual cells as deformations of these templates (Grenander, 1993; Grenander and Miller, 1994).

We propose to use a deformable template to model the outline of each cell–type, using the parameters of the deformation as the marks for each cell within the marked point process framework. This requires forming the mark distribution as a mixture distribution of the different template types m and the deformations $\boldsymbol{s}$ given a particular template type:

$$\pi(\boldsymbol{s}, m) = \pi(\boldsymbol{s}|m)\pi(m).$$

The deformed template is located in the data window by a point, which in this case, is the location of a designated vertex rather than its centroid. The standard MCMC move types would then be births and deaths of cells, and changes to the locations or the marks of any particular cell. Notice that the numbers of parameters defining the different templates may not be constant (they may be polygons with different numbers of vertices) and so these last updates either alter $\boldsymbol{s}$ alone (altering just the shape) or both $\boldsymbol{s}$ and m together (altering the object's label and shape). Unfortunately, this collection of moves may lead to slow movement around the sample space. In particular proposed label changes may be accepted infrequently because existing updates may have already matched the current polygon rather well to the data so that even if the proposed cell–type is more appropriate, the proposed fit will not be as good. For this reason, we extend the template classes, allowing each template to have a stochastic number of vertices in order to generate good moves between the different object–types. If the number of vertices is denoted n, the mark mixture distribution then becomes

$$\pi(\boldsymbol{s}, n, m) = \pi(\boldsymbol{s}|n, m)\pi(n|m)\pi(m).$$

Changing the resolution n of a template requires a dimension change in the model (Green, 1995). Label changes are then only proposed between templates with the same number of vertices. This means that the label can be changed without altering the current shape, preserving existing edge-fitting. The acceptance probability then depends on the ratio of the density of the observed polygon under the two model interpretations. See Rue and Hurn (1997) for more details.

Another important class of move types would be to allow nearby cells to merge together forming a single cell, or conversely to split apart again into two cells. This improves the MCMC mixing over the varying dimension of the number of cells present (the same changes would be very slow working with birth and death steps alone). Using a reparameterisation of the marked point process from locations c and deformations s to the positions of the vertices themselves, it is possible to devise efficient transitions which again preserve much of the existing edge-fitting (Rue and Hurn, 1997).

We will consider a small illustrative data set, shown in Figure 3(d). The image consists of three different types of cells, each crudely described by either a triangle, or an ellipse, or a circle. The three basic templates are shown in Figures 3(a)-(c) using the minimum number of vertices allowed for each shape. Increasing the number of vertices for any template does not change its shape; the additional vertex is placed at the midpoint of the currently longest side.

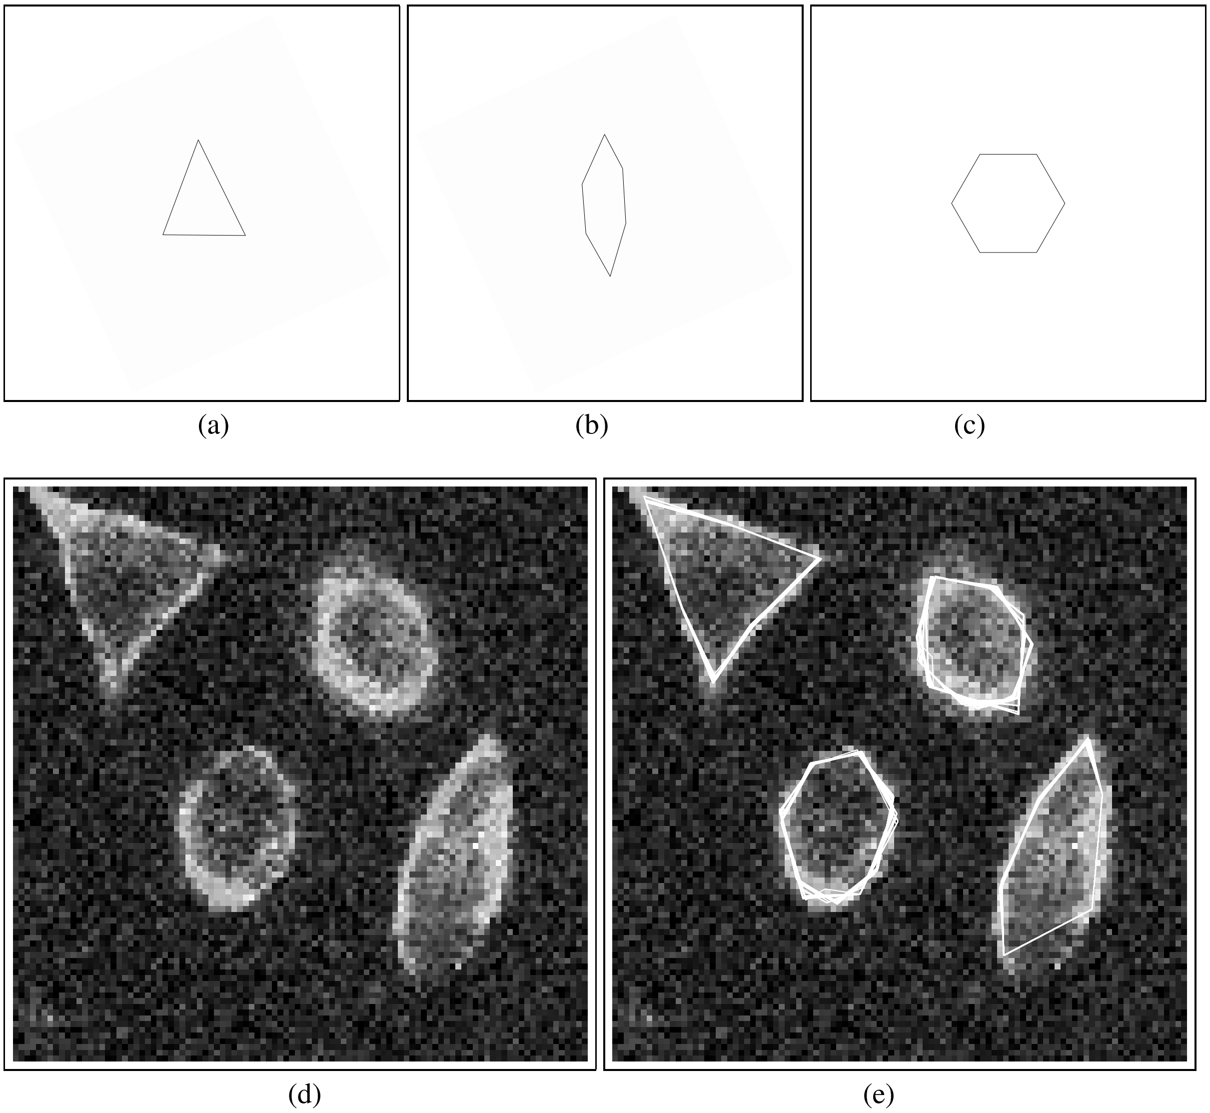

Figure 3: The three template types, each shown with the minimum allowable number of vertices (a) triangle; (b) ellipse; (c) circle; (d) the small data set; (e) some samples of the polygons shown overlaid on the data.

The likelihood model required in this application will differ from that in the previous application because here only the cell edges are fluorescing (hence the bright rims and darker interiors to the objects). The data are assumed to be independently Normally distributed with a common variance and with means which depends on whether the pixel in question lies outside the polygon, on a two–pixel–wide strip lying on the polygon, or within its interior. These model parameters are held fixed during the simulation. This model is a fairly crude approximation, however as already stated we are primarily interested in cell location and labelling in this application rather than assessing the quality of segmentation.

Using the varying resolution families of templates and the move types described, in addition to cell births, deaths and fixed resolution shape updates, the MCMC algorithm is initialised with a configuration containing no cells. After discarding an initial burn–in run, Figure 3(d) shows several well–spaced realisations from the model overlaid on the data. The algorithm has located the four quite clearly defined cells. For each of the displayed configurations of polygons, the corresponding labels are as a triangle, two circles and an ellipse. If a point estimate is required rather than a sequence of typical realisations, Rue and Hurn (1997) describe a point estimator for the number of objects together with their locations, shapes and labels, based on Baddeley's delta metric (Baddeley, 1992).

Finally, notice that it is possible to extend the template models to deal with the problem observed in fitting the ellipse in the lower right corner. The difficulty here is that the fluorescent material is not uniformly distributed over the surface of the cell; as a result, part of the cell's rim appears to be missing in the image. Rue and Husby (1997) demonstrate how the destructive deformation idea of Grenander (1993) can easily be applied to correct for this effect.

ACKNOWLEDGEMENTS

The authors are grateful to the Research Council of Norway and the British Council for funding under their scholarship scheme for collaborative research. MAH also acknowledges the financial support of the Nuffield Foundation. Thanks to Nick White and the Plant Science Department at Oxford University for describing the application and providing the data used in our first example, and to Mari Hjelstuen of the Biophysics group at NTNU for providing the data used in our second example.

REFERENCES

Baddeley, A. J. (1992). An error metric for binary images. In Förstner, W. and Ruwiedel, S., editors, *Robust Computer Vision: Quality of Vision Algorithms*, pages 59–78. Wichmann Verlag.

Baddeley, A. J. and Van Lieshout, M. N. M. (1993). Stochastic geometry models in high–level vision. In Mardia, K. V. and Kanji, G. K., editors, *Statistics and Images*, volume 20, chapter 11, pages 235–256. Abingdon: Carfax Publishing.

Green, P. J. (1995). Reversible jump MCMC computation and Bayesian model determination. *Biometrika*, 82(4):711–732.

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

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.

Hastings, W. K. (1970). Monte Carlo simulation methods using Markov chains and their applications. *Biometrica*, 57:97–109.

Hurn, M. A. (1996). Bayesian image analysis in confocal fluorescence microscopy. Statistical Research Report 96:01, School of Mathematical Sciences, University of Bath.

Qian, W. and Mardia, K. V. (1995). Recognition of multiple objects with occlusions. Stat 95/01, University of Leeds.

Qian, W., Titterington, D. M., and Chapman, J. N. (1996). An image analysis problem in electron microscopy. *Journal of the American Statistical Association*, 191(435):944–952.

Rue, H. and Hurn, M. A. (1997). Bayesian object identification. Statistics, Department of Mathematical Sciences, Norwegian University of Science and Technology, Trondheim, Norway.

Rue, H. and Husby, O. K. (1997). Identification of partly destroyed objects using deformable templates. Statistics 1, Department of Mathematical Sciences, Norwegian University of Science and Technology, Trondheim, Norway.

Rue, H. and Syversveen, A. R. (1995). Bayesian object recognition with Baddeley's delta loss. Statistics 8, Departments of Mathematics, The University of Trondheim, Trondheim, Norway.