Statistical Snakes: Active Region Models†

Jim Ivins†† and John Porrill

Artificial Intelligence Vision Research Unit
University Of Sheffield
England S10 2TN

Abstract

This paper describes a new region-growing technique that uses a closed snake driven by a pressure force that is a function of the statistical characteristics of image data. This statistical snake expands until its elements encounter pixels that lie outside user-defined limits relative to a seed region; when these limits are violated the pressure force is reversed to make the model contract. Tension and stiffness forces keep the boundary of the region model smooth, and a repulsion force prevents self-intersection. Boundary elements can be added and removed in response to complexity changes, and the tension, stiffness and pressure parameters can be adjusted to preserve the energy balance of the changing model. Statistical snakes have been used to segment a variety of images including composite textures and NMR data volumes.

1 Introduction

Active contour models (snakes) – see Kass et al (1987) – are energy-minimising splines constrained by internal *tension* (stretching) and *stiffness* (bending) forces. A *pressure* force can also be added so the models attempt to maximise their areas – see Sullivan et al (1990).

The mechanical properties of a pressurised snake $\mathbf{u}(\lambda) = (\, x(\lambda), y(\lambda)\,)$ are specified by an energy functional E containing four terms weighted by constants α, β and ρ:

$$E = \underbrace{\frac{\alpha}{2}\oint\left|\frac{\partial\mathbf{u}}{\partial\lambda}\right|^2 d\lambda}_{\text{Tension}} + \underbrace{\frac{\beta}{2}\oint\left|\frac{\partial^2\mathbf{u}}{\partial\lambda^2}\right|^2 d\lambda}_{\text{Stiffness}} - \underbrace{\frac{\rho}{2}\oint\frac{\partial\mathbf{u}}{\partial\lambda}\times\mathbf{u}\, d\lambda}_{\text{Pressure}} + \underbrace{\oint P(\, I(\mathbf{u})\,)\, d\lambda}_{\text{Potential}}$$

The *potential* P is usually generated by processing the image I to highlight edges.

Unfortunately, an image feature must produce quite strong edges to overcome the pressure and allow the snake to reach equilibrium. As a result, it is not always possible to segment an image using pressurised snakes. *Statistical snakes* – see Ivins and Porrill (1993) – overcome this difficulty and can model image features that cannot be extracted using ordinary snakes.

---

† This work was carried out in collaboration with IBM UK Scientific Centre as part of AIM Project A2003: Computer Vision in Radiology (COVIRA).

†† E-mail J.Ivins@aivru.sheffield.ac.uk. JI was funded by the SERC.

BMVC 1994 doi:10.5244/C.8.37

A statistical snake is an *active region model* driven by an energy term that links the pressure force to the image data as described in sections 2 and 3. The new model can expand and contract considerably during energy minimisation; section 4 therefore introduces a boundary-tracking mechanism that prevents self-intersection, and section 5 describes how boundary elements can be added and removed in response to complexity changes. (An appendix describes how the tension, stiffness and pressure parameters can be adjusted to preserve the energy balance of a snake during re-parameterisation.) Finally, the versatility of statistical snakes is illustrated in section 6 which concerns texture segmentation.

2 Region Energy

Let G be a functional that measures the *goodness* of pixels within the region R of an image $I(x, y)$ enclosed by a statistical snake $\mathbf{u}(\lambda)$; if desirable features produce large values of G then a region energy E_{region} can be defined as an area integral:

$$E_{region} \quad = \quad -\rho \iint_R G(\, I(x, y)\,)\; dx\, dy$$

The region energy is the total goodness of the pixels enclosed by the snake. The change in area δA that arises from a small arbitrary change $\delta\mathbf{u}$ to the snake is:

$$\delta A(\mathbf{u}) \quad = \quad \frac{1}{2} \oint \left(\frac{\partial \mathbf{u}}{\partial \lambda} \right)^{\perp} \cdot \delta\mathbf{u}\; d\lambda$$

Since $\delta\mathbf{u}$ is small, the corresponding change in region energy δE_{region} can be approximated by multiplying the goodness value $G(\, I(\mathbf{u})\,)$ at each element by the local change in area:†

$$\delta E_{region} \quad \approx \quad -\rho \oint G(\, I(\mathbf{u})\,)\; \delta A(\mathbf{u})\; d\lambda \quad = \quad -\frac{\rho}{2} \oint G(\, I(\mathbf{u})\,) \left(\frac{\partial \mathbf{u}}{\partial \lambda} \right)^{\perp} \cdot \delta\mathbf{u}\; d\lambda$$

The equations of motion for energy minimisation by iterative gradient descent are obtained using variational calculus:

$$\frac{\partial \mathbf{u}}{\partial t} \quad = \quad \underbrace{\alpha \frac{\partial^2 \mathbf{u}}{\partial \lambda^2}}_{\text{Tension}} \quad - \quad \underbrace{\beta \frac{\partial^4 \mathbf{u}}{\partial \lambda^4}}_{\text{Stiffness}} \quad + \quad \underbrace{\frac{\rho}{2}\, G(\, I(\mathbf{u})\,) \left(\frac{\partial \mathbf{u}}{\partial \lambda} \right)^{\perp}}_{\text{Pressure}}$$

The three terms in these equations are the tension, stiffness and region (pressure) forces. The pressure parameter is usually kept at unity while the tension and stiffness parameters vary between 0.1 and 1.0. The smoothing forces, although often small in comparison with the pressure, are very important because they hold the region model together during expansion and contraction. Pressure forces are considered in more detail in the next section.

---

† This approximation can be improved by computing G from the mean pixel intensity along the two boundary sections that incorporate the element.

3 Region (Pressure) Forces And Goodness Functionals

According to the equations of motion the pressure force **f** arising from the region energy is simply the normal to the snake boundary, weighted by the local G value:

$$E_{region} = -\rho \iint_R G(\, I(x,y)\,)\; dx\, dy \qquad \mathbf{f}(\mathbf{u}) = \frac{\rho}{2}\, G(\, I(\mathbf{u})\,) \left(\frac{\partial \mathbf{u}}{\partial \lambda}\right)^{\perp}$$

Because the goodness functional takes the same form in both the region energy and the pressure force, *any* functional can be used to evaluate the pixels inside the region model. For example, the pressure term used by Sullivan et al (1990) is equivalent to a region energy that measures the area of the snake:

$$G(\, I(x,y)\,) = +1 \qquad \mathbf{f}(\mathbf{u}) = \frac{\rho}{2}\left(\frac{\partial \mathbf{u}}{\partial \lambda}\right)^{\perp}$$

All pixels are considered good enough for inclusion in the model, and the boundary expands until either it is trapped by strong edges, or the smoothing forces become prohibitive. Three alternative goodness functionals are defined below using first-order statistics.

Binary Pressure. A binary functional links the pressure force to the image data and makes the region model contract if user-defined statistical limits are violated:

$$G(\, I(x,y)\,) = +1 \qquad (\, |I(x,y) - \mu| \le k\,\sigma\,) \qquad \mathbf{f}(\mathbf{u}) = +\frac{\rho}{2}\left(\frac{\partial \mathbf{u}}{\partial \lambda}\right)^{\perp}$$

$$G(\, I(x,y)\,) = -1 \qquad (\, |I(x,y) - \mu| > k\,\sigma\,) \qquad \mathbf{f}(\mathbf{u}) = -\frac{\rho}{2}\left(\frac{\partial \mathbf{u}}{\partial \lambda}\right)^{\perp}$$

The mean intensity μ and standard deviation σ are computed from a seed region;† k is a user-defined constant (typically 2 or 3).

Linear Pressure. The direction of the pressure force is continually changing as the boundary elements move between pixels on either side of the statistical threshold for inclusion in the model. Stability can therefore be improved by using a normalised goodness functional:

$$G(\, I(x,y)\,) = 1 - \frac{|I(x,y) - \mu|}{k\,\sigma} \qquad \mathbf{f}(\mathbf{u}) = \rho \left(\frac{\partial \mathbf{u}}{\partial \lambda}\right)^{\perp}\left(1 - \frac{|I(\mathbf{u}) - \mu|}{k\,\sigma}\right)$$

This linear pressure force is maximal – making the model expand rapidly – when $I(\mathbf{u}) - \mu = 0$. Conversely, when $|\, I(\mathbf{u}) - \mu\,|$ is very large the model will contract very quickly. Of course, when $|\, I(\mathbf{u}) - \mu\,| = k\,\sigma$ the pressure force will be zero, allowing the model to stabilise (figure 1).

---

† Standard algorithms exist for region filling – for example, see Hechbert (1990).

Mahalanobis Pressure. The binary and linear pressure terms can be extended to deal with a multi-dimensional feature vector $\mathbf{I}(\mathbf{u})$ instead of the scalar $I(\mathbf{u})$; for example:

$$G(\,\mathbf{I}(x,y)\,) \;=\; 1-\frac{1}{k}\sqrt{\chi^2} \qquad\qquad \mathbf{f}(\mathbf{u}) \;=\; \rho\left(\frac{\partial\mathbf{u}}{\partial\lambda}\right)^{\perp}\left(1-\frac{1}{k}\sqrt{\chi^2}\,\right)$$

$$\chi^2 \;=\; (\mathbf{I}(\mathbf{u})-\boldsymbol{\mu}_{\mathbf{I}})^T\mathbf{S}_{\mathbf{I}}^{-1}(\mathbf{I}(\mathbf{u})-\boldsymbol{\mu}_{\mathbf{I}})$$

The limit k is chosen from standard χ^2 tables according to the number of images. The mean vector $\boldsymbol{\mu}_{\mathbf{I}}$ and covariance matrix $\mathbf{S}_{\mathbf{I}}$ are calculated from a seed region in each image. Using the Mahalanobis pressure force a statistical snake can combine information from multiple images such as the texture potentials described in section 6.

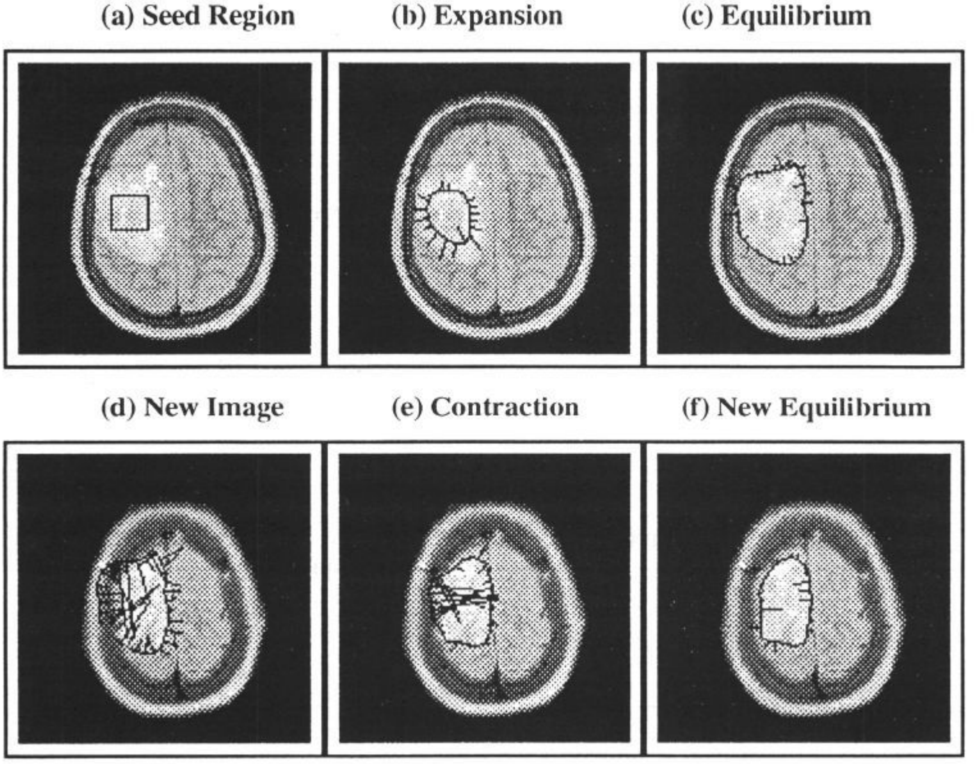

Figure 1 shows two different slices of NMR data. (**a**) The user defines a seed region (inside the pale area of inflamed tissue) from which the pixels are automatically analysed for their mean and standard deviation. (**b**) The boundary of the seed region is then converted into a statistical snake that expands until its elements encounter pixels that lie outside pre-defined limits relative to the seed statistics. (**c**) At equilibrium the model oscillates slightly – as shown by the linear pressure forces normal to the boundary. (**d**) The model is copied to another image where it is exposed to new pressure forces. (**e**) The model deforms to fit the new data – boundary sections that lie outside the pale region contract. (**f**) Within 50 iterations the model fits the new data ($\alpha = \beta = \rho = 1$; $k = 2$).

Detail in an image is only represented to pixel accuracy; steps of less than one pixel are therefore unnecessary except in the final stages of energy minimisation. Steps greater than unity are undesirable because of the possibility of overlooking detail. The energy of a snake can therefore be minimised by altering the position of each element *in turn* according to the equations of motion. The component forces in these equations are approximated using finite differences, and the total force is normalised to give one-pixel steps. This method avoids the need for implicit (matrix inversion) schemes to guarantee stability, and leads to efficient implementation of boundary tracking (section 4) and re-parameterisation (section 5).

4 Boundary Tracking

If two sections of an active region model are allowed to cross then it becomes unstable because the direction of the pressure force is reversed relative to the interior of the model (figure 2). The model is therefore tracked by calculating the discrete coordinates of every point on its boundary, and incrementing the corresponding cells in a two-dimensional accumulator. Intersection can then be avoided by testing the appropriate accumulator cells each time an element is moved. The existing boundary sections are first subtracted from the accumulator, and the new locations are then tested for non-empty cells which indicate that an intersection will occur if the element is moved as intended.†

Although there are N elements, each of which can pass through any of the $N-2$ non-adjacent boundary sections, the accumulator mechanism retains $O(N)$ complexity and will deal with multiple snakes.

If moving a boundary element will produce an intersection then the simplest response is to cancel the move. However, this will inhibit the dynamic properties of the model because portions of the boundary are frozen until nearby elements move away. A repulsion force is therefore calculated from the tension, stiffness, and *reversed* pressure. If this alternative force also produces an intersection then – as a last resort – no move is made.

† Because the model is segmenting a discrete image it is only necessary to record the boundary to pixel accuracy. Points between elements can therefore be obtained using a line plotting algorithm; for example, see Bresenham (1965). Boundary *testing* must be done with *four-way* line connections; however, *plotting* can be done using *eight-way* connections.

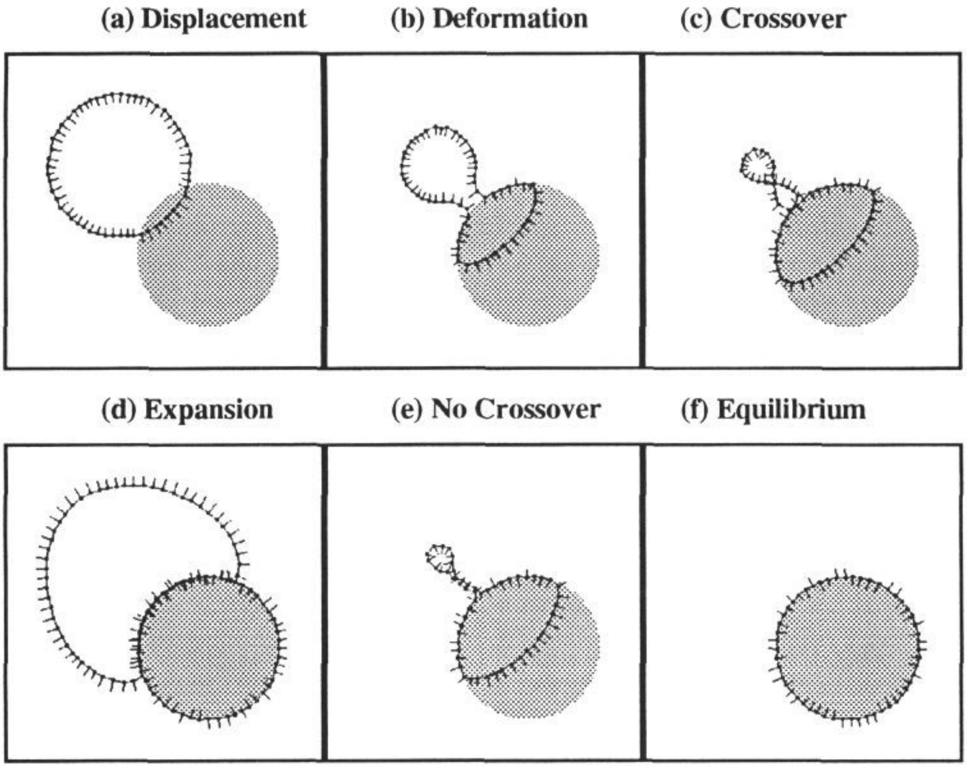

Figure 2. (**a**) An active region model manually displaced from the synthetic grey circle. (**b**) The model begins to deform under the influence of the binary pressure forces shown as normals to the boundary. (**c**) Opposing boundary sections cross over as the model collapses. (**d**) The sections that crossed now expand uncontrollably because the direction of the pressure force is reversed relative to the interior of the model. (**e**) Self-intersection cannot occur if the region boundary is tracked (compare with c). (**f**) The impenetrable model regains equilibrium within 100 iterations ($\alpha = \beta = 0.1$; $\rho = 1$; $k = 2$).

5 Re-parameterisation

Adjacent boundary elements become separated as a statistical snake expands; it is therefore necessary to insert extra elements if a suitable amount of detail is to be incorporated into the model. Similar considerations apply to a collapsing model – closely packed elements can be removed to reduce processing requirements. Boundary elements are therefore inserted or deleted whenever the length of a boundary section violates maximum or minimum limits. (Good results are obtained with a minimum around 10 pixels and a maximum about twice this.)

The behaviour of the model is generally easier to predict if the energy balance is preserved during re-parameterisation. This is achieved using constants α_0, β_0 and ρ_0 to calculate α, β and ρ according to the number of snake elements N at the current iteration:

$$\alpha = \alpha_0 N \qquad \beta = \beta_0 N^3 \qquad \rho = \rho_0$$

These formulae are derived in the appendix. Experience suggests that the energy parameters of a typical snake containing 32 elements can all be set to unity. The initial constants are therefore set as follows:

$$\alpha_0 = 1 / 32 \qquad \beta_0 = 1 / 32768 \qquad \rho_0 = 1$$

6 Example: Texture Energy

Laws (1979) introduced the concept of texture energy by convolving images with small masks and replacing the pixels with local sums of squared convolution values. Eight energy masks are shown in table 1; they perform combinations of smoothing (Level-detection), Edge-detection, and Spot-detection.

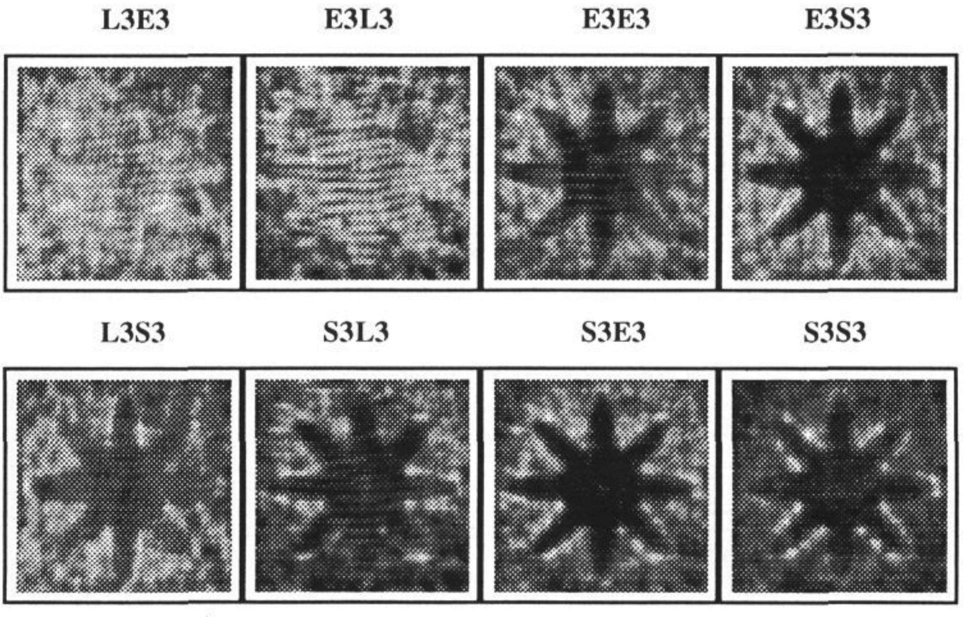

Figure 3 shows the potentials generated by applying energy masks to figure 4a; each energy field has been smoothed using a linear sequential filter, and scaled for display. Note that the L3E3 and E3L3 masks, which together constitute the Sobel edge operator, do not separate the textures because the edge density is approximately uniform throughout the original image.

A statistical snake can be driven using a combination of texture potentials such as those shown in figure 3.† First, a seed region is marked inside the texture of interest; the portions of the energy fields enclosed by the model are then analysed

---

† The performance of texture segmentation algorithms is often assessed by processing test images composed of textures from the photographic album by Brodatz (1966). However, these images are used as much for convenience as for the unofficial standard that they provide. All of the textures in this paper come from this source.

to obtain the mean vector and the associated covariance matrix. A region model can then be grown using the Mahalanobis pressure functional (figure 4).

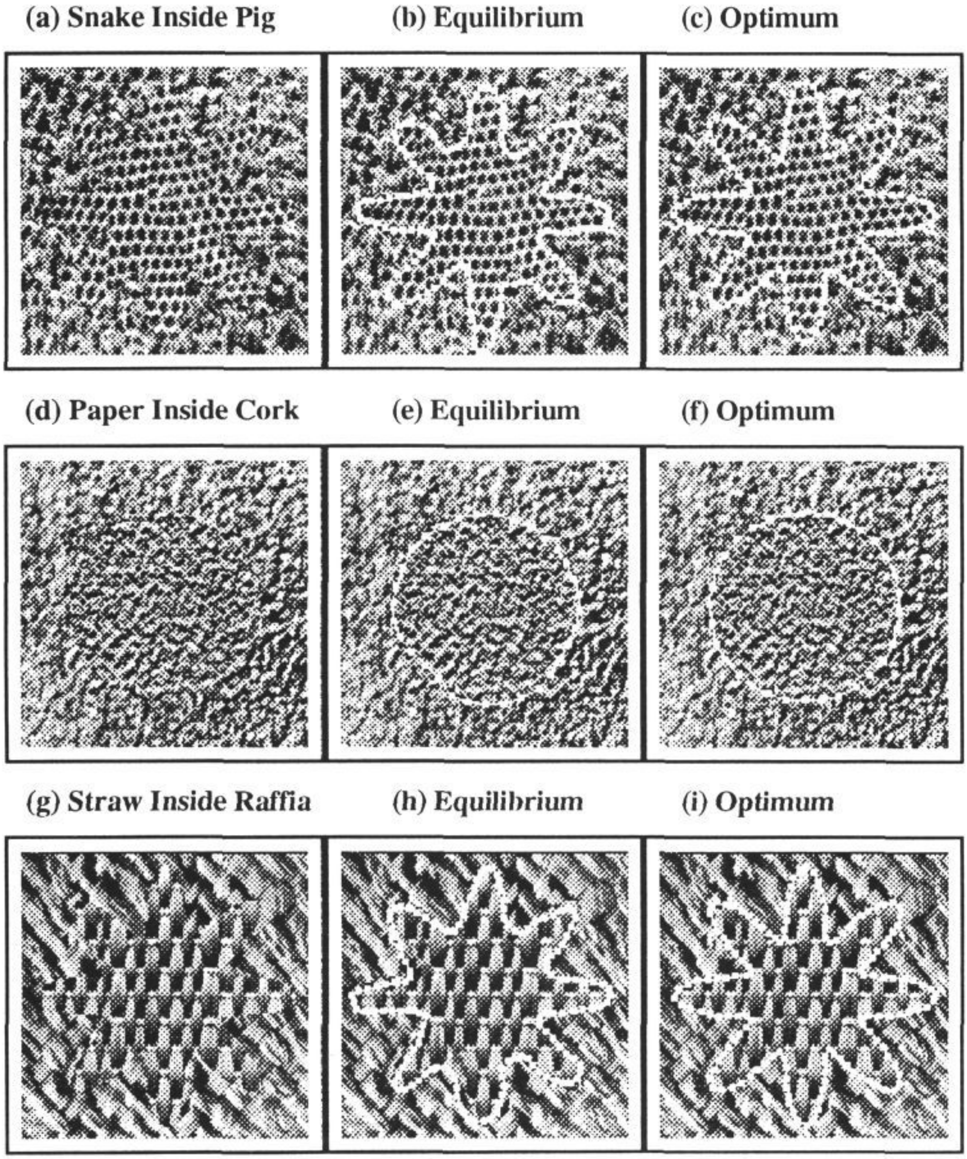

Figure 4 shows three different composite 256-by-256 pixel images. Each texture has been histogram equalised to minimise first-order statistical differences. (**a**) Snake skin surrounded by pig skin. (**b**) A model generated using the Mahalanobis pressure functional to combine the energy fields shown in figure 3. (**c**) An ideal segmentation for comparison with b. (**d-f**) These images show paper surrounded by cork; the sequence follows the same order as the previous three images. (**g-i**) A similar sequence for straw and raffia. In each example the parameters were $\alpha = \beta = \rho = 1$ and $k = 5$ for eight images (degrees of freedom). Each region model was grown from a 64-by-64 pixel seed region at the centre of the image.

L3E3			E3L3			E3E3			E3S3		
−1	0	+1	−1	−2	−1	+1	0	−1	+1	−2	+1
−2	0	+2	0	0	0	0	0	0	0	0	0
−1	0	+1	+1	+2	+1	−1	0	+1	−1	+2	−1

L3S3			S3L3			S3E3			S3S3		
−1	+2	−1	−1	−2	−1	+1	0	−1	+1	−2	+1
−2	+4	−2	+2	+4	+2	−2	0	+2	−2	+4	−2
−1	+2	−1	−1	−2	−1	+1	0	−1	+1	−2	+1

Table 1 shows the eight energy masks used to generate the texture potentials shown in figure 3.

7 Conclusion

A pressurised snake is a semi-global *edge detector* with internal smoothing constraints. In contrast, a statistical snake is a *region grower*, also with internal smoothing, that incorporates three new features: region energy, boundary tracking, and re-parameterisation. Statistical snakes have been implemented in a tool that allows trained users to segment anatomical objects in NMR and CT data cubes. By using region models to extract features such as tumours it is possible to produce three-dimensional images that are accurate enough to use when planning surgery and radiotherapy – see Porrill and Ivins (1994). Current work is exploring the use of active region models to track colour in real-time.

References

Bresenham, J E: 1965. Algorithm For Computer Control Of A Digital Plotter. *IBM Systems Journal: vol 4, no 1, pp 25-30.*

Brodatz, P: 1966. Textures – A Photographic Album For Artists And Designers. *Dover, New York.*

Hechbert, P: 1990. Concave Polygon Scan Conversion. *Glassner, A S (Editor) – Graphics GEMS (London, Academic Press Limited): pp 681-684.*

Ivins, J; Porrill, J: 1993. Statistical Snakes: Active Region Models. *AIVRU Memo #89.*

Kass, M; Witkin, A; Terzopoulos, D: 1987. Snakes: Active Contour Models. *First International Conference on Computer Vision: pp 259-268.*

Laws, K I: 1979. Texture Energy Measures. *DARPA Image Understanding Workshop (Los Angeles): pp 47-51.*

Porrill, J; Ivins, J: 1994. A Semiautomatic Tool For 3-D Medical Image Analysis Using Active Contour Models. *Medical Informatics (In Press).*

Sullivan, G D; Worrall, A D; Hockney, R W; Baker, K D: 1990. Active Contours In Medical Image Processing Using A Networked SIMD Array Processor. *First BMVC: pp 395-400.*

Appendix: Energy Balance

This appendix examines the energy changes that arise from inserting and deleting the boundary elements of a statistical snake. Consider a snake $\mathbf{u}(\lambda)$ – where $\lambda = 0, 1, ..., N{-}1$ – with energy E defined as:

$$E \;=\; \frac{\alpha}{2}\oint\left|\frac{\partial\mathbf{u}}{\partial\lambda}\right|^2 d\lambda \;+\; \frac{\beta}{2}\oint\left|\frac{\partial^2\mathbf{u}}{\partial\lambda^2}\right|^2 d\lambda \;-\; \rho\iint_R G(\,I(x,y)\,)\,dx\,dy$$

An affine re-parameterisation $\mu = k\,\lambda + c$ gives rise to the following relationships:

$$\hat{\mathbf{u}}(\mu) = \mathbf{u}(k\,\lambda + c) \qquad \frac{\partial}{\partial\mu} = \frac{1}{k}\frac{\partial}{\partial\lambda} \qquad d\mu = k\,d\lambda$$

$$N = \oint d\lambda \qquad \hat{N} = \oint d\mu = \oint k\,d\lambda = k\,N \qquad k = \frac{\hat{N}}{N}$$

After re-parameterisation, the new energy of the snake can be written:

$$\hat{E} \;=\; \frac{\hat{\alpha}}{2}\oint\left|\frac{\partial\hat{\mathbf{u}}}{\partial\mu}\right|^2 d\mu \;+\; \frac{\hat{\beta}}{2}\oint\left|\frac{\partial^2\hat{\mathbf{u}}}{\partial\mu^2}\right|^2 d\mu \;-\; \hat{\rho}\iint_R G(\,I(x,y)\,)\,dx\,dy$$

This is equivalent to:

$$\hat{E} \;=\; \frac{\hat{\alpha}}{2}\oint\frac{1}{k}\left|\frac{\partial\mathbf{u}}{\partial\lambda}\right|^2 d\lambda \;+\; \frac{\hat{\beta}}{2}\oint\frac{1}{k^3}\left|\frac{\partial^2\mathbf{u}}{\partial\lambda^2}\right|^2 d\lambda \;-\; \hat{\rho}\iint_R G(\,I(x,y)\,)\,dx\,dy$$

Re-parameterisation will cause the energy balance of the snake to change unless the following relationships are maintained:

$$\frac{\alpha}{2}\oint\left|\frac{\partial\mathbf{u}}{\partial\lambda}\right|^2 d\lambda \;=\; \frac{\hat{\alpha}}{2k}\oint\left|\frac{\partial\mathbf{u}}{\partial\lambda}\right|^2 d\lambda \quad\Rightarrow\quad \alpha = \frac{\hat{\alpha}}{k} = \hat{\alpha}\,\frac{N}{\hat{N}}$$

$$\frac{\beta}{2}\oint\left|\frac{\partial^2\mathbf{u}}{\partial\lambda^2}\right|^2 d\lambda \;=\; \frac{\hat{\beta}}{2k^3}\oint\left|\frac{\partial^2\mathbf{u}}{\partial\lambda^2}\right|^2 d\lambda \quad\Rightarrow\quad \beta = \frac{\hat{\beta}}{k^3} = \hat{\beta}\,\frac{N^3}{\hat{N}^3}$$

$$\rho\iint_R G(\,I(x,y)\,)\,dx\,dy \;=\; \hat{\rho}\iint_R G(\,I(x,y)\,)\,dx\,dy \quad\Rightarrow\quad \rho = \hat{\rho}$$

The overall energy balance of the snake can therefore be made invariant to re-parameterisation by using constants α_0, β_0 and ρ_0 to calculate α, β and ρ according to the number of elements N at each iteration:

$$\alpha = \alpha_0\,N \qquad \beta = \beta_0\,N^3 \qquad \rho = \rho_0$$

These formulae can also deal with local (geometric) re-parameterisations if the model is allowed to stabilise; this is possible because the tension in the boundary tends to space elements uniformly at equilibrium.