Robust Boundary Delineation Using Random-Phase-Shift Active Contours

Astrit Rexhepi and Farzin Mokhtarian

Centre for Vision, Speech, and Signal Processing
School of Electronics and Physical Sciences
University of Surrey
Guildford GU2 7XH
United Kingdom
`a.rexhepi@surrey.ac.uk`,
`f.mokhtarian@surrey.ac.uk`
`http://www.surrey.ac.uk`

Abstract. When an active contour is applied to a noisy image, the contour is sometimes attracted to a local energy minimum, since the noise gives rise to high rates of change of the image gray levels. In this paper we will describe a novel method of overcoming this problem by using a sparse set of points to represent the active contour C and randomly varying the positions of these points.

1 Introduction

Active contours, also known as "snakes" or deformable models, have proven to be an effective method of boundary delineation. Since the original work by Kass, Witkin, and Terzopoulos (KWT) in 1988 [1], extensive research has been carried out on such models [2]. An active contour locates a boundary in an image by minimizing an energy function. This function includes "internal" terms that depend on the length and curvature of the contour; these terms are small when the contour is short and smooth. It also includes "external" terms that depend on the image gray levels at or near the points of the contour. For example, if the inverse rate of change of the image gray level is used as the external term, it will be small when the contour lies close to a strong boundary in the image.

In the KWT active contour model, the energy function is defined by an expression of the form

$$E(C) = \int_C [w_1|v_s|^2 + w_2|v_{ss}|^2 + \xi(v)]ds$$

where C is a curve in the image plane; s is a parameter representing a point on C; $v = (x, y)$ is the position of the point in the plane; v_s and v_{ss} are the first and second derivatives of v with respect to s; w_1 and w_2 are (possibly variable) weights; and $\xi(v)$ is a function of the image values in a neighborhood of v.

In this expression for $E(C)$, the first two terms are called "internal" energy terms because they depend only on the geometry of C itself, while the third

term is called an "external" energy term because it depends on the image gray levels at or near C. For example, suppose the value of ξ at a point (x, y) is the inverse rate of change of the image gray level at (x, y); then ξ has low values on or near image boundaries and high values elsewhere.

It is well known that at a minimum of $E(C)$, the coordinates x, y of the points of C must satisfy the Euler equations

$$e_x \equiv w_1 x_{ss} + w_2 x_{ssss} + (\partial \xi / \partial x) = 0$$

$$e_y \equiv w_1 y_{ss} + w_2 y_{ssss} + (\partial \xi / \partial y) = 0$$

In the KWT model these equations are solved by an iterative process in which C is approximated by a discrete set of points, and at each iteration, the positions of the points are adjusted so as to reduce e_x and e_y.

It has long been realized that this basic active contour model must be modified in order to enable it to detect a distant object boundary, avoid local energy minima due to noise, and conform to the details of a boundary's shape. Since $\xi(x, y)$ depends only on the image gray levels near (x, y), it has an influence on $E(C)$ only near image features. For example, the inverse rate of change of the image gray level is low only near an object boundary; elsewhere, the boundary has no effect on $E(C)$, and C has no tendency to get closer to the boundary. On the other hand, since noise gives rise to high rates of change of the image gray level, the inverse is low at noise points, and $E(C)$ may have a local minimum when C passes through or near noise points. If C does succeed in approaching an object boundary, it may have difficulty conforming to parts of the boundary that have high curvature, since the internal energy of C is high on such parts. Methods of overcoming these difficulties will be discussed in the next sections.

In the remainder of this section we review selected publications on active contours, emphasizing papers that made contributions to the representation of the contour; the definitions of the internal and external energy terms; the methods of initialization and energy minimization; and the methods of handling noise.

To improve the performance of active contours, alternative representations for the contour have been proposed. Kass et al. [1] represented an active contour by a digital curve. Delagnes et al. [3] defined adjustable polygons: sets of active line segments that can approximate any object shape; this representation gives good results if the object to be delineated is noise-free. Wang et al. [4] used a spline representation; this resulted in some improvement in accuracy and convergence speed over the KWT model. Wong et al. [5] proposed a segmented snake model; this converted the problem of global optimization of a closed curve into local optimization of a number of open arcs. Their approach was able to locate convex, concave and high-curvature parts of an object boundary; its performance was similar to that of Wang's spline representation. Chesnaud et al. [6] proposed a region snake model rather than using a boundary-based representation. Their model was based on Maximum Likelihood estimation of the statistics of the inner and outer regions defined by the boundary. This approach works well if we can use a priori assumptions about the statistics of the regions, and if these statistics are invariant or at least easy to classify. Ray et al. [7] proposed

a multiresolution approach in which the snake algorithm is applied at an initial scale, and after the snake stabilizes, a higher resolution is used to adjust it. A potential disadvantage of this model is that high-curvature parts of the boundary and thick curves may be eliminated at the coarser resolution. Velasco and Marroquin [8] proposed Sandwich Snakes, which can detect contours that have complex shapes and reject false minima (up to some level) due to noise. Their model consists of two snakes, one inside and the other outside the boundary; it requires a one-to-one correspondence between the two snakes.

The performance of active contour models can be improved by modifying the definition of their internal energy. Cohen [9] proposed the use of a pressure force which inflates the active contour in the normal direction until it conforms to the boundary of the object. This model gives good results in noise-free images, but improper selection of the pressure force yields poor results. Davatzikos and Prince [10] proposed spatio-temporal variation of the internal energy terms as functions of position in the image and the number of iterations. This allows the model to handle high-curvature parts of the boundary more effectively than fixed-parameter algorithms. Xu *et al.* [11] proposed a method of compensating for the normal internal force so as to make it independent of shape. The resulting model works well, with no need to fine-tune internal parameters, and can conform to high-curvature parts of a boundary such as corners; however, its ability to overcome noise is reduced. Wang *et al.* [4] divided the energy minimization process into multiple stages. The first stage was designed to optimize the convergence speed in order to allow the snake to quickly approach the minimum-energy state. The second stage was devoted to snake refinement and local minimization of the energy, thereby driving the snake to a quasi-minimum-energy state. Finally, the third stage used the Bellman optimality principle to fine-tune the snake to a global minimum-energy state. Metaxas and Kakadiaris [12] presented a technique for the automatic adaptation of a deformable model's elastic parameters in a Kalman filter framework. The parameters are initialized and are subsequently modified, depending on the data and the noise distribution, until the contour conforms to the boundary; this works well if the spatial variations of the data are smooth. Mokhtarian and Mohanna [20] proposed an active contour model in which the smoothness internal evergy term is replaced by the output of a Curvature Scale Space filtering process.

Other authors have modified the definition of the external energy to increase the capture range of a snake and thus make the snake robust to noise. Kass *et al.* [1] used a scale space approach to guide a snake toward the boundary of an object. Xu and Prince [13] proposed a new external energy term called the "Gradient Vector Flow Field" computed by diffusion of the gradient vectors of a gray-level or binary edge map derived from the image. This force field is insensitive to initialization of the snake and allows the snake to move into concave boundary regions in noise-free images. However, using the diffusion of gradient vectors to develop this field may increase the effect of noise. Peterfreund [15] used spatio-velocity space (a combination of optical flow and image forces) to track boundaries of nonrigid objects on cluttered backgrounds.

An active contour model can be semi- or fully automatic, depending on how it is initialized. Kass *et al.* [1] initialized the snake near a boundary. Cohen [9] initialized the snake inside or outside an object and used a pressure force to push the contour outward or inward until it reached the boundary. Neuenschwander *et al.* [16] presented a model in which the user has to specify only the two endpoints of the contour rather than a polygonal approximation. The snake converges from this minimal initialization by propagating image information along the contour from both endpoints. The Gradient Vector Flow Field used by Xu and Prince [13] made the snake insensitive to initialization.

Alternative algorithms for minimizing the energy of an active contour have also been used. Kass *et al.* [1] minimized the energy by solving the Euler equations. Amini *et al.* [17] used dynamic programming to optimize an active contour; their approach was more stable than the original KWT approach, but it was time-consuming. Williams and Shah [18] proposed a greedy algorithm, which gave results comparable to those of Amini's method but was much faster. A common disadvantage of both methods was that they are local and hence are relatively sensitive to noise. Caselles *et al.* [19] proposed Geodesic Active Contours, which combined a geometric contour model with energy function minimization. The performance of this approach is comparable to that of conventional active contour models up to a constant that depends on the initial parameterization of the contour. The Geodesic Active Contour model combined with level set methods can be used to delineate the boundaries of multiple objects. This model has advantages over the original active contour model, but it has the drawback of being nonlinear.

The convergence speed and accuracy of active contours depends greatly on the level of noise in the image. Filtering techniques can be used to reduce the noise to some degree, but it is almost impossible to eliminate it completely. As a result, a snake may get stuck at energy minima caused by noise before it reaches a boundary. To avoid this situation, Davatzikos and Prince [10] proposed an algorithm in which the internal energy varies spatially. By giving high weight to the internal energy in noisy parts of the image they were able to overcome local minima. His model worked well when the object to be delineated had smooth boundaries. Delagnes *et al.* [3] proposed a new energy function based on textural characteristics of objects to resolve conflict situations when tracking objects on a cluttered background. Their method worked well when the textures were easily distinguishable. Other active contour models that were designed to overcome noise include those proposed by Metaxas and Kakadiaris [12], Chesnaud *et al.* [6], and Velasco and Marroquin [8]; these models were described above.

The organization of this paper is as follows: In Section 2 we show how local minima in the energy of an active contour, due to noise in the image, can be avoided by perturbing the contour representation during the energy minimization process. Section 3 and Section 4 describes our methods. Section 5 describes an experiment in which an active contour is used to delineate the boundary of a moving hand in an image sequence. Section 6 summarizes our methods.

2   Overcoming Local Energy Minima Due to Noise

When an active contour is applied to a noisy image, the contour sometimes is attracted to a local energy minimum, since the noise gives rise to high rates of change of the image gray levels. In this section we will describe a method of overcoming this problem by using a sparse set of points to represent the active contour C and randomly varying the positions of these points.

The number of points chosen to represent C must be a fraction of the total number of pixels on the (digital) contour, so there will be room to vary the position of the points. (We can roughly estimate the number of points on the contour by examining the output of the boundary extraction process [21]) On the other hand, the fraction cannot be too small, since it must be possible to closely approximate the shape of the contour by interpolating between the points. In the experiments described in this section, C was several hundred pixels long, and we represented it by about 60 points. Note that when we use scale-space methods to detect object boundaries at a distance, we are reducing (and afterwards increasing) the number of pixels on the contour, and the number of points used to represent the contour must be reduced or increased correspondingly.

In an active contour algorithm that uses a sparse representation, the contour C is represented by (say) n points. At each iteration of the algorithm, the points shift slightly; the new n points represent a new contour C'. This process is repeated at each iteration.

We have investigated a method of incorporating random variation into the points that represent the contour. In our method the number n of points remains constant. At each iteration of the active contour algorithm, we interpolate a smooth digital curve C' through the points. We choose n equally spaced points of C', one of which coincides with one of the original points on C'. We then randomly shift the new points along C' by an amount less than the spacing between the points. The points all shift together; their spacing remains the same. We refer to this method as "phase perturbation".

If many of the points that represent the contour coincide with noise points in the image, the external energy of the contour will be low, since external energy is inversely related to the gradient of the image gray level, which is high at noise points. Thus a contour configuration in which many of the points coincide with noise points may give rise to a local minimum in the energy, and the active contour algorithm may not be able to leave this minimum. However, if we perturb the points that represent the contour using phase perturbation, the perturbed points will no longer coincide with the noise points, so the contour has a chance of escaping the local energy minimum.

When phase perturbation is used, the number n of points that represent the contour remains constant, so the expected number of coincidences between these points and noise points is also constant. Thus the random displacement of the n points can be chosen from a uniform distribution over an interval.

We will now verify experimentally that using a uniform distribution is preferable in phase perturbation.

Figure 1(a) shows a 100×100 image that contains a blurred hollow square on a noisy background; the outer boundary of the square is represented by a solid curve S, and an initial active contour C surrounding S is overlaid on the image. We chose n equally spaced points to initially represent C, and performed 500 iterations of a sparse version of the KWT algorithm. (Randomness can be introduced into any active contour algorithm, but in our experiments we used the KWT algorithm.) After each iteration we also had the option of performing a phase perturbation by shifting the n points by an integer number of steps along the contour, where the number never exceeded the spacing s between the points. We performed six versions of this experiment:

a) No shift.

b-e) A shift chosen randomly in the range [0,2], [0,3], [0,4], or [0,s].

f) The shift that resulted in the highest external energy of C' (to maximize the likelihood of C escaping from the local energy minimum).

where (a to f) stands for (top-down). Figure 1(b) shows plots of the area (in pixel units) contained between C and S as a function of iteration number, averaged over 500 instances of the noise. We see that in all cases, the area at first decreases rapidly from its initial value of about 5000 as C shrinks toward S, but it then levels off. When no shifts were used (version (a)) the area levels off at about 2700. When random shifts were used (versions (b-e)), the area continues to drop; the larger the range of the shifts, the greater the drop, because there are more possibilities for increasing the energy. For shifts in the range [0,s] the area drops
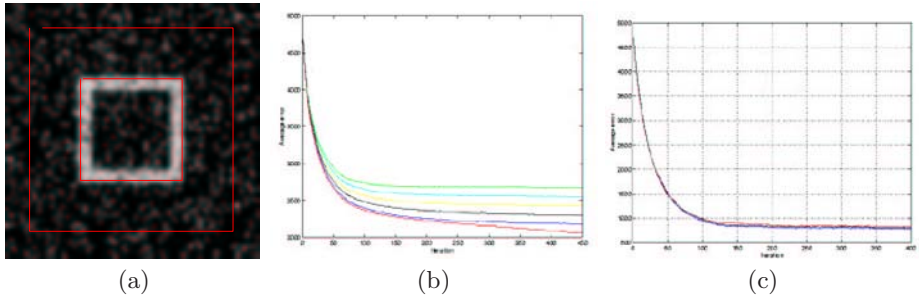

(a)   (b)   (c)

Fig. 1. (a)- A square in a noisy image. The initial active contour and the boundary of the square are overlaid on the image. (b)- Phase perturbation. Each curve shows the area between the active contour and the boundary of the square as a function of iteration number. The curves are averages over 500 instances of the image noise. Top to bottom: (a) No perturbation. (b-e) Shifts chosen randomly in the ranges [0,2], [0,3], [0,4], and [0,s], where s is the spacing between the points that represent the active contour. (f) The shifts that resulted in the highest external energy of the active contour. (c)- Comparison of two algorithms on a non-noisy version of Figure 1(a): a sparse version of the KWT algorithm, and algorithm that incorporated phase perturbation.

to less than 2200; this is nearly as good as when we use the shift that results in the highest external energy of C (version (f)).

Our algorithms represent a contour by a discrete set of points, and interpolate a smooth digital curve on these points each time a perturbation is applied. It might be thought that this repeated interpolation would result in a contour evolution process, but in fact this did not happen. To demonstrate this, we applied two algorithms to a version of Figure 1(a) that contained no noise: a sparse version of the KWT algorithm, and versions that incorporated phase perturbation. As we see in Figure 1(C), the plots of the area between C and S are virtually indistinguishable for both algorithms. This demonstrates that our use of repeated interpolation did not result in contour evolution.

In the experiments described in this section, the active contour C never penetrated the boundary of the square S; we could therefore use the area contained between C and S as an error measure. In the real examples described in the next section it is possible for C to penetrate the object boundary B. We will therefore use a more general error measure: the area of the symmetric difference between the regions surrounded by C and B.

3 Delineating an Object Boundary

Active contour performance can be improved by dividing the energy minimization process into stages [4] and allowing the energy function to vary during the process [12,13]. In Section 4 we will describe how such an adaptive active contour algorithm can be used to detect an object boundary at a distance and then locate details of the boundary's shape.

An active contour can be used to track the boundary of a moving object in an image sequence. This is usually done by locating the boundary (by minimizing the energy of the contour) in each frame of the sequence, and then using the result to initialize the contour in the next frame. In Section 5 we will use an active contour to locate the boundary of a moving object in an image sequence, using the moving boundary extraction process described in [21]. We will describe an experiment in which an active contour is used to locate the boundary of a hand moving against a complex background.

4 Detecting and Conforming to the Boundary

Since the external energy ξ of C depends only on the image values in the vicinity of C, distant object boundaries have no effect on ξ. Thus if C is initialized far from an image boundary, minimization of $E(C)$ does not attract C toward the boundary. This problem can be overcome by blurring the image before initializing C; but blurring the image may destroy details of the shapes of object boundaries. To achieve both detection at a distance and accurate location of shape details, we can vary the amount of image blur during the minimization process [1,12,13]. The blur remains high until ξ becomes low, indicating that C is approaching a boundary; the blur can then be gradually reduced so that C

can accurately conform to the boundary shape. The spacing of the points that we use to represent the contour does not exceed the amount of the blur.

The need to conform to boundary parts that have high curvatures introduces another problem: the internal energy of C is high when its curvature is high. This problem can be overcome [13] by gradually reducing the weight given to the curvature term of $E(C)$ as C approaches the boundary.

5   An Application: Delineating the Boundary of a Moving Hand in an Image Sequence

In this section we use an active contour to delineate the boundary of a moving object; the boundary is initially extracted by the method described in [21].

Figure 2(a) shows part of a frame of a 20-frame video sequence of a hand and arm moving in an indoor scene. A number of final boundaries located by the active contour after energy minimization is overlaid on the image. Figure 2(b) shows the moving boundary points extracted from that frame by the method described in [21]. Most of these points are concentrated near the hand and arm boundaries.

C was initialized on a large square which was close to the image border. Gaussian blur was initially applied to the moving boundaries that were extracted
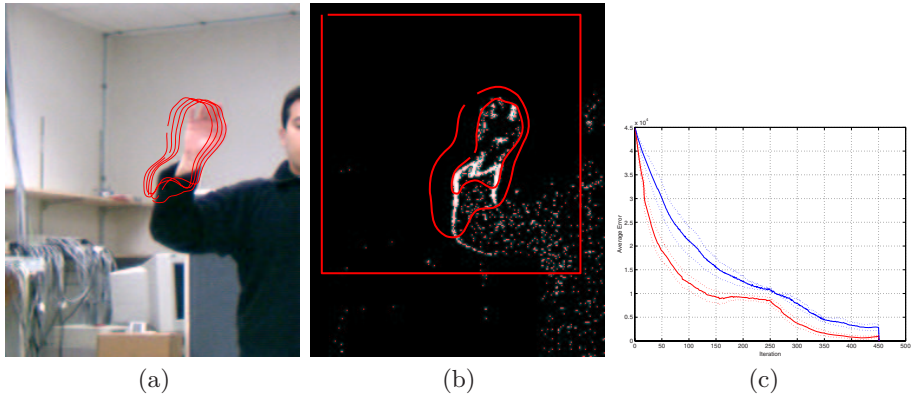

(a)   (b)   (c)

Fig. 2. Delineation of a moving hand in a video. (a) One frame of the video, with the final contours overlaid. (b) Initial contour, final Stage 1 contour, and final Stage 2 contour overlaid on the output of the moving boundary points that were extracted from the image. (c) Area (in pixels) of the symmetric difference between the hand and the interior of the contour, as a function of iteration number, averaged over 20 frames. The solid curve is the mean; the dotted curves are one standard deviation above and below the mean. Upper curves: Algorithm without perturbation. Lower curves: Algorithm with perturbation. Note the shoulder in the curves at iteration 250, when the first stage ended.

from that image (Figure 2(b)). The rate of change of the gray level in this blurred image was used as the external energy term of C. This rate of change is a maximum at the inflections of the Gaussian; hence its inverse is a minimum. Minimization of $E(C)$ thus causes C to shrink and to approach the hand and arm boundary until it reaches the inflections. In Figure 2(b) the location of C after 250 iterations overlaid on the image.

By this time the external energy of C was quite low. A second stage of energy minimization was then initiated, in which the amount of image blur and the weight given to curvature in the internal energy of C were both progressively decreased, as described in Section 4. This allowed C to approach the boundary closely and to conform to its shape. The location of C after 200 iterations of the second-stage process is also overlaid on the Figure 2(b).

To reach its final location, C must cross noisy parts of the image background. As discussed in Section 2, it is possible for C to be trapped by a local energy minimum caused by the noise, but this can be avoided by applying phase perturbation to the points that define C.

To study how this perturbation improves performance, we applied two versions of our active contour algorithm to the 20 frames of our image sequence; the first version perturbed the points that represent the contour and the second version did not. Figure 2(c) compares the average performance of the two versions on the 20 frames; the upper curves are for the second version and the lower curves for the first version. In each frame, we computed the area of the symmetric difference between the hand/arm region and the interior of the active contour, as a function of iteration number. (The solid curve is the 20-frame average; the dotted curves are one standard deviation above and below the average.) We see that the first version of the algorithm converged more quickly and approximated the boundary more accurately. The lower curve comes close to a minimum after 150 iterations in the first stage of the process, and after 450 iterations it is less than 20% as high as the upper curve.

6   Concluding Remarks

This paper has made the following contributions: First; an active contour can be trapped by local energy minima when too many of its points are influenced by image noise. We have shown that this situation can be prevented by incorporating randomness to the points that represents the contour. Second; to speed up the process of convergence and to conform to boundary shape we developed a two stage algorithm. Our model is a mixture of deterministic and random components that made it robust with respect to noise. We have shown that our model performs equally well as other sparse models when there is no noise, the robustness becomes visible only when the image contains noise. Using this active contour model, we were able to locate and track a moving boundary in a sequence of images in the presence of noise.

References

1. Kass, M., Witkin, A.P., Terzopoulos, D.: Snakes: Active Contour Models. International Journal of Computer Vision 1, 321–331 (1988)
2. Blake, A., Isard, M.: Active Contours. Springer, Heidelberg (1998)
3. Delanges, P., Benois, J., Barba, D.: Active Contours Approach to Object Tracking in Image Sequences with Complex Background. Pattern Recognition Letters 16, 171–178 (1995)
4. Wang, M., Evans, J., Hassebrook, L., Knapp, C., Multistage, A.: Optimal Active Contour Model. IEEE Trans. on Image Processing 5, 1586–1591 (1996)
5. Wong, Y.Y., Yuen, P.C., Tong, C.S.: Segmented Snake for Contour Detection. Pattern Recognition 31, 1669–1679 (1998)
6. Chesnaud, C., Refregier, P., Boulet, V.: Statistical Region Snake-Based Segmentation Adapted to Different Physical Noise Models. IEEE Trans. on Pattern Analysis and Machine Intelligence 21, 1145–1157 (1999)
7. Ray, N., Chanda, B., Das, J.: A Fast and Flexible Multiresolution Snake with a Definite Termination Criterion. Pattern Recognition 34, 1483–1490 (2001)
8. Velasco, F.A., Marroquin, J.L.: Robust Parametric Active Contours: The Sandwich Snakes. Machine Vision and Applications 12, 238–242 (2001)
9. Cohen, L.D.: On Active Contour Models and Balloons. Computer Vision, Graphics, and Image Processing 53, 211–218 (1991)
10. Davatzikos, C., Prince, J.L.: Adaptive Active Contour Algorithms for Extracting and Mapping Thick Curves. Computer Vision and Pattern Recognition, pp. 524-529 (1993)
11. Xu, G., Segawa, E., Tsuji, S.: Robust Active Contours with Insensitive Parameters. Pattern Recognition 27, 879–884 (1994)
12. Metaxas, D., Kakadiaris, I.A.: Elastically Adaptive Deformable Models, European Conference on Computer Vision, pp. 550–559 (1996)
13. Xu, C., Prince, J.L.: Gradient Vector Flow: A New External Force for Snakes. In: Conference on Computer Vision and Pattern Recognition, pp. 66–71 (1997)
14. Davatzikos, C., Prince, J.L.: Convexity Analysis of Active Contour Problems. Image and Vision Computing 17, 27–36 (1999)
15. Peterfreund, N.: The Velocity Snake: Deformable Contour for Tracking in Spatio-Velocity Space. Computer Vision and Image Understanding 73, 346–356 (1999)
16. Neuenschwander, W., Fua, P., Szekely, G., Kubler, O.: Making Snakes Converge from Minimal Initialization. International Conference on Pattern Recognition, pp. 613–615 (1994)
17. Amini, A., Tehrani, S., Weymouth, T.E.: Using Dynamic Programming for Minimizing the Energy of Active Contours in the Presence of Hard Constraints. In: International Conference on Computer Vision, pp. 95–99 (1988)
18. Williams, D.J., Shah, M.: A Fast Algorithm for Active Contours. In: International Conference on Computer Vision, pp. 592–598 (1990)
19. Caselles, V., Kimmel, R., Sapiro, G.: Geodesic Active Contours. International Journal on Computer Vision 22, 61–79 (1997)
20. Mokhtarian, F., Mohanna, F.: Fast Active Contour Convergence through Curvature Scale Space Filtering. Image and Vision Computing, pp. 157–162 (2001)
21. Rexhepi, A., Rosenfeld, A., Mokhtarian, F.: Extracting Boundaries from Images by Comparing Cooccurrence Matrices. Digital Image Computing Techniques and Applications DICTA2003 (2003)