Globally Optimal 3D Image Reconstruction and Segmentation via Energy Minimisation Techniques

Brian C. Lovell

Intelligent Real-Time Imaging and Sensing Group, EMI
The School of Information Technology and Electrical Engineering
The University of Queensland
`lovell@itee.uq.edu.au`,
WWW home page: `http://www.itee.uq.edu.au/~lovell`

Abstract. This paper provides an overview of a number of techniques developed within our group to perform 3D reconstruction and image segmentation based of the application of energy minimisation concepts. We begin with classical snake techniques and show how similar energy minimisation concepts can be extended to derive globally optimal segmentation methods. Then we discuss more recent work based on geodesic active contours that can lead to globally optimal segmentations and reconstructions in 2D. Finally we extend the work to 3D by introducing continuous flow globally minimal surfaces. Several applications are discussed to show the wide applicability and suitability of these techniques to several difficult image analysis problems.

1 Introduction and Time-Line

In 1992 we began a research project to automatically segment cell images from Pap Smear slides for the detection of cancer of the cervix. We investigated simple techniques based on edge detection, grayscale thresholding, and grayscale morphology (*e.g.,* watersheds), but could only achieve accurate segmentation on about 60% of cell images. In 1997 we started looking at dual snake energy minimisation techniques as proposed by Gunn [1], but this method suffered from poor robustness. However, Gunn did suggest a fast globally optimal method based on converting the circular contour finding problem into a linear trellis and then applying the Viterbi shortest path finding algorithm. This approach worked extremely well as reported by Bamford and Lovell [2] and yielded 99.5% correct segmentation on a cell database of nearly 20,000 cell images.

As this method was so remarkably effective on cell images there was little incentive to improve the method for the Pap Smear problem itself, but we still held a desire to develop energy minimisation techniques which were more general. In 2002, Appleton and Sun [3] put the problem of representing closed contours on a linear trellis on to a firm mathematical basis. Then, in 2003, Appleton and Talbot [4, 5] extended and generalized the energy minimisation approach to handle

the optimal segmentation of planar concave objects as well as convex images such as cells. This extension avoided dependance on a coarse discretization grid so that grid-bias could be removed. The extension to 3D was achieved in late 2003 by Appleton and Talbot [6] by converting the shortest path techniques into an equivalent continuous maximum flow/minimal surface problem.

In this paper we outline the various energy minimisation segmentation techniques and show how they can be applied to solve quite difficult segmentation and reconstruction problems from diverse domains such as volumetric medical imaging to multiview reconstruction.

2 Cell Image Segmentation using the Viterbi Shortest Path Method

Although the use of active contours [7] is well established, it is well known that these methods tend to suffer from local minima, initialisation, and stopping criteria problems. Fortunately global minimum energy, or equivalently shortest-path, searching methods have been found which are particularly effective in avoiding such local minima problems due to the presence of the many artefacts often associated with medical images [8–10].

An energy minimization method employed was based on a suggestion in [1]. A circular search space is first defined within the image, bounded by two concentric circles centralised upon the approximate centre of the nucleus found by an initial rough segmentation technique (*e.g.,* converging squares algorithm). This search space is sampled to form a circular trellis by discretising both the circles and a grid of evenly-spaced radial lines joining them (figure 1). This circular trellis is then unwrapped in a polar to rectangular transformation yielding a conventional linear trellis.
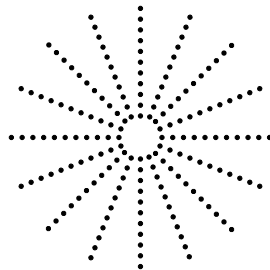

Fig. 1. Discrete search space

Every possible contour that lies upon the nodes of the search space is then evaluated and an associated energy or cost function is calculated. This cost is a function of both the contour's smoothness and how closely it follows image

edges. The relative weighting of the cost components is controlled by a single regularization parameter, $\lambda \in [0, 1]$. By choosing a high value of λ, the smoothness term dominates, which may lead to contours that tend to ignore important image edges. On the other hand, low values of λ allow contours to develop sharp corners as they attempt to follow all high gradient edges, even those which may not necessarily be on the desired object's edge. Once every contour has been evaluated, the single contour with least cost is chosen as the global solution. The well-known Viterbi algorithm provides an efficient method to find this global solution as described in [2].

A data set of 19946 Pap stained cervical cell images was available for testing. The single parameter λ was empirically chosen to be 0.7 after trial runs on a small sub-set of the images. The effect of the choice of λ on segmentation accuracy on this trial set is shown by the graph of figure 2. Figure 3 shows a value of $\lambda = 0.7$ as being the most suitable for these particular images. Every image in the data set was then segmented at $\lambda = 0.7$ and the results verified by eye. Of the 19946 images, 99.47% were found to be correctly segmented.
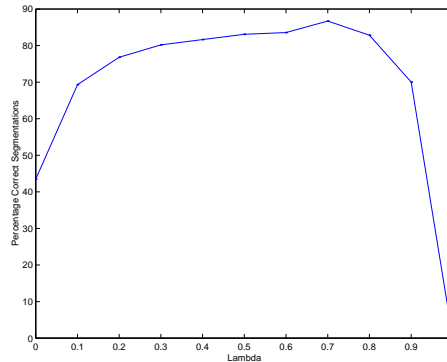
Fig. 2. Plot of percentage of correct segmentations against λ for a set of images consisting of known 'difficult' images and randomly selected images.

With λ set at 0.0, the smoothness constraint is completely ignored and the point of greatest gradient is chosen along each search space radius. Previous studies [11] have shown that for approximately 65% of images, all points of greatest gradient actually lie upon the nucleus cytoplasm border, so these "easy" cell images will be correctly segmented. For the remaining 35% of images, a large gradient due to an artefact or darkly stained chromatin will draw the contour away from the desired border. As λ increases, the large curvatures present in these configurations become less probable (figure 3).

Comments: We show in [12] that the above segmentation method can be viewed as the application of hidden Markov model techniques where the transition ma-
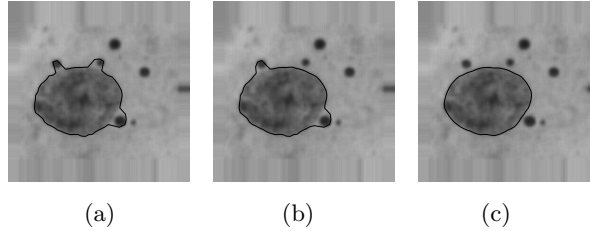
Fig. 3. The effect of increasing λ. a) $\lambda = 0.1$, b) $\lambda = 0.2$, c) $\lambda = 0.5$.

trix is determined by the curvature constraints and the observation matrix is determined by the gradient image. Conceptually the Viterbi algorithm progresses like a planar wavefront through the linear trellis in a manner not unlike an electromagnetic wave passing through a waveguide. Later when we look at the fast marching [13, 14] and Dijkstra's[15] shortest path algorithm, we see that these two algorithms are more akin to spherical electromagnetic wavefronts propagating in free space.

A rather unsatisfactory step in the above development is the formation of the linear trellis from the circular domain and the actual determination of the shortest path. A simple way to find shortest paths on the linear trellis that correspond to closed contours in the image domain is to replicate the M nodes where we unwrap the circular domain such that the last column of nodes in the trellis ia a copy of the first column. Then if there are M such nodes in a column, we would need to evaluate each of the M paths starting and finishing on the same node $i \in [0 \ldots M - 1]$. This would require M evaluations of the Viterbi algorithm.

Gunn [1] suggests a heuristic whereby we choose node i where the image gradient is maximal (*i.e.,* likely to be on the cell boundary), then we find the shortest path beginning and ending at i. Even if the first node i is not actually on the cell boundary, halfway around the cell image the nodes on the shortest path are very likely to be on the boundary. So we rotate the wrapping point 180° and rerun the Viterbi algorithm starting from this node. In our case we selected $M=30$, so we can replace 30 Viterbi evaluations with just 2. Although this heuristic works very well in practice, in theory there are clearly situations where it could fail.

Appleton and Sun [3] investigated this general problem of circular shortest paths. Their method is guaranteed to find the shortest circular path and uses a branch and bound technique to quickly locate it. The root of the branch and bound search tree consists of the entire first column of nodes. The shortest path to the other end of the trellis is found and this forms a lower bound on the circular shortest path. The root node is then split in two and the shortest path algorithm run again keeping only paths that are circular. By applying this process recursively the node on the circular shortest path is quickly identified.

However despite this improvement a major shortcoming of all methods based on a polar to rectangular mapping is the inability to handle concave contours, thus severely limiting their application domain.

3   Globally Optimal Geodesic Active Contours

The classic active contour or snake model proposed by Kass [7] modelled a segmentation boundary by a series of point masses connected by springs. This explicit view of curves as a polygon was replaced by an implicit view of curves as the level set of some 3D surface by Osher and Sethian [16]. Level sets offer significant advantages over traditional snakes including improved stability and much better handling of topology (*e.g.,* segmentation of multiple objects with just one contour). Another advance came in the form of geodesic active contours as proposed by Caselles *et al* [17]. They demonstrated the equivalence of their energy function to the length of a geodesic (*i.e.,* path of least cost, path of least time) in an isotropic space. A problem with traditional geodesic active contours is that they are a gradient descent method and thus have all the usual problems of initialisation, termination, and local maxima associated with such methods. They just do not have the stability and simplicity of application of globally optimal segmentation methods.

The globally optimal method we outline here finds closed contours in the image domain itself rather than unwrapping the image through polar to rectangular transformation. Working in the image domain means that we cannot find simple shortest paths as that would cause a bias towards small contours. Instead we use a contour energy of the form [18]

$$E[C] = \oint_C \frac{g}{r}\,ds \tag{1}$$

where g is a measure of probability of being on the boundary (*e.g.,* gradient) and r is the radius of the contour C. Thus all circles centred on the origin would have the same contour energy. Another innovation is to find the minimal contour as an open path in an augmented helicoidal representation that allows us to represent concave contours as illustrated in figure 4.

We use the Fast Marching Algorithm [13] to find the surface of minimal action, also known as the distance function, whose gradient curves form minimal geodesics. This algorithm finds the surface of minimal action by considering it as the first time-of-arrival of a wavefront emanating from the starting point and travelling with speed $\frac{1}{g}$, where g is usually the image gradient as before. The algorithm is identical to Dijkstra's shortest distance algorithm [15] apart from the need to update g. Figures 5 and 6 show some segmentation results from the globally optimal geodesic active contours method.

4   Globally Minimal Surfaces

The planar segmentation technique outlined in the last section cannot be extended to higher dimensions so we need an entirely new approach. Minimum
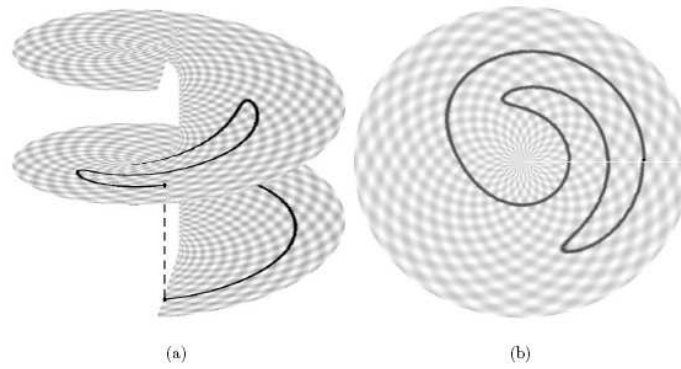
**Fig. 4.** The helicoidal representation of the cut-concave shape from [18]. (a) open curve (b) closed curve.
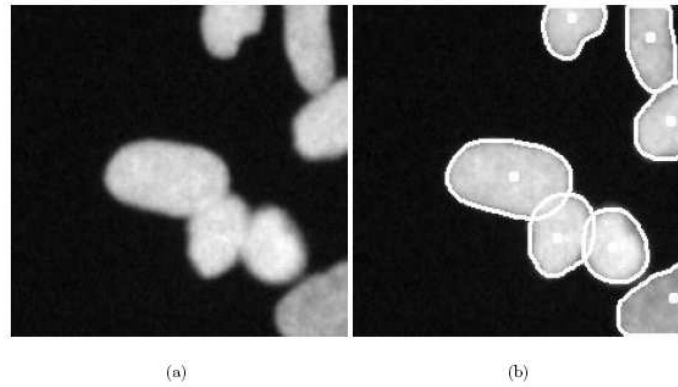
Fig. 5. Globally optimal geodesic active contours applied to overlapping objects from [18]. The cells (a) are separated despite the weak intesity gradient between them (b).

cuts and Maximum flow techniques are naturally suited to globally optimal segmentation in higher dimensions. Although this has been done in the past with discrete approximations, Appleton and Talbot [6] have developed a method based on continuous maximal flows by solving a system of partial differential equations. It is shown [18] that this method gives identical results to the previous globally optimal geodesic active contour method in the case of planar images. Figure 7 shows the segmentation of volumetric MRI data and figure 8 shows the 3D reconstruction from multiview images using minimal surfaces as a postprocessor.
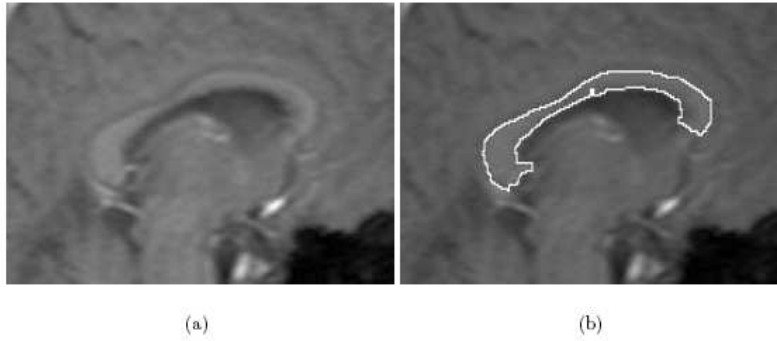
Fig. 6. Segmentation of MRI image of a concave contour, the corpus callosum in a human brain, from [18]. (a) is the original Image and (b) is the segmentation via globally optimal geodesic active contours.
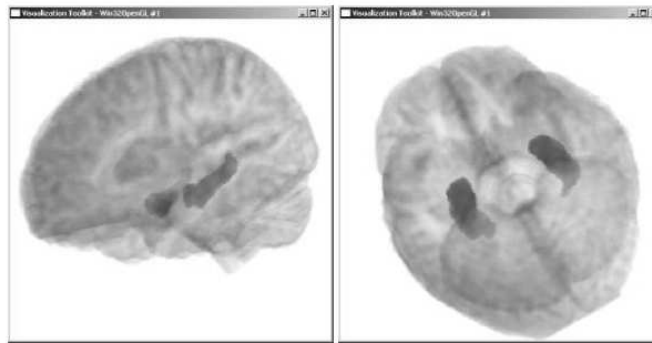

Fig. 7. Segmentation of the hippocampi from an MRI dataset from [18].

5   Conclusions

These globally optimal energy minimisation methods are fast, very easy to apply, and tend to yield robust solutions. Future work is focussed on integrating these techniques with statistical shape models to develop an 3D EM-like algorithm incorporating prior knowledge for detection and segmentation of known shapes.

6   Acknowledgements
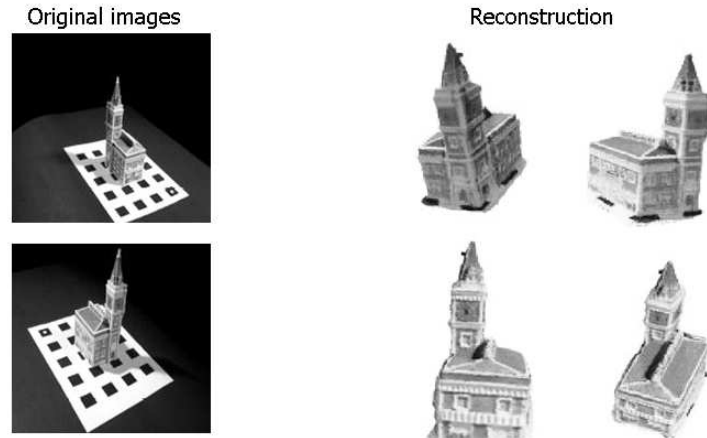
Fig. 8. Multiview reconstruction using globally minimal surfaces for post-processing based on [19].

current and former PhD students Ben Appleton, Carlos Leung, David McKinnon, Pascal Bamford, and Christian Walder.

References

1. Gunn, S.R.: Dual Active Contour Models for Image Feature Extraction. University of Southampton (1996) PhD Thesis.
2. Bamford, P., Lovell, B.: Unsupervised cell nucleus segmentation with active contours. Signal Processing Special Issue: Deformable Models and Techniques for Image and Signal Processing **71** (1998) 203–213
3. Appleton, B., Sun, C.: Circular shortest paths by branch and bound. Pattern Recognition **36** (2003) 2513–2520
4. Appleton, B.: Optimal geodesic active contours: application to heart segmentation. In Lovell, B.C., Maeder, A.J., eds.: APRS Workshop on Digital Image Computing. Volume 1., Brisbane, APRS (2003) 27–32
5. Appleton, B., Talbot, H.: Globally optimal geodesic active contours. Journal of Mathematical Imaging and Vision (2005)
6. Appleton, B., Talbot, H.: Globally optimal surfaces by continuous maximal flows. In Sun, C., Talbot, H., Ourselin, S., Adriaansen, T., eds.: Digital Image Computing: Techniques and Applications. Volume 2., Sydney, CSIRO Publishing (2003) 987–996
7. Kass, M., Witten, A., Terzopoulos, D.: Snakes: Active contour models. International Journal of Computer Vision **1** (1987) 321–331
8. Cohen, L.D., Cohen, I.: Finite-element methods for active contour models and balloons for 2-D and 3-D images. IEEE Transactions on Pattern Analysis and Machine Intelligence **15** (1993) 1131–1147
9. Davatzikos, C.A., Prince, J.L.: An active contour model for mapping the cortex. IEEE Transactions on Medical Imaging **14** (1995) 65–80

10. Geiger, D., Gupta, A., Costa, L., Vlontzos, J.: Dynamic programming for detecting, tracking, and matching deformable contours. IEEE Transactions on Pattern Analysis and Machine Intelligence **17** (1995) 294–302
11. Bamford, P., Lovell, B.: Improving the robustness of cell nucleus segmentation. In Lewis, P.H., Nixon, M.S., eds.: Proceedings of the Ninth British Machine Vision Conference, BMVC '98, University of Southampton (1998) 518–524
12. Lovell, B.C.: Hidden markov models for spatio-temporal pattern recognition and image segmentation. In Mukherjee, D.P., Pal, S., eds.: International Conference on Advances in Pattern Recognition. Volume 1., Calcutta (2003) 60–65
13. Sethian, J.A.: A fast marching level set method for monotonically advancing fronts. Proceedings of the National Academy of Sciences **93** (1996) 1591–1595
14. Sethian, J.A.: Level Set Methods and Fast Marching Methods — Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science. Cambridge University Press (1999)
15. Dijkstra, E.: A note on two problems in connexion with graphs. Numerische Mathematik **1** (1959) 269–271
16. Osher, S., Sethian, J.A.: Fronts propagating with curvature dependent speeed: Algorithms based on hamilton-jacobi formulations. Journal of Computational Physics **79** (1988) 12–49
17. Caselles, V., Kimmel, R., Sapiro, G.: Geodesic active contours. International Journal of Computer Vision **22** (1997) 61–79
18. Appleton, B.C.: Globally Minimal Contours and Surfaces for Image Segmentation. The University of Queensland (2004)
19. Leung, C., Appleton, B., Lovell, B.C., Sun, C.: An energy minimisation approach to stereo-temporal dense reconstruction. In: International Conference on Pattern Recognition. Volume 1., Cambridge (2004) 72–75