

Globally Optimal Surfaces by Continuous Maximal Flows

Ben Appleton¹ and Hugues Talbot²

¹ Intelligent Real-Time Imaging and Sensing Group, ITEE
The University of Queensland, Brisbane, QLD 4072, Australia
`appleton@itee.uq.edu.au`

² CSIRO Mathematical and Information Sciences,
Locked Bag 17, North Ryde, NSW 1670, Australia
`hugues.talbot@csiro.au`

Abstract. In this paper we solve the problem of computing exact continuous optimal curves and surfaces for image segmentation and 3D reconstruction, using a maximal flow approach expressed by means of a PDE model. Previously existing techniques yield either grid-biased (graph-based approaches) or sub-optimal answers (active contours and surfaces).

The proposed algorithm simulates the flow of an ideal fluid with spatially varying velocity constraint. A proof is given that the algorithm gives the globally maximal flow at convergence, along with an implementation method. The globally minimal surface may be obtained trivially from its output.

The new algorithm is applied to segmentation in 2D and 3D medical images and to 3D reconstruction from a stereo image pair. The results in 2D agree remarkably well with an existing planar minimal surface algorithm and the results in 3D segmentation and reconstruction demonstrate that the new algorithm does not exhibit grid bias.

1 Introduction

Geometric optimisation techniques have been applied for some time to reconstruction and to image segmentation. These techniques are concerned with the extraction of curves or surfaces which are optimal according to a measure dictated by the application domain. They have the distinct advantage over many other analysis techniques that practitioners need only define an appropriate measure of 'goodness' and then optimise accordingly.

Geometric optimisation techniques include on the one hand active contour methods such as snakes [1], level sets [2, 3] and geodesic active contours [4, 5]; and on the other hand graph-theoretic methods such as shortest paths and minimal cuts.

Graph-theoretic [8] methods have enjoyed success in 3D reconstruction. Stereo matching was first cast as a shortest path problem in the mid 1980's by Ohta and Kanade [9] and independently by Lloyd [10]. The approach remains much

the same in current research [11], where the path through the correlation matrix of maximum sum is obtained. In recent years minimal cuts have been applied to stereo matching yielding improved spatial consistency at the cost of additional computation [12]. These methods have also been successfully applied to image segmentation [13].

Recently several advances have been made in extending optimal methods from discrete graphs to continuous space. Dijkstra's classic shortest path algorithm [14] was extended by Tsitsiklis [15] and later Sethian [16] to computing minimal geodesics and continuous distance functions. These have found broad application to optimal control, wave propagation and computer vision.

In this paper we present an algorithm to compute optimal curves and surfaces in arbitrary Riemannian spaces.

2 Previous work

2.1 Geodesic Active Contours and Surfaces

Geodesic Active Contours and Surfaces were introduced by Caselles et al [4, 5] for segmentation in 2D and 3D images. They are manifolds of co-dimension one which minimise the integral $E[S] = \int_S g(S) dS$, where $E[S]$ is defined as the energy of the surface S . In object segmentation $g : \mathbb{R}^N \rightarrow \mathbb{R}^+$ is typically a decreasing function of edge strength.

Geodesic Active Contours and Surfaces evolve an initial surface via a gradient descent flow toward a local minimum of the energy functional. This evolution is implemented using a level-set embedding due to Osher and Sethian [2]. For a function $\phi : \mathbb{R}^N \rightarrow \mathbb{R}$ whose zero level set is $S = \{\mathbf{x} | \phi(\mathbf{x}) = 0\}$, we may evolve ϕ so as to implement a gradient descent flow, itself derived from variational calculus: $\frac{\partial \phi}{\partial t} = -(g\kappa - \nabla g \cdot \mathbf{N}) |\nabla \phi|$. A fast implicit scheme has also been presented by Goldenberg et al [17]. Unfortunately these gradient descent flows usually stop at local minima. Numerous schemes have been introduced in an attempt to increase robustness [6, 7]. Nonetheless active contours often require additional user interaction, limiting their effectiveness for many image analysis problems.

2.2 Weighted graphs and Riemannian spaces

A number of geometric optimisation techniques have been proposed for computer vision based on discrete graphs [13, 12] and later continuous Riemannian spaces [18, 19]. Due to limited space, we only introduce the most important concepts and results.

Graphs and minimal paths, manifolds and geodesics A path on a positively weighted graph between two points s and t is minimal if there exists no connected path of lower length. Such paths may be computed using Dijkstra's shortest path algorithm [14].

A Riemannian space R can be seen as the continuous equivalent of a weighted graph. It consists of a N -manifold Ω and an associated metric $g : \Omega \rightarrow \mathbb{R}$ (we only consider positive scalar metrics). A simple curve between two points s and t is a minimal geodesic if there exists no curve of lower length. Minimal geodesics may be computed using Sethian's Fast Marching Method [16].

Minimal Cuts and Minimal Surfaces A partitioning of a graph G decomposes its vertex set into a collection $\Gamma_G = \{V_1, V_2, \dots\}$ of disjoint subsets. To each partition Γ_G we associate a cost $C(\Gamma_G)$ which is the total cost of the edges whose endpoints lie in different partitions: $C(\Gamma_G) = \sum_{e \in E^*} C_E(e)$, where $E^* \subseteq E$ denotes the set of edges crossing the partition. The s - t minimal cut problem seeks the partition of minimal cost such that the disjoint vertex sets $s, t \subseteq V$ lie in different partitions. Sedgewick [8] has a good introduction to algorithms solving this problem.

Similarly, a partitioning of a Riemannian space R decomposes the space into a collection $\Gamma_R = \{\Omega_1, \Omega_2, \dots\}$ of disjoint subsets: To each partition Γ_R we associate a cost $C(\Gamma_R)$ which is the integral of the metric g over the partition surfaces $\partial\Omega_i$. In the continuous case, the s - t minimal cut problem seeks the partition Γ_R of minimal total cost such that the disjoint point sets $s, t \subseteq \Omega$ fall in different partitions. Until now no algorithm had been proposed to solve this general problem.

2.3 Discrete and continuous maximal flows

Let G be a graph with edge costs C_E now reinterpreted as *capacities*. A flow $F_G : E \rightarrow \mathbb{R}$ from a *source* $s \subseteq V$ to a *sink* $t \subseteq V$ has the following properties:

- Conservation of flow: The total (signed) flow in and out of any vertex is zero.
- Capacity constraint: The flow along any edge is less than or equal to its capacity: $\forall e \in E, \quad F(e) \leq C_E(e)$.

An edge along which the flow is equal to the capacity is *saturated*. In this and all future formulations, we implicitly add a directed edge connecting $t \rightarrow s$ of infinite capacity to conserve flow uniformly throughout G . A maximal flow in a weighted graph G maximises the flow through the $t \rightarrow s$ edge. Ford and Fulkerson [20] showed that the maximal flow equals the minimal cut, with the flow saturated uniformly on the cut.

Strang [21] explored the extension of maximal flows to continuous domains, and showed that a maximal flow saturates the minimal surface. A continuous maximal flow has the following properties:

- Conservation of flow: $\nabla \cdot \mathbf{F} = 0$
- Capacity constraint: $|\mathbf{F}| \leq g$

The duality between maximal flows/minimal cuts and surfaces has a simple interpretation. Every cut forms a bottleneck for the flow, limiting the maximal flow to be less than the minimal cut. Furthermore the maximal flow is indeed equal to the minimal cut, and on the minimal cut the maximal flow must be saturated uniformly.

2.4 The planar case

In the 2-D case some special properties hold. Given a planar graph G , it is possible to define a dual graph G^* such that each vertex in G^* corresponds to an open region bounded by edges in G (a face). Then minimal cut/maximal flow problem in G corresponds to a shortest path problem in G^* . Similar, slightly more complex properties can be expressed in the continuous case as well.

These dualities are important in the design of planar minimal cut algorithms because the computation of shortest paths is efficient, compared to more general maximal flow methods. They are used in Weihe's discrete maximal flow algorithm [22] and in Strang's continuous maximal flow algorithm [21].

Minimal Cycles The problem of obtaining minimal cycles or closed contours is more challenging than that of obtaining minimal paths and geodesics. The discrete case has been investigated by Sun and Pallottino [23] and Appleton and Sun [24]. The continuous case has been more recently studied by Appleton and Talbot [19] resulting in the Globally Optimal Geodesic Active Contour (GOGAC) algorithm.

3 Minimal surfaces in arbitrary dimensions

Unfortunately the duality between minimal paths and minimal surfaces breaks down in higher dimensions, in part because they have different dimensions. Here we present an algorithm for obtaining continuous maximal flows in arbitrary spaces with scalar metric.

3.1 A continuous maximal flow algorithm

Consider the following system of partial differential equations:

$$\frac{\partial P}{\partial t} = -\nabla \cdot \mathbf{F} \quad (1)$$

$$\frac{\partial \mathbf{F}}{\partial t} = -\nabla P \quad (2)$$

subject to

$$|\mathbf{F}| \leq g \quad (3)$$

Eq. 2 introduces coupling such that gradients in the scalar field P drive the flow $\mathbf{F}$. Eq. 1 and 2 form a simple system of wave equations. They may be recognised as a linearised form of the Navier-Stokes equations describing the dynamics of an ideal fluid with pressure P and velocity $\mathbf{F}$. Eq. 3 constitutes a harsh constraint on the magnitude of the flow velocity $\mathbf{F}$. It is unique to the maximal flow problem and does not appear to have an immediate physical analogy.

For boundary conditions we fix the scalar field P at the source s and sink t : $P_s = 1$ and $P_t = 0$. These values are chosen arbitrarily without loss of generality.

3.2 Properties of the continuous maximal flow algorithm

Conservation of potential P Let $P_A = \int_A P dA$ denote the total integral of P in a given region A not including s, t .

Then

$$\frac{\partial P_A}{\partial t} = \int_A \frac{\partial P}{\partial t} dA = - \int_A (\nabla \cdot \mathbf{F}) dA = - \oint_{\partial A} \mathbf{F} \cdot \mathbf{N}_{\partial A} d(\partial A) \quad (4)$$

So P is conserved in the interior of any *sourceless* region A (any region not including the source s or sink t).

Monotonic reduction of ‘energy’ $P^2 + \mathbf{F}^2$ Consider the temporal rate of change of the total quantity of $P^2 + \mathbf{F}^2$ in a given region A not including s, t :

$$\frac{\partial}{\partial t} \int_A P^2 + \mathbf{F}^2 dA = -2 \int_A P \nabla \cdot \mathbf{F} + \nabla P \cdot \mathbf{F} dA = -2 \oint_{\partial A} P \mathbf{F} \cdot \mathbf{N}_{\partial A} d(\partial A) \quad (5)$$

Note that we have momentarily ignored the effect of the magnitude constraint (Eq. 3). Consequently $P^2 + \mathbf{F}^2$ is conserved in the interior of any sourceless region A . Including the magnitude constraint may only decrease $|\mathbf{F}|$ and hence the energy $P^2 + \mathbf{F}^2$ must monotonically decrease in the interior of a sourceless region. Since the energy is positive it must converge. To ensure convergence of P and $\mathbf{F}$ independently, a dissipative term can be added to the equations. In practice this term is not necessary.

3.3 Correctness at convergence

At convergence P is an indicator function for the interior of the (globally) minimal surface S_{min} separating the point sets s and t . Setting temporal derivatives to zero at convergence, we may restate the system (Eq. 1, 2, 3):

$$\nabla \cdot \mathbf{F} = 0 \quad (6)$$

$$\begin{aligned} \nabla P &= 0 & \text{if } |\mathbf{F}| < g \\ \nabla P &= -\lambda \mathbf{F} & \text{where } \lambda \geq 0 \quad \text{if } |\mathbf{F}| = g \end{aligned} \quad (7)$$

Isosurfaces of P must occur in areas where ∇P is non-zero, i.e. where the flow $\mathbf{F}$ is saturated. By choice of boundary conditions we know that there must exist at least one isosurface separating the source s and sink t , because they have fixed different values of P .

Consider a sourceless region A of constant P with locally extremal value. Then $\mathbf{F}$ must be directed uniformly outward or inward on the boundary ∂A , so $\frac{\partial P_A}{\partial t} \neq 0$. This contradicts the assumed convergence of the system and hence there can be no local extrema of P except on the boundaries s, t .

Consequently every isosurface of P is a simple closed curve containing the source s and hence all isosurfaces have flows directed uniformly outward in the direction of decreasing P . By conservation of flow, the flow into any region must

equal the flow out of that region. Applying this observation to the region between two sequential isosurfaces shows that they must have the same total flow, and hence that all isosurfaces have the same total flow. Isosurfaces are saturated, so at convergence all isosurfaces have the same total capacity. In the usual case of a unique minimal surface, S_{min} will be the only isosurface at convergence and hence P will be an indicator function for its interior.

3.4 Implementation

Eq. 1, 2 are discretised using an explicit first-order scheme in time and space. The scalar field P is stored on grid points while the vector field F is stored by component on grid edges. The system of equations is iterated sequentially with the flow magnitude constraint (Eq. 3) enforced after each timestep. Timesteps are limited by the Courant-Freidrichs-Levi (CFL) stability condition to $\Delta t < 1/N$. In general this simulation may be replaced by any suitable iteration scheme for the linearised Navier-Stokes equations. The fundamental iteration is simple enough that a single implementation may handle input data of arbitrary dimension.

Several heuristics have been found to greatly increase the speed of convergence. The fields P and F are rapidly initialised using the pre-flow push discrete maximal flow algorithm of Goldberg and Tarjan with both global and gap relabelling [8]. A multiscale approach is then applied recursively for rapid convergence at the finest grid resolution from a coarse grid estimate. Computation may be avoided in the interior of the source s and sink t , yielding great savings when they occupy a significant portion of the space.

At convergence, the pressure field P is theoretically perfectly binary with value 1 within the volume bounded by the minimal surface, and 0 outside. In practice convergence is deemed to be attained if the sum of the relative areas of pressure $A_{P \geq 0.97}$ and $A_{P \leq 0.03}$ is greater than 0.99.

4 Applications

In this section we present results for the application of minimal surfaces to the segmentation of 2D and 3D medical images and 3D reconstruction from stereo images.

All tests were performed on a 700MHz Toshiba P-III laptop with 192MB of RAM under the Linux operating system. The algorithm presented here has been implemented in C and has not been optimised significantly.

4.1 2D Image Segmentation

Here we apply discrete minimal cuts, GOGAC [19], and the algorithm presented in this paper to segment a microscope image of a cluster of cells (Figure 1(a)) and compare the results. In spite of its apparent simplicity this problem outlines the challenge of delineating faint boundaries between cells without leaking.




Fig. 1. Segmentation of a microscope image of a cell cluster. (a) The microscope image. (b) Segmentation via a discrete minimal cut. (c) Segmentation via Globally Optimal Geodesic Active Contours. (d) Segmentation via continuous maximal flows.

The segmentation of each cell is performed independently in sequence. The source sets are depicted in (Figure 1(b-d)) while the sink is the image boundary. The discrete minimal cut solves a discretised minimal surface problem, resulting in a clear grid bias and a poor segmentation. GOGAC and the continuous maximal flow algorithm solve the same continuous optimisation problem and are in clear agreement. Note that the continuous segmentations follow the perceived cell contours even in the absence of local cues.

The image depicted in Figure 1(a) has dimensions 231×221 . We reduce the amount of computation required by expanding the sink to include only the cells of interest, a region of size 150×100 .

The discrete minimal cuts required 12.9 seconds to compute in total. GOGAC required 1.48 seconds to compute in total. The continuous minimal surface algorithm presented here required 27.3 seconds in total to converge.

4.2 3D Image Segmentation

We apply the algorithm proposed in this paper to segment the lungs in the Computed Tomography image of a chest depicted in Figure 2, and compare this to a segmentation using a discrete minimal cut. The same metric is used for both the discrete and continuous minimal cut computations. The sources are small spheres inside each lung and the sink is the boundary of the volume, not including the uppermost face. The lungs are segmented separately in turn.

Observe the directional bias along the grid in the discrete minimal cut. This is particularly evident at the flat boundaries in the interior surfaces at the top of the lungs (Figure 2(b)). By contrast, the continuous minimal surface does not exhibit such a directional bias.

The data shown in Figure 2 has dimensions $200 \times 160 \times 90$. The discrete minimal cut required 14 minutes to compute using Goldberg and Tarjan's classic pre-flow push algorithm to compute a maximal flow [8]. The continuous minimal




Fig. 2. Segmentation of the lungs in a chest CT image. (a) The CT image. (b) Segmentation using a discrete maximal flow algorithm. Observe the directional bias due to the grid. (c) Segmentation from identical input using continuous maximal flows.

surface algorithm required 1 hour and 52 minutes to converge, including all initialisation.

4.3 3D Scene Reconstruction from Stereo Images

We apply our maximal flow algorithm to the reconstruction of a scene from a stereo image pair. The metric used here is based on the Zero-mean Normalised Cross Correlation (ZNCC) area based matching score, popular for its good statistical basis and efficient computation [11]. Here ZNCC scores are computed using a window of size 5×5 . Scores are computed for integer disparities in the range $[-15, 0]$. We set $g = 1 - ZNCC$ throughout the disparity volume to convert the problem of finding a maximal surface into that of finding a minimal surface. Following Roy and Cox [12] the source and sink are the first and last disparity layer respectively.

Figure 3 depicts the results of the reconstructions using both a discrete and a continuous minimal cut. Both disparity maps and the corresponding textured surfaces are shown. The results computed from the discrete minimal cut show very distinct flat zones due to the small number of disparities considered, while the continuous minimal surface is able to capture detail below the discrete level. As a result additional features are visible in the continuous minimal surface, including the third parking meter and the large scale surface texture of the bushes. Observe the absence of bias in the shape of the frame of the car.

The stereo image pair used here has dimensions 256×240 . The discrete minimal cut required 3 minutes and 50 seconds to compute. The continuous minimal surface algorithm required 16 minutes to converge including initialisation.




Fig. 3. Reconstruction of a scene from two views. (a) Left view only. (b) Reconstruction from a discrete maximal flow. (c) Reconstruction from identical input using continuous maximal flows.

5 Conclusions

In this paper we have described a novel algorithm to compute continuous maximal flows in Riemannian spaces with scalar metric.

We have modelled this flow using a system of PDEs simulating the flow of an ideal fluid with spatially varying velocity constraints. The computation is implemented by a first order finite differences scheme. For efficiency reasons a discrete maximal flow result is used for initialization and a multi-resolution scheme is also used. In spite of this the total computational cost is relatively high, especially in 3-D at high resolution.

At convergence, the solution exhibits globally maximal flow, trivially yielding the expected minimal surface. However a full proof of convergence is not given in this paper and requires more work.

The new algorithm has been applied to segmentation in 2D and 3D medical images and to 3D reconstruction from a stereo image pair. The results in 2D agreed remarkably well with an existing planar minimal surface algorithm. The results in 3D segmentation and reconstruction demonstrated that, in contrast to existing discrete optimisation algorithms, the new algorithm computes surfaces which are not biased by the choice of computational grid. We believe these benefits may outweigh the computational cost of the method in cases where accuracy is paramount. The proposed algorithm also provides an accuracy benchmark for faster methods.

References

1. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. *International Journal of Computer Vision* **1** (1998) 321–331

2. Osher, S., Sethian, J.A.: Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics* **79** (1988) 12–49
3. 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)
4. Caselles, V., Kimmel, R., Sapiro, G.: Geodesic active contours. *IJCV* **22** (1997) 61–79
5. Caselles, V., Kimmel, R., Sapiro, G., Sbert, C.: Minimal surfaces based object segmentation. *IEEE Trans. on PAMI* **19** (1997) 394–398
6. Cohen, L.D.: On active contour models and balloons. *Computer Vision, Graphics, and Image Processing. Image Understanding* **53** (1991) 211–218
7. Xu, C., Prince, J.L.: Gradient vector flow: A new external force for snakes. In: *Proc. IEEE Conf. on Comp. Vis. Pat. Rec. (CVPR)*, Los Alamitos: Comp. Soc. Press (1997) 66–71
8. Sedgewick, R.: *Algorithms in C*. Third edn. Number 5. Addison-Wesley (2002)
9. Ohta, Y., Kanade, T.: Stereo by intra- and inter-scanline search using dynamic programming. *IEEE Trans. Pattern Anal. Mach. Intell.* **7**(2) (1985) 139–154
10. Lloyd, S.: Stereo matching using intra- and inter-row dynamic programming. *Pattern Recognition Letters* **4** (1986) 273–277
11. Sun, C.: Fast stereo matching using rectangular subregioning and 3D maximum-surface techniques. *International Journal of Computer Vision* **47** (2002) 99–117
12. Roy, S., Cox, I.J.: A maximum-flow formulation of the n-camera stereo correspondence problem. In: *Int. Conf. on Computer Vision (ICCV'98)*, Bombay, India (1998) 492–499
13. 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** (1988) 203–213
14. Dijkstra, E.: A note on two problems in connexion with graphs. *Numerische Mathematik* **1** (1959) 269–271
15. Tsitsiklis, J.N.: Efficient algorithms for globally optimal trajectories. *IEEE Transactions on Automatic Control* **40** (1995) 1528–1538
16. Sethian, J.: A fast marching level set method for monotonically advancing fronts. In: *Proceedings of the National Academy of Sciences*. Volume 93(4). (1996) 1591–1595
17. Goldenberg, R., Kimmel, R., Rivlin, E., Rudzsky, M.: Fast geodesic active contours. *IEEE Trans. On Image Processing* **10** (2001) 1467–1475
18. Cohen, L.D., Kimmel, R.: Global minimum for active contour models: A minimal path approach. *International Journal of Computer Vision* **24** (1997) 57–78
19. Appleton, B., Talbot, H.: Globally optimal geodesic active contours. *Journal of Mathematical Imaging and Vision* (2002) Submitted.
20. L. R. Ford, J., Fulkerson, D.R.: *Flows in Networks*. Princeton University Press, Princeton, NJ (1962)
21. Strang, G.: Maximal flow through a domain. *Mathematical Programming* **26** (1983) 123–143
22. Weihe, K.: Maximum (s, t) -flows in planar networks in $O(|V|\log|V|)$ time. *Journal of Computer and System Sciences* **55** (1997) 454–475
23. Sun, C., Pallottino, S.: Circular shortest path in images. *Pattern Recognition* **36** (2003) 709–719
24. Appleton, B., Sun, C.: Circular shortest paths by branch and bound. *Pattern Recognition* (2003) Accepted.