

Real-time Dynamic Deformable Meshes for Volumetric Segmentation and Visualisation

Bowden, R. Mitchell, T. A. Sahardi, M.
Vision and VR Group, Dept M & ES
Brunel University
Uxbridge, Middlesex
UB8 3PH
richard.bowden@brunel.ac.uk

Abstract

This paper presents a surface segmentation method which uses a simulated inflating balloon model to segment surface structure from volumetric data using a triangular mesh. The model employs simulated surface tension and an inflationary force to grow from within an object and find its boundary. Mechanisms are described which allow both evenly spaced and minimal polygonal count surfaces to be generated. The work is based on inflating balloon models by Terzopolous[8]. Simplifications are made to the model, and an approach proposed which provides a technique robust to noise regardless of the feature detection scheme used. The proposed technique uses no explicit attraction to data features, and as such is less dependent on the initialisation of the model and parameters. The model grows under its own forces, and is never anchored to boundaries, but instead constrained to remain inside the desired object. Results are presented which demonstrate the technique's ability and speed at the segmentation of a complex, concave object with narrow features.

1 Introduction

The popularity of medical imaging as a source of 3D datasets has fuelled a huge interest in the processing and segmentation of volumetric data. The problems of understanding 3D structure from a discretely sampled volume have shown the benefit of visualisation techniques. Surface approximations, such as isosurfacing, allow surface to be extracted that which, rendered and shaded, provide an invaluable insight into a volume's internal structure.

The reconstruction of multi model datasets from different sources of volumetric data is greatly simplified by the successful segmentation of surface topology. Surfaces that directly correspond to a volume can be matched far more simply than the original volumes[6].

In addition to structural insight, surface approximations are invaluable in reducing the processing time needed for traditional image processing techniques, as processing can be localised to a contour boundary. Further more, these surfaces can provide a mathematical representation of shape which can then be used statistically to model and classify shape and deformation [3][1].

Since the original formulation of Active Contour Models (Snakes) [4] a significant interest has been shown in extending the technique to dynamic 3D models. Snakes have shown to be useful in contour reconstruction, but require large amounts of user intervention to successfully segment complex objects. Point Distribution Models [3] can simplify the problem of object recognition and segmentation by statistically constraining the shape of the model within suitable bounds, through the analysis of a training set of shapes. However, in 3D, where models become too large to manufacture by hand, another means of generating training sets for statistical analysis must be found. The purpose of this work is to produce a generic technique that can be used to produce 3D surface approximations from volumetric data which can be subsequently used as the basis for the generation of 3D PDMs.

Terzopolous and Vasilescu [8] extended the snake model to include an inflation force that helps remove the need for initial contour placement and thus avoid convergence on local minima. The inflation force drives the snake model outwards towards the object boundary like an inflating balloon. Terzopolous and Vasilescu formulated the model as a finite element mesh and later extended the model to a thin plate spline, demonstrating successful results in the reconstruction of range data and volumetric CT data surface representations[5].

This paper presents an iterative dynamic mesh model which uses simulated physical forces to segment desired surface approximations from volumetric datasets. The work is based on the work of Chen and Medioni[2] which is itself a continuation of the work on dynamic balloon models by Terzopolous and Vasilescu[8]. Chen and Medioni simplified Terzopolous and Vasilescu's model and applied the work to the constrained problem of reconstruction from pre-registered range images. Here we make further simplifications and extend the problem back into the domain of volumetric data.

1.1 Overview of the Dynamic Mesh Model

The mesh structure consists of a triangular mesh which can vary in size, shape and connectivity. Each vertex is connected to other vertices in the model by the edges of the polygonal facets. These interconnections are used to simulate springs that connect the mesh mathematically. The force of these springs gives a resulting surface tension to the model which attempts to keep the surface as smooth as possible. An inflation force is used at each vertex to inflate the overall model, while surface tension attempts to keep the mesh spherical. A simple local feature detection scheme is used at each vertex to remove the inflation force as nodes reach the boundaries of desired structures. A dynamic mesh subdivision scheme is used to subdivide polygons locally if they exceed set size or curvature criteria. This allows the mesh to inflate and grow until a boundary is located. Once the mesh has converged on a solution, a good local edge detection scheme can be used to lock vertex points to the boundary. The process starts with a small minimal polygon object which is inflated from within a volumetric image with the inflation force driving the surface towards the object boundary. The mesh grows in size and complexity

to fill the object like an inflating balloon until the mesh vertices lie close to the true object boundary (See Fig 1). This technique requires no user intervention after the initial placement and provides a simple, fast method for object segmentation, which produces surfaces with a minimal polygon count.

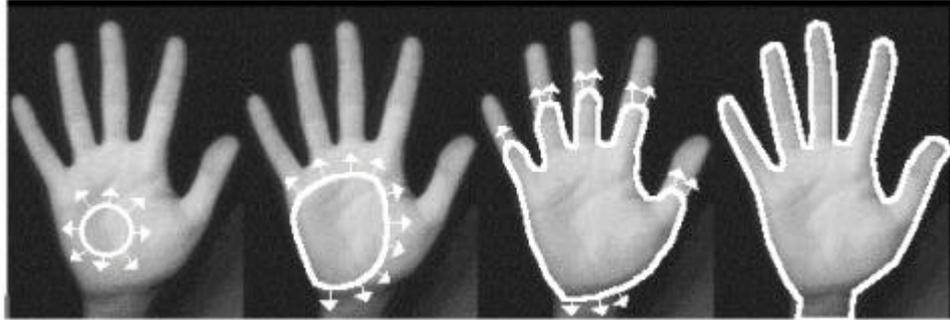


Figure 1 - Simple 2D Contour Inflating Towards the Object Boundary

The remainder of this paper is organised as follows, Section 2 formulates the dynamic mesh structure and subdivision mechanisms. Section 3 shows the resulting model applied to an MRI image of the human hand. Finally conclusions and further work are discussed.

2. The Balloon Model

The balloon model consists of a mesh of triangular facets or patches. The initial triangulated surface can be any shape or size allowing the re-application of a segmented surface to a new dataset. Each node (vertex) has two forces acting upon it. The spring force derived from the sum of the vectors of the interconnections of the mesh, and the inflation force, derived from the weighted normal direction of the surface at each node.

The operation of the inflating balloon model can be encapsulated by the following algorithm.

Algorithm 1.

for a given closed form polygonal model,

build a connected mesh of vertices

while number of polygons is not constant do

 compute the normal at each node

 for each node do

 compute the elastic force using eq (4) (See Section 2.2)

 test node position in dataset using feature detection scheme.

 If feature not found calculate the inflation force using eq (5) (See

Section 2.3)

 add this to the elastic force

 compute the new node position v_i^{t+1} using eq (3) and update node

 perform dynamic subdivision using algorithm (2) (See Section 2.4)

2.1 A Simple Dynamic Model.

The motion of any element i on a finite element mesh model can be described by the set of coupled second order differential equations [8]

$$(1) \quad m_i \frac{d^2 x_i}{dt^2} + \gamma_i \frac{dx_i}{dt} + g_i = f_i \quad i = 1, \dots, n.$$

Here, x is the location of the element i , m is its mass, g is the surface tension, generated by the interconnections of the elastic mesh, f is the inflation force and γ is the velocity-dependent damping coefficient that controls the rate of dissipation of kinetic energy. Giving the mesh these simulated physical properties provides a robust model that performs well but at a computational cost. The main rationale for the momentum term $m_i \frac{d^2 x_i}{dt^2}$ is its ability to reduce the mesh's susceptibility to noise. Due to the momentum

of nodes the damping term γ is necessary to bring the model to rest. The mesh reaches an equilibrium state when $\frac{dx_i}{dt} = 0$ which can take some time [2]. Chen and Medioni simplify this model by making $m=0$ and $\gamma=1$ for all i reducing eq 1 to,

$$(2) \quad \frac{dx_i}{dt} + f_i = g_i \quad i = 1, \dots, n.$$

Due to this simplification the equation (2) has a very simple explicit integration [2]

$$(3) \quad x^{t+\Delta t} = (f_i + g_i) \Delta t + x^t$$

2.2 Simplified Spring Force

The spring force exerted on node i by the spring linking node i and j of natural length l_{ij} can be expressed as [8],

$$s_{ij} = \frac{c_{ij} e_{ij}}{\|r_{ij}\|} r_{ij}$$

where c_{ij} is the stiffness, $r_{ij} = x_j - x_i$ the vector separation of the nodes, $\|r_{ij}\|$ is the length of the spring and $e_{ij} = \|r_{ij}\| - l_{ij}$ is the deformation of the spring.

In order to generate a generic technique for the segmentation of objects, and due to the large nature of 3D objects it is not feasible to assign values to c_{ij} and l_{ij} for each node. Further simplifications can therefore be made by setting all stiffness coefficients to a constant value with a minimum spring length of zero, $c_{ij} = c$ and $l_{ij} = 0$.

The total elastic force on a node i is therefore,

$$(4) \quad g_i = \frac{c}{n} \sum_{j=0}^n r_{ij}$$

2.3 Inflation Force

The inflation force applied to each node i is,

$$(5) \quad f = kn_i$$

where n_i is the normal at node i , and k is the amplitude of the inflation force.

Node normals are calculated as the average normal of the surrounding polygons sharing the node i . Other, more complicated schemes as used by Chen and Medioni [2], give little benefit as errors in this normal estimation technique are reduced by the surface smoothing properties of the surface tension (elastic force) This also gives a significant performance increase as normals must be recalculated at least once every iteration of the algorithm.

2.4 Dynamic Subdivision

As the inflation force increases the surface area of the mesh, individual polygons grow in size. As the elastic force is directly proportional to the size of polygons, there becomes a point where the elastic force will not allow the mesh to increase in size further unless the inflation amplitude is increased accordingly. Dynamic subdivision can be used to subdivide polygons which exceed set size criteria and keep polygons within a suitable limit. Each edge of the mesh is checked at each iteration to see if it exceeds the subdivision threshold. Fig 2 demonstrates how the process works.

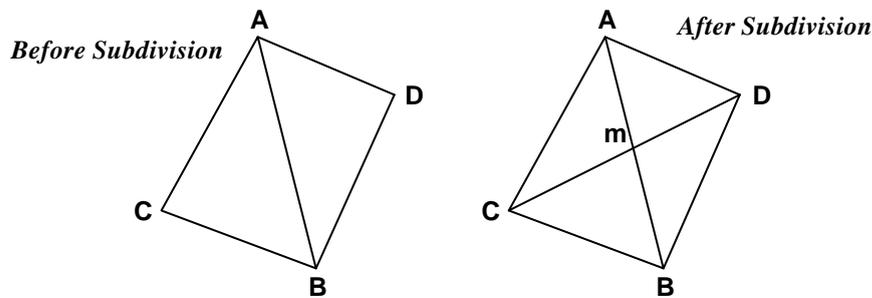


Fig 2 - Dynamic Subdivision

When the length of a node connection **AB** exceeds set criteria, distance or curvature, the two triangles that contain this edge are located (**ABC, ADB**) and removed from the polygon list. The mid point **m** of **AB** is calculated and four new polygons constructed **AMC, CMB, ADM, and MDB**. The internal connectivity of the mesh is also altered to reflect this new local structure. Long thin triangles are undesirable as they do not model local surface properties well. This technique ensures that they never occur as any edge that exceeds a distance threshold is immediately subdivided.

2.4.1. Dynamic Subdivision Algorithm

The dynamic subdivision procedure can be encapsulated by the following algorithm.

Algorithm 2.

```

for each node (V1) do
    for each connection to another node (V2) do
        if the connection (V1V2) matches the subdivision criteria do
            remove connection (V1V2)
            remove the two polygons that share this edge
            find the mid point m of V1V2
            construct four polygons using m as a common node
            update the connections of the mesh
recalculate the normal at each node
    
```

2.5 Subdivision Criteria

Using a distance threshold for subdivision produces an evenly spaced mesh which can alter its structure locally to fit any dataset. In this work we use the length of a triangle's edge to determine if that edge (and the corresponding triangles) needs to be subdivided, this produces an evenly spaced mesh which aids in the construction of the PDM. It is also possible to use other criteria to provide a more flexible approach. As the normal at each node is known for use with the inflation force, the dot product of two adjacent vertices' normals represents local surface curvature. This can be used to further subdivide the mesh if the dot product drops below a certain threshold value, i.e. the area has a high degree of curvature, allowing more vertices to be placed in these areas of high curvature. This is useful where long narrow features are present in the dataset.

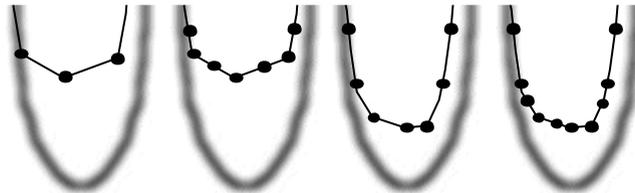


Figure 3 - Curvature Based Subdivision

Figure 3 demonstrates an image boundary and an inflating balloon front. The boundary shown has found an equilibrium state in the narrow feature. By subdividing the mesh on a curvature basis, in addition to distance, extra vertices are added to the front of the model providing the inflation force needed to successfully segment the long narrow feature.

Both subdivision criteria can be used in conjunction to minimise the polygon count of a mesh, removing the need for post-processing techniques such as Delaunay Triangulation [7]. An edge is subdivided only if it exceeds both a distance and a curvature threshold. Polygons on parts of the surface with low curvature grow beyond the threshold keeping polygon counts to a minimum. Therefore, areas of high curvature have larger numbers of small polygons that better model the surface features.

2.6 Feature Detection

Edge features within an image are typically identified as a change in intensity from one range to another via an isointensity which depicts the boundary of these two regions.

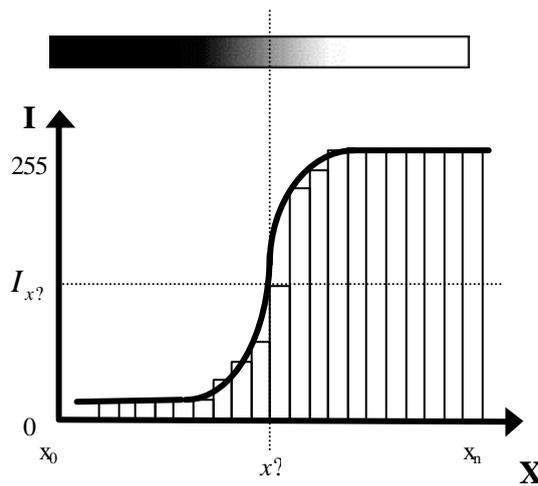


Figure 4 - The Boundary between Light and Dark

Figure 4 shows a cross section through an image depicting a sharp boundary between light and dark. The intensity x_i depicts the threshold that would generate an isointensity boundary for this feature within an image. Providing scanning starts within the model boundary, it can be said the boundary (x_i) has been passed when either

$$I_{x_i} \geq I_{x_i} \quad \text{or} \quad I_{x_i} \leq I_{x_i} \quad \text{where } i = 0, 1, \dots, n$$

depending on the direction of the intensity gradient along the isosurface boundary normal.

This simple thresholding mechanism can be used to detect when the balloon has just passed through a possible isosurface boundary, at which point the inflation force can be removed for that node. Due to the noisy nature and simplicity of this mechanism, many false boundary points are detected. However, elasticity is a constant force and as such provides the function of a simple momentum term which pulls the nodes away from false boundary points.

Where complicated internal structures are required this approach may not provide adequate results. In this situation, other more sophisticated feature detection schemes can be employed. However as we are primarily concerned with the external boundary of the model where a distinct boundary is present, this approach provides an efficient and simple solution.

2.7 Robustness to Noise

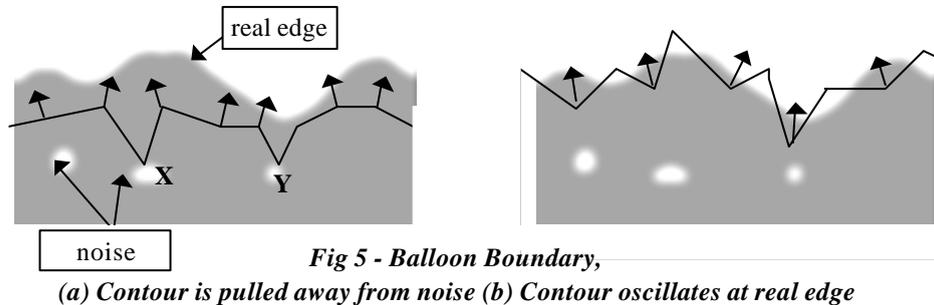


Fig 5 - Balloon Boundary,

(a) Contour is pulled away from noise (b) Contour oscillates at real edge

Figure 5 demonstrates this invulnerability to noise spikes. In fig 5a the boundary moves towards the true boundary through the influence of the inflationary force. Pts X and Y are located on noisy areas of the image. Where these false edges are located the inflation force is removed. However, as the remainder of the contour progresses forward under the inflation force the elasticity pulls these points away from the noise. Once a sufficient distance from the noise has been reached the edge detection criteria no longer apply and the inflation force is reapplied. Elasticity then helps smooth these features as the process iterates. Figure 5b demonstrates what happens when the contour approaches the true boundary. As points are inflated beyond the boundary their inflation force is removed and elasticity pulls the point back within the model, where the inflation force is then re applied. This causes the contour to oscillate around the true edge. As points oscillate back and forth chaotically their overall movement is at a minimum and therefore mesh subdivision approaches zero. At this point a good local edge detection scheme can be used to clamp nodes onto their closest edge. This creates an evenly spaced mesh that is a good surface approximation to the desired object.

3 Results

The balloon model is now presented working on a raw MRI scan of a human hand. The volume is 256x256x20 voxels in size. This is rescaled by 1x1x2 to reconstruct a cuberville and trilinear interpolation used to estimate values within the volume. Fig 6 shows the development of the balloon mesh when applied to this dataset. Initially, a seed balloon is placed within the volumetric dataset.

The seed consists of a simple diamond shape with 8 polygons and 6 vertices. Forces are applied to the model and after 10 iterations it has grown to 307 polygons. The almost spherical shape is due to the surface tension of the model. Its non-spherical symmetry demonstrates that positive features have been detected early on in the process and thus the inflation force has not been applied evenly. This demonstrates the algorithms robustness to false boundaries and noise.

As the process iterates further the final shape very quickly starts to take form. Although mesh subdivision continues we can see that its is starting to decrease in rate considerably after 40 iterations. Although the model will finally converge on a stable solution, it is sufficiently complete at around 70 iterations which takes approximately 35 seconds on a single MIPS R4400 200MHz processor, including render time. This is significantly faster than previous researchers techniques, the most comparable being the

work of Chen and Medioni[2], where a comparable complexity model takes approximately 30 mins to iterate on a SUN Sparc-10 machine. This can also be compared with a standard isosurface of the external hand boundary which generates a surface of some 235000 polygons.

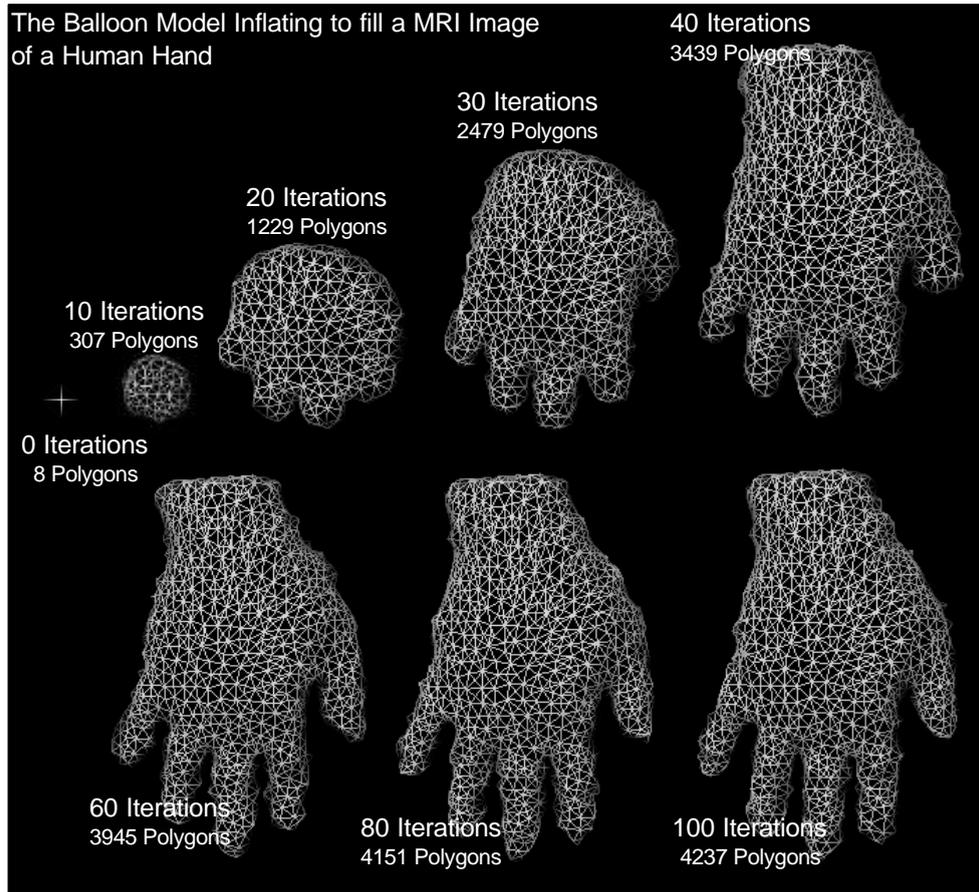


Figure 6 - Segmentation of an MRI dataset of the Human Hand

The hand dataset is a good example of the effectiveness of the technique, demonstrating its ability to work with complex noisy images which contain an object with convex, concave and long narrow features.

4 Conclusions

This paper has presented a surface segmentation method which uses a simulated inflating balloon model to estimate structure from volumetric data using a triangular mesh. The model uses simulated surface tension and an inflationary force to grow from within an object and find its boundary. Mechanisms have been described that allow either evenly spaced or minimal polygonal count surfaces to be generated. Unlike previous work by researchers, the technique uses no explicit attraction to data features and as such is less

dependent on the initialisation of parameters and local minima. Instead, the model grows under its own forces, never anchored to boundaries but constrained to remain inside the desired object. Results have been presented that demonstrate the technique's ability and speed at the segmentation of a complex, concave object with narrow features, while keeping model complexity within acceptable limits.

5 Future Work

This work is ongoing, with constant refinements being made to the procedure. As stated earlier, the primary rationale for this work is the ability to produce low level polygonal surface approximations to allow 3D Point Distribution Models to be built for automatic recognition, segmentation and analysis of volumetric data. Work has also been done in the area of mesh self-intersection. A set of criteria have been developed which allow the detection of mesh self-intersection. Future work includes allowing this criteria to be used to detect intersections, and re-join the mesh at these points to allow more complex torus like shapes to be successfully extracted.

6 Acknowledgements

This work is supported by EPSERC and Foster Findlay Associates Ltd. The authors would like to thank Tony Heap and the Centre for Medical Imaging Research (CoMIR) at the University of Leeds for access to the MRI data of the hand model.

References

- [1] Bowden, R., Heap, A. J., and Hart, C., "Virtual Datagloves: Interacting with Virtual Environments Through Computer Vision", in Proc. 3rd UK VR-Sig Conference, DeMontfort University, Leicester, UK, July 1996.
- [2] Chen, Y., and Medioni, G., "Description of Complex Objects from Multiple Range Images Using an Inflating Balloon Model.", *Computer Vision and Image Understanding*, 61(3), 1995, 325-334.
- [3] Cootes, T. F., Taylor, C. J., Cooper, D. H., and Graham, J., "Active Shape Models - Their Training and Application.", *Computer Vision and Image Understanding*, 61(1), 1995, 38-59.
- [4] Kass, M., Withkin, A., and Terzopoulos, D., "Snakes: Active Contour Models.", *International Journal of Computer Vision*, 1988, 321-331.
- [5] McInery, T., and Terzopoulos, D., "A finite element model for 3D shape reconstruction and nonrigid motion tracking", in Proc. International Conference on Computer Vision, Berlin, Germany, May 1993, 518-523.
- [6] Moshfeghi, M., Ranganath, S., and Nawyn, K., "Three-Dimensional Elastic Matching of Volumes. *IEEE Transactions on Image Processing*", 3(2), 1994, 128-137.
- [7] Soucy, M., and Laurendeau, D., "Multiresolution Surface Modeling Based on Hierarchical Triangulation.", *Computer Vision and Image Understanding*, 63(1), 1996, 1-14.
- [8] Terzopoulos, D., and Vasilescu, M., "Sampling and Reconstruction with Adaptive Meshes", in Proc. Conference on Computer Vision and Pattern Recognition, Maui, HI, June 1991, 70-75.