Tracking the interframe deformation of structures in 3D ultrasound imaging

M. Syn[†], J.P.M. Gosling[†‡], R.W. Prager[†], L. Berman[‡], J. Crowley[‡]

[†]Cambridge University Engineering Department
[‡]University Department of Radiology, Addenbrookes Hospital

Abstract

Three dimensional ultrasound imaging with a freehand probe[1] allows a flexible approach to medical visualisation and diagnosis. Given the imperfect accuracy of proprioceptive devices used to log the position and tilt of the probe, it is important to utilise the position constraints provided by image evidence. This is also important if we wish to consider the visualisation of structures which move significantly during acquisition, such as a heart or foetus. We present here an initial approach to more robust segmentation and shape recovery in a particularly noisy modality. We consider 2D segmentation based on edge evidence, using first an **active contour**, then finding an optimal segmentation using **simulated annealing**.

Correspondence between contours in adjacent frames can only be solved in general cases by use of a 3D prior model. Dynamic physics-based mesh models as used by Pentland [20] and Nastar [17], allow for **shape** modelling, and over-constrained 3D shape recovery can then be performed using the intrinsic **vibration modes** of the model.

1 Motivation

One of the skills of the clinical ultrasonographer is the ability to interpret two-dimensional slices through complex three-dimensional structures. However, the presentation of true three-dimensional information (even after rendering or slicing to produce a viewable two-dimensional image) is of definite benefit to other practitioners (e.g. surgeons) who may need to use the data. It is also of likely benefit to the ultrasonographer in easing the interpretation of particularly ambiguous or complex structures.

The most important clinical application of segmentation in three dimensional imaging is the ability to accurately estimate volumes for various anatomical structures (e.g. prostate, thyroid or lesion volumes), whether in the planning of treatment, or when estimating treatment effectiveness. Automatic techniques for segmentation have the very significant advantage of repeatability over traditional hand-segmentation approaches, and eliminate the problems of inter-operator variation.

A second benefit of automatic segmentation is that it allows much less computationally intensive approaches to three-dimensional visualisation. The number of operations required to map a comparatively small number of control points describing a segmentation curve into three dimensions is significantly lower than that required to map every pixel within each acquired image. A third benefit that may be derived from such a segmentation, even in the event that one *does* map all of the original image data into three

[1]EPSRC project GR/H74032 (Stradivarius)

dimensions, is that the superposition of the segmentation data on the rendered image data can make interpretation of this data easier.

A number of problems arise from the flexibility of the freehand approach, the most important of which is the accuracy of position data returned by the proprioceptive device attached to the probe [11]. Errors in magnetic proprioception can be caused by the presence of large metallic objects, by passing near the origin of the probe's coordinate system, or by random noise in position data returned by the probe.

Other problems specific to our setup arise from the use of video tape to record images during scanning, and the audio track to log position information. Logging errors give rise to spurious frame positions, emphasising the need for feedback confirmation using image evidence. We are currently implementing a system which allows direct logging of image and position data to workstation memory.

2 Review

Approaches to 2D boundary detection in ultrasound images are briefly reviewed by Feng *et al* [8] who apply fuzzy reasoning techniques to left ventricular boundary detection. The referenced approaches utilise assorted optimisation methods, and general or application-specific heuristics in segmenting structures based on edge information.

Friedland and Adam [9] apply Simulated Annealing (SA) as an optimisation technique [10, 25] to find an appropriate segmentation based on edge evidence and continuity considerations. We apply this formulation, using a general initial contour and modified energy functionals. We also examine segmentation using a closed "active contour"[2] formulation [13].

In simple cases, the segmented boundaries can be linked between adjacent frames using simple heuristics of volume minimisation or continuity [4, 14, 18] or computing a 3D Delaunay triangulation between adjacent contours [3]. Correspondence between contours in adjacent frames can only be solved in general cases by use of a 3D prior model of the structure under investigation, since there are frequently a number of plausible permutations when more than one contour exists in a scan window (Figure 1).

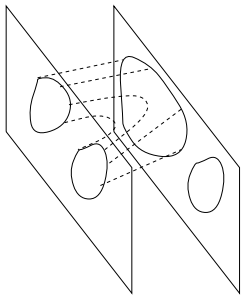

Figure 1: Possible erroneous interpolation

[2] An active contour both searches and segments during its evolution.

This approach frees us from the constraints in our sequential 2D segmentation that

- we acquire most of the structure's *closed* boundary within the scan window,
- the imaged structure has simple topology with simple and well-behaved branching behaviour,
- during the scan pass a single direction is used.

Further advantages in using a 3D prior model are that

- for volume (as opposed to surface) models, structural information from internal boundaries can be incorporated to limit torsional degrees of freedom in contour correspondence,
- boundary initialisations are no longer derived from adjacent frames; this would otherwise be unsatisfactory if frames are undersampled compared to boundary variation,
- positional constraints established within a frame can be propagated throughout the entire structure in a consistent way.

The 2D "active contour" can be extended into 3D as an "active balloon". Cohen and Cohen [5], and McInerney and Terzopoulos [15] utilise an *inflating* active balloon to segment medical structures. The additional constraint of surface normal direction given by a 3D edge, makes a 3D edge segmentation attractive. Given that 2D edges in ultrasound are unreliable however, 3D edges will be even more so, unless (preferably 3D) texture discriminators are incorporated. Applying SA in 3D segmentation would be computationally expensive, and would be very much dependant on a good initialisation to give a small search region within which annealing occurs.

We use 3D mesh models whose dynamics can be solved using the Finite Element Method (FEM). Nastar and Ayache [17] have applied such deformable models in segmentation and tracking in ultrasound and magnetic resonance images. Pentland *et al* [19, 20], and Metaxas and Terzopoulos [16] have applied superquadric meshes to 3D model fitting and tracking, as well as other applications.

We aim to use the intrinsic vibration modes of the dynamic model as an ordered set of deformations, a number of which can be used to over-constrain the shape recovery procedure, given the number of correspondences between model and data. This global segmentation can then be fine tuned locally within each 2D slice, since the limited set of deformations used cannot encompass all local variations. In any case, our image data is acquired along a set of scan planes, and does not constitute a true volumetric image.

3 2D Segmentation

2D images acquired from the ultrasound scanner have been interpolated from polar scans into cartesian form by internal circuitry[3]. The uncertainty in each pixel value therefore varies over the image, and the scale-space and preprocessing operations that are described below ought to take this into account. We would ideally be able to operate directly on the radial data from the probe, but since we cannot yet do so, we work under the assumption that our data is uniformly sampled in cartesian form.

[3] An unknown amount of image preprocessing also occurs in scanner, digitiser and VCR circuitry.

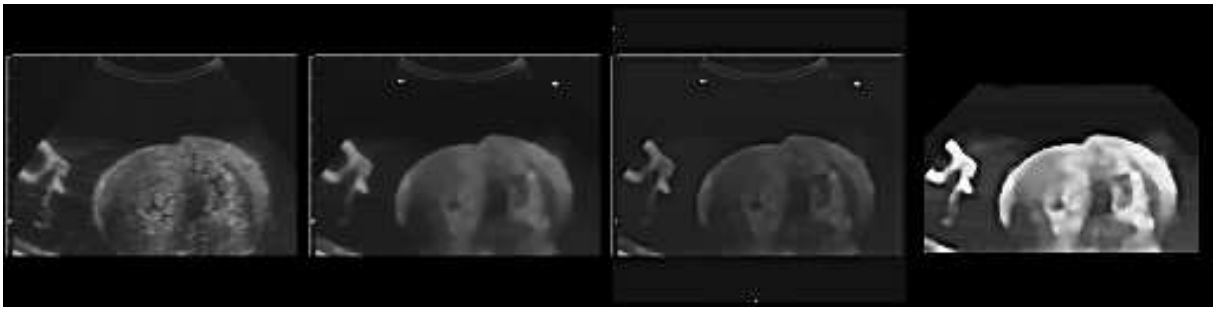

(a) (b) (c) (d)

Figure 2: (a) original (b) anisotropic diffusion (c) high-pass filter (d) histogram equalisation

We preprocess the images to facilitate boundary detection. The use of anisotropic diffusion followed by high-pass filtering and histogram equalisation, has been found to be effective in promoting edge features [7], provided appropriate scale and filter parameters have been established (Figure 2).

3.1 Active contour

We implement four functions weighted relative to each other, which constitute the internal and external "forces"[4] acting upon the active contour (Figure 3). One function describes the attraction to local edges[5], another minimises the local curvature at the current control point. This second function tends to maintain a circular contour in the absence of other information, but we can incorporate a more specific bias if a prior model is available (see section 4).

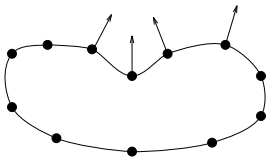

Figure 3: Active contour with radial forces acting on control points

The third function "inflates" the contour so that it seeks outward from its initial placement, and allows coarse control over the noisy local minima which the contour should ignore. The final function biases the contour in favour of a configuration similar to that in the previous frame. This function is made redundant if we implement a prior model as noted above.

The disadvantage of this formulation is that the contour must be initialised sufficiently close to the required boundary. Weight parameters defining the relative importance of each "force" need to be set so that the correct choice is made as the contour evolves over a number of possible candidate boundaries. These parameters will vary from image to image, making this an awkward basis for a semi-automated segmentation procedure.

[4] They serve to deform the active contour by assigning radial displacements to each control point on the contour. They are *not* true forces since no record is kept of accumulated momentum.

[5] A global isotropic edge map of the image is built during anisotropic diffusion.

3.2 Simulated annealing

SA is an optimisation technique that guarantees convergence to a global minimum provided an appropriate cooling schedule [25] is adhered to. In practice, faster schedules need to be used for tractable computation, and since the ideal cooling schedule is only a sufficient condition, global optima can still be achieved.

We use a closed contour as before, but reformulate the "forces" from section 3.1 as Gibbs functions. This defines a Markov dependency over adjacent control point triplets, allowing the use of SA to optimise over a search range defined on both sides of the initial contour. See Friedland and Adam [9] for a more detailed explanation.

Such a formulation is attractive since it enables prior knowledge to be explicitly stated in probability distributions, and allows flexibility in the form of additional constraints or discriminators imposed on the segmentation process. For example, the probability distribution of intensity levels around each control point[6] is used by Cootes *et al* in an active contour [6] which restricts deformation modes to the most statistically likely ones, derived by principal components analysis (PCA) of a sample of typical boundaries. The use of 3D deformable meshes in the next section allows physics-based mesh constraints to be imposed on the position of each mesh node (control point). Texture measures can also be incorporated as boundary discriminators.

Figure 4 shows a sequence of detected boundaries from scans of a pig's heart suspended in a low noise background. The sequence has been truncated at both ends where the segmentation process does not discriminate fatty tissue from the actual heart boundary.

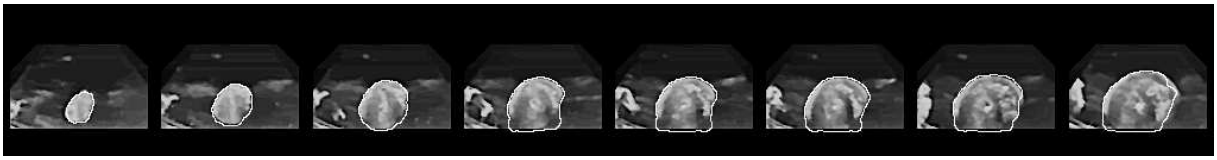

Figure 4: Detecting boundaries in image sequence of pig's heart

For reconstruction purposes, contours in adjacent frames are connected by first corresponding control points from one contour to another. This set of correspondences is then refined by optimising over an energy function comprising curvature match, radial distance between corresponding points and smoothness between adjacent correspondences. Figure 5 shows a surface reconstruction of the pig's heart.

4 3D Segmentation

During manual segmentation, we find that the consultant can sometimes do so based very much on expected position. Figure 6(a) from a thyroid scan has hardly any texture or edge segmentation cues defining the thyroid[7].

[6] "Point Set Distribution"

[7] Bottom of the top left quadrant.

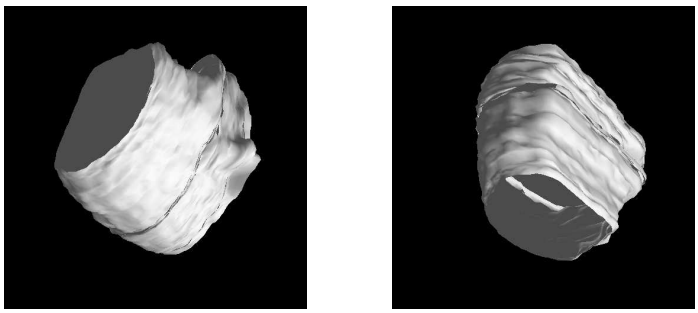

Figure 5: Reconstruction of pig's heart from sequence of detected boundaries

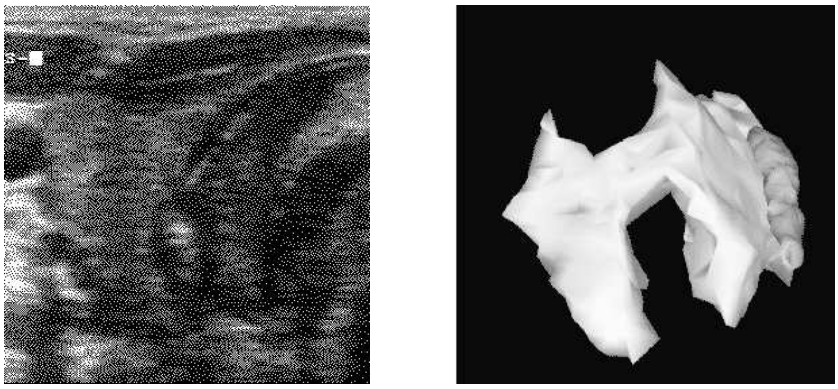

(a) (b)

Figure 6: (a) frame from thyroid scan; (b) thyroid surface model (with carotid on right)

The consultant has a mental model of the structure being imaged, and targets possible modes of variation from model to image, allowing a rough visualisation from the sequence of 2D ultrasound images. This is achieved by matching landmarks between model and image, and orientating the model by rigid-body translation and rotation, as well as performing large-scale (i.e. roughly whole body) deformations of the model. Finer segmentation in the absence of further landmark correspondences is then achieved by deforming the model locally until boundaries "match" in some sense.

We use 3D mesh models consisting of nodal *masses* connected by links with *damping* and *linear elasticity*. The dynamic equations for such a model can be diagonalised, so that nodal quantities (e.g. displacements, forces) in cartesian space, are transformed into modal quantities varying along a set of basis functions representing the vibration modes of the mesh model. This is known as a *modal decomposition* [1].

These dynamic equations can be solved using the *finite element method* (FEM), which defines interpolation functions over the domain of each constituent element of the mesh model. Compared to the technique of *finite differences*, which discretises the mesh and poses sampling problems since it cannot allow for forces to act anywhere other than at the discretisation nodes, FEM can model continuous parameters such as mass, stiffness, and displacement, over the entire model.

The low-frequency[8] vibration modes tend to represent whole-body deformations of the model such as "bending", "pinching", "tapering", as well as the six rigid-body translation and rotation modes in 3D. See Pentland in [20, 21] for further examples. Reasons are also given in [20] why high-frequency modes can be discarded in favour of low-frequency ones, if we wish to restrict the number of degrees of freedom the model has in deformation.

In the absence of a statistical sample, without which we cannot derive a mean model and its statistically important deformation modes using principal components analysis (PCA) [6], the low-frequency vibration modes intrinsic to a mesh model form an interesting alternative set of deformations. We are currently conducting experiments which aim to show that a significant amount of the variance in the sample population can be accounted for using the low-frequency vibration modes of a mesh model of the structure in question [24]. The advantage of being able to derive such a set of deformations from a mesh model, is that we are no longer required to acquire, segment and normalise a statistical sample for PCA. This leads to a significantly greater generality of application, since we only need to build a reasonable mesh approximation to the structure in question[9].

4.1 Model construction

Boissonat and Geiger [3] present a method for interpolating between contours in adjacent parallel planes using Delaunay triangulation. Boissonat *et al* [2] provide a triangulation algorithm which can do this with any number of arbitrarily orientated planes, leading to an efficient method of constructing mesh models from manually segmented contours in 3D ultrasound or MRI. We are currently applying a general 3D Delaunay triangulation algorithm[10] to construct the convex hull of the segmented points, after which smoothing is applied to create a reasonable prior model (Figure 6(b)) [12]. An alternative approach would be to use geometric primitives such as ellipsoids and cylinders to construct a model. Such a model can then be deformed to form a reasonable approximation for use as a prior model (see Footnote 9).

We will also need the ability to build mesh models from segmented volume images or artificially constructed volume models. Schroeder and Shephard [22, 23] apply octree decomposition and 3D Delaunay triangulation to fully automated mesh generation. We are working on a variation of this algorithm which results in a well-formed mesh of either surface triangles or volumetric tetrahedra.

4.2 Segmentation

We are currently investigating the design of a graphical user interface which gives the clinician a volume rendered view of a set of acquired 2D slices. This system would allow the clinician to establish correspondences between the volume image and a mesh model. This set of correspondences can then over-constrain[11] an initial shape recovery procedure by allowing fewer deformation modes than correspondences, leading to a method of volume estimation[12].

---

[8] Low *temporal* frequency during dynamic simulation.

[9] Low order modes are determined by the low-order moments of inertia, and vary little between compact objects with roughly the same overall shape [20].

[10] This does not exploit the fact that segmented points lie on planes, and tends to be computationally inefficient.

[11] Obviating the need for smoothness or symmetry heuristics.

[12] Using Gauss' theorem and numerical integration [24].

Given this initial segmentation, we can refine it by applying 2D segmentation methods along each of the acquired 2D slices, initialising the segmenting contour(s) as the cross sectional boundary of the 3D model intersected with each scan plane. This is a more computationally tractable approach than a true 3D segmentation, and a more valid one since we acquire image data along a set of scan planes and do not have a true volumetric image.

5 Conclusion

We have presented 2D segmentation algorithms which are driven by edge cues alone. These will need to be extended to include texture cues for better and more robust performance in a modality which is characterised by significant noise problems in both imaging and proprioception. The algorithms are highly parallelisable, and the SA algorithm in particular allows for flexible incorporation of segmentation cues and constraints. We also aim to incorporate 3D position constraints by using prior models of the structure being imaged, allowing a measure of robustness against particularly noisy local patches (e.g. Figure 6(a)) which provide no segmentation cues.

Acknowledgements

Many thanks to Roberto Cipolla for insightful comments. J.P.M. Gosling is supported by EPSRC project Stradivarius (GR/H74032). M. Syn is supported by an EPSRC Research Studentship. Both gratefully acknowledge the funding and support of the Engineering and Physical Sciences Research Council.

References

[1] K.J. Bathe. *Finite element procedures in engineering analysis.* Prentice-Hall, 1982.

[2] J-D. Boissonat, A. Cerezo, O. Devillers, and M. Teillaud. Output sensitive construction of the 3d delaunay triangulation of constrained sets of points. Technical report, INRIA, France.

[3] J-D. Boissonat and B. Geiger. Three-dimensional reconstruction of complex shapes based on the delaunay triangulation. Technical Report 1697, INRIA, BP 93–06902 Sophia Antipolis Cedex (France), April 1992.

[4] S.Y. Chen, W.C. Lin, C.C. Liang, and C.T. Chen. Improvement on dynamic elastic interpolation technique for reconstructing 3d objects from serial cross sections. 9(1):71–83, 1990.

[5] L. Cohen and I. Cohen. Finite-element methods for active contour models and balloons for 2-d and 3-d images. 15(11):1131–1147, 1993.

[6] T.F. Cootes, A. Hill, C.J. Taylor, and J. Haslam. The use of active shape models for locating structures in medical images. In H.H. Barrett and A.F. Gmitro, editors, *Proc. Information Processing in Medical Imaging*, volume 511 of *Lecture Notes in Computer Science*. Springer-Verlag, 1991.

[7] R. Cromartie and S.M. Pizer. Edge-affected context for adaptive contrast enhancement. In *Lecture Notes in Computer Science: IPMI91 (511)*, 1991.

[8] J. Feng, W-C. Lin, and C-T. Chen. Automatic left ventricular boundary detection in digital two-dimensional echocardiography using fuzzy reasoning techniques. In *Proc. Biomedical Image Processing*, volume 1245. SPIE, 1990.

[9] N. Friedland and D. Adam. Automatic ventricular cavity boundary detection from sequential ultrasound images using simulated annealing. 8(4):344–353, 1989.

[10] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. 6(6):721–741, 1984.

[11] J.P.M. Gosling, R.W. Prager, and L.A. Berman. Proprioception accuracy in free-hand three-dimensional ultrasound imaging. Technical Report CUED/F-INFENG/TR169, Cambridge University Engineering Department, Trumpington Street, CB2 1PZ.

[12] J.P.M. Gosling, M. Syn, R. Prager, and L.A. Berman. Estimation for thyroid volumes using three-dimensional free-hand ultrasound. In *12th Biosignal 94*, Brno, Czech Republic, June 1994.

[13] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. In *Proc. First International Conference on Computer Vision*, pages 259–268, 1987.

[14] Y. Kim and R.C. Luo. Surface reconstruction based on descriptions of cross-sectional contours. In *SPIE 1260: Sensing and Reconstruction of Three-Dimensional Objects and Scenes (1990)*, pages 191–198, 1990.

[15] T. McInerney and D. Terzopoulos. A finite element model for 3d shape reconstruction and nonrigid motion tracking. In *Proc. Fourth ICCV*, May 1993.

[16] D. Metaxas and D. Terzopoulos. Shape and nonrigid motion estimation through physics-based synthesis. 15(6):580–591, 1993.

[17] C. Nastar and N. Ayache. Non-rigid motion analysis in medical images: a physically based approach. In *IPMI93*, volume 687 of *Lecture Notes in Computer Science*, 1993.

[18] O.S. Odesanya, W.N. Waggenspack, and D.E. Thompson. Construction of biological surface models from cross-sections. *IEEE Trans. Biomedical Eng.*, 40(4):329–334, 1993.

[19] A. Pentland and B. Horowitz. Recovery of nonrigid motion and structure. 13(7):730–742, 1991.

[20] A. Pentland and S. Sclaroff. Closed-form solutions for physically based shape modeling and recognition. 13(7):715–729, 1991.

[21] A. Pentland and J. Williams. Good vibrations: modal dynamics for graphics and animation. *Computer Graphics*, 23(3):215–222, 1989.

[22] W.J. Schroeder and M.S. Shephard. Geometry-based fully automatic mesh generation and the delaunay triangulation. 26:2503–2515, 1988.

[23] W.J. Schroeder and M.S. Shephard. A combined octree/delaunay method for fully automatic 3d mesh generation. 29:37–55, 1990.

[24] M. Syn and R.W. Prager. Model based three-dimensional ultrasound imaging. Technical Report CUED/F-INFENG/TR180, Cambridge University Engineering Department, Trumpington Street, CB2 1PZ, 1994.

[25] P. van Laarhoven and E. Aarts. *Simulated Annealing: Theory and Applications*. Reidel, Dordrecht, The Netherlands, 1987.