Surface interpolation for sparse cross-sections using region correspondence

G.M. Treece*, R.W. Prager, A.H. Gee[a] and L. Berman[b]

[a]Department of Engineering, University of Cambridge, Trumpington Street, Cambridge, UK, CB2 1PZ
[b]Department of Radiology, University of Cambridge, Addenbrooke's Hospital, Cambridge, UK, CB2 2QQ

Abstract. The ability to estimate a surface from a set of cross-sections allows calculation of the enclosed volume and display of the surface in three-dimensions (3-D). However, extracting the cross-sections (segmenting) can be very difficult, and automatic segmentation methods are not sufficiently robust to deal with all situations. Hence, it is an advantage if the surface reconstruction algorithm can work effectively on a small number of cross-sections. An algorithm is presented which can interpolate a surface through sparse, complex cross-sections. This is an extension of *maximal disc guided interpolation* [1], which is itself based on *shape based interpolation* [2]. The performance of this algorithm is demonstrated on various types of medical data (X-ray Computed Tomography, Magnetic Resonance Imaging and three-dimensional ultrasound). Although the correspondence problem in general remains unsolved, it is demonstrated that correct surfaces can be estimated from limited real data, through the use of *region* rather than *object* correspondence.

1 Introduction

Three-dimensional (3-D) data is now available in many medical imaging modalities. One of the ways of displaying such data is by rendering a surface within it. A surface can be estimated by segmenting a set of cross-sections from the data and, if they are sufficiently dense, rendering directly from them. However, in practice, there are many situations where this is not the case. In X-ray Computed Tomography (CT) data, for instance, the resolution of the pixels is typically greater than the inter-plane spacing, and it is desirable to interpolate new data between the planes. In order to minimise interpolation artifacts in subsequent surface renderings, object based interpolation methods are often used, which interpolate the object cross-section rather than the grey scale values.

Segmentation of medical data is non-trivial, and the numerous automatic segmentation techniques are not robust enough for use on generic data. As a result of this, manual segmentation, or computer assisted manual segmentation, is often used instead. Typical applications can require manual segmentation of several hundred planes. Hence, even where the data is sufficiently dense that surface interpolation is not strictly necessary, it can be helpful in reducing the number of planes to be segmented. Freehand 3-D ultrasound is a particularly difficult case, since the original data is acquired on planes which are not parallel, and ultrasound data is notoriously hard to segment.

2 Related work

Object based interpolation methods use the shape of objects as a guide for interpolation of either grey-scale or binary data. In shape-based interpolation [2], a 3-D distance function is created by linearly interpolating the 2-D minimum distance functions for each cross-section. This is thresholded at zero to reveal the interpolated surface. Shape-based interpolation gives good results for similar cross-sections, but performs poorly when the cross-sections are dissimilar, or do not have significant overlap, as in Figure 1(d). This problem has been addressed, for tree-like shapes, by moving the cross-sections to align the centroids of each object [3], as in Figure 1(e). This works well for scans transverse to the main axis, like the top three cross-sections in Figure 1(e), but not when the shape of each cross-section changes dramatically. Scaling each cross-section, such that the bounding rectangles are the same size, can improve the result in some cases [4] — but not in the case of Figure 1(f). A similar approach based on mathematical morphology, which aligns objects using an iterative minimisation of distance transform values, is presented in [5].

The connectivity in all of these algorithms is based on whole contours, but it can be seen from Figure 1(e) and (f) that this is not always appropriate. An alternative approach, dynamic elastic interpolation [6], uses a force field to iteratively deform one cross-section towards the other. The initial formulation produced similar results to centroid-guided interpolation. However, this behaviour was improved by imposing an upper threshold on the forces at each point on the contour, thus preventing far portions of the cross-section having a detrimental effect on the

*Corresponding author: gmt11@eng.cam.ac.uk

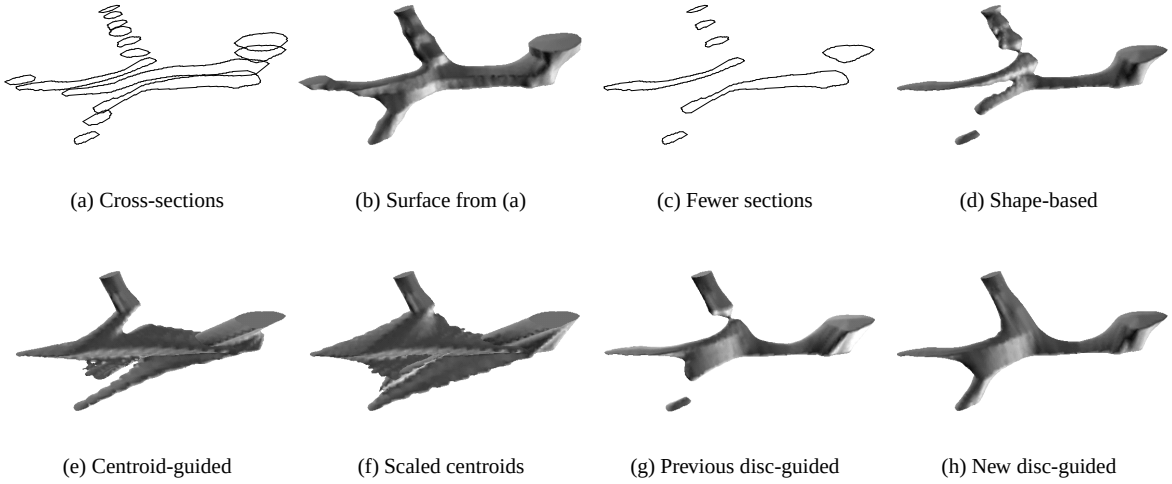

(a) Cross-sections (b) Surface from (a) (c) Fewer sections (d) Shape-based

(e) Centroid-guided (f) Scaled centroids (g) Previous disc-guided (h) New disc-guided

Figure 1. Part of the portal venous system. The cross-sections were manually outlined on the original B-scans of a freehand 3-D ultrasound investigation. (a) and (b) show the original data, (d) to (h) show the surface created from the selected cross-sections shown in (c).

deformation, but also adding further complexity. The result of maximal disc guided interpolation [1] is shown in Figure 1(g). This, although an improvement, is still not a faithful representation of the actual anatomy, Figure 1(b). In addition, in common with the centroid based techniques, it can introduce artifacts. In contrast, the extension presented in this paper gives the surface of Figure 1(h), which correctly represents the actual anatomy.

Algorithms which can create *surfaces from scattered point sets* can also be applied to cross-sections — although this increases the complexity of the problem. Hoppe [7] developed a distance function approach, formed by combining surface normal contributions from each point. This is a very general solution, but as a result it is not well constrained, and many points are required to ensure a correct reconstruction. In freehand 3-D ultrasound, the cross-sections are already non-parallel, and therefore this approach becomes more attractive [8].

3 Maximal disc guided interpolation

The overall strategy of the surface interpolation algorithm is reviewed in Figure 2. Steps (a), (b) and (f) are identical to maximal disc guided interpolation, in [1] (which also examines the additional processing for non-parallel cross-sections). The novel steps (c), (d) and (e) are described in detail in [9]. Maximal discs can be used to exactly represent a binary image: the disc centres are the ridges of the distance transform. We use a sub-set of these discs to loosely represent the object shape. External discs (representing troughs, not ridges) are also extracted — this allows the additional correspondence of holes, concavities and gaps between objects, which improves the results for concave or multiple contours. The set of discs for the cross-sections in Figure 2(a) is shown in Figure 2(c).

The region correspondence of neighbouring cross-sections (i.e. which regions should be connected, and in what direction) is determined by considering the correspondence between each pair of discs. The calculation of correspondence likelihood for each pair of discs is the most important step in this process. It is estimated by comparing, in each plane, the difference in distance transform values at the centre of each disc with the planar distance between the disc centres. This gives an error which, if it is small compared to the radii of each of the discs, is used as the likelihood estimate for this pair of discs. The calculation of this estimate is examined in detail in [9].

The approach has a variety of features which make it attractive. Firstly, discs will only correspond with the nearest discs on the neighbouring cross-section if the nearness relationship is reciprocal. This allows regions to be left unconnected, if appropriate. Secondly, it is *not* necessary for regions (or discs) to overlap in order to correspond. Thirdly, small discs will only tend to have a local effect on correspondence, unless the contour they represent is itself small. Larger discs on the same contour will take priority in the far field. The only limiting assumption is that at least one contour must be connected to one other contour, i.e. at least one object spans the cross-sections. The result of this operation on the cross-sections of Figure 2(a) is shown in Figure 2(d) — note that in this case the external discs have no partners on neighbouring planes, and hence do not contribute to the correspondence.

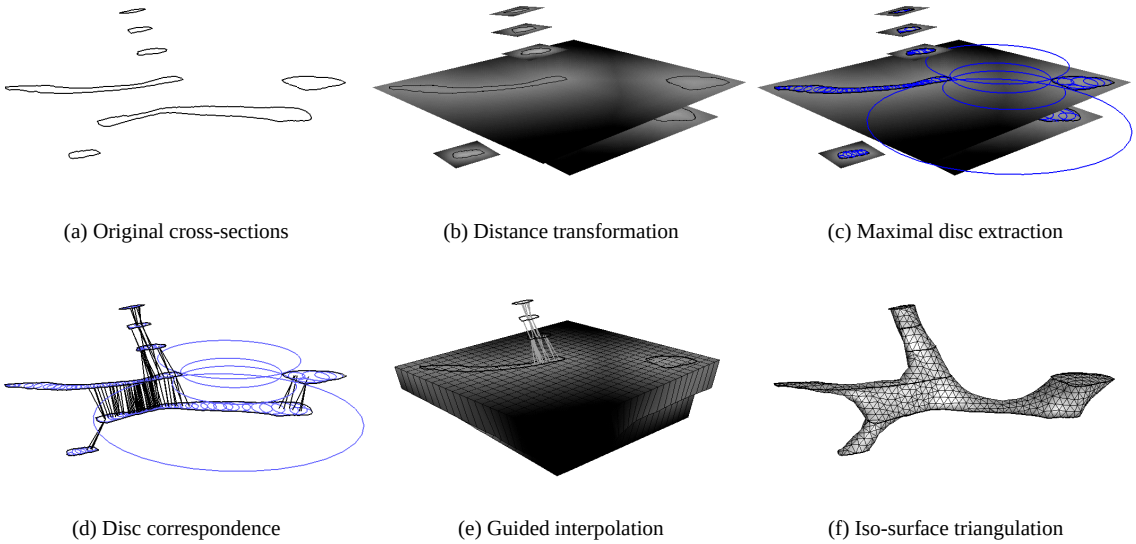

(a) Original cross-sections (b) Distance transformation (c) Maximal disc extraction

(d) Disc correspondence (e) Guided interpolation (f) Iso-surface triangulation

Figure 2. Surface interpolation algorithm. Distance transformation is first performed on each cross-section (b), then a set of representative discs is extracted from these transformations (c). Region correspondence is estimated for each disc (d), and this is combined to give a correspondence direction at any point. This provides the interpolation direction for the distance field (e), which is then thresholded and triangulated to reveal the surface (f).

4 Results

Surfaces were estimated by shape-based interpolation and three derivatives: centroid-guided (a combination of techniques described in [3] and [4]), the previous maximal disc algorithm and the new disc-guided algorithm. Surface visualisation is used here, rather than more rigorous quantitative measures, since it gives more insight into the relative behaviour of the algorithms. All the surfaces were generated by triangulating the zero iso-surface of the interpolated distance fields [9]. Processing was done on a Silicon Graphics Indigo 2 workstation. *In vivo* ultrasound data was recorded using a Toshiba Powervision 7000 ultrasound machine with a Polhemus FASTRAK magnetic field position sensor mounted on the probe. All processing was performed using Stradx v5.3[1] software [10].

The child's skull in Figure 3 is from CT data provided with the 3DViewnix[2] visualisation package. Every sixth scan has been thresholded with the appropriate coefficient for bone to give the cross-sections. The data is not appropriate for the centroid-guided technique in this case, since there are multiple contours which are neither simple nor tree-like. The cross-sections in the upper part of the skull are far from overlapping, which results in the separate rings constructed by shape-based interpolation. The previous disc-guided method improves on this, but is still limited by the condition that the discs must overlap in order for them to correspond. There is no such constraint in the new disc-guided interpolation, and hence the skull is correctly reconstructed by this algorithm.

The human liver in Figure 3 is from MRI data of the male cadaver from the Visible Human Project. Five cross-sections were manually segmented, which is enough information to give a reasonable idea of the shape and volume of the liver. However, it is not enough information for shape-based interpolation, as in Figure 3(f). The disc-guided interpolation performs slightly better than centroid-guided (particularly on the thin upper edge), since the calculation of centroid position is dominated by the bulk of the object: small features on large objects do not contribute and may not be reconstructed correctly. In contrast, disc-guided interpolation uses local information and can handle localised small features.

5 Conclusions

The method of surface interpolation presented in this paper is capable of improving on the results of shape-based interpolation in many cases where there are few cross-sections or the cross-sections vary significantly between planes. In addition, and unlike previous methods, it can be relied upon not to generate additional artifacts in cases

[1] `http://svr-www.eng.cam.ac.uk/~rwp/stradx/`

[2] From `CHILD.IM0`, 3DViewnix v1.1.1 (c) 1993-1996 MIPG Univ. of Pennsylvania, All Rights Reserved.

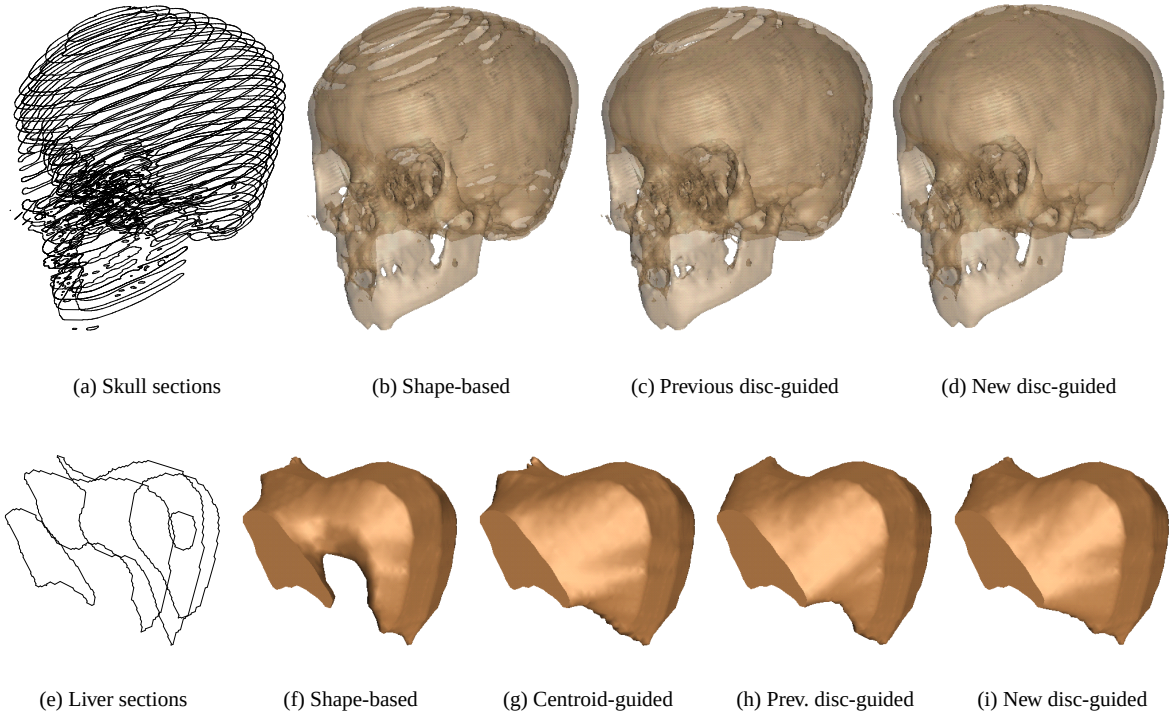

(a) Skull sections (b) Shape-based (c) Previous disc-guided (d) New disc-guided

(e) Liver sections (f) Shape-based (g) Centroid-guided (h) Prev. disc-guided (i) New disc-guided

Figure 3. Results: skull and liver. The skull was segmented by thresholding CT data, and has been rendered slightly transparent, so that both surfaces can be seen. The liver was manually segmented from coronal MRI scans.

where shape-based interpolation is already a good surface estimator. This is an important result, since it indicates that the technique can be safely used on many different types of data. There is a processing overhead (see [9]), but this overhead is small for many practical cases. Even where the overhead is large, the total processing time is still small compared to the time required to segment the data (times varied from 2 seconds for the liver to 2.5 minutes for the skull). In addition, the ability of this algorithm to interpolate sparse cross-sections reduces the number which need to be segmented, thus potentially *reducing* the time from scanning to 3-D display.

References

1. G. M. Treece, R. W. Prager, A. H. Gee et al. "Fast surface and volume estimation from non-parallel cross-sections, for freehand 3-D ultrasound." *Medical Image Analysis* **3(2)**, pp. 141–173, 1999.
2. G. T. Herman, J. Zheng & C. A. Bucholtz. "Shape-based interpolation." *IEEE Computer Graphics and Applications* pp. 69–79, May 1992.
3. W. E. Higgins, C. Morice & E. L. Ritman. "Shape-based interpolation of tree-like structures in three-dimensional images." *IEEE Transactions on Medical Imaging* **12(3)**, pp. 439–450, 1993.
4. M. Loughlin, I. Carlbom, C. Busch et al. "Three-dimensional modelling of biopsy protocols for localized prostate cancer." *Computerized Medical Imaging and Graphics* **22**, pp. 229–238, 1998.
5. Y.-H. Liu, Y.-N. Sun, C.-W. Mao et al. "Edge-shrinking interpolation for medical images." *Computerized Medical Imaging and Graphics* **21(2)**, pp. 91–101, 1997.
6. S.-Y. Chen, W.-C. Lin, C.-C. Liang et al. "Improvement on dynamic elastic interpolation technique for reconstructing 3-D objects from serial cross sections." *IEEE Transactions on Medical Imaging* **9(1)**, pp. 71–83, 1990.
7. H. Hoppe, T. DeRose, T. Duchamp et al. "Surface reconstruction from unorganized points." *Computer Graphics* **26(2)**, pp. 71–78, July 1992.
8. M. E. Legget, D. F. Leotta, E. L. Bolson et al. "System for quantitative three-dimensional echocariography of the left ventricle based on a magnetic-field position and orientation sensing system." *IEEE Transactions on Biomedical Engineering* **45(4)**, pp. 494–504, April 1998.
9. G. M. Treece, R. W. Prager, A. H. Gee et al. "Surface interpolation from sparse cross-sections using region correspondence." Technical Report CUED/F-INFENG/TR 342, Cambridge University Engineering Dept, March 1999.
10. R. W. Prager, A. H. Gee & L. Berman. "Stradx: real-time acquisition and visualisation of freehand 3D ultrasound." *Medical Image Analysis* **3(2)**, pp. 129–140, 1999.