
**NARROW-BAND
VOLUME RENDERING FOR
FREEHAND 3D ULTRASOUND**

A. H. Gee, R. W. Prager,
G. M. Treece and L. Berman

CUED/F-INFENG/TR 392

September 2000

University of Cambridge
Department of Engineering
Trumpington Street
Cambridge CB2 1PZ
England

Email: ahg/rwp/gmt11@eng.cam.ac.uk, lb@radiol.cam.ac.uk

Narrow-Band Volume Rendering for Freehand 3D Ultrasound

Andrew Gee, Richard Prager, Graham Treece and Laurence Berman*
University of Cambridge

Department of Engineering
Trumpington Street
Cambridge CB2 1PZ

*Department of Radiology
Addenbrooke's Hospital
Cambridge CB2 2QQ

Abstract

Volume rendering, the projection of volumetric intensity data into a 2D image, is finding an increasing number of applications across a diverse range of scientific disciplines. A key ingredient in the volume rendering recipe is a procedure to allow the user to specify a limited volume of interest, since rendering the entire volume is often too slow and also leaves the image full of clutter that obscures genuine structures of interest: this is particularly important for 3D ultrasound, where speckle noise has traditionally frustrated the successful application of volume rendering. This paper describes novel techniques, tuned to the specific needs of 3D ultrasound, for specifying these target volumes. The basic idea is to extend the familiar planar and non-planar reslicing tools, such that the reslice plane or surface is endowed with a certain thickness, producing a *narrow-band* volume. The attraction of such volumes is that they are simple to set up, limit the amount of clutter in the final renderings and, when thin, produce reslice-like images that are far less sensitive to the precise position of the reslice plane (or surface) than those generated by standard reslice algorithms. The paper gives full implementation details of both planar and non-planar narrow-band volume rendering, and describes several practical applications of the resulting tools.

1 Introduction

Conventional diagnostic ultrasound imaging is performed with a hand-held probe which transmits ultrasound pulses into the body and receives the echoes. The magnitude and timing of the echoes are used to create a 2D greyscale image (B-scan) of a cross-section of the body in the scan plane. **Freehand 3D ultrasound imaging** [11, 14] extends this technique by constructing 3D data sets from a series of 2D B-scans — see Figure 1.

A conventional 3D freehand examination can be broken into three stages: scanning, reconstruction and visualisation. Before scanning, some sort of position sensor is attached to the probe. This is typically the receiver of an electromagnetic device [2, 5, etc.], as illustrated in Figure 1, although alternatives include acoustic spark gaps [6], mechanical

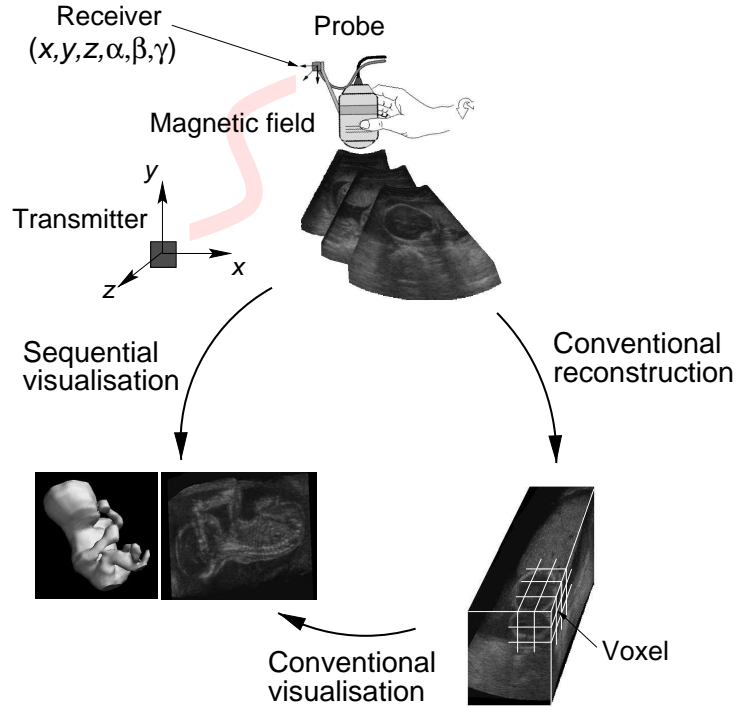


Figure 1: **Freehand 3D ultrasound imaging.** The conventional approach is a three-stage process, involving scanning, reconstruction and visualisation. The alternative, sequential approach allows visualisation directly from the raw B-scans and positions.

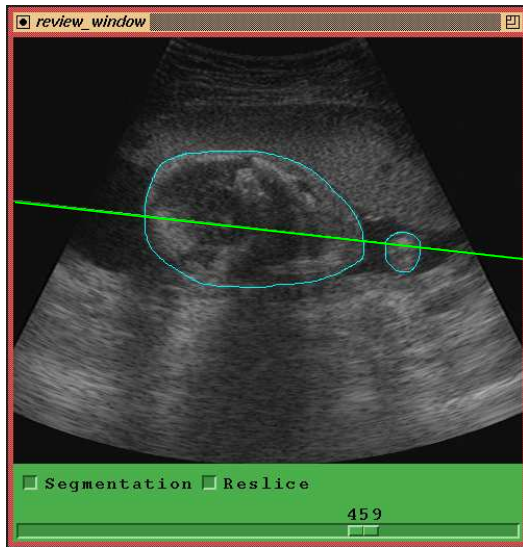
arms [7] and optical sensors [17]¹. Measurements from the position sensor are used to determine the positions and orientations of the B-scans with respect to a fixed datum, usually the transmitter of the electromagnetic device. In the next stage, the set of recorded B-scans and their relative positions are used to fill a regular voxel array. Finally, the array is visualised using, for example, planar reslicing or surface rendering (after segmentation).

Recently, we have proposed an alternative approach to freehand 3D ultrasound, which bypasses the voxel array stage. In **sequential** freehand 3D ultrasound, the data is visualised and analysed directly from the raw B-scans and positions — see Figure 1. Planar and non-planar reslicing, panoramic imaging, volume estimation and surface rendering can be performed without the use of an intermediate voxel representation [4, 8, 9, 10, 15, 16]: all of these facilities are implemented in the Stradx freehand 3D ultrasound system². The sequential approach offers many, significant advantages, as discussed in [4, 16, 18].


Figure 2 illustrates a typical obstetrics application of the Stradx system. The ‘Review’ window shows a selected B-scan, onto which the user has drawn two segmentation contours; the ‘Reslice’ window shows a particular planar reslice of the data set; the ‘Manifold’ window

¹More recently, some progress has been made in determining the relative positions of the B-scans from the ultrasound images themselves [19], thus obviating the need for a position sensor.

²The software can be downloaded from <http://svr-www.eng.cam.ac.uk/~rwp/stradx/>.




'Review' window



'Reslice' window



'Manifold' window



'Outline' window

Figure 2: **3D ultrasound visualisation using Stradx.** This example shows an examination of a 22-week foetus. The 'Review' window displays a selected B-scan, onto which the user has drawn two segmentation contours. The 'Reslice' window shows a planar reslice, which intersects the selected B-scan along the straight line in the 'Review' window. The vertical line in the 'Reslice' window corresponds to the B-scan in the 'Review' window. The user has drawn a non-planar reslice contour onto the 'Reslice' window. The 'Manifold' window shows the resulting non-planar reslice, which runs along the foetus's spine. The 'Outline' window shows the outlines of all the recorded B-scans and also the positions of the segmentation contours, reslice plane and non-planar reslice surface.

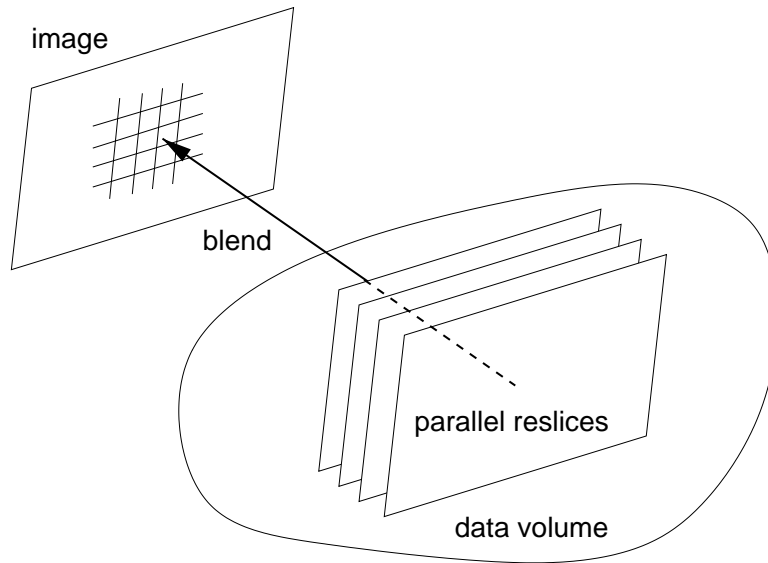


Figure 3: **Volume rendering.** A volume rendering can be generated by blending together a number of parallel reslices. Different blending functions produce different effects in the final image: for example, simple averaging produces a semi-transparent rendering with all data elements given equal prominence, while maximum intensity compounding (projecting only the brightest data into the image) highlights bright reflectors like bone. Changing the orientation of the reslice and image planes produces a rendering from a different viewpoint. A number of renderings from neighbouring viewpoints can be spliced together to make a movie: this is a particularly good way to appreciate 3D structure within the data volume.

shows a particular non-planar reslice of the data set; while the ‘Outline’ window shows the outlines of all the recorded B-scans and also the positions of the segmentation contours, reslice plane and non-planar reslice surface.

Clinical users of the Stradx system have identified two shortcomings of the visualisation interface in Figure 2. First, it is often necessary to position the reslice plane (or surface) very precisely in order to see small structures of interest. Second, the inherent two-dimensional nature of a reslice makes it difficult to appreciate the relative positions and orientations of key structures: it is not always the case that the anatomy intersected by the reslice plane (or surface) provides sufficient geometrical context.

The latter shortcoming is traditionally addressed by **volume rendering** [3], a visualisation technique which projects the entire data volume into a two-dimensional image: see [12, 20] for recent surveys. One way to construct a volume rendering is to take a set of parallel reslices and blend them into a single image, as illustrated in Figure 3. While volume rendering is fairly commonplace for visualising CT and MRI data (especially now that graphics hardware can generate renderings within acceptable timescales [1]), its application to 3D ultrasound is less widespread. This is because of the low signal-to-noise ratio of ultrasound images: the data volume is dominated by speckle which obscures key

anatomical features in typical volume renderings. The speckle cannot be easily removed by global filtering without also removing small-scale structures of interest: instead, the user must somehow define a limited region of interest [12], perhaps using a ‘virtual scalpel’ to cut away portions of the data volume, leaving only a small, pertinent region for rendering.

In this paper, we present two techniques with a virtual scalpel flavour tuned to the specific requirements of 3D ultrasound. In essence, the new techniques are conceptually straightforward extensions of planar and non-planar reslicing. The reslice plane or surface is given some thickness, producing a narrow, wafer-like volume, which is then volume rendered in the usual way. For modest thicknesses, this solves the aforementioned problem of having to position the reslice surface precisely in order to see small structures of interest. For thicker, planar reslices, the result is a more conventional volume rendering, but with a rapid and straightforward setup procedure. The practical implementation of narrow-band, non-planar volume rendering is challenging, and forms the main theoretical contribution of this paper.

2 Narrow-band volume rendering

Planar volume rendering

Narrow-band, planar volume rendering is a straightforward, though effective, extension of planar reslicing. A number of parallel, planar reslices are computed and subsequently blended together to produce the volume rendering, as illustrated in Figure 3. The spacing between the individual reslice planes is set to the resolution cell size of the B-scans (the width of one pixel), while the number of reslices depends on the user-specified thickness of the volume rendering. The user can also choose between maximum intensity compounding, minimum intensity compounding and simple averaging. The maximum and minimum intensity options are good for highlighting bright reflectors and fluid-filled cavities respectively, while simple averaging produces X-ray like images of the selected volume. Figures 4 and 5 show the effects of maximum intensity and average compounding respectively.

When compounding colour Doppler data, we need to define what we mean by the ‘brightest’ or ‘darkest’ pixel. For this purpose, we consider any coloured pixel (red or blue) to be brighter than any grey pixel. Where we need to choose between red and blue pixels of the same intensity, we arbitrarily consider the red pixel to be brighter. Figure 6 shows a maximum intensity volume rendering of a colour Doppler data set.

When average compounding is applied to Doppler data, and at least one of the pixels to be averaged is coloured, then the averaging is applied to the underlying fluid velocities and not the pixel intensities. So grey, very dark red and very dark blue pixels are all associated with zero velocity fluid, bright red pixels are associated with high, positive velocity fluid, and bright blue pixels are associated with high, negative velocity fluid. The velocities are averaged and then converted back to an appropriate colour for display on the screen.

The time taken for narrow-band planar volume rendering depends on the number of thin reslices compounded together, and is therefore directly proportional to the selected

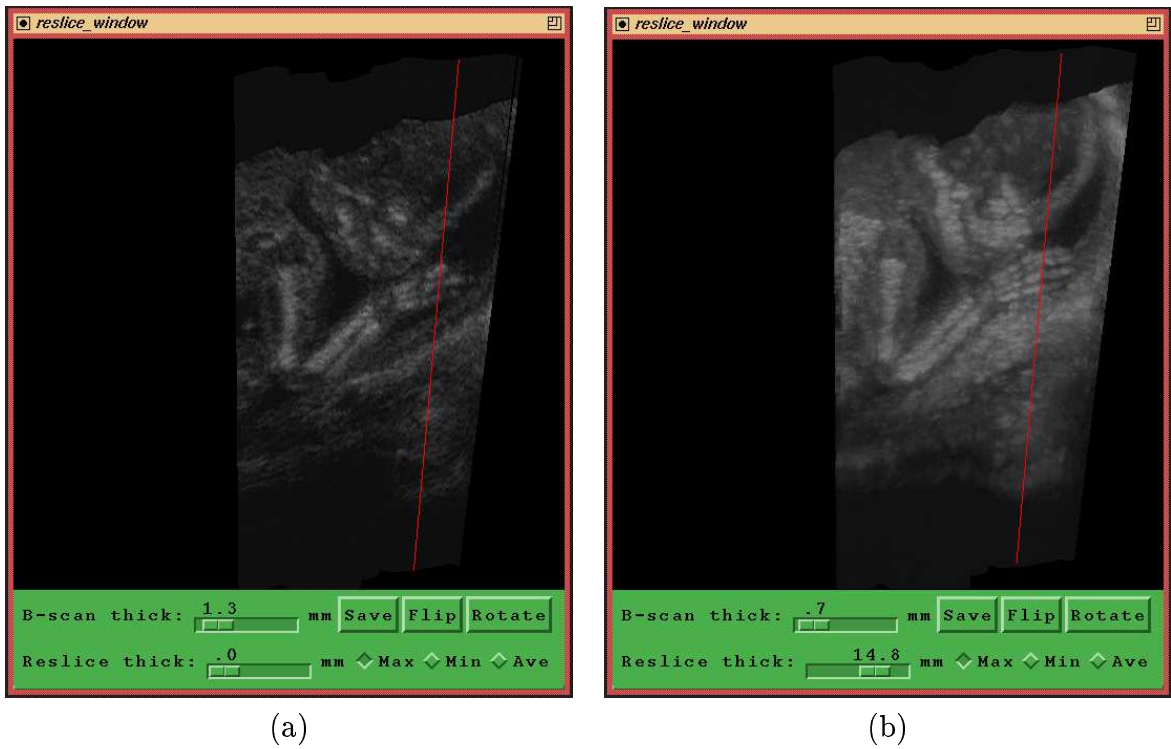


Figure 4: **Maximum intensity compounding.** Compare the standard, thin reslice (a) with the 14mm thick reslice (b). The bone structure of the foetus's hand is much more clearly visible in (b). Moreover, it was not necessary to take particular care in positioning the reslice plane.






Figure 5: **Average compounding.** Compare the standard, thin reslice (a) with the 13mm thick reslice (b). The averaging has reduced the level of speckle noise and revealed the structure of the foetus's spine and ribcage.



(a)



(b)

Figure 6: **Maximum intensity Doppler compounding.** The ‘Outline’ window in (a) shows the outlines of the B-scans, a set of automatically-generated segmentation contours around the blood vessels, and the limits of the 20mm narrow-band volume rendering. The volume rendering itself is in (b): note the bifurcations in the artery (red), and the vein (blue) hidden behind the artery. Also note the white artefact caused by the Doppler region-of-interest displayed on the ultrasound machine’s screen. To give some idea of the scale of the reconstruction, the volume of the visible blood vessels is approximately 0.7ml.




Figure 7: **Non-planar reslicing.** The non-planar reslice image shows the intersection of the 3D data with a special class of developable surface, unrolled for display on a flat screen.

thickness of the band. Using the efficient, sequential reslicing algorithm described in [8, 10], volume renderings like those in Figures 4–6 can be produced in about 10 seconds on modest computer hardware. Since no voxel array is required, renderings like these are available immediately after scanning.

Non-planar volume rendering

Far more challenging from a theoretical perspective is narrow-band, non-planar volume rendering. We motivate this topic by first reviewing non-planar *reslicing*. Figure 7 illustrates the general idea: a **developable** (unrollable) surface is specified by the user and its intersection with the 3D ultrasound data is computed, so that the surface can be ‘painted’ with the data it intersects. The painted surface is then unrolled for display on a flat screen. The particular class of developable surface considered here is defined by the plane π with unit normal \mathbf{n} and the plane curve $\mathbf{c}(s)$ lying in π — see Figure 7. The surface is swept out by the set of **rulings** \mathcal{R} which are all of length $2r$, are all parallel to \mathbf{n} and intersect $\mathbf{c}(s)$ at their midpoints. The distance between the rulings is set to the resolution cell size of the B-scans (the width of one pixel), so each ruling maps to a row of pixels in the unrolled reslice image.

The Stradx implementation of non-planar reslicing is illustrated in Figure 2. The current reslice plane corresponds to π . The user draws the curve $\mathbf{c}(s)$ in the ‘Reslice’ window and can adjust the distance r by resizing the resulting ‘Manifold’ window. Details of the sequential algorithm used by Stradx to compute the non-planar reslice can be found in [4].




Figure 8: **Generating non-planar narrow-band volume renderings.** (a) shows the original curve $\mathbf{c}(s)$ and its immediate neighbours, $\mathbf{c}_u(s)$ and $\mathbf{c}_d(s)$. The distance d between adjacent curves is set to the resolution cell size of the B-scans (the width of one pixel). The rulings associated with $\mathbf{c}(s)$ (which come out of the page towards the reader) are marked as filled circles. For each ruling, we find the closest points on $\mathbf{c}_u(s)$ and $\mathbf{c}_d(s)$: these points define new rulings which will be compounded with those on $\mathbf{c}(s)$ to create the narrow-band volume rendering. The correspondences between rulings are shown as narrow lines in (a). (b) shows the surfaces defined by the three sets of rulings, while (c) shows the compounding process used to produce the final volume rendering.

Here, we extend this work to give the developable surface some thickness. The resulting narrow-band volume is unrolled and rendered using one of the standard compounding options. As with the planar case, the motivation for doing this is to make the precise positioning of π and $\mathbf{c}(s)$ less critical, and to reveal some of the surrounding anatomy which is not intersected by the reslice surface.

Non-planar narrow-band volume rendering

The general strategy for non-planar narrow-band volume rendering is to define a series of adjacent developable surfaces, each with the same number of rulings, and calculate a non-planar reslice for each surface. The resulting set of reslice images can then be compounded together in the usual way — see Figure 8. At this point, we should note that defining a suitable set of surfaces is equivalent to finding a suitable set of plane curves $\mathbf{c}(s)$ in π , since the curves, along with the dimension r , define the surfaces. Given the family of curves, nearest neighbour search can be used to obtain a set of rulings on the new curves in direct correspondence with those on the original curve, as illustrated in Figure 8(a). The families




Figure 9: **Contour propagation strategies.** (a) shows a portion of the curve $\mathbf{c}(s)$ that defines the non-planar reslice. (b) illustrates a naive strategy for generating narrow-band volume renderings: the curve is sampled at regular intervals, lines are drawn normal to $\mathbf{c}(s)$ through the sample points, and new curves $\mathbf{c}_u(s)$ and $\mathbf{c}_d(s)$ constructed by joining the end points of the lines. Unfortunately, this simple strategy results in self-intersecting curves, which in turn lead to self intersecting surfaces and confusing narrow-band volume renderings. A better solution is shown in (c): $\mathbf{c}(s)$ is propagated using the normal flow equations to form $\mathbf{c}_u(s)$ and $\mathbf{c}_d(s)$. Note the creation of a ‘shock’ in $\mathbf{c}_u(s)$ and two ‘rarefaction fans’ in $\mathbf{c}_d(s)$.

of rulings are painted with the data they intersect, then compounded together to form the narrow-band volume rendering. Note that, as a result of the curvature of $\mathbf{c}(s)$ and the nearest neighbour compounding, the rulings on $\mathbf{c}_d(s)$ are spaced further apart than those on $\mathbf{c}(s)$, which in turn are spaced further apart than those on $\mathbf{c}_u(s)$. Since the rulings map to rows of pixels in the non-planar reslice images, it follows that anatomical structures lying on $\mathbf{c}_d(s)$ will appear squashed compared with those lying on $\mathbf{c}_u(s)$. This is an inevitable side effect of nearest neighbour compounding of curved surfaces.

We have thus far omitted a discussion of how to generate the curves $\mathbf{c}_u(s)$ and $\mathbf{c}_d(s)$ given $\mathbf{c}(s)$. Figure 8(a) suggests a simple strategy: for each ruling on $\mathbf{c}(s)$, we draw a line perpendicular to $\mathbf{c}(s)$ with length $2d$ and midpoint lying on $\mathbf{c}(s)$. The two paths swept out by the endpoints of this line define $\mathbf{c}_u(s)$ and $\mathbf{c}_d(s)$. Figure 9(b) reveals the flaw in this appealingly simple strategy. At tight corners in $\mathbf{c}(s)$, the line perpendicular to $\mathbf{c}(s)$ will cut back across its own path and $\mathbf{c}_u(s)$ will be self-intersecting. Consequently, structures intersected by $\mathbf{c}_u(s)$ will be repeated at disjoint locations in the narrow-band volume rendering, producing an extremely confusing image.

Figure 9(c) illustrates a more principled way to construct $\mathbf{c}_u(s)$ and $\mathbf{c}_d(s)$. We consider $\mathbf{c}(s)$ to be a wavefront propagating under normal flow, so that each point on the wavefront moves normal to the front with constant speed s . We generate the family of curves $\mathbf{c}_u(s)$ and $\mathbf{c}_d(s)$ by considering wave propagation in both the ‘upwind’ and ‘downwind’ directions.

Robust algorithms to compute $\mathbf{c}_u(s)$ and $\mathbf{c}_d(s)$ are nontrivial but have been developed to maturity by researchers in fluid dynamics, combustion, image processing and computer graphics: [13] offers an excellent description of the state-of-the-art. The entropy-satisfying weak solution to the normal flow problem can be computed using either level set methods or fast marching methods [13]. The entropy-satisfying weak solution avoids ‘swallowtails’, where the front passes through itself as in Figure 9(b), creating instead ‘shocks’, as in Figure 9(c). For the downwind flow, sharp corners spawn circular arcs known as ‘rarefaction fans’.

Observe in Figure 9(c) how the same ruling on $\mathbf{c}_u(s)$ is repeated several times in the narrow-band volume rendering. This is a direct consequence of shocks in the entropy-satisfying weak solution to the normal flow equations. Consequently, anatomical structures lying on this ruling appear smeared out in the rendering³. This is an unavoidable artefact which the user must learn to recognise and interpret, as with the countless other artefacts inherent to both 2D and 3D ultrasonic imaging. The alternative, swallowtail solution in Figure 9(b) results in the far more confusing artefact of a structure being *repeated* at several, disjoint locations in the rendering.

A fast algorithm for approximate normal flow

For this particular application, we do not need the full entropy-satisfying weak solution, since we are less interested in the wavefronts themselves than in the set of rulings lying on the wavefronts. Recall that these rulings are the nearest neighbours on $\mathbf{c}_u(s)$ and $\mathbf{c}_d(s)$ to the original set of rulings on $\mathbf{c}(s)$. Consider now the downwind contour $\mathbf{c}_d(s)$ in Figure 9(c), and in particular the two rarefaction fans. Since points on the rarefaction fans are the nearest neighbours to the sharp corners of $\mathbf{c}(s)$, and there is a vanishingly small probability that a ruling will be exactly coincident with a corner, we do not need to go to the trouble of computing the rarefaction fans. In contrast, however, we certainly do need to compute the shocks, as demonstrated in Figure 9(c).

The fast algorithm used by Stradx to compute sufficient approximations to $\mathbf{c}_u(s)$ and $\mathbf{c}_d(s)$ is illustrated in Figure 10. For each vertex of $\mathbf{c}(s)$, we start by computing the vertex normals, which are the perpendicular bisectors of the angles at the vertices: these are shown as dotted lines in Figure 10. The vertex normals of the two end vertices are simply the normals to the single edge incident at each vertex. As the wavefront advances, its vertices move along the corresponding vertex normals, which can be thought of as a ‘route map’ for the flow. We next locate any points where the vertex normals intersect (the filled circles in Figure 10): at these points, the vertices of the propagating contour meet to form a new shock, which subsequently moves along the path defined by the new vertex normal at this point, as shown in Figure 10.

Once the route map of vertex normals has been computed, the upwind and downwind contours can be discovered by stepping the vertices of $\mathbf{c}(s)$ along the map. To generate

³In understanding this point, the reader may find it helpful to imagine a ‘V’-shaped narrow-band volume made out of rubber and straightened out by force. If the centerline of the ‘V’ is to remain undistorted, then the point at the apex of the concavity must be stretched out over a finite distance.




Figure 10: **Fast, approximate normal contour propagation.** The diagram shows the original contour $c(s)$ (thick, solid line), the ‘route map’ of vertex normals (dotted lines) and the shock creation points (filled circles). Also shown are two propagated contours in both the upwind and downwind directions (thin, solid lines), one pair immediately adjacent to $c(s)$, the other pair more distant. Note the incorrect propagation of an exaggerated protrusion to the right of the diagram (and also at every downwind vertex, though less obviously so). These errors do not, however, affect the eventual sets of rulings that define the non-planar reslice surfaces.

the next contour in the sequence, the vertices are advanced a distance $d/\sin\theta$ along the dotted lines, where d is the resolution cell size of the B-scans (the width of one pixel) and θ is the angle between the vertex normal and the incident edges. This algorithm generates a correct solution to the normal wave propagation equations, except that rarefaction fans are replaced by sharp corners, an extreme example of which is seen at the right of Figure 10. These errors have no effect on the final volume rendering, since the nearest neighbour procedure used to generate rulings along $c_u(s)$ and $c_d(s)$ will ignore all points on the offending corners.

Example

A typical application is illustrated in Figures 11–13. The non-planar narrow-band volume rendering in Figure 12(b) shows a remarkable image of a 16-week foetus’s unrolled skeleton. For comparison, Figure 11 shows the corresponding thin reslices, and Figure 13 shows one

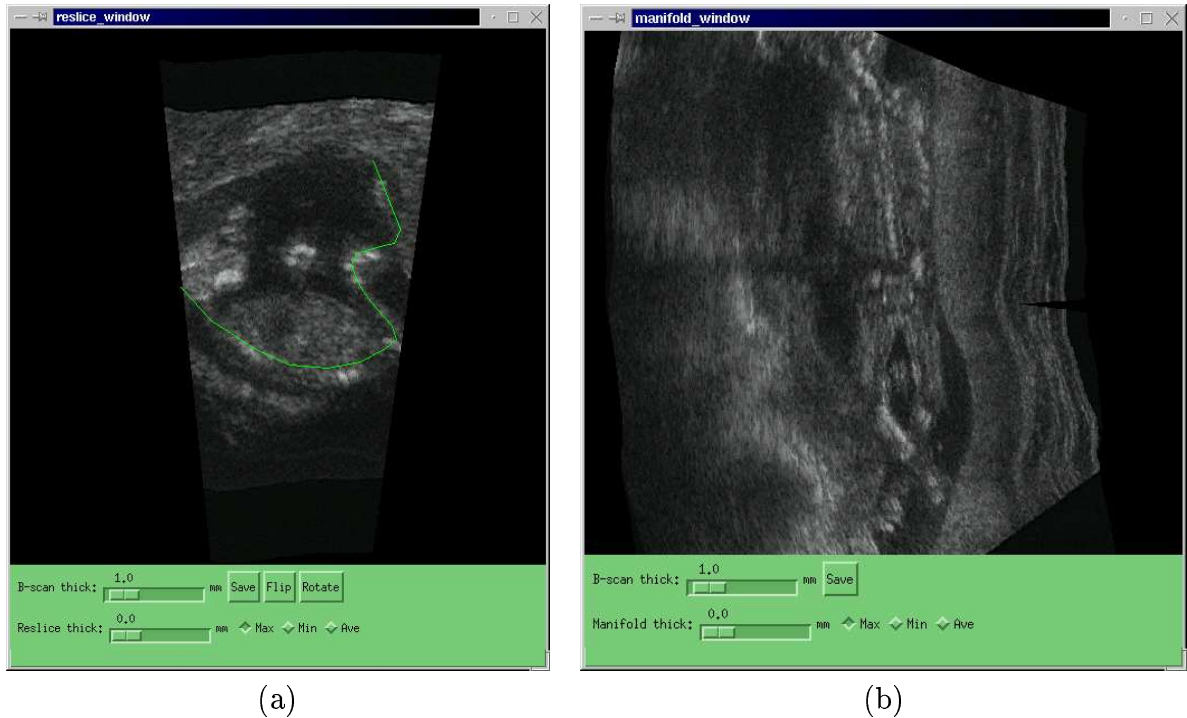


Figure 11: **Planar and non-planar reslicing.** The ‘Reslice’ window (a) shows a planar reslice through a 16-week fetus’s torso and legs. The user has drawn the contour $c(s)$ through the spine, femur, tibia and fibula, and foot. The resulting non-planar reslice, showing the entire skeleton unrolled, is displayed in the ‘Manifold’ window (b). Note, however, that many of the bones are incomplete or missing, since the reslice surface does not intersect them precisely.

of the raw B-scans used to construct the renderings. The smearing artefact previously referred to is evident in Figure 12(b) at locations where the surface curvature is high, but its effect is minor and the rendering is by no means unintelligible.

The renderings in Figure 12 were available approximately two minutes after completion of the ultrasound examination. This includes the time taken by the user to define the reslice plane and non-planar reslice contour. The time required to compute the non-planar narrow-band volume rendering was approximately 10 seconds on a modest laptop PC. The rendering time is, of course, proportional to the thickness of the narrow-band volume.

3 Conclusions and further work

This paper has described conceptually simple but technically challenging extensions to planar and non-planar reslicing. What emerged was a narrow-band volume-rendering technique which significantly facilitates the review of 3D ultrasound data sets: reslice planes and surfaces no longer need to be positioned with pinpoint accuracy, and more anatomical

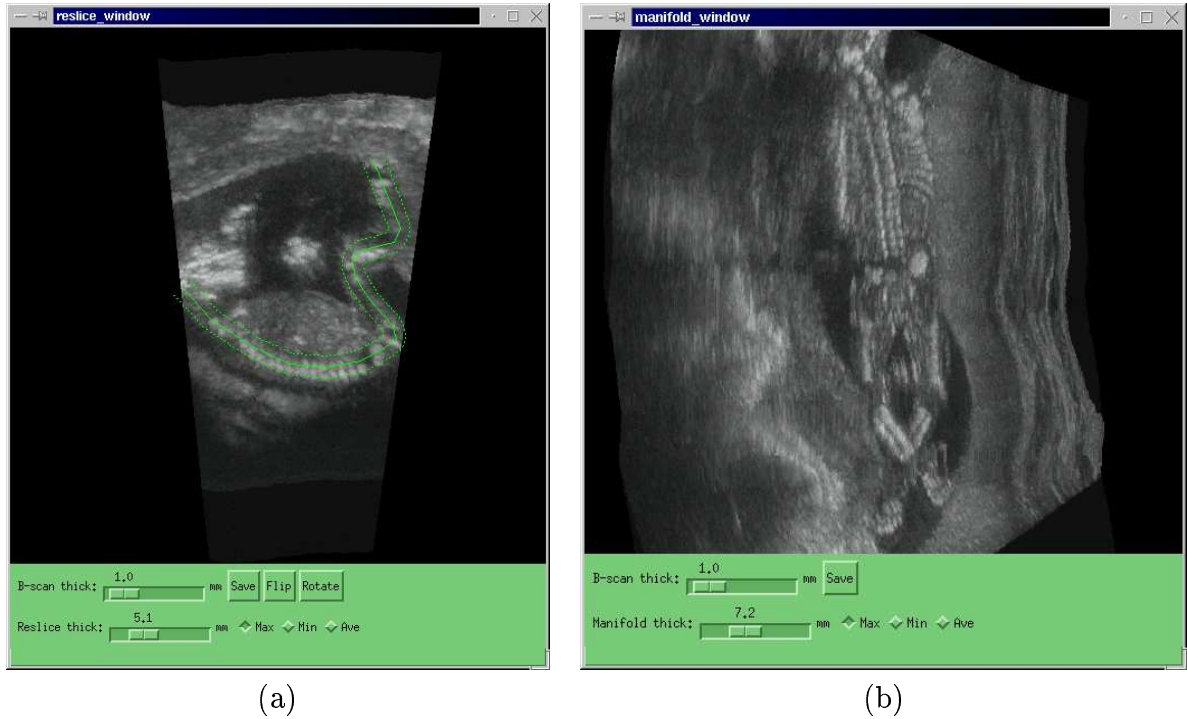



Figure 12: Narrow-band planar and non-planar volume rendering. This example is identical to Figure 11, except the planar reslice thickness has been set to 5.1mm and the non-planar reslice thickness has been set to 7.2mm. Maximum intensity compounding was used to highlight bone in both the resulting narrow-band volume renderings. Compared with Figure 11(a), note how much more bone is visible in (a) to help the user draw the contour $c(s)$. The non-planar narrow-band volume rendering in (b) shows an almost complete picture of the unrolled skeleton, including the spine, ribcage, pelvis, both femurs, both fibulas and tibulas, and both feet. A small section of umbilical cord is visible between the femurs, though neither arm is contained within the narrow-band volume. Note also the extreme upwind and downwind contours $c_u(s)$ and $c_d(s)$ displayed in (a).



(a)



(b)

Figure 13: **Raw data used for Figures 11 and 12.** The ‘Review’ window (a) shows one of the raw B-scans used to construct the renderings in Figures 11 and 12. For navigation purposes, landmarks, numbered 1–3, have been placed on one foot, lower leg and near the base of the spine. For cross-reference, these same landmarks are also shown in the thick planar reslice (b).

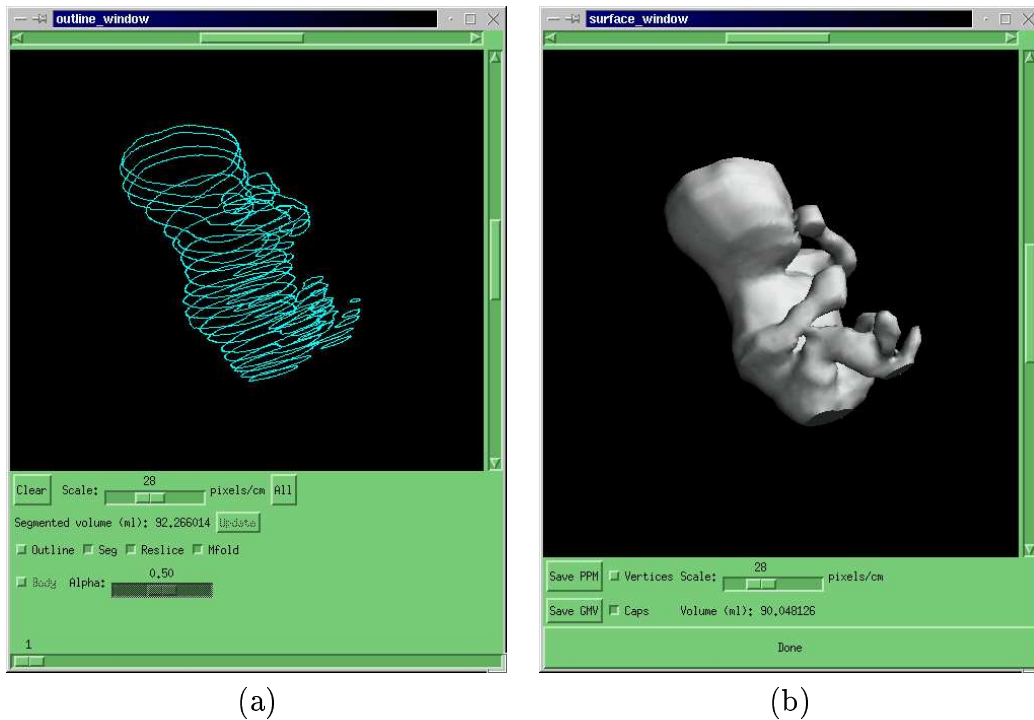


Figure 14: **Contour-based segmentation.** The user has outlined a sparse set of contours on the original B-scans (a), which is interpolated to produce a surface rendering (b). Using these tools, the volume of the fetus can be estimated to within $\pm 5\%$ [15].

structure can be observed in a single image.

Future work in this area will seek to generate volume renderings from contour-based segmentations. For example, Figure 14 shows existing Stradx facilities for segmentation, surface rendering and volume estimation. A natural extension would be to develop efficient, sequential volume rendering algorithms to visualise structures *within* the segmented object. The user would be able to interactively change the viewpoint of the volume rendering, in order to better appreciate the rendered anatomy.

Acknowledgements

The authors would like to thank Diana Galletly and Patrick Gosling for permission to publish the data in Figures 2, 4, 5, 11, 12, 13 and 14.

References

- [1] B. Cabral, N. Cam, and J. Foran. Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. In *Proceedings of the 1994 Symposium*

- on *Volume Visualization*, pages 91–98, Washington, DC, October 1994.
- [2] P. R. Detmer, G. Bashein, T. Hodges, K. W. Beach, E. P. Filer, D. H. Burns, and D.E. Strandness Jr. 3D ultrasonic image feature localization based on magnetic scanhead tracking: in vitro calibration and validation. *Ultrasound in Medicine and Biology*, 20(9):923–936, 1994.
 - [3] R. Drebin, L. Carpenter, and P. Hanrahan. Volume rendering. *Computer Graphics*, 22(4):65–74, August 1988.
 - [4] A. Gee, R. Prager, and L. Berman. Non-planar reslicing for freehand 3D ultrasound. In *Medical Image Computing and Computer-Assisted Intervention — MICCAI'99*, pages 716–725, Cambridge UK, September 1999. LNCS 1679, Springer.
 - [5] S. W. Hughes, T. J. D'Arcy, D. J. Maxwell, W. Chiu, A. Milner, J. E. Saunders, and R. J. Sheppard. Volume estimation from multiplanar 2D ultrasound images using a remote electromagnetic position and orientation sensor. *Ultrasound in Medicine and Biology*, 22(5):561–572, 1996.
 - [6] D. L. King, D. L. King Jr., and M. Y. Shao. Evaluation of in vitro measurement accuracy of a three-dimensional ultrasound scanner. *Journal of Ultrasound in Medicine*, 10:77–82, 1991.
 - [7] R. Ohbuchi, D. Chen, and H. Fuchs. Incremental volume reconstruction and rendering for 3D ultrasound imaging. In R. A. Robb, editor, *Proceedings of Visualization in Biomedical Computing*, SPIE 1808, pages 312–323. International Society of Optical Engineering, Bellingham, WA, USA, 1992.
 - [8] R. W. Prager, A. H. Gee, and L. Berman. 3D ultrasound without voxels. In *Proceedings of Medical Image Understanding and Analysis*, pages 93–96, Leeds, July 1998.
 - [9] R. W. Prager, A. H. Gee, and L. Berman. Real-time tools for freehand 3D ultrasound. In *Medical Image Computing and Computer-Assisted Intervention — MICCAI'98*, pages 1016–1023, Cambridge MA, October 1998. LNCS 1496, Springer.
 - [10] R. W. Prager, A. H. Gee, and L. Berman. Stradx: real-time acquisition and visualization of freehand three-dimensional ultrasound. *Medical Image Analysis*, 3(2):129–140, 1999.
 - [11] R. N. Rankin, A. Fenster, D. B. Downey, P. L. Munk, M. F. Levin, and A. D. Vellet. Three-dimensional sonographic reconstruction: techniques and diagnostic applications. *American Journal of Roentgenology*, 161(4):695–702, 1993.
 - [12] W. Schroeder, K. Martin, and B. Lorensen. *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*. Prentice Hall, London, 2nd edition, 1998.

- [13] J. A. Sethian. *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision and Materials Science*. Cambridge University Press, Cambridge, UK, 1999.
- [14] H. Steiner, A. Staudach, D. Spitzer, and H. Schaffer. Three-dimensional ultrasound in obstetrics and gynaecology: technique, possibilities and limitations. *Human Reproduction*, 9(9):1773–1778, 1994.
- [15] G. M. Treece, R. W. Prager, A. H. Gee, and L. Berman. Fast surface and volume estimation from non-parallel cross-sections, for freehand 3-D ultrasound. *Medical Image Analysis*, 3(2):141–173, 1999.
- [16] G. M. Treece, R. W. Prager, A. H. Gee, and L. Berman. Volume measurement of large organs with 3D ultrasound. In *Proceedings of SPIE Vol. 3982, Medical Imaging 2000 (Ultrasonic Imaging and Signal Processing)*, pages 2–13, San Diego, CA, February 2000.
- [17] J. W. Trobaugh, D. J. Trobaugh, and W. D. Richard. Three-dimensional imaging with stereotactic ultrasonography. *Computerized Medical Imaging and Graphics*, 18(5):315–323, 1994.
- [18] P. M. Tuomola, A. H. Gee, R. W. Prager, and L. Berman. Body-centered visualisation for freehand 3D ultrasound. *Ultrasound in Medicine and Biology*, 26(4):539–550, June 2000.
- [19] T. A. Tuthill, J. F. Krücker, J. B. Fowlkes, , and P. L. Carson. Automated three-dimensional US frame positioning computed from elevational speckle decorrelation. *Radiology*, 209(2):575–582, 1998.
- [20] J. K. Udupa. Three-dimensional visualization and analysis methodologies: a current perspective. *Radiographics*, 19(3):783–806, May 1999.