Surface Reconstruction
in
3D Medical Imaging

Jonathan Carr BE (Hons 1)

A thesis presented for the degree of
Doctor of Philosophy
in
Electrical and Electronic Engineering
at the
University of Canterbury,
Christchurch, New Zealand.

February 1996

ABSTRACT

This thesis addresses two problems in medical imaging, the development of a system for 3D imaging with ultrasound and a system for making titanium prostheses for cranioplasty. Central to both problems is the construction and depiction of surfaces from volume data where the data is not acquired on a regular grid or is incomplete.

A system for acquiring 3D pulse-echo ultrasound data using a conventional 2D ultrasound scanner equipped with an electro-magnetic spatial locator is described. The non-parallel nature of 2D B-scan slices acquired by the system requires the development of new visualisation algorithms to depict three dimensional structures. Two methods for visualising iso-valued surfaces from the ultrasound data are presented. One forms an intermediate volume reconstruction suitable for conventional ray-casting while the second method renders surfaces directly from the slice data. *In vivo* imaging of human anatomy is used to demonstrate reconstructions of tissue surfaces. Filtering and spatial compounding of scan data are used to reduce speckle.

Pulse-echo ultrasound primarily depicts tissue boundaries. These are characterised by incomplete acoustic interfaces contaminated by noise. The problem of reconstructing tissue interfaces from ultrasound data is viewed as an example of the general problem of reconstructing an object's shape from unorganised surface data. A novel method for reconstructing surfaces in the absence of *a priori* knowledge of the object's shape, is described and applied to 3D ultrasound data. The method uses projections through the surface data taken from many viewpoints to reconstruct surfaces. Aspects of the method are similar to work in computer vision concerning the determination of the shape of 3D objects from their silhouettes. This work is extended by considering the reconstruction of incomplete objects in the presence of noise and through the development of practical algorithms for pixel and voxel data. The reconstruction of realistic, non-convex objects is considered. 2D and 3D ultrasound data derived from phantoms, as well as artificial data, are used to demonstrate reconstructions.

The second problem studied in this thesis concerns designing cranial implants to repair defects in the skull. Skull surfaces are extracted from X-ray CT data by ray-casting iso-valued surfaces. A tensor product B-spline interpolant is used in the ray-caster to reduce ripples in the surface data due to partial voluming and the large spacing between CT slices. The associated surface depth-maps are characterised by large irregular holes which correspond to the defect regions requiring repair. Defects are graphically identified by a user in surface-rendered images. Radial basis function approximation is introduced as a method of interpolating the surface of the skull across these defect regions. The fitted surface is used to produce CNC milling instructions to machine a mould in the shape of the surface from a block of hard plastic resin. A cranial implant is then formed by pressing flat titanium plate into the mould under high pressure in a hydraulic press. The system improves upon current treatment procedures by avoiding the manual aspects of fashioning an implant. It is also suitable when other techniques which use symmetry to reconstruct the skull are inadequate or not possible. The system has been successfully used to treat patients at Christchurch Hospital. Radial basis function (RBF) approximation has previously been restricted to problems where the number of interpolation centres is small. The use of newly developed fast methods for evaluating radial basis interpolants in the surface interpolation software results in a computationally efficient system for designing cranial implants and demonstrates that RBFs are potentially of wide interest in medical imaging and engineering problems where data does not lie on a regular grid.

ACKNOWLEDGEMENTS

I would like to thank John and Molly Andreae. John for his valued wisdom throughout the course of my research and particularly for his hard work and criticism during the writing of this thesis. Molly for her relief from the aforementioned hard work and criticism. Dr. Richard Fright for introducing all these interesting problems and trying to make an engineer out of me — hopefully some of his 'can do' attitude has rubbed off. Dr. Kathy Garden for initially co-supervising me and introducing me to the Medical Physics group at Christchurch Hospital. Dr. Rick Beatson, initially unappreciated when I was an undergraduate student, but greatly valued in the course of my post-graduate research when all that maths suddenly appeared useful. Dr. Philip Bones for keeping medical imaging alive in the department and overseeing the completion of this thesis.

I am especially grateful to Dr. Nigel Anderson and Richard Allen in the Radiology Department of Christchurch Hospital for their cooperation and enthusiasm for the 3D ultrasound work.

I am indebted to the late John Hinton who pioneered titanium cranioplasty at Christchurch Hospital and without whom the work in chapter 6 would not have come about. Also, I wish to acknowledge Michael Parker who made the titanium plates and neurosurgeons Dr's McFarlane and Bonkowski. Furthermore, the staff of the Medical Physics Department of Christchurch Hospital who provided a happy and informative working environment and supported the work presented in this thesis.

I would like to acknowledge the financial contribution of Telecom Corporation of New Zealand in the form of a post-graduate scholarship and New Zealand Lottery Health Research who funded computer equipment for the 3D ultrasound project. Also, the Acuson corporation who provided additional equipment and assistance.

A general thank-you goes to the staff, technicians and students of the Department of Electrical and Electronic Engineering at Canterbury University for being both able colleagues and sources of entertainment and amusement. In particular, Helen Devereux and for her organisation of departmental rafting trips and climbing expeditions. Robert Van Nobelen (Dr. Bob) a fellow colleague and originator of many (as yet unsuccessful) get-rich-quick schemes.

Research does not occur in a vacuum. Thus, a final thanks goes to the members of the Canterbury University Tramping Club for many memorable trips, in particular, Catherine Brown for motivating the pursuit of 3000m peaks and carrying the rope.

CONTENTS

PREFACE

A successful final year project with Dr. John Andreae initiated my interest in postgraduate research and introduced me to applications of engineering in medicine [Carr 1991]. The project involved replicating the KARDIO expert system developed by Ivan Bratko [Bratko *et al.* 1988] for diagnosing cardiac arrhythmia, disorders of the heart's electrical system. However, rather than being a study of the heart's electrical system, the aim of the work was to investigate Bratko's model-based approach to constructing an expert system. This system was interesting because it used simulation of the heart's electrical system to automatically derive a knowledge database and machine-learning to solve the diagnostic problem.

This thesis describes research carried out by the author between 1992-95. The path which led to the final content of this thesis began with a study of the human visual system. This work, under Dr. John Andreae's guidance, involved a computational study of the lower visual cortex. In particular, investigating log-polar models of retino-cortical mapping [Wilson 1983, Wilson 1982], Marr's theory of vision [Marr 1982, Marr and Hildreth 1980] and later, Michael Leyton's theory of shape [Leyton 1992].

The search for a link between the pixel world and high-level visual concepts proved elusive and Marr's framework was found to be unsatisfying. However, Michael Leyton's theory of shape was stimulating and provocative. A suitable application was sought to focus the work. At the time, Dr. Kathy Garden's recently acquired CT scanner provided a convenient application for biologically-based image processing techniques: the analysis of X-ray CT data.

Searching for a means to display the results of 3D segmentation procedures applied to CT data sets proved an enjoyable distraction. Initially wire meshes and shaded polygonal meshes were used, but the desire for more realistic depictions, with better detail, led to the development of ray-tracing and volume rendering software. The problem of interpolating ray-traced depth-maps containing discontinuities provided further distraction and proved to be an addictive problem. The algorithms resulting from this work are not reported in this thesis but aspects of this research manifest themselves in Chapter 6 where interpolation of incomplete depth-maps is used to design titanium cranial implants.

A common interest in 3D graphics led to a meeting with Dr. Richard Fright at the Department of Medical Physics at Christchurch Hospital. This proved to be a turning point in my research. Richard introduced me to the problems of 3D ultrasound and titanium cranioplasty which, in hindsight, have unified many aspects of my diverse interests.

A seminar by Dr. Rick Beatson from the Department of Statistics and Mathematics at Canterbury University introduced me to radial basis function approximation. His fast evaluation methods were recognised as being relevant to a number of problems in medical imaging. The titanium cranioplasty problem seemed ideally suited and a happy and successful collaboration with Rick resulted.

A fascination with vision, the process of seeing and its role in cognition has persisted throughout my research. Remnants of this interest can be seen in Chapter 5 in the algorithms for reconstructing an object from incomplete surface data. Here multiple projections through the data are used to reconstruct an object. The inspiration for this novel approach was the apparent ability of people to mentally reconstruct 3D objects from incomplete data by viewing the data from multiple viewpoints. For example, structures which appear ambiguous when depicted in a single static view using semi-transparent rendering methods often become clear when multiple views are presented in an animated sequence. The algorithms presented in this chapter can be seen as a computational mechanism which attempts to mimic this ability. Consequently the method shares limitations with those of the human viewer concerning how far 3D shapes can be understood from 2D views. Although reconstructions from real ultrasound data of phantoms are presented, the work presented in Chapter 5 might have more significance when viewed in this light rather than as a practical method for reconstructing tissue boundaries from clinical ultrasound data. Due to multiple objects, noise and artifacts, the clinical ultrasound reconstruction problem is a very difficult one representing a worst case scenario when compared to analogous problems in computer vision where reconstructions of 3D objects are required from stereo range data or laser surface scanners.

The construction of a *clinically useful* 3D ultrasound system involves many aspects including development of acquisition hardware, spatial reconstruction algorithms, signal processing, artifact reduction, visualisation methods and tools for quantitative exploration of the data. In this thesis it has not been possible to give each of these problems the attention they deserve, consequently, the thesis has attempted to focus on a few of these problems independent of the others, in particular, shape reconstruction for quantitative measurements and surface visualisation.

The visualisation software, shape reconstruction and ultrasound acquisition software were all written by the author. Software for the solution and evaluation of radial basis functions (RBFs) were also written by the author and code for the fast evaluation of RBF interpolants was provided by Dr. Rick Beatson.

During the course of this research the following papers have been either published, submitted or presented at conferences:

[1] K. E. Surman, J. C. Carr, Y. Brichieri-Colombi, W. R. Fright and K. L. Garden, A System for Three Dimensional Graphics from Ultrasound, *Proceedings of the First New Zealand Conference on Image and Vision Computing*, Auckland, New Zealand, pp367-373, August 1993.

[2] W. R. Fright, N. A. Oien, R. N. Tremewan, J. W. Hinton, J. C. Carr and R. K. Beatson, Automated Production of Titanium Prostheses for Cranioplasty at Christchurch Hospital, *Joint Conference of Engineering & Physics in Medicine and IEEE Engineering in Medicine and Biology Society*, Queenstown, New Zealand, November 20-24, 1995.

[3] J. C. Carr and W. R. Fright, Shape reconstruction for 3D ultrasound, *IEEE Transactions Pattern Analysis and Machine Intelligence*, Submitted March 1995.

[4] J. C. Carr, W. R. Fright and R. K. Beatson, Surface Interpolation with Radial Basis Functions for Medical Imaging, *IEEE Transactions on Medical Imaging*, provisionally accepted February 1996.

GLOSSARY OF TERMS

The terms in this glossary are defined in the context of volume visualisation and diagnostic ultrasound. The general and formal definition is not necessarily reflected in the definition given here.

A-scan: a pulse-echo scan where the ultrasound beam points in one direction and received echoes are filtered, envelope-detected and displayed as vertical deflections of the display trace at locations corresponding to the interfaces which produced the echoes.

B-scan: a pulse-echo scan in which the ultrasound beam is swept across a plane and the received echoes are filtered, envelope-detected and presented as grey tone spots on a display screen at positions corresponding to those of interfaces in the scan-plane which produced the echoes.

real-time B-scanning: the acquisition and display of B-scan images at frame-rates greater than 5 per second.

depth-map: see *Z-buffer*.

depth-only shading: a shading method in which the shading of a surface varies as a function of distance from the image plane.

dynamic focusing: rapidly moving the focus along the ultrasound beam axis to coincide with the range of received echoes.

M-scan: a pulse-echo scan where the ultrasound beam points in one direction and the echoes received from interfaces are plotted against time to show variations of the positions of the interfaces.

partial voluming effect: the case where a voxel value represents a weighted integral over the voxel volume. The value reflects the average property of the material mixture within the voxel. Chapter 3, Section 3.1.1.2 (X-ray computed tomography); Chapter 4, Section 4.6.2 (ultrasound).

ray-casting: a viewing algorithm in which sight rays are cast from the image plane through the volume. The tracing of rays stops when the visible voxels are determined by accumulating or encountering an opaque value.

ray-tracing: a viewing algorithm in which light behaviour is simulated by recursively tracing imaginary rays of light through the scene.

resectioning: a simple viewing algorithm which displays 2D cross-sections taken along a user-defined axis through the 3D data volume.

resonance: occurs when the vibration of a structure is enhanced by constructive interference of internally reflected waves.

reverberation: occurs when multiple reflections arising from a single interface are falsely reconstructed as interfaces at distances corresponding to the times of the received secondary echoes.

shadowing: signal-drop out arises when a high density material such as bone totally absorbs transmitted ultrasound waves, leaving structures further from the transducer in shadow.

splatting: a viewing algorithm in which each visible voxel is projected on to an image plane as a cloud of points (a footprint) that spreads the voxel's contribution [Westover 1989]. It is used by Ohbuchi and Fuchs [1990] and Nelson and Elvins [1993] for semi-transparent rendering of 3D ultrasound data. Chapter 3, Section 3.1.1.2; Chapter 4 Section 4.2.

surface rendering: the indirect visualisation of volume data by first converting them into an intermediate surface representation and then using conventional computer graphics techniques to render them [Kaufman 1991].

ultrasound: vibrations of frequency greater than the upper limit of the human audible range. Longitudinal travelling waves of frequencies in the 0.5MHz-30MHz range are used in diagnostic ultrasound.

ultrasound pulse: a burst of ultrasonic vibrations.

volume rendering: the direct visualisation of volume data without any intermediate conversion of the volumetric dataset to a surface representation.

Z-buffer: the table of distances from image pixels to visible voxels in the volume dataset.

Z-buffer gradient shading: a volume shading technique in which the surface gradient is approximated from the depth (distance) values of the visible voxels in a small 2D neighbourhood.

Chapter 1

INTRODUCTION

This thesis is concerned with two practical problems in medical imaging: three dimensional (3D) imaging with ultrasound and titanium cranioplasty — the process of forming cranial implants to repair defects in the skull. Common to both these problems is the reconstruction and depiction of surfaces from two dimensional (2D) images. In the cranioplasty problem the images are X-ray computed tomography (CT) scans which depict X-ray attenuation in a plane through the body. The images are acquired as a stack of parallel slices and collectively represent a 3D distribution which is related to tissue density. In the ultrasound problem the images are pulse-echo B-scans which represent the scattering of ultrasound by tissue in a plane through the body. Unlike the CT scan slices, the ultrasound slices are typically not parallel.

Both problems involve fitting surfaces to incomplete surface data. The parallel nature of CT slices mean that conventional ray-tracing techniques can be applied to render an iso-valued surface corresponding to the skull. The surfaces extracted from the CT data are considered incomplete where a defect or a hole occurs in the skull. The problem is to reconstruct the natural shape of the skull across the defect region. Tissue surfaces depicted in the ultrasound data are incomplete on two accounts. Firstly, the non-parallel nature of slices means that some areas are undersampled relative to others. Secondly, the acoustic interfaces depicted in the B-scan images are contaminated by noise and artifacts. Constructive and destructive interference effects result in noise known as 'speckle' while shadowing and attenuation artifacts result in signal dropout. Tissue boundaries therefore need to be reconstructed from incomplete acoustic interfaces which have been sampled irregularly.

The primary contributions of this thesis are methods for displaying iso-value surfaces from non-parallel slices (Chapter 4), a new method for reconstructing an object from incomplete surface data (Chapter 5) and a new method for interpolating across irregularly shaped holes in bone surfaces (Chapter 6). In addition, the novel application of B-splines to the ray-casting viewing algorithm is described in Chapter 3 and used to extract cranial bone surfaces from CT data. Chapter 2 describes the system for acquiring 3D ultrasound data, independently developed in the course of this research.

Although the surface reconstruction and interpolation algorithms in Chapters 4, 5 and 6 are illustrated with *in vivo* ultrasound and CT data, the problems of fitting surfaces to non-uniform and incomplete data are of wider interest, particularly in the fields of scientific visualisation of multidimensional data and computer vision. The ultrasound and cranioplasty problems are among many applications, not necessarily associated with medical imaging, which share common characteristics. The algorithms developed in Chapter 5 are particularly relevant to recent theoretical work in computer vision concerning the reconstruction of objects from silhouettes and shadows [Laurentini 1995, Yuan 1995, Murch 1990].

1.1 ORGANISATION OF THESIS

The thesis is organised as follows. The remainder of this chapter introduces the motivation for 3D ultrasound imaging and describes basic principles of 2D diagnostic ultrasound imaging. Some knowledge concerning basic computer graphics, wave phenomena and computed tomography is assumed. However, a detailed understanding of diagnostic ultrasound is not necessary to understand the main contributions presented in this thesis. Section 1.3 describes the general characteristics and limitations of conventional 2D pulse-echo imaging which are relevant to the development of the 3D imaging system. A glossary of relevant terms used throughout this thesis is also provided.

Chapter 2 reviews 3D ultrasound systems and describes the approach of spatially registered B-scans. It details the system built in the course of this research and describes its limitations. Conclusions are drawn concerning the advantages and disadvantages of this approach.

Chapter 3 introduces the ray-casting paradigm for rendering iso-valued surfaces from regular 3D arrays of data. It distinguishes between volume rendering techniques, commonly used for visualising ultrasound data, and the surface rendering techniques used in this thesis. Tensor product B-splines are introduced as an improved method for interpolating scalar data lying on a regular grid which is especially suited to rendering cranial bone surfaces from CT data.

Chapter 4 reviews methods for rendering 3D ultrasound data and describes the method developed for mapping non-parallel slice data on to a regular grid so that conventional ray-casting methods may be applied. A direct method for rendering surfaces from non-parallel slices as they are acquired in real-time is also presented. Filtering, necessary for depicting meaningful iso-valued surfaces from ultrasound data, is detailed and results are presented for *in vivo* imaging of human anatomy. 3D artifacts are discussed.

The depiction of iso-valued surfaces is limited to imaging particular structures. In Chapter 5 the problem of reconstructing incomplete tissue interfaces from ultrasound data is considered. The problem is viewed as an example of the general problem of reconstructing an object from incomplete surface data. An original algorithm for reconstructing isolated objects, based on projections and shadows, is introduced and used to reconstruct ultrasound phantoms.

Chapter 6 introduces the problem of titanium cranioplasty and describes the system developed at Christchurch Hospital for making titanium implants. Radial basis function approximation is introduced as a novel solution to the problem of interpolating the surface of the skull across holes and defect regions. Implants made in the shape of the interpolant are used to restore the shape of the skull. The approach is illustrated with reconstructions from real CT data and newly developed fast evaluation methods are employed to improve the computational aspects of the method.

Chapter 7 presents conclusions and discusses future work.

1.2 DIAGNOSTIC ULTRASOUND

Ultrasound is the term used to describe the transmission of mechanical vibrations through matter when the frequency of oscillations is greater than 18kHz, the highest audible frequency. For medical applications ultrasound is taken to mean longitudinal waves of frequency ranging from 0.5MHz up to 30MHz. Although ultrasound has medical applications in surgery and therapy, this thesis is concerned only with diagnostic ultrasound, in particular pulse-echo systems. In these systems, images of soft tissues are produced by transmitting ultrasound into the body and detecting echoes produced by reflection at tissue boundaries.

Diagnostic ultrasound has several favourable properties compared to other medical imaging modalities. It is regarded as being relatively safe [Woodcock 1979, Dunn 1991], it involves no ionising radiation and most examinations are non-invasive and do not distress the patient [McDicken 1991]. It is also safe for the operator, consequently examinations may be readily repeated. Ultrasound does not require the special facilities which X-ray CT and magnetic resonance imaging (MRI) require and portable instruments are available. Ultrasound allows soft tissues, which are difficult to depict by conventional X-ray techniques, to be imaged in detail. Unlike other tomographic techniques, ultrasound offers interactive visualisation of underlying anatomy in real time and has the ability to image dynamic structures within the body. Blood flow can be recorded using the Doppler effect and, in some instances, measured quantitatively. Needles and catheters can also be directed under ultrasonic guidance.

1.2.1 Motivation for 3D ultrasound

Currently, much pathology is readily diagnosed with conventional 2D ultrasound equipment. However, complex 3D anatomy is difficult even for specialists to visualise. The development of 3D ultrasound offers several additional benefits over conventional 2D ultrasound imaging. These include the visualisation of complicated anatomy, the ability to reformat data to obtain views from orientations previously not possible, and the ability to make quantitative measurements concerning organ size and volume. There is also the potential to reduce some of the artifacts present in 2D scans, such as *shadow*, *speckle* and *reverberation* (cf. Section 1.3). Improvement of the signal to noise ratio in the data through spatial compounding is also possible with the 3D imaging technique described in Chapter 2.

The vast majority of '3D ultrasound' systems, including the one presented in this thesis, use directed, focused ultrasound beams to generate 3D images. Because these systems do not insonify the whole volume it might be argued that this is not true 3D imaging and the term '3D imaging with ultrasound' is more appropriate. Despite this distinction, the term '3D ultrasound' is used to describe these systems in the literature. However, it is important to note that the same B-scan imaging principle, described in Section 1.3.1 for conventional 2D imaging, is being used to generate 3D images.

1.3 2D PULSE-ECHO ULTRASOUND

In Chapter 2 3D images are formed from an ensemble of 2D *B-scans* obtained from a real-time sector scanner. The input to the system is the grey-level video image produced by the ultrasound scanner. Because it is not possible to access the signals from the scanner at a lower level of processing, the performance of the 3D system is limited by the quality of the 2D scans. Artifacts which affect the 2D scans are therefore relevant to subsequent 3D reconstructions.

The working details of conventional pulse-echo ultrasound scanners vary from system to system. Although these affect the resolution, slice thickness, focal zones and general system performance, these variations are of limited importance to the algorithms and reconstruction methods presented in this thesis. More important are the basic characteristics and artifacts which scanners share in common. These are described in depth elsewhere: Wells [1977] provides a general reference covering the basic physics, signal processing and safety of pulse-echo systems while McDicken [1991] presents material relevant to the clinical application of modern 2D pulse-echo systems. A detailed reference to the interactions between ultrasound and biological tissues can be found in Hussey [1975], and Dunn [1991] surveys studies on the safety of diagnostic ultrasound.

1.3.1 Principle of pulse-echo imaging

In a pulse-echo imaging system, a burst of vibrations is transmitted into the body via a piezoelectric transducer which converts applied electrical voltages to mechanical vibrations. Because the reverse effect is also observed in piezoelectric materials the same device can be operated as a pressure transducer. In most systems the tasks of transmitting and receiving are implemented by the same transducer. The transducer is housed in a hand-held probe. A short burst of ultrasound is transmitted into the body through physical contact with the probe. After transmitting a short burst of oscillations, the transducer operates in receive mode. After a period of time, sufficient for returning echoes to die out, a second burst can be transmitted. Echoes arising from inhomogeneities in the medium are recorded by the transducer and then filtered and envelope-detected by the scanner's electronics. The time of flight for received echoes is used to reconstruct the position of acoustic interfaces. In most systems, a constant speed of sound is assumed and used to calculate the distance of interfaces from the transducer. Invariably a narrow, focused beam is directed into the body. If refraction and diffraction effects are neglected, and a straight ray path is assumed for ultrasound propagation, the direction and range of a structure relative to the transducer can be determined.

The received echo information can be displayed in a number of ways. The three commonly used display methods are the A-scan, B-scan and time-position scan or M-scan.

1.3.2 A-scan

In the A-scan display an ultrasonic beam is directed in a single direction into the body. Reflections are envelope detected, rectified and displayed as vertical deflections along a horizontal trace. The size of the vertical displacement is a measure of the echo amplitude, hence the term amplitude-modulated timebase (A-scan). The position of the deflection along the horizontal trace is a measure of the time taken for the echo to return. A-scans are useful for accurate measurement of dimensions.

1.3.3 B-scan

In the B-scan display a brightness-modulated timebase is used to display echoes. The ultrasound beam is swept across a plane through the body and the received echoes are filtered, envelope-detected and presented as grey tone spots on a display screen at positions corresponding to those of interfaces in the scan-plane which produced the echoes. Early B-scan devices used a mechanical arm to determine the position and orientation of the probe as it was manually manipulated in the scan-plane. Real-time B-scanners have now superseded these systems.

1.3.3.1 Real-time B-scan

In a real-time B-scan an ultrasonic beam is swept across the scan-plane by mechanical or electrical means from a single probe position. A complete sweep across the imaging plane produces a frame. Each frame consists of an ensemble of regularly spaced lines, displayed relative to the probe. Interpolation and compounding of successive sweeps are usually used to fill-in gaps between adjacent lines in the image. The term 'real-time' refers to the capability to generate 2D B-scan images at frame-rates greater than 5/s [McDicken 1991]. Due to the high velocity of sound in soft tissue, it is possible to image at frame-rates up to 40/s with a real-time probe.

The region in the scan-plane imaged by the probe is termed the *field of view*. The field of view is determined by the ultrasound penetration depth and by the number, and spacing between, lines in the image. The field of view, the number of lines in a single frame and the frame-rate are interrelated. The product of these three terms is a constant. Thus, imaging a greater number of lines reduces the frame-rate for the same field of view. Similarly, increasing the penetration depth reduces the frame-rate since it takes longer for pulses to propagate from and return to the transducer. Other factors such as compounding of multiple scans, dynamic focusing and 2D image filtering also reduce the frame-rate.

Real-time B-scan probes can be realised in a number of ways. Electro-mechanical devices move a single or multiple transducers to produce each line in the image. Transducers which use arrays of piezoelectric elements either excite groups of elements in the array to produce particular lines in the image or use the principle of a phased array to electronically steer a beam through the use of delay lines on the array elements.

This thesis is concerned with real-time B-scans. The term 'B-scan' is used throughout to refer to a real-time B-scan image independent of its field of view (whether it be a sector or linear scanner) and whether it be an electronic or mechanical device. The B-scan is regarded as a 2D slice through the imaged body and as such, the method by which the B-scan image is formed is not of great importance.

1.3.4 M-scan

An M-scan or time-position scan is a pulse-echo scan where the ultrasound beam points in one direction and the echoes received from interfaces are plotted as brightness-modulated lines against time. Lines from successive pulses are plotted side by side to show how the positions of interfaces vary with time.

1.3.5 B-scan resolution

The range or axial resolution of a system is determined by the length of the ultrasound burst which is transmitted. The quoted frequency for a pulse-echo system refers to the carrier frequency. The shorter the burst, the wider its bandwidth is. For example, a 2.5MHz pulse of length 7 cycles has a bandwidth of about 1MHz, while a 3 cycle pulse has a bandwidth of about 2MHz. Practical range or axial resolution is normally limited to 2-3 times the carrier wavelength. Because the pulse has a number of frequency components, wave phenomena such as interference are more complicated than for narrow band, continuous wave, systems. The shape of the pulse is altered in an unknown way as it propagates since the high-frequency components are absorbed preferentially in tissue. However, the many frequencies present in a pulse result in a more uniform intensity distribution in its ultrasonic field [McDicken 1991].

The lateral resolution of a B-scan is governed by the beam width. A conventional scanner will show the echo of a structure within the reception zone of the beam, that is the region in which ultrasound can be detected, as if it were on the beam axis. Lateral resolution is affected by the focusing of the beam and thus it can vary greatly along the beam axis. Lateral resolution is typically several wavelengths and lower than the range resolution [Wells 1977]. Modern multi-element transducers improve lateral resolution through dynamic focusing. Dynamic focusing involves rapidly sweeping the focus along the beam axis to coincide with the range of received echoes for a given instant in time.

1.3.6 Attenuation and penetration depth

Attenuation of ultrasound results from several phenomena including absorption, scattering, reflection and beam divergence. Despite the complexity of the phenomena, the attenuation property of soft tissue increases approximately linearly with frequency over the frequency range used by diagnostic ultrasound. Thus, a given power and frequency establishes a maximum depth of penetration beyond which echoes are too weak to be detected. External scanning frequencies of 3-7MHz are common. Higher frequencies, up to 30MHz, are usually used in catheter or transoesophageal imaging systems where the region of interest is closer to the transducer. For a given transmitted power a tradeoff arises between axial resolution and penetration depth. The transmitted power associated with a pulse is limited by safety considerations. The biological effects of ultrasound are also a function of the carrier frequency, the length of pulses, the pulse repetition frequency, and the maximum pulse amplitude [Wells 1977].

1.3.7 Time gain compensation

Time gain compensation (TGC) compensates for soft tissue attenuation by increasing the amplification of echo signals received at greater depths. Simple automatic systems are designed to correct for attenuation in homogeneous tissue and so the gain is purely a function of depth. Newer methods, however, adjust the gain across fluid-filled regions which exhibit lower attenuation than soft tissue. For example, with 5MHz ultrasound the attenuation compensation for a 1cm layer of liquid should be around 10 dB less than for a 1cm layer of soft tissue [Evans 1986]. These fluid-filled regions can be detected by an absence of backscatter. Despite more sophisticated systems, manual variation of the TGC characteristic is required during most examinations. TGC directly affects grey-levels in the B-scan image.

1.3.8 Scattering

Scattering of ultrasonic waves is caused by spatial variations in acoustic impedance in the human body. Two kinds of waves are present, specular reflections and diffuse scattering.

Specular reflections appear between layers of tissues with different acoustic impedances when the boundary is smooth. Specular reflections are only registered by the transducer if the layers are approximately normal to the direction of sound propagation — a direct consequence of the laws of reflection. Thus, received echo signals from structures lying parallel to the propagation direction can be very weak and thus difficult to render, whereas similar structures lying normal to the propagation direction give much stronger echoes. Specular reflections appear as bright boundaries blurred by a point spread function.

High frequency spatial variations in acoustic impedance give rise to diffuse scattering. Low intensity echoes are scattered over a wide angular range in the beam plane due to inhomogeneities within the medium which are small with respect to the ultrasound wavelength. Because wide-angle back-scattered echoes are nearly independent of angle of incidence, they form the basis of pulse-echo imaging systems.

1.4 ARTIFACTS

Many artifacts affect B-scan images. Some, such as beam width artifacts, grating lobes and focusing effects, are a characteristic of how the system is implemented. Others are inherent to the nature of pulse-echo imaging. McDicken [1991] discusses and illustrates both types of artifact in detail. It is the second class, however, which is of particular interest since these will always be present in a pulse-echo imaging system, whether it be 2D or 3D. Among these are shadowing, specular reflection, reverberation, refraction and variations in the speed of sound.

When a pulse of ultrasound passes through a highly absorbing layer it is greatly attenuated and is capable of producing only weak echoes at subsequent interfaces. A high density material such as bone may totally absorb the transmitted ultrasound waves, leaving structures further from the transducer in shadow. Reverberation may occur between a strong reflector and the transducer face or between interfaces in the body. Reverberation artifact occurs when multiple reflections arising from a single interface are falsely reconstructed as interfaces at distances corresponding to the times of the received secondary echoes. Resonance occurs when the vibration of a structure, due to its dimensions, is enhanced by constructive interference of internally reflected waves.

Geometric distortion of tissue morphology can arise from refraction and variations in the speed of sound. Refraction occurs at an interface between two tissues, in which the velocity of ultrasound differs, and results in deviation of the beam. Subsequent echoes are falsely reconstructed on a straight line path. A beam may also be dispersed by refraction and diffraction if it passes along a tissue boundary.

The speed of sound in soft tissue is not constant as assumed by pulse-echo systems. Excluding air and bone, it ranges between 1450m/s for fat to 1750m/s for tendon and is on average 1540m/s [McDicken 1991]. It also varies with temperature and tissue condition. Variation in the speed of sound results in range measurement errors when the average value is assumed. However, these errors are generally small [McDicken 1991]. In air or gas the speed of sound is approximately 330m/s and 3500m/s for bone. Both gas and bone pose additional problems for conventional imaging systems. Bone exhibits high attenuation and a large change of acoustic impedance from soft tissue to bone means that 40-50% of normal incident wave energy is reflected. Consequently, it is difficult to image through (calcified) bone because echoes from interfaces behind the bone are likely to be very weak and may not be detected.

1.5 ULTRASOUND COMPUTED TOMOGRAPHY (UCT)

An alternative to conventional pulse-echo ultrasound worth noting is ultrasound computed tomography (UCT). UCT is so named for its similarity to X-ray computed tomography. It is of interest because it offers another method by which 3D ultrasound imaging could be achieved and it has the potential to reduce some of the artifacts inherent in conventional pulse-echo imaging [Anderson 1991]. There are two types of ultrasound CT: reflection-mode UCT and transmission-mode UCT.

In reflection-mode UCT a transducer is used to transmit unfocused pulses into a target and receive the reflected or backscattered echoes. If the speed of sound is assumed constant, then echoes from scatterers at a given range will arrive at the same time at the receiver. This sum of echoes can be considered as the integral of echoes, or object reflectivity, at that particular range. The output signal from the receiver as a function of time can be considered as a 1D projection of the object reflectivity as a function of range. By rotating the transducer around the object in discrete steps, projections from many different angles can be obtained and algorithms developed for X-ray CT reconstructions can be applied to reconstruct images of ultrasonic reflectivity. The principle generalises to three dimensions. Reflection-mode UCT can potentially improve upon the resolution of conventional B-scan devices, particularly lateral resolution. It can also improve the speckle nature of pulse-echo images and the visualisation of specular reflecting interfaces [Jago and Whittingham 1993].

Transmission-mode UCT is of interest because the assumptions that ultrasonic beams travel in straight lines and that the speed of sound is constant can be avoided. Consequently, it has the potential to overcome geometric distortions associated with conventional pulse-echo ultrasound. Transmission-mode UCT requires two transducers, a transmitter and a receiver located located either side of the region to be imaged. The amplitude and time-of-flight of transmitted pulses are recorded at the receiver. The former can be used to reconstruct images of ultrasonic attenuation using X-ray CT techniques. The time-of-flight (TOF) measurements can be used to reconstruct images of the refractive index within the imaged region and consequently the speed of sound distribution within the target. Speed of sound images have clinical value in their own right and offer the possibility of correcting registration artifacts in reflection-mode UCT images of the same region. Minimum time-of-flight measurements, which correspond to the minimum ray path taken through the object by an ultrasonic beam, are used for performing reconstructions.

The problem that ultrasonic beams follow curved ray paths through the body is known as bent-ray CT. Iterative approaches have been developed for the reconstruction of a refractive index distribution involving ray-tracing [Enright 1992]. In these techniques the path of rays through an initial estimate of the refractive index distribution are computed and used to generate a better estimate of the distribution. Enright *et al.* [1992a] proposes an algorithm involving the iterative solution of a set of linear equations which does not require ultrasound ray paths to be known, thereby avoiding the computationally intensive task of ray-tracing. The method has been shown to cope with appreciable spatial variations in the refractive index in simulations of a circularly symmetric sinusoidal distribution, but has not been demonstrated in a practical system. Bent-ray ultrasound CT scanners are not expected in the near future.

Aside from the theoretical difficulties with transmission UCT, the development of a practical system is hindered by the awkward requirements of scanning. Transmission UCT typically requires that the body to be imaged is immersed in a water bath. The relatively long acquisition time of current experimental systems means that internal organ movement is an additional problem. The ergonomic problems associated with producing a practical and clinically useful system are significant. Furthermore, such a system is likely to be relatively expensive and the traditional benefits of ultrasound scanning, such as real-time imaging, relatively cheap operation, portability and the convenience of hand-held scanning will be negated.

1.6 SUMMARY

The interaction of ultrasound with tissue involves non-linear and anisotropic effects: absorption, diffraction, refraction, reflection and altered speed of propagation. B-scan imaging relies on coherently detecting the backscattered waves from inhomogeneities within the body due to incident pulses of ultrasound. Parameters such as tissue attenuation, frequency, pulse-length, resolution, bandwidth, pulse amplitude, the shape of transmitted pulses and the repetition frequency are all interrelated. Consequently, system designers are faced with a number of compromises.

B-scan images are distorted by several noise sources including thermal noise from amplifying circuits, acoustic noise, phase effects (speckle noise and phase aberrations) and, depending on the probe type, side and grating lobes. McDicken [1991] describes these and other practical artifacts at length. Gas, bone and fat are particularly difficult to image. Despite the numerous source of errors and artifacts, ultrasound has proven to be a useful clinical tool. In 2D B-scan imaging some artifacts can even provide useful diagnostic information, for example, shadowing can help identify gallstones or some types of tumour masses.

The aim of 3D imaging with ultrasound is to assist the comprehension of complex anatomy in 3D and to assist the extraction of quantitative information concerning organ size, morphology and volume. 3D ultrasound has advantages over CT and MRI in its cost, safety for both patient and operator and its ability to depict soft tissue. However, 3D reconstructions from 2D B-scans will be affected by the assumptions which underly conventional B-scan imaging. These will have implications for the interpretation of 3D structures and the accuracy of quantitative measurements.

Chapter 2

A SYSTEM FOR 3D GRAPHICS FROM ULTRASOUND

The reconstruction of surfaces from 3D ultrasound data is affected by the accuracy and other limitations of the acquisition system. The purpose of this chapter is to describe the system developed at Christchurch Hospital and its limitations. Section 2.1 presents an overview of techniques for acquiring 3D ultrasound data. Section 2.2 describes the Christchurch Hospital system. Section 2.3 summarises the performance and limitations of the system and the consequent scanning protocols for acquiring 3D data. Technological limitations are distinguished from intrinsic physical limitations, and improvements for future development are suggested.

2.1 REVIEW OF 3D ULTRASOUND SYSTEMS

A wide variety of systems have been proposed for the acquisition of 3D ultrasonic data [Rankin *et al.* 1993]. Some techniques have been developed to image specific anatomy, some involve transmission rather than pulse-echo methods, and some involve the development of new 3D probes, while others propose acquisition with conventional 2D equipment using spatial locators. Techniques also differ between those which are invasive and those which are non-invasive. This section presents an overview of the range of 3D acquisition systems which have been developed. 3D ultrasound research has been focused in three main areas; vascular imaging, imaging the heart and foetal imaging. It is difficult to discriminate clearly between systems, because these different objectives have tended to produce systems with common features.

2.1.1 Vascular imaging

3D ultrasound systems have been built to image arterial walls for the purposes of detecting arterial disease and surgical planning. Kitney *et al.* [Kitney *et al.* 1991, Kitney *et al.* 1990, Burrell *et al.* 1989] have developed an intra-arterial system consisting of a catheter-mounted probe. The probe is a multi-element phased-array device which acquires 2D slices at high speed as the catheter moves through the artery. Consecutive 2D slices, described in polar coordinates, are approximated as being parallel and reconstructed in a 3D Cartesian voxel space. Herrington *et al.* [1991] have developed a similar system for acquiring 3D images of the coronary artery. A 64 element, 20MHz transducer mounted on the tip of an intra-coronary catheter allows atherosclerotic plaque to be imaged over several centimetres of an artery. Mintz *et al.* [1993] have also developed a motorised catheter system for 3D intravascular imaging. A 25MHz transducer is rotated at 1800rpm as it is withdrawn through the artery at 0.5mm/s. 15 cross-sectional planar images/mm are acquired.

Intra-arterial techniques allow the use of relatively high frequencies in the range 20-30MHz. The resulting image quality is higher than conventional, external B-scan imaging where frequencies in the range 3-7MHz are commonly used. The greater attenuation experienced at these frequencies is not a difficulty since the transducer is close to the region of interest, namely, the wall of the artery. However, the development of catheter-mounted probes is expensive and subject to the general criticism of invasive techniques that they magnify the risk of infection. A major advantage of medical ultrasound has been that it is inherently non-invasive. For these reasons Nixon *et al.* [Nixon and Hames 1993, Nixon *et al.* 1992] have proposed a non-invasive technique for 3D artery visualisation using external B-scans. Transverse slices taken normal to the artery are acquired with careful manual scanning using a 5MHz phased-array transducer. The scanning head is moved such that images are grabbed at approximately constant spacing at the peak of the systolic cycle. The result is a stack of parallel slices. The artery is initially modelled as circular and the Hough transform is used to extract artery walls within each slice. Consequently, the slices must be normal to the artery so that the artery cross-section is circular rather than elliptical. Nixon *et al.* propose the development of a 3D probe to meet these scanning requirements in a clinical system. Pretorius *et al.* [1992] have developed mechanical device which translates a 3.5MHz sector probe along a 5cm path to generate 3D colour Doppler images of the carotid artery. An intravenous bag filled with acoustic gel is placed between the transducer and the subject's neck to maintain acoustic coupling.

2.1.2 Cardiac imaging

Imaging the heart, in order to visualise its dynamics and make quantitative measurements, has been the motivation for much research into 3D ultrasound. The morphology of the heart depicts its functional status. 3D ultrasound offers the possibility of detailed evaluation of morphology without the assumption of a simple geometric model. This is particularly important in the case of the cardiac chambers, whose size and shape are complex and directly reflect their performance as muscular pumps [McCann *et al.* 1988].

Verlande *et al.* [1991] have built a 3D probe for transoesophageal echocardiography. Their 5MHz transducer is capable of scanning any angle between 0° and 180° (from transverse through to longitudinal to inverse transverse) from a single transducer position. A multi-planar view of the heart is built up by scanning cross-sections from 0° to 180° in 15° increments. A four dimensional spatio-temporal dataset is assembled over a number of cardiac cycles. Martin *et al.* [1989] have also developed a similar 3D transoesophageal probe which differs in the axis about which the transducer is rotated. From a single position of the probe, transverse slices through the heart are obtained by mechanically tilting the transducer. Typically a total of 13 transverse slices through the heart are acquired. The heart is scanned from the base to the apex with the imaging planes separated by 3° on average. Scanning is synchronised to the end-expiration of the respiratory cycle and to the cardiac cycle via the electrocardiogram (ECG). Belohlavek *et al.* [1991] have developed yet another transoesophageal probe which acquires 20 2D fan-shaped scans. In their system, images are obtained in the longitudinal plane while the transducer is mechanically rotated in 1.8° increments.

Transoesophageal echocardiology provides access to cardiac structures from a short distance without interference from bone or lung interposition. However, other researchers have favoured non-invasive external scanning methods. Pini *et al.* [1991] have developed a 3.5MHz dynamically focused annular array transducer that allows acquisition of 50 standard fan-shaped 2D echocardiograms at 3.6° increments of rotation about the central axis of the scan. A transducer rotation of 180° permits the 3D visualisation of a solid cone encompassing the heart through the chest wall. The first 2D scan plane (0°) is compared with the last (180°) scan to provide a check on the stability of the rotation axis during acquisition. Acquisition is synchronised to the ECG and takes 60-100 seconds to acquire an entire cardiac cycle in normal sinus rhythm. McCann *et al.* [1988] have also developed a similar 3D probe which acquires a cone of data by rotating a sector scanning transducer. 2D slices are reconstructed in a 3D voxel array. Due to the rotational nature of scanning, data points are not uniformly distributed and 'holes' arise in the voxel array. Local averaging and smoothing are applied to fill-in holes on a slice by slice basis through the voxel array.

2.1.3 Techniques for imaging general anatomy

Anderson [1991] proposes a novel, theoretical technique which offers the possibility of real-time 3D ultrasonic imaging. This is not possible with approaches discussed thus far since these have used a steered, focused beam. The speed of sound in tissue severely limits the size of the volume which can be imaged in real-time. Anderson's technique involves taking time-of-flight measurements between multiple receivers and point source transmitters. Received echos are back-projected on ellipsoids to reconstruct a 3D volume. The technique is a type of transmission-mode ultrasound computed tomography

(UCT) and assumes wave fronts propagate along straight ray paths. Simulations of 3D point images have been successful; however, the development of a clinically useful system would suffer from the same difficulties discussed in Section 1.5 for general UCT methods. When transmitters are placed within the body, the method would additionally suffer from being invasive.

Many researchers have proposed general-purpose, external techniques for acquiring 3D data. Most systems fall into one of two categories, those advocating the development of 3D probes, where a 3D volume is imaged from a single probe position, and those where a 3D volume is constructed from a set of registered 2D slices, obtained using a conventional real-time B-mode scanner in conjunction with a remote localiser. Those developing 3D probes argue that the use of a single observation site simplifies the reconstruction and visualisation of the 3D data since the geometry of acquired slices is precisely known and no external frame of reference is required. However, 3D probes are generally suited to imaging smaller volumes while the registered B-scan approach allows larger volumes to be imaged. In either case, it is critical that the geometry of the slices be accurately known to reduce distortions in the reconstructed volume.

Some groups are developing 3D probes which utilise phased 2D arrays of piezo-electric elements to electronically focus and steer a beam, so that it scans through a volume. These operate in a manner directly analogous to current multi-element 2D B-scan transducers. However, they are difficult to build; if a conventional 2D B-scan transducer uses 32 elements then an analogous 3D transducer will require 1024 elements. Consequently, most 3D probes are mechanical in nature and house a conventional 2D transducer which is driven to acquire a set of slices through a volume. The 3D probes developed by Pini and McCann for imaging the heart are examples of this approach. Rickey *et al.* [1991] propose a 3D probe where the transducer is translated rather than rotated, so that parallel slices are acquired. Steen and Olstad [1994] place a 2D probe in a holder which uses a stepper motor to either tilt, rotate or translate the scan plane. The stepper motor is synchronised with the scanner. A typical acquisition takes approximately two seconds and consists of 95 scans. Tilting the probe covers an angle of 51° while rotating the probe covers a full revolution. The data is directly scan-converted into a regular voxel array and trilinear interpolation is used to fill gaps between adjacent scan-converted rays. The axis of rotation or translation is used as the frame of reference for the resulting 3D image. Consequently, 3D probes must be held still during acquisition. This time may range from a few seconds to over a minute for ECG-gated cardiac studies. In the system developed by McCann *et al.* the probe is mounted on a fixed stand and attached to the body via a specially designed body brace.

Rather than developing new 3D probes, and consequently new scanners, the registered B-scan approach has the advantage of utilising current technology and offers the ability to upgrade a conventional 2D scanner to 3D capability, avoiding the need to purchase a new scanner. In this approach the conventional scanner is augmented by a graphics workstation equipped with a spatial locator and appropriate software. The convenience of a hand-held probe is preserved and this appeals to clinicians. Imaging over a wide field of view with high resolution is possible.

A consequence of the registered B-scan approach is more complicated reconstruction and rendering algorithms. Unlike other 3D medical imaging modalities such as CT or MRI, where images are usually built up from stacks of parallel 2D tomographic slices, reconstructions must be performed from a collection of arbitrarily oriented slices where the anatomy is non-uniformly sampled. Although slices acquired by mechanical 3D probes are not usually parallel, simpler algorithms can be developed by taking advantage of the fixed geometry of acquired slices.

Although the advent of 3D probes appears inevitable, the technology and algorithms required to reconstruct 3D volumes from registered 2D scans is not likely to be made redundant. 3D probes, augmented with spatial locators, would provide the possibility of compounding multiple 3D scans. This would allow larger volumes to be imaged where the resolution of the scan might otherwise be compromised to improve the probe's field of view for the same number of image lines. Compounding of 3D scans from multiple transducer locations would also offer the possibility of improving the signal to noise ratio in the image and removing shadow and specular artifacts (Section 4.4.2.1).

2.1.4 3D reconstructions from registered B-scan slices

Some researchers have relied on careful manual scanning to image volumes using conventional B-scan probes, where the position of the probe is estimated by translating the probe so that a set of parallel planes is acquired at regular intervals [Cheng *et al.* 1993, Balen *et al.* 1993] or by tilting the probe from a fixed location on the surface of the body [Coppini *et al.* 1995, Thune and Olstad 1991]. However, manual localisation of the probe is inadequate for accurate imaging and clinically impractical.

A number of different technologies have been developed for six degrees-of-freedom tracking. These include optical, acoustic, electromagnetic and mechanical methods. Meyer *et al.* [1992] present a survey of current position-tracking technologies and their relative performance with respect to resolution, accuracy, responsiveness, robustness and working volume. The principle requirements of a 3D ultrasound system are that the tracking method does not interfere with the natural handling of the probe, that accurate, high resolution position and orientation measurements are achievable over a working volume of at least a 1m^3, and that measurements can be updated at a rate at least equivalent to the frame-rate of the scanner (>25Hz) with minimal sensing time and a total latency time (the time between requesting a measurement and it appearing at the output) which allows the acquisition frame-rate to be maintained (<40ms). The working volume estimate is based on the typical size of volumes to be imaged and the closest practical distance at which a fixed reference point (i.e. a transmitter) can be located near a patient.

Prior to the advent of scanners with mechanically driven transducers and multi-element phased arrays, mechanical arms were used to track the position and orientation of transducers in early 2D B-scans. Potentiometers in the joints of the arm determined angular displacements. Various groups have since used these to localise B-scan probes for 3D studies. Some arm designs restrict the degrees-of-freedom in which the probe can be moved. Ohbuchi and Fuchs [1990] describe the use of an arm which restricts scanning to three degrees-of-freedom (x, y, θ). Nikravesh *et al.* [1984] use a mechanical arm which permits a full six degrees-of-freedom.

Mills and Fuchs [1990] developed an optical tracking system to determine the position and orientation of their probe. Two video cameras view the position of a crown of lights attached to the probe. Stereo triangulation then yields six degrees-of-freedom tracking information. Over a small field of view an accuracy of 2mm of translation and up to $0.1°$ of rotational resolution can be obtained. However, they report that, in practice, a tradeoff between low tracking resolution and ease of implementation was made. The large distance by which the tracking beacons were separated from the probe (over 300mm) resulted in a probe which many clinicians would find awkward to work with. Advantages of this scheme included that it was relatively insensitive to environmentally induced distortion, that it had a large working volume and had a fast response time. Although they have subsequently improved the performance of their system, reducing the size of LED beacons on the probe and improving accuracy, optical occlusion of the probe from camera view must be avoided while scanning.

Moritz *et al.* [1983] used an acoustic method to locate an ultrasound probe in their system. Three spark-gap transmitters were mounted on the probe and an array of microphones, tuned to 100kHz, was located nearby, within 2m. With this method the probe can be located to within $\pm$1mm with an orientation calculated to within $1°$ at a distance of 50cm from the microphones. However, an unobstructed line of sight must be maintained between the transmitters and receivers. This somewhat restricts the scanning freedom of the operator. Variation of speed of sound in air with temperature and humidity is significant and so temperature correction is required. This system takes at least 30ms to update the position of the probe and possibly as long as 120ms if inconsistency in the measurements is detected.

More recently the advent of electromagnetic spatial locators has seen these used in 3D ultrasound systems [Krieg and Sanders 1984, Selzer *et al.* 1989, Gardener 1991, Nelson and Elvins 1993, Pretorius and Nelson 1994, Kelly *et al.* 1994, Kelly *et al.* 1992, Surman *et al.* 1993]. A spatially varying electromagnetic field is generated by a transmitter located near the patient and a receiving sensor is mounted on the probe. This type of localiser permits six degrees-of-freedom but is subject to the presence of ferro-magnetic and conductive (metallic) objects which may distort the field depending on the type of device and whether it uses ac or pulsed dc fields (Section 2.2.2.1).

Of the various technologies available, electromagnetic spatial locators currently offer the best resolution and accuracy, up to 1mm position accuracy and $0.15°$ orientation accuracy. Optical systems and six-degrees-of-freedom mechanical arms are also worthy of consideration. However, these are relatively expensive and restrict or interfere with normal hand-held scanning. An important advantage of electromagnetic systems is that they do not require unoccluded line of sight between the transmitter and receiver, as do optical systems. Traditionally ultrasound has been seen as portable and flexible since it does not require dedicated enclosures as do MRI and CT scanners. However, the use of fixed cameras in the examination room by optical systems means that this quality is lost.

The aim of the system being developed at Christchurch Hospital is to achieve a non-invasive, general purpose and clinically useful system for acquiring 3D ultrasound data. The registered B-scan approach, employing an electromagnetic spatial localiser, meets these objectives.

2.2 A SYSTEM FOR ACQUIRING 3D ULTRASOUND DATA

A simplified block diagram of the system built for acquiring 3D ultrasonic data is illustrated in Figure 2.1. It consists of three principal components, the Acuson XP10 scanner, the Polhemus 3Space spatial locator and a graphics workstation. The workstation consists of an IBM-compatible 486 PC which hosts two transputer boards, a Transtech TMB08 single transputer board and a Quintek MosaiQ frame grabber. Position and orientation measurements determined by the Polhemus 3Space localiser are synchronised with the acquisition of 2D B-scans from the Acuson scanner. The latter are digitised from the video output of the Acuson scanner via a Quintek MosaiQ frame grabber. Figure 2.2 is a photograph of the acquisition hardware located in the Department of Radiology at Christchurch Hospital.

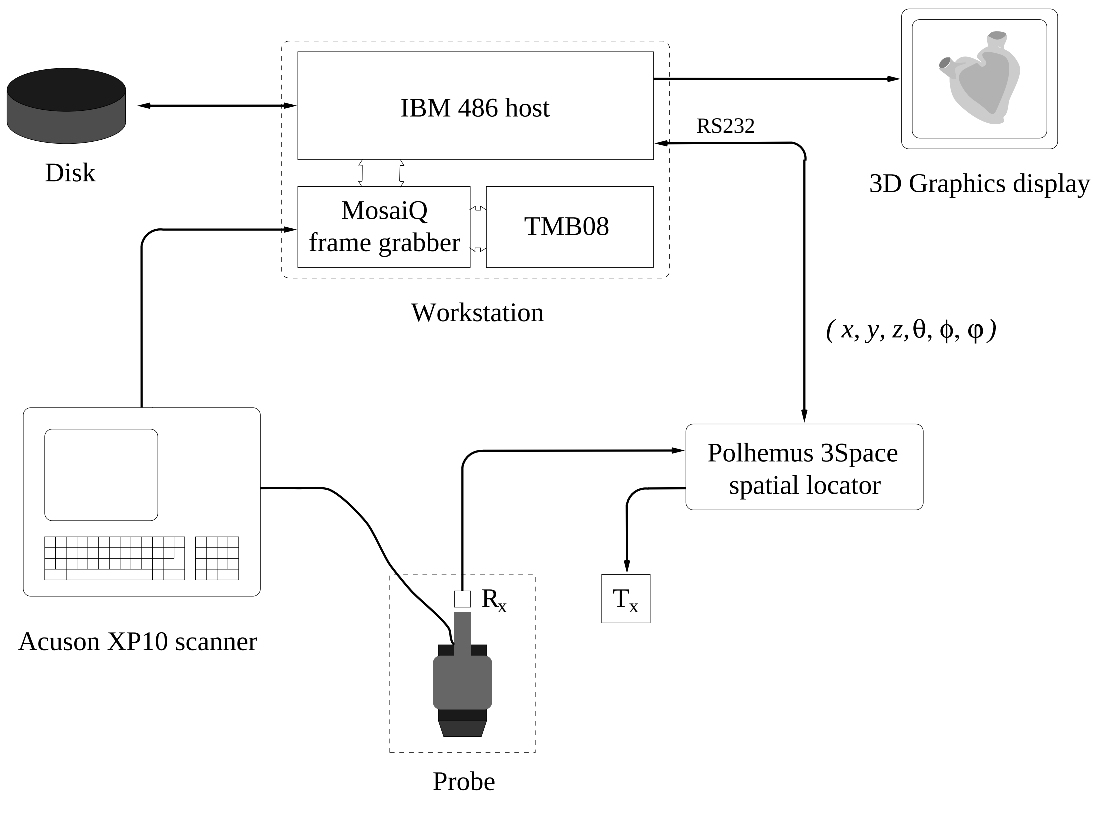

Figure 2.1 Block diagram of the 3D ultrasound data acquisition system.

2.2.1 The Acuson XP10 scanner

The resolution and accuracy of a 3D reconstruction can be no better than the constituent 2D slices from which it is assembled. The reconstructions presented in this thesis were obtained using an Acuson XP10 ultrasound scanner equipped with an L558

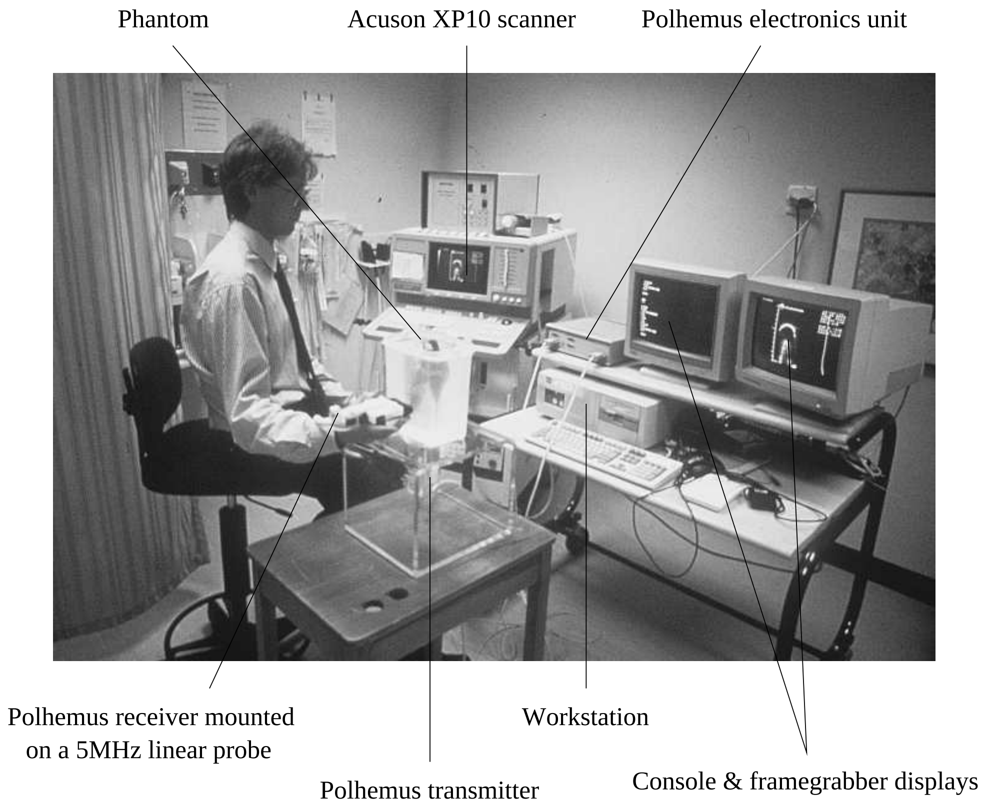

Figure 2.2 The 3D ultrasound data acquisition system at Christchurch Hospital. An ultrasound phantom is shown being scanned with a 5MHz linear probe.

5MHz linear phased-array probe and a 3.5MHz phased-array sector probe. The performance of the 5MHz probe was independently evaluated with an RMI ultrasound test phantom [Fright and Allan 1994]. The principal findings are summarised in table 2.1.

The term "registration" in table 2.1 refers to the lateral accuracy of the sector scan. The separation of two wires of known separation is measured to determine this. Calibration refers to the axial accuracy and is determined by measuring the separation between two wires in the axial direction (along the line of sound propagation). The axial resolution is the smallest axial separation between two objects before they can no longer be distinguished. Similarly, lateral resolution is the closest spacing which can be resolved in the lateral direction. The dead zone refers to the minimum depth at which tissue can be imaged, while sensitivity refers to the maximum depth at which tissue can be imaged.

The Acuson scanner provides no digital interface to scan data nor access to internal signal processing; only an RGB PAL video output is provided. The rate at which scans are updated depends on the probe used, the field of view, resolution and the number of focal zones, as well as the type of filtering selected. Typically scans are updated at rates ranging from 10 to 28 frames per second.

Test		Result
Registration		0.7%
Calibration		1.3%
Axial resolution	near-field	0.5mm
	mid-field	0.6mm
	far-field	-
Lateral resolution	near-field	1.0mm
	mid-field	2.0mm
	far-field	-
Dead zone		3.6mm
Sensitivity		120.6mm

Table 2.1 Test results for 5MHz linear probe, from Fright and Allan 1994.

The Acuson XP10 has an extensive operator interface which allows the user to modify many parameters and significantly alter the appearance of an imaged slice. Among these is the ability to manually vary the focal zones, the axial gain and the field of view. In addition, contrast enhancement tools and filters are available. Although it is feasible to monitor some of these parameter settings from the video output of the scanner, most cannot be readily monitored during acquisition. In order to reduce the initial complexity of reconstructing and visualising volumes, the current scanning protocol adopted for 3D acquisition relies on a constant field of view being maintained during scanning and a fixed axial gain characteristic. However, variation of the focal zone is permitted, even though this information is not currently utilised by the reconstruction and rendering algorithms.

2.2.2 The Polhemus *3Space* spatial locator

The Polhemus 3Space Fastrak is an electromagnetic six degrees-of-freedom measurement device manufactured by Polhemus Incorporated [Polhemus 1992]. The device consists of an electromagnetic transmitter connected to an electronics unit via a 3m cable and a receiver also connected with a 3m cable. The transmitter approximates a cube with dimensions 53x53x58mm. The receiver has dimensions 28.3x22.9x15.2mm and weighs 17.0gm. The physical characteristics of the receiver are important since it is attached to the ultrasound probe and should not interfere with its handling.

The single receiver Polhemus system consists of a stationary magnetic dipole transmitting antenna and a mobile receiving antenna. The position of the receiver is described by its relationship to a fixed Cartesian coordinate system (x, y, z) centred on the transmitter. A direction, or orientation, in relation to the transmitter is described by three angles; azimuth (yaw), elevation (pitch), and roll. Both antennas consist of three mutually orthogonal coils, or loops. The diameter of the loops is small compared to the distance separating the transmitter and receiver so that each loop may

Specification	Polhemus	Bird
Operation mode	ac field	pulsed dc field
Translation range	0.76m*	0.91m
Translation resolution	0.5mm/m of range	0.076mm
Translation accuracy	0.138mm	0.254mm RMS
Angular resolution	0.025°	0.1° at 0.3m
Angular pointing accuracy	0.21° RMS	0.5° RMS
Maximum update rate	120Hz**	144Hz

Table 2.2 Manufacturer's specification for Polhemus and Bird spatial locators as of 1993.
* Polhemus range for the accuracy figures quoted in this table.
** Update rate for a single receiver.

be regarded as a point or infinitesimal dipole. Exciting the loop antenna produces a far-field component and a near-field component. The far-field intensity decreases with the inverse of distance from the antenna ($1/r$) while the near-field, also referred to as the "quasi-static" component in the literature, decreases as the inverse cube of the distance from the antenna ($1/r^3$). The quasi-static field is not detectable at long distances but dominates at short distances such that the far-field component is negligible. The 3Space uses near-field measurements to calculate position and orientation. Due to the symmetry of the magnetic fields produced by the transmitter, there are two mathematical solutions to each set of receiver data processed. Consequently, only half the spatial sphere enclosing the transmitter can be practically used at any one time to avoid ambiguity. The zenith of the hemisphere for valid measurements can be defined by the user through software commands.

The transmitter radiates three mutually perpendicular magnetic fields when each of the coils is excited in sequence by an alternating current. The driving signals for each coil are identical in frequency (8013Hz for a single receiver) and phase. Each excitation produces a single axis source dipole which induces currents in the receiver coils; three currents are sensed in the receiver for each transmitter dipole. The current induced in each coil varies with distance from the transmitter. The position and orientation of the receiver are calculated from the nine currents sensed by the receiver. Small changes in the sensed coordinates are used to update previous measurements. Imperfections in the antenna loops, which create distortions in the magnetic field, are measured and compensated for in software and hardware at the time of manufacture. A more detailed discussion on the principles behind the 3Space's operation can be found in Raab *et al.* [1979].

The proximity of conductive materials can alter the 3Space's performance since the ac magnetic field radiated by the transmitter induces eddy currents and distorts the transmitter field. The effect of eddy currents increases with the size of the metallic object. The larger the cross-section area through which eddy currents can flow, the lower the eddy current resistance. Hence, large conductive objects are to be avoided in the local environment. A family of electromagnetic trackers produced by *Ascension*, known as the Bird, Big Bird and Flock of Birds, radiate a sequence of dc pulses. Because the 3Space transmitter radiates a continuously changing field, it is continuously generating eddy currents. The Bird radiates magnetic pulses that only generate eddy currents at the beginning and end of a pulse. Measurements are taken after a delay which is sufficient to ensure that the field measured is approximately dc. However, the Bird devices are affected by background dc magnetism, of which the Earth's magnetic field is a significant and ever-present component. The devices overcome this problem by measuring the background magnetism and subtracting it from subsequent detected values. Consequently, the Bird family of devices claims to be less susceptible to conductive-metal interference. However, both the 3Space Fastrak and the Bird remain sensitive to the permeability of ferromagnetic objects and Ascension recommends that both the Bird's transmitter and receiver be located 150mm from ferromagnetic objects. Both systems are also vulnerable to environmentally induced noise from power cables, transformers, CRTs, and other devices which generate noise in the 8-1000Hz range. Table 2.2 compares manufacturers' specifications for the Bird and Polhemus 3Space Fastrak as of 1993. The 3Space Fastrak was chosen in preference to the Bird for its higher resolution and accuracy at the time of development. Subsequent development of both devices means that this may no longer be the case. The 3Space localiser is also used in the following ultrasound systems, Krieg and Sanders [1984], Selzer *et al.* [1989], Nelson and Elvins [1993] and Kelly *et al.* [1994].

The 3Space unit provides software which allows the unit to be calibrated to the particular magnetic shape of a specific working volume. However, this feature is not used in the current ultrasound system since it is only useful if the environment is static. The metallic content of the mobile probe, to which the receiver is mounted, is of particular concern in this application since this potential source of interference cannot be readily eliminated. Every attempt is made to remove all other metallic objects from the scanning locality. Section 2.2.2.1 summarises investigations into the effects of the probe on the performance of the 3Space unit.

2.2.2.1 Receiver-transducer interactions

Krieg and Sanders [1984] used a Polhemus *Tracker*, a predecessor of the 3Space Fastrak, to locate the position of an ultrasound probe and reported the following findings concerning the interaction between the probe and the sensed magnetic field. They found that a separation of 100mm provided adequate accuracy, typically 1.25mm translation and $0.1°$ orientation, provided that the ultrasound transducer did not lie between the transmitter and the receiver, and that the receiver was no further than 0.76m from the transmitter.

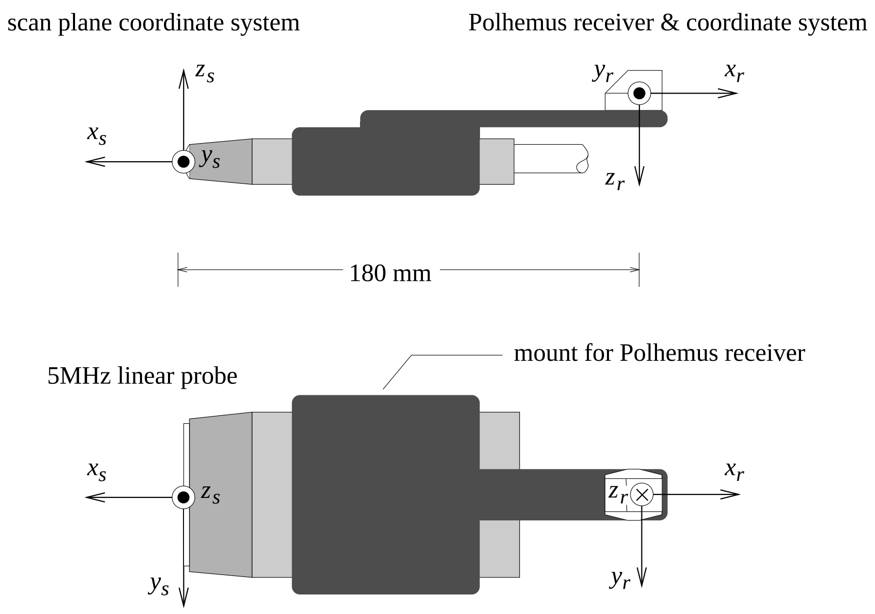

Figure 2.3 5MHz probe with Polhemus receiver mounting.

Predicting the interference caused by a metallic object on an electromagnetic field is a difficult problem. It depends on the conductivity of the metal and its magnetic permeability, which in turn depend on the frequency of the applied field. Also, residual magnetism and the shape and thickness of the object are important. In this case the internal construction of the Acuson probes was not known.

A series of static tests were therefore conducted at Christchurch Hospital to determine the effect of the proximity of an Acuson 5MHz linear probe on the receiver [Brichieri Colombi 1993]. Various receiver-transmitter separations were recorded along a radial path from the transmitter. The variation in the recorded measurements as the probe was moved closer to the receiver were noted. The tests were not a complete study of the 12 degrees of freedom, i.e. all possible combinations of receiver-probe positions and orientations relative to the transmitter. The tests assumed that the transmitter field was dependent only on the distance from the transmitter and that interference due to the probe was independent of the orientation of the receiver. Furthermore, only a limited number of probe positions, relative to the transmitter and receiver, were considered. At close proximity, the variation in the 3Space tracking data due to the presence of the probe was significant [Brichieri Colombi 1993]. However, adequate accuracy was obtained with a receiver separation of 150mm from the transducer face, provided that the probe did not lie between the receiver and transmitter. Consequently a plastic receiver mounting, illustrated in Figure 2.3, was designed in consultation with sonographers. Mounting the receiver at the rear of the probe provided adequate accuracy and least interfered with the handling of the probe. A similar holder was designed for the 3.5MHz sector probe, which also uses a phased array to generate a steerable scan beam.

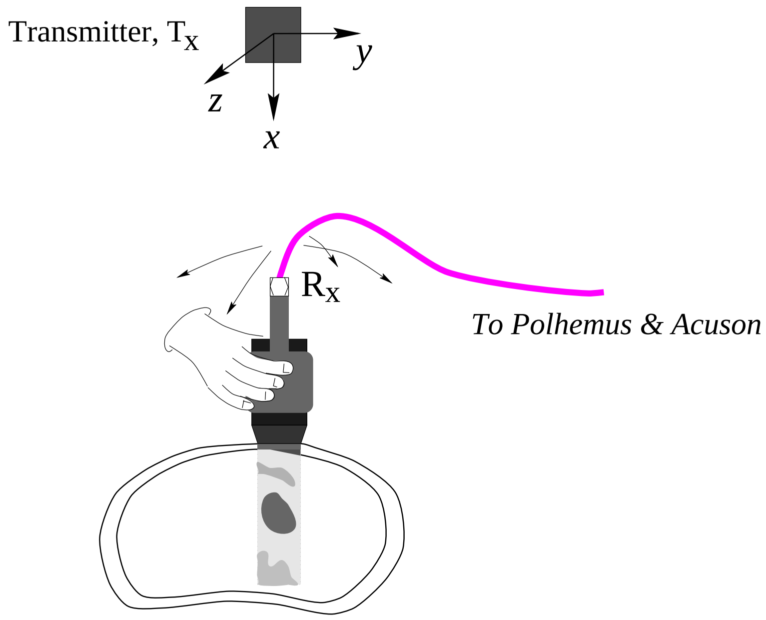

Figure 2.4 Overhead transmitter arrangement minimises positioning errors during scanning.

The scanning geometry is illustrated in Figure 2.4. The transmitter is positioned above the examination region to ensure that the probe will not obscure the receiver from the transmitter and will guarantee that the receiver operates within the valid measurement hemisphere of the transmitter. With this arrangement the receiver cable is conveniently bundled with the probe cable. Further tests have found that the presence of the probe cable has no adverse effect on receiver measurements.

2.2.2.2 Accuracy

The purpose of the spatial locator is to determine the position in transmitter space of points imaged in the probe's scan plane. The error associated with determining the location in transmitter-space of a point, p_s, in scan space can be considered in two parts: the error associated with determining the position of the receiver relative to the transmitter, ε_{tr}, and the position error of the imaged point relative to the receiver, ε_{rs} (Figure 2.5). The magnitude of ε_{rs} is a function of the uncertainty in the receiver's orientation, $\Delta\theta$, and the separation between the imaged point and the receiver, r_{rs},

$$|\varepsilon_{rs}| = 2|r_{rs}| \sin \frac{\Delta\theta}{2} \tag{2.1}$$

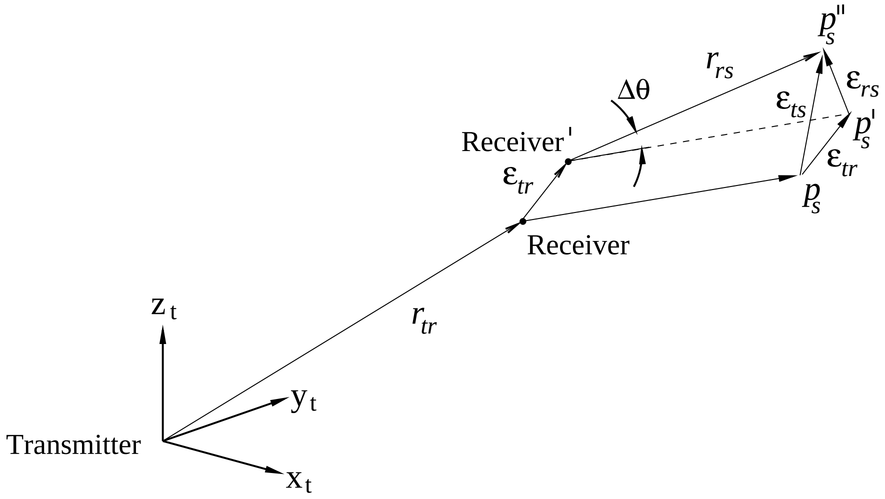

Figure 2.5 Spatial location error for determining the location of a point in the ultrasound scan plane, p_s, in transmitter space (x_t, y_t, z_t).

Figure 2.6 is a contour plot of $|\varepsilon_{rs}|$ as a function of the $\Delta\theta$ and r_{rs}. Imaging tissue at the maximum range of the 5MHz probe corresponds to a receiver separation of 300mm for the receiver-probe mounting illustrated in Figure 2.3. It is apparent from Figure 2.6 that accurate registration of tissue imaged at this range requires an orientation accuracy better than 0.2° for a position error better than 1mm.

The total position error, ε_{ts}, is the vector sum of ε_{tr} and ε_{rs}. If both ε_{tr} and ε_{rs} are independent random variables with zero mean, then the RMS position error is,

$$\begin{aligned} |\varepsilon_{ts}|_{RMS} &= \sqrt{|\varepsilon_{tr}|^2 + |\varepsilon_{rs}|^2} \\ &= \sqrt{|\varepsilon_{tr}|^2 + 4|r_{rs}^2| \sin^2 \frac{\Delta\theta}{2}} \end{aligned} \quad (2.2)$$

This is the position accuracy for the ultrasound imaging system. Polhemus rates the static translation accuracy of the receiver as 0.8mm RMS for x, y and z, and the angular accuracy as 0.15° RMS for θ_x, θ_y and θ_z, when the receiver is less than 0.76m from the transmitter [Polhemus 1992]. These measurements were determined with the aid of a precise mechanical positioning instrument with a precision gimble [Polhemus 1992]. Although a maximum transmitter-receiver operating distance of 3m is specified, the quoted accuracies degrade over this extended range.

The position error for an imaged point as a function of its distance from the receiver is displayed in Figure 2.7. The figure used to compute the receiver translation accuracy is 1.38mm RMS, the vector sum of the quoted RMS uncertainty in the x, y and z

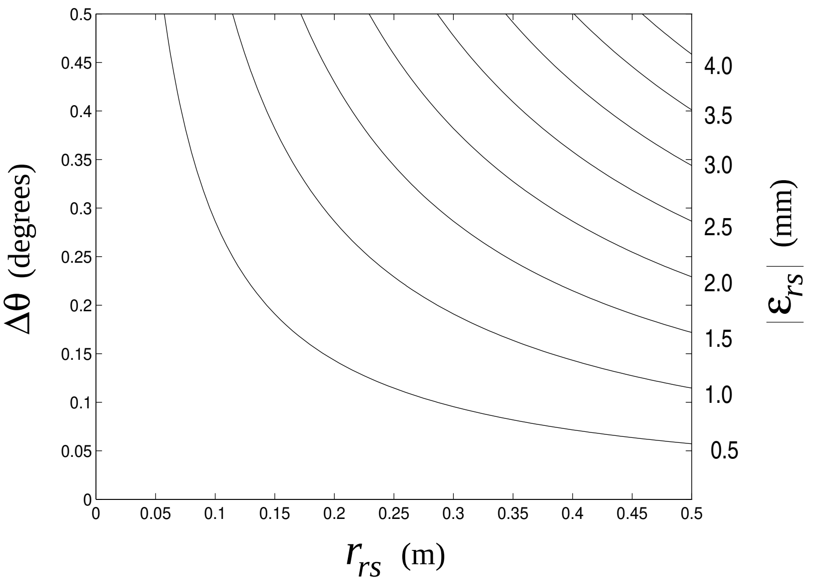

Figure 2.6 Magnitude of the position error for an imaged point relative to the receiver as a function of receiver orientation uncertainty $\Delta\theta$ and the distance of the imaged point from the receiver, r_{rs}.

directions. Similarly, the orientation accuracy, 0.260° RMS, is vector sum of the angular uncertainties θ_x, θ_y and θ_z. These uncertainties result in an RMS error at 300mm of 1.9mm. Figure 2.7 assumes that uncertainties are independent and have zero mean. The distribution of uncertainties is unknown and how these error distributions vary spatially is also unknown.

Two important assumptions are relevant to this discussion. The first is the assumption that the Polhemus specifications are accurate and that variations in the receiver position and orientation are uniformly distributed throughout the 0.76m operating hemisphere. If this is not the case, then the error in some regions, or at certain orientations, may be larger than predicted. The second assumption is that the uncertainties in position and orientation behave as random variables. Without the benefit of a precision positioning instrument, independent verification of the Polhemus specifications is impossible. A simple experiment, however, has indicated that there is reason to query the specified accuracies and the modelling of variations as random variables. In this experiment two receivers were mounted on a rigid support, approximately 0.3m apart. The transform relating one receiver to the other was not known precisely, so measurements from both receivers were used to estimate the separation and rotation between the two. The receivers were moved within the 0.76m hemisphere about the transmitter in a metal-free environment without the presence of the probe. The expected variation in the distance measured between the two receivers was $\sqrt{1.38^2 + 1.38^2} = 1.95$mm RMS. However, the measured separation varied as much as 5mm between certain positions and orientations. These variations were repeatable measurements rather than random. The experiment suggested that it might be more appropriate to consider the position

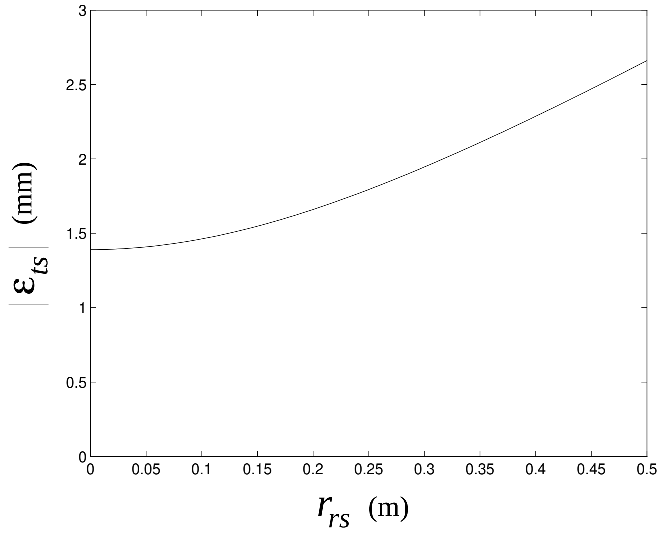

Figure 2.7 Spatial location error, $|\varepsilon_{ts}|$, as a function of the distance of an imaged point from the receiver for $\Delta\theta = 0.26°$ RMS and $|\varepsilon_{tr}| = 1.39$ RMS.

error as a warping of the true measurement rather than a random variable. If this is the case, then it is possible to have high accuracy for relative position measurements taken in a small neighbourhood, while errors in absolute measurements are comparatively large.

This might explain why discrepancies have not been reported by other users of the system for 3D ultrasound imaging and why reasonable results have been obtained with this system. Many of the 3D reconstructions reported in the literature have been from short sweeps of a probe where semi-parallel slices were acquired [Pretorius and Nelson 1994, Balen *et al.* 1993, Selzer *et al.* 1989]. Relative position and orientation accuracy over a short temporal period appeared to be tested by these scanning procedures rather than absolute position accuracy.

Training a neural network to compensate for distorted measurements was considered during the course of this research. However, systematic training of the network is a difficult task due to the six degrees of freedom to be considered and the need for an independent positioning system.

As a consequence of these studies, confidence in the Polhemus has been restricted to relative measurements. In Section 4.2.1, 3D reconstructions are restricted to slices acquired during short sweeps of the probe. This has not proved to be a significant limitation in practice and has simplified aspects of the reconstruction algorithm. Despite the doubts concerning the current system, Chapter 5 considers the problem of reconstructions from completely arbitrary slices where absolute position accuracy is important, since the current limitations of spatial locators are anticipated to be temporary.

2.2.3 Computing platform

2D ultrasound slices are digitised from the PAL video output of the Acuson scanner with a Quintek MosaiQ frame grabber. The MosaiQ frame grabber is a 24-bit full colour image capture board which incorporates an IMS T805-25MHz transputer. In addition to a 24-bit image plane, the MosaiQ provides an 8-bit overlay plane, 4 MBytes VRAM, 8 MBytes DRAM and four transputer expansion slots. The latter allow parallel processing to be implemented through links to other transputers.

The transputer is an electronic chip designed by Inmos for general-purpose computing. Contained on the chip is a fast 32-bit processor, a double precision arithmetic unit, a scheduler and four built-in communications links which allow it to talk to four neighbouring transputers. The transputer platform is an ideal basis for parallel processing; it allows large processor arrays to be constructed with ease. These can be organised in a wide range of topologies, provided interconnections are limited to a maximum of four per processor.

A major benefit of adopting the transputer architecture is that a system can be developed with an initial low-cost configuration and then expanded to extend its power and facilities through the addition of extra transputer modules, once the system is proven. The processing tasks associated with acquiring ultrasound data, filtering the data, then performing 3D reconstructions and graphical renderings, are naturally suited to parallel implementation. Parallel processing through the addition of further processors allows real-time speed to be achieved. A parallel architecture also facilitates the development and maintenance of software.

In the current system an IBM-compatible 486 PC hosts the MosaiQ frame grabber which is linked to a TMB08 transputer board, also equipped with 8MBytes of DRAM. A 768x576 monochrome image is grabbed with 8-bits resolution from the Acuson video output. The Acuson's pulse-echo display is monochrome. The colour Doppler imaging capabilities of the scanner are not utilised. The MosaiQ on-board transputer acts as the master processor and the IBM PC host runs in a server mode. Communication with the Polhemus 3Space unit is via the IBM PC bus and a 19200 baud RS232 interface. This system has been used for the acquisition of scan data while the development of the reconstruction algorithms, described in Chapters 4 and 5, has been implemented on Sun workstations. Eventually, these too will be implemented on the transputer network.

2.2.3.1 Data Acquisition

Figure 2.8 summarises the data acquisition cycle. Upon entering the data acquisition mode, a frame is immediately digitised from the Acuson video output. From this initial frame the system identifies the type of probe being used, the field of view, the resolution

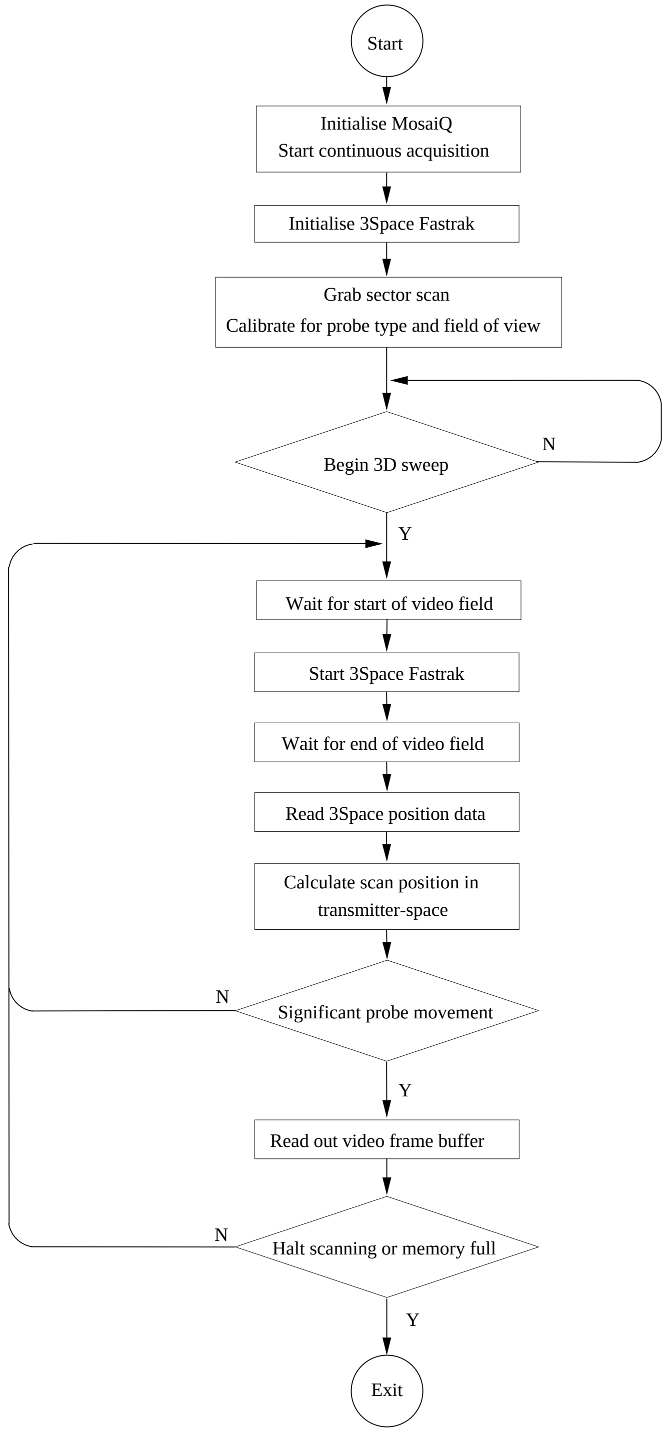

Figure 2.8 Flow diagram for the ultrasound data acquisition cycle.

and the bounding region containing the ultrasound image. The system lies in a dormant state until the acquisition of a sequence of slices is instigated by the operator. This is done via the workstation keyboard or by a foot switch. The latter is more convenient when performing an examination, and its use is familiar to sonographers. The system continuously displays real-time B-scans from the output of the scanner when in the dormant state. Once acquisition commences, a sequence of frames is grabbed until the process is halted either by the operator or by the system running out of memory. Sweeps of up to 60 slices at a time, depending on the type of probe and field of view, are typically acquired with the current configuration.

Acquisition frame-rate

Two acquisition methods have been implemented in the current system. One acquires frames at the constant video rate; the other, illustrated in Figure 2.8, acquires frames at a variable rate pending the detection of significant probe movement. In the latter method, the extrema points for the probe's field of view are calculated in transmitter space using the Polhemus position data. When probe movement exceeding a user-specified threshold is detected, the next available frame from the Acuson output is acquired. This is done to compensate for different operator scanning speeds. Ensuring that a minimum amount of movement occurs between scans means that an approximately uniform spacing between slices is achieved, independent of the operator's scanning speed and variation during a sweep of the probe. Manually achieving a constant scanning speed is an otherwise difficult skill.

Although the 3Space unit can sample position data at a rate of 120Hz in continuous output mode, it takes 15.6ms to transmit a full position record for a single receiver. The slow speed of the RS232 link between the 3Space unit and the PC, combined with the limited size of the 3Space output buffer, mean that erratic skipping of position readings will occur if the unit is operated in the continuous output mode. Consequently, the 3Space unit is run in the single transmission mode where position data is provided upon receiving a request. The maximum position update rate possible with the current system is 46.3Hz. A 21.6ms delay occurs between requesting a position measurement and receiving it. This delay consists of 0.5ms to transmit the request, 3.5ms to sample the magnetic field, 2ms computation time and 15.6ms to transmit the data record and two control bytes.

Probe velocity

The delay between when tissue is imaged and when the position of the probe is determined is a potential source of registration error and imposes restrictions on probe velocity during scanning.

The video output of the Acuson scanner is an interlaced signal consisting of two consecutive fields of 20ms duration, each beginning with a synchronisation pulse. Requests for 3Space position data are synchronised to the beginning of each frame to be grabbed. The 3Space unit operates in parallel with the frame grabber, so the tasks of acquiring a frame and measuring probe position can be executed simultaneously. The time to transmit a request command to the 3Space is 0.5ms and the time for the 3Space to sample the magnetic field at the receiver is 3.5ms [Polhemus 1992], giving a total delay of 4.0ms. Although the video frame-rate of the scanner is 25Hz, the actual imaging rate may be higher or lower depending on the probe, the field of view, the number of focal zones and the filtering specified by the operator. Furthermore, the actual scan rate may vary dynamically with the tissue being imaged. The details of how and when the video output is updated by the scanner are unknown. If the video output is buffered, then the critical delay between when tissue is probed by ultrasound to when a position measurement is made could be increased by as much as the length of a full frame, more if the actual imaging rate is slower than 25Hz. If a full frame delay is taken as a worst case, then a 40ms + 4.0ms = 44ms delay occurs between when tissue is first probed by ultrasound and a position measurement is made. If a spatial registration error of 1mm is tolerated, then the maximum allowed velocity of a pixel in the scan plane is 23mm/s. Velocities of this order correspond to relatively slow sweeps and twists of the probe, but not uncomfortably so. The scan velocity can be estimated by comparing the position of extrema points in consecutive scan planes. A large movement between slices is indicative of excessive probe movement and probable registration error. Consequently, excessive scanning velocity is monitored to prevent misregistration of slices. Alerting the operator and repeating the scan is preferable to compensating for motion in following reconstruction algorithms.

Placing a limit on the probe velocity constrains the spacing between acquired slices, the time to image a volume therefore becomes a function of the probe's field of view and the video frame-rate.

2.3 CONCLUSION

This chapter has described the construction of a system to acquire 3D ultrasound data. Registered 2D B-scans are used to build up 3D volumes with the aid of an electromagnetic spatial locator. This approach has been adopted by a number of researchers [Krieg and Sanders 1984, Selzer *et al.* 1989, Nelson and Elvins 1993, Kelly *et al.* 1992]. It allows conventional ultrasound equipment to be upgraded to provide 3D imaging capability, rather than the development and purchase of a new scanner. The use of a magnetic spatial locator maintains the freedom and convenience of a hand-held probe but requires the development of new reconstruction and visualisation algorithms, due to the arbitrary orientation of constituent slices.

The resolution of the imaged volume is limited primarily by the performance of the spatial locator rather than the scanner resolution. The analysis in Section 2.2.2.2 indicated that the accuracy of the spatial locator needs to be better than the figures quoted for the Polhemus and Bird devices in order to achieve tissue registration better than 1mm. Despite this analysis, reliable reconstructions have been achieved with the current system when the probe is restricted to smooth sweeps where near-parallel slices are grabbed. These rely on accurate relative measurements over a small region while reconstructions from arbitrary oriented slices require accurate absolute measurements. Fortunately, acquisitions of probe sweeps are compatible with the natural manner of performing an examination. Reconstructions with 0.5mm^3 resolution are presented in Chapter 4.

Interference from metallic objects, including the Acuson probe, has a measurable effect on the performance of the Polhemus spatial locator. Consequently, the receiver should be located far from the ultrasound probe; however, the further this distance, the greater the effect that inaccuracies in orientation have on the registration of imaged tissue. The current separation is a compromise between these two requirements. A filtering option is provided in the 3Space system to compensate for variations in position measurements. However, the time required to accumulate data and compute the results compromises the responsiveness of the system and so this facility is not used. The current position update rate, 46.3Hz, is limited by the RS232 baud rate but is adequate given the 25Hz video frame-rate.

Spatial registration errors also originate from probe movement and the delay between probing tissue and making a position measurement. The limited video output of the Acuson scanner means that optimal synchronisation of position measurements with scan acquisition is not possible. The estimated maximum scan-plane velocity of 20mm/s is not troublesome.

Spatial location technology is continually improving and is not seen as an intrinsic limitation of the system. Scanner manufacturers are beginning to provide digital output interfaces with their scanners as well as video. This will simplify the acquisition of data and improve the ability to monitor important scanner settings. Despite limitations, the current system has proved to be useful for 3D studies of a wide range of anatomy, particularly where internal tissue movement is minimal. The fundamental limitation of this approach, in terms of the volume which can be imaged in a given period and the maximum scanning velocity, is determined by the speed of sound in tissue and the fact that volumes are effectively imaged with a single swept beam.

Chapter 3

VISUALISATION OF SURFACES FROM VOLUMETRIC DATA

The depiction of surfaces from ultrasound and X-ray CT data is directly relevant to the work presented in Chapters 4, 5 and 6 concerning surface reconstruction. This chapter describes how images are produced from 3D volume data in this thesis. The visualisation of 3D medical data differs from traditional computer graphics in several aspects. Firstly, the data from which images are formed is captured from real objects rather than synthesised by a designer or generated mathematically from a computer model. Secondly, accurate depiction is more important than pleasing images — although this does not imply that good image quality need conflict with clinical accuracy.

Scientific visualisation is an area of active research and is populated by a large number of algorithms. Section 3.1 provides an overview of visualisation techniques and establishes the focus of the work presented in this chapter, which is the extraction and visualisation of opaque surfaces from scalar volumetric data. The ray-casting paradigm in particular is described and variations on interpolation and shading methods are discussed with respect to the tasks of rendering surfaces from ultrasound data and rendering bone surfaces from CT data. A new 3D B-spline interpolant is formulated which extracts smooth iso-value surfaces from voxel data. In this technique discretization and partial voluming artifacts are reduced by approximating sampled data at voxel centres. Furthermore, surface normals, necessary for realistic shading of surfaces, are explicitly defined by the B-spline surface rather than estimated by an *ad hoc* method.

3.1 VOLUME VISUALISATION

The purpose of scientific visualisation is to extract meaningful information from complex datasets through the use of interactive graphics. Volume visualisation is concerned with the representation, manipulation, and rendering of 3D volumetric data. In the context of medical imaging this usually involves revealing underlying anatomy and organ morphology for diagnostic or treatment planning purposes.

Volume visualisation involves three distinguishable aspects, data representation, data classification or surface extraction, and data viewing. The primary source of volume data in medical imaging is empirical. A discrete sampling of an object, usually as a sequence of cross-sectional scans, is generated by a medical scanner. Most viewing algorithms use the 3D voxel array representation. The volumetric dataset resides in a discrete voxel space which is a 3D integer grid of unit volume cells or elements called *voxels*.

A voxel is the 3D counterpart of the 2D pixel. Each voxel is a quantum unit of volume and has a numeric value (or values) associated with it that represents some measurable properties of the sampled objects (e.g. density, refractive index, acoustic impedance, velocity). Visualising 3D Doppler ultrasound data [Rickey *et al.* 1991, Burrell *et al.* 1989] is an example of visualising a vector field. This thesis, however, is concerned with the visualisation of scalar quantities, in particular, density in the case of CT data and changes in acoustic impedance in the case of pulse-echo ultrasound data.

Most visualisation methods require that the voxels lie on a uniform orthogonal grid. A few methods have been developed for rectangular lattices of uneven spacing and connected lattices of irregularly shaped cells [Giertsen 1992, Shirley and Neeman 1989]. The uniform voxel representation is convenient for the storage and manipulation of volumetric data within a computer, and also for interpolating between voxel centres. A voxel is either represented as a rectangular prism centred at a grid point or, interchangeably, as a zero-dimensional point located at the grid coordinates. The aggregate of voxels tessellating the volume forms the volumetric dataset. CT and MRI datasets are conveniently *reconstructed* into a regular 3D voxel dataset by stacking parallel cross-sections. However, non-parallel slices, such as those acquired by a six degrees-of-freedom hand-held ultrasound probe, are not readily suited to this representation. These require the development of new viewing algorithms or alternatively, reconstruction techniques which generate a regular array of voxels from the data and therefore allow conventional methods to be used.

3.1.1 Viewing algorithms

Viewing algorithms generate 2D images from 3D data by projecting the 3D data in a specified direction on to a 2D viewing plane. There is a necessary reduction of information, or filtering of information, and this depends on the projection method. Resectioning of the data volume and display of the resulting 2D cross-sections is the simplest viewing method, and is popular among radiologists. However, the 3D nature and structure of imaged objects is lost in doing so. An alternative is to render structures or "objects" within the data in a way which is similar to how 3D objects are naturally viewed in the physical world. There are many viewing algorithms and variations which attempt to do this. Most can be classified into two broad categories: surface rendering techniques and volume rendering techniques.

3.1.1.1 Surface rendering

Surface rendering implicitly assumes that the data possesses tangible surfaces that can be extracted and visualised. In effect a binary classification of the data volume is made. Traditional surface rendering algorithms convert the volume data into geometric primitives in a process variously known as iso-surfacing, iso-contouring, surface extraction, or border following. The geometric primitives (e.g., a polygon mesh, a set of contours) are then rendered to the screen by conventional computer graphics algorithms [Foley *et al.* 1990, Shuey *et al.* 1986]. Some well-known examples are the *cuberille* technique [Chen *et al.* 1985, Herman and Liu 1979], the *marching cubes* algorithm [Lorensen and Cline 1987, Cline *et al.* 1988] and the *dividing cubes* algorithm [Cline *et al.* 1988].

The cuberille technique generates rectangular polygons at every voxel, thus representing the volumetric dataset as a binary array of opaque cubes. Linear interpolation or shape-based interpolation [Herman *et al.* 1992] between adjacent 2D data slices is often used to generate intermediate slices so that an array of approximately cubic voxels is achieved. The *marching cubes* algorithm generates a fine polygon mesh by sequential tessellation of the boundary with tiny triangles. The sampled values are placed at the corners of a voxel and each voxel which contains a boundary transition is identified as belonging to one of 15 topologically distinct patterns. A look-up table generates the appropriate structure of triangular patches. The *dividing cubes* algorithm estimates the path of a surface through boundary voxels by subdividing the voxels into smaller cubes of similar size to the pixels in the image plane. New sample values are interpolated at the sub-voxel centres, those sub-voxels which lie on the surface boundary being identified and directly projected on to the image plane.

Purely geometric surface representations often lack realism and detail. The use of normals from the constituent facets to shade the image results in "flat shading". The shading of geometric representations can be substantially improved by using surface gradient estimates which are closer to those of the sampled object, assuming that the actual surface is probably smooth and curved [Hohne *et al.* 1990]. Chen *et al.* [1985] use a contextual shading technique which considers the adjacent facets to better estimate the actual surface gradient. Other researchers use volume gradients calculated at voxel centres, also known as grey-level gradient shading [Hohne and Bernstein 1986]. The gradients are interpolated across facets according to the well-known Gouraud and Phong shading methods [Foley *et al.* 1990] to provide smooth, realistic shading of pixels in the image plane.

3.1.1.2 Volume rendering

Volume rendering techniques use opacity (attenuation), colour, scattering and emissivity to generate images [Dreblin *et al.* 1988, Levoy 1988, Ney *et al.* 1990, Yoo *et al.* 1992, Sabella 1988, Westover 1990]. A binary classification of the data is avoided as light is modelled as being partially attenuated and reflected as it passes through the data volume. Consequently, volume rendering is dependent on a process called *classification* which involves the assignment of properties such as opacity (attenuation coefficients), colour, luminescence, etc, to the data. This contrasts with surface rendering which is dependent on the surface extraction process, or the binary segmentation of the data. Levoy [1988] and others argue that in cases where binary segmentation is difficult, a semi-transparent modelling of an object can improve the visualisation. For example, in the case of visualising bone from CT data, thresholding may fail to detect a thin bone surface, resulting in apparent holes in the rendered surface which do not actually exist. This is because a voxel grey level represents the average density of the tissue within it; thus the grey level may no longer be representative of the membership of a particular tissue class, an effect known as "partial voluming". Such fine bone structures occur in the orbits and nasal septum. Consequently, Levoy [1988] uses both grey-level value and grey-level gradient to assign transparency. Although transparency methods can sometimes compensate for partial voluming effects and generate smoother images [Hohne *et al.* 1990], the assignment of opacities and related parameters is arbitrary and a broad range of visual impressions can be obtained for the same object. It is difficult to know which rendering is the most accurate and how precisely the rendering represents the actual surface. Tiede *et al.* [1990] have attempted to quantify image quality in their study of rendering methods. They conclude, with particular reference to rendering bone surfaces from CT data, that if exact surfaces can be determined then non-transparent rendering often yields the best perception.

Rendering volumetric properties such as transparency usually requires the use of image-space methods, where the image plane is mapped on to the data by ray-casting. Because the volume is sampled along each viewing ray, the volume data must be stored and re-sampled in order to compute each new view. The computing overheads for these algorithms is especially great when the whole data volume does not fit into the memory of the hardware. In contrast, the majority of surface rendering techniques use object-space algorithms where the object is mapped on to the image plane. The object surface is stored in a data structure suitable for subsequent transformations and consequently it is often faster to compute new views and there is no need to store the complete data volume in memory. Typically, interactive rendering in real-time is possible on a moderate-sized serial machine.

Udupa and Odhner [1993] have proposed a technique for rendering semi-transparent surfaces which partly overcomes the high computational cost and storage requirements of ray-casting methods. They define a data structure called the *shell*, whereby a surface is identified as a subset of voxels. The shell is effectively a fuzzy set of boundary points which can be projected directly on to the image plane. *Splatting* is another object-oriented method where every voxel is projected on to the image plane. Each voxel is 'splatted' on to the image plane in the direction of the viewer and makes a weighted contribution to the shading of a neighbourhood of pixels determined by an appropriate 'spreading' function. The general method is described further in [Westover 1990, Hanrahan and Laur 1991, Yoo *et al.* 1992]. Splatting is suitable for general volume rendering and has been used to render CT and ultrasound data [Nelson and Elvins 1993]. An interesting and conceptually satisfying approach is the use of the Fourier Projection-Slice theorem to generate volume rendered images [Totsuka and Levoy 1993, Malzbender 1993]. A view from any direction and orientation can be obtained by inverse transforming the appropriate data slice through the 3D Fourier-transform of the volume data. However, images lack depth information similar to a conventional X-ray image. Totsuka and Levoy [1993] describe frequency domain shading and depth cueing techniques but the main difficulty with Fourier methods is the interpolation required in Fourier space. Although progressive refinement to final image is more straightforward than with spatial domain volume rendering, and spatial filtering is simple to implement, Fourier methods are not widely used.

3.1.2 Summary

Volume and surface rendering strategies can be complementary depending on the goal — whether it be qualitative or quantitative. Levoy [1990] has even developed a hybrid ray-tracer which combines polygon surface rendering with transparent volume rendering. This system allows opaque surfaces and transparent volumes to be visualised simultaneously. Despite being topical in the computer graphics literature, semi-transparent rendering is regarded as being of limited clinical value. Irrespective of computational considerations and image quality, fuzzy or transparent renderings are not suitable for making quantitative measurements or for fitting surfaces due to the ambiguous depiction of anatomy. In Chapter 5 of this thesis the aim is to make quantitative volume measurements from 3D ultrasound data. In Chapter 6 the problem of fashioning cranial implants to fit the skull is considered and computer graphics are used to extract surfaces from CT volume data to which an implant is later fitted. Consequently, this thesis focuses on surface rendering. The ray-casting viewing algorithm is used to depict opaque surfaces rather than semi-transparent volumes.

3.2 RAY-CASTING

Ray-casting is an image-oriented method which can be used to render opaque or transparent surfaces. Tuy and Tuy [1984] were among the first researchers to apply ray-casting to the display of binary volumetric data. Subsequent implementations have varied primarily in their use of transparency, shading and in coding details which improve the computational efficiency of the method. This discussion is restricted to rendering opaque, iso-value surfaces from regular voxel array data.

The terms ray-casting and ray-tracing are sometimes used interchangeably in the literature. However, Kaufman [1991] makes a distinction between the two, with ray-casting being the process of casting sight rays from a viewing plane through the volume data and ray-tracing being the recursive process of tracing rays from a viewing plane to multiple light sources. Ray-casting is therefore a special case of ray-tracing by this definition.

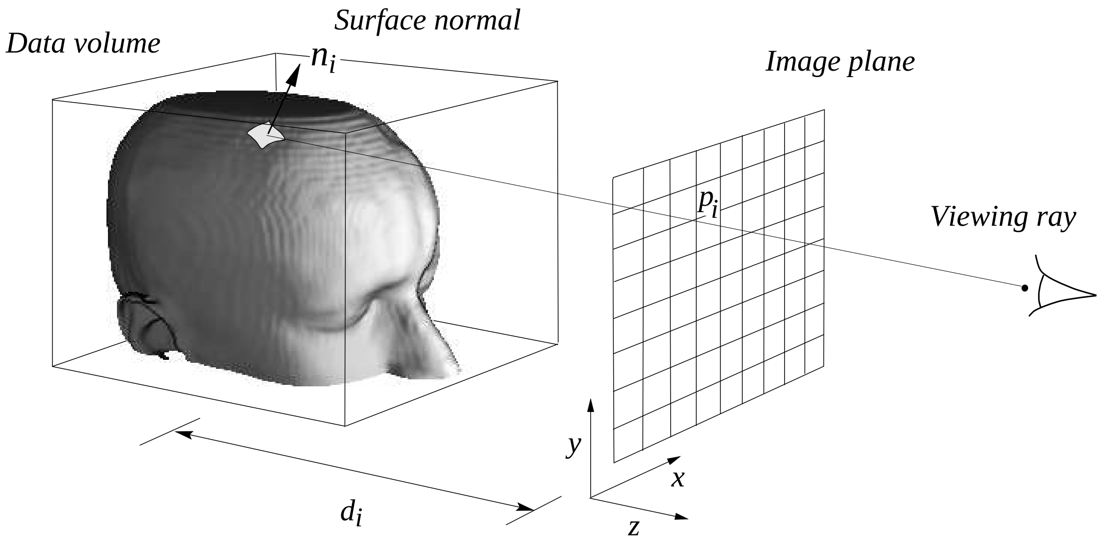

Figure 3.1 The ray-casting algorithm used in this thesis. For each pixel p_i in the image plane a ray is cast into the data volume until a surface, determined by some function of the intensity profile, is encountered. The distance d_i from the image plane to the surface is stored in a depth-buffer or Z-buffer together with the grey-level gradient computed locally at the surface point. Images are rendered from depth and gradient information in the Z-buffer. (The skin surface in this figure was rendered from CT data by thresholding the intensity data to segment regions denser than air).

Figure 3.1 illustrates the principle of the ray-casting method used to render surfaces in this thesis. Rays emanate from a viewpoint, pass through pixels in an image plane, and then intersect the data volume. For each pixel in the image plane a ray is tracked through the volume until either a surface is detected or the volume is exhausted. A surface is determined by thresholding the scalar intensity distribution. Consequently, an iso-valued surface is rendered. Figure 3.1 suggests a perspective projection, but parallel rays are often used in medical imaging.

Intensity values are interpolated between voxel centres. Tri-linear interpolation, where an intermediate value is computed from the linear combination of the 8 nearest voxels (Figure 3.2), is commonly used. Here the interpolation neighbourhood forms a rectangular prism whose vertices are defined by the eight nearest voxels. A ray can propagate efficiently through the volume data by determining when it enters a new voxel neighbourhood. If no value in the neighbourhood exceeds the threshold then the next neighbourhood is considered. Otherwise the ray is traversed in fine increments, interpolating a new data value at each increment, until a threshold transition is detected or the neighbourhood is exited. In this way a surface transition is determined with sub-voxel resolution and aliasing is avoided.

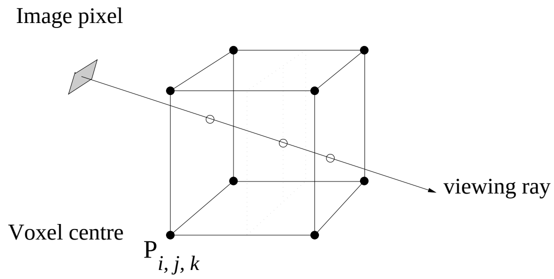

Figure 3.2 Tri-linear interpolation between voxel centres.

Whether a threshold transition can be determined analytically along a ray depends on the interpolation strategy used. Incremental evaluation along a ray to detect threshold transitions is commonly used and is often unavoidable. Consequently, it is conceivable that fine structure might be missed if the spacing between steps is coarse. The associated computation cost increases linearly with the number of steps. However, prior processing of the volume data to restrict the number of voxels through which rays are traced substantially reduces the rendering time, particularly when a more computationally intensive interpolation strategy is used such as the cubic B-spline interpolatant described in Section 3.3.3.

An approximation to tri-linear interpolation which is sometimes used avoids incremental evaluation along a ray. It uses bi-linear interpolation to evaluate the scalar intensity at a ray's entry and exit points from a voxel neighbourhood. The intensity profile along the ray is then linearly interpolated between these two points. Using this method, a point in the data volume will generally be evaluated with slightly different values depending on the ray path taken through the point, i.e. for different views, the scalar distribution will appear to be multi-valued. However, this effect is not usually noticeable when voxels are comparable in size to pixels in the image plane.

Some authors apply viewing transformations to the data, rather than the rays, with the advantage that a single coordinate transform is performed per viewing direction. Tri-linear interpolation is used to assign grey levels to voxels in the rotated volume.

Each row of voxels in the transformed data volume aligns with a single pixel in the image plane so that rays then propagate along the voxel rows. This simplifies the implementation of the ray-caster and avoids excessive paging of the data set into and out of memory, which may otherwise occur when rays are cast with arbitrary orientations through the data. This approach comes at the cost of increased storage requirements and tends to break even in terms of computation when properties such as transparency are rendered, i.e. when the image is composited from every voxel in the data volume. It is therefore of limited advantage when binary surfaces are rendered.

The basic ray-caster developed in the course of this thesis is complicated by provisions for cut-away views, multiple light sources, multiple thresholds and distributed processing. It employs tri-linear and cubic B-spline interpolation to super-sample the data. The processes of ray-casting and rendering images are implemented by separate packages. The ray-caster outputs a linked list of *depth-maps* which the rendering package uses to form images.

The depth-map structure stores the Z-buffer (the table of distances from image plane pixels to surfaces) for a particular view together with the associated surface gradients calculated in object-space. Gradient information is encoded into three bytes per image pixel. The viewing parameters are also stored in the structure along with the Z-buffer information. This means that the position where each ray terminates on a surface can be computed in object space. The significance of this is that it allows ray-casting to be recommenced from the point on the surface where it left off. This means that a higher threshold can be efficiently rendered from the same viewpoint by continuing ray-tracing along the existing ray paths, or that surface gradients can be estimated in a different manner, or that ray-tracing can be continued to another light source.

3.3 SURFACE SHADING

The viewing process, in this case ray-casting, determines the visible parts of a 3D scene. Shading determines the intensity value that reaches the viewer's eye from each point on the visible surface. The goal of shading is optimal visual recognition or discernment of the displayed objects and to show the natural appearance of real objects.

Depth-only rendering is the simplest method for producing images from a depth-map. The intensity at each pixel in the image plane is determined solely by the distance, d, from the pixel to the surface; surfaces which are far from the viewer appear darker than those which are closer. In Figure 3.6(a) grey-levels are assigned inversely proportional to the distance from the image plane squared. The square term models the natural fall off in intensity of light originating from a point source. The response of the eye to changes in intensity is logarithmic over a wide range of intensities [Pratt 1991], hence a logarithmic assignment of grey-levels to a linear variation in depth is perceived as a linear progression in intensity. Other authors [Tiede *et al.* 1990, Gordon 1985, Chen *et al.* 1985] drop the square factor and assign grey-levels inversely proportional to distance, arguing that natural lighting has properties intermediate between a point source and an extended light source which emits parallel rays. The grey-level range is normalised to make optimal use of the 8-bit dynamic range typically available on grey-scale displays. The scale is normalised so that the illumination range, from maximum to zero intensity, corresponds to the diameter of the sphere which circumscribes the data volume. This avoids apparent fluctuating illumination when the data volume is rotated in an animated sequence. The default normalisation parameter values are chosen empirically after extensive experimentation.

Depth-rendered images generally appear smooth because small changes in the distance from the observer result in small changes in grey-level. The eye's ability to discern small changes in intensity is very poor since the human perceptual system is edge-detection oriented [Marr and Hildreth 1980, Frisby 1979]. Psychological tests reveal that the eye is actually capable of distinguishing little more than 32 grey levels [Foley and Van Dam 1982]. Consequently, there is a loss of textural information and although this smoothes out any sampling noise, it also suppresses minute details such as discontinuities in the surface that can be of great significance in medical imaging. For example, little detail is evident in the cranial surface depicted in Figure 3.6(a). The addition of shading based on the surface gradient substantially increases the information revealed. The following section investigates more sophisticated techniques based on modelling the way natural objects are illuminated.

3.3.1 Shading models

The shading observed in scenes originating from the physical world around us are the result of complex interactions between light and objects. Grimson [1981] provides a detailed discussion of how these images are formed when a simple scene consisting of opaque objects is considered.

The apparent "brightness" recorded by an imaging device is a measurement of image *irradiance*, the radiant flux striking a pixel in the image plane. Provided that the imaging system is properly focused, the flux reaching a small part of the receptive field will be a function of a corresponding small surface element. The irradiance value associated with a particular image point is a function of three factors, the amount of radiant flux striking the surface, the percentage of incident flux reflected by the surface material and the distribution of reflected flux as a function of direction.

The amount of flux absorbed depends on the surface material. The behaviour of reflected flux is determined by the surface microstructure, that is those features too small to be resolved by the imaging system. If the surface structure is flat and homogeneous, the reflected ray will lie in the same plane as that formed by the incident ray and the surface. The surface normal will coincide with the normal of the incident and reflected rays. Such a material is referred to as "specular", "metallic" or "dielectric". If the surface is not perfectly flat on a microscopic scale, it will scatter parallel incident rays in many directions. For small deviations of the local surface normal, most rays will lie in the direction for ideal specular reflection. The perceived surface is referred to as "shiny" or "glossy". If the surface layer is not homogeneous on a microscopic scale, the light rays which penetrate the surface will be scattered by refraction and reflection at boundaries between regions with differing refractive indices. The scattered rays may re-emerge in many directions. Such surfaces are referred to as having a "diffuse" or "matte" reflection. Thus, the distribution of reflected light depends on the direction of incident rays and the details of the surface microstructure.

In order to specify the effects of surface microstructure on the reflected light, a reflectance function, known as the Bidirectional Reflectance-Distribution Function (BRDF), is associated with any real surface. The BRDF specifies reflectance in terms of incident- and reflected-beam geometry, it specifies how bright a surface element will appear when illuminated from a given direction and viewed from a second given direction. If a surface possesses the property that the reflectance of a surface element is invariant to rotation about its normal, then the beam geometry of the situation is greatly simplified and only three angles are needed to determine reflectance (Figure 3.3(b)). Such surfaces are termed *isotropic*.

In the field of medical graphics the primary intention is not to accurately model the interaction of light so that rendered objects appear as they would in the physical world, but to achieve a representation which readily, and faithfully, communicates to the viewer the 3D nature of structures present in the data. Computational reasons alone make the former an unrealistic proposition. However, effective shading can be implemented efficiently by ignoring absorption and modelling the light reflected from a surface as having diffuse, specular and ambient components.

Ideal diffuse reflecting surfaces appear to have the same brightness from all viewing directions; the brightness depends upon the illumination which is given by

$$I = I_{max} \cos\theta \tag{3.1}$$

where θ is the angle between the incident light and the normal to the surface (Figure 3.3(a)). I_{max} is the maximum illumination which results when $\theta = 0$. Equation 3.1 can therefore be written as

$$I = I_{max} \mathbf{N} \cdot \mathbf{L} \tag{3.2}$$

where $\mathbf{N}$ is the unit normal to the surface and $\mathbf{L}$ is the unit vector along the line from the surface to the light source. Diffuse reflectance can be combined with depth rendering to enhance the 3D perception. Unless otherwise stated, the light source is assumed to coincide with the viewpoint in the ray-casting paradigm, thus the complication of shadows is avoided. Depth shading is also beneficial to the comprehension of a 3D scene when a projection of parallel rays is taken, even though a fall off in illumination is inconsistent with the physics of a planar source.

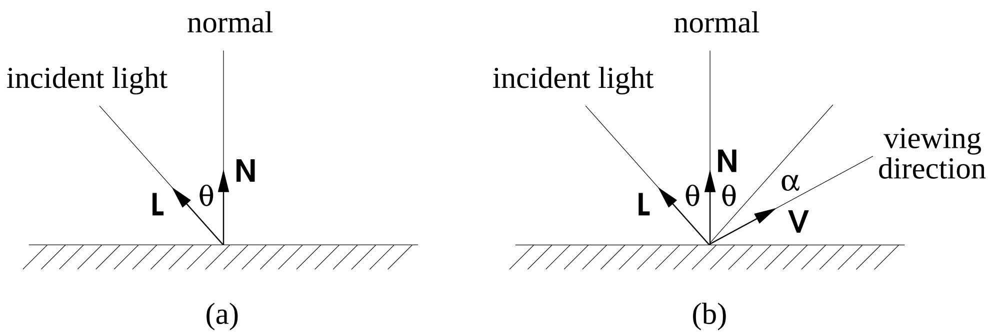

Figure 3.3 (a) Diffuse reflection and (b) specular reflection.

Some researchers find that the $\cos\theta$ term in (3.1) creates too harsh an angular dependence and replace it by $(\cos\theta)^p$ where p is an empirically determined parameter. However, this is usually done to smooth out variations of θ which occur when a noisy estimate of the surface normal is available. A rapidly changing surface normal will manifest itself as a series of light and dark bands in an otherwise smooth surface. Throughout this thesis $p = 1$ is used.

Specular reflection (observed as a highlight on shiny surfaces) depends on the direction of the incident light **L** and the viewing direction **V** (Figure 3.3(b)). It can be approximated by Phong's empirically derived model,

$$I = I_{max} W(\theta) \cos^n \alpha \tag{3.3}$$

where n is the exponent used to approximate highlights, α is the angle between the reflected ray and the viewing direction, and $W(\theta)$ represents the fraction of specularly reflected light as a function of the incident angle. The latter is often set to a constant value $W(\theta) = k_s$.

Many researchers also include an ambient light component in their shading model so that surfaces shielded from the main light source are still visible. The complete shading model, known as the *Phong shading model* [Foley and Van Dam 1982], can then be expressed as a combination of three components,

$$I = I_a k_a + I_p F(d)[k_d L \cdot N) + k_s (R \cdot V)^n] \tag{3.4}$$

where

$$\begin{aligned}
I &= \text{observed brightness} \\
I_a &= \text{ambient light} \\
I_p &= \text{point light source} \\
n &= \text{exponent used to approximate highlights} \\
\mathbf{V} &= \text{normalised vector in the direction of the viewer} \\
\mathbf{L} &= \text{normalised vector in the direction of the light source} \\
\mathbf{R} &= \text{normalised vector in the direction of the reflected ray} \\
k_a &= \text{fraction of ambient reflection} \\
k_d &= \text{fraction of diffuse reflection} \\
k_s &= \text{fraction of specular reflection} \\
F(d) &= \text{factor determining illumination fall off with depth}
\end{aligned}$$

This model is based on empirical observation. Theoretically-based models have also been developed. These model surfaces as collections of microscopic facets, each being a perfect reflector, with orientations determined by a Gaussian probability distribution. The Phong model, however, is more efficient to evaluate. Models similar to (3.4) are described in [Kaufman 1991, Tiede *et al.* 1990, Levoy 1988].

Many researchers have found diffuse reflection combined with depth-rendering is adequate for medical imaging purposes [Chen *et al.* 1985]. Figure 3.4 compares a purely diffuse rendering with that produced by (3.4) using parameters suggested by Levoy [1990] $k_a = 0, k_d = 0.8, k_s = 0.2$ and $n = 5.0$. Here the viewpoint and a single point-light source coincide. Slightly more detail is apparent in the Phong rendered skull on surfaces which are almost perpendicular to the direction the light source, in this case the forehead above the right eye. This is where the addition of a specular reflection component has the most effect. Establishing the accuracy of a rendered object is difficult and has been investigated by Tiede *et al.* [1990] using artificial geometric objects and cadavers (for bone surfaces). In the interest of minimising the number of parameters and facilitating comparison between reconstruction methods, the simpler diffuse reflectance model was chosen and used throughout the following figures.

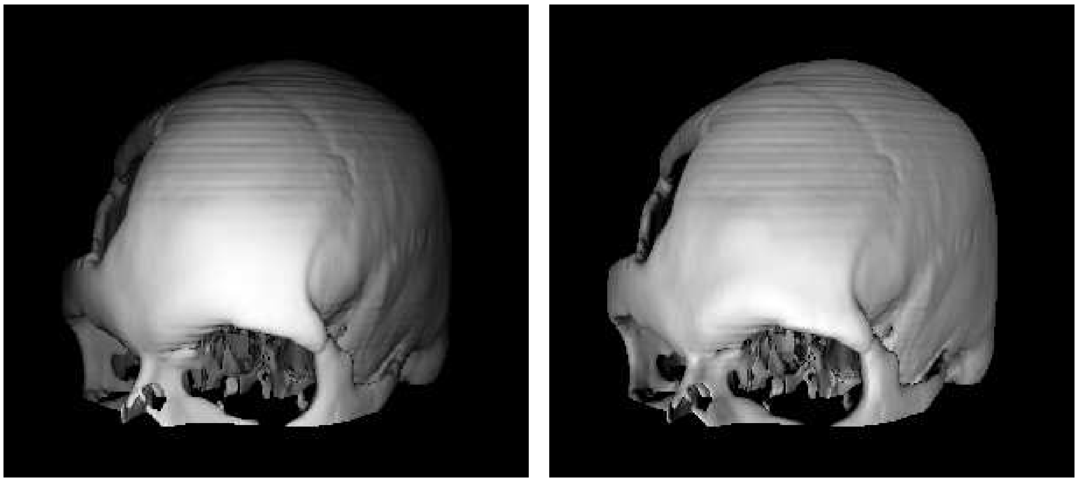

(a) (b)

Figure 3.4 Comparison of diffuse-only shading and Phong shading with a specular reflection component. (a) Diffuse shading, (b) Phong shading.

The important observation to be made concerning lighting models is that they are fundamentally dependent on the determination of surface normals. Other parameters are usually either known or can be easily determined. Consequently, estimating surface normals is the main graphics issue.

Surface normals can either be determined by object-space or image-space methods. Object-based shading uses voxel intensities in the 3D data to estimate surface gradients whereas image-space methods are computed from the Z-buffer alone. The former usually provide better results and exhibit a greater dynamic range of surface normals. However, these can not always be used. An example is the direct method for rendering non-parallel ultrasound slices described in Chapter 4, Section 4.3.

3.3.1.1 Object-space gradient estimation

A consequence of the tomographic data acquisition process is that voxel grey-level values reflect the relative average of various tissue types in the immediate voxel neighbourhood. These relative volumes are related to the surface inclination. Thus, the grey-level gradient at a surface voxel can be considered as a measure of surface inclination.

A grey-level gradient can be estimated at a voxel by considering the 6 central neighbours in a 3x3x3 voxel neighbourhood. Denoting the grey level of a voxel at location i, j, k by $f_{i,j,k}$, the grey level gradient is computed as:

$$\begin{aligned} F_x &= f_{i+1,j,k} - f_{i-1,j,k} \\ F_y &= f_{i,j+1,k} - f_{i,j-1,k} \\ F_z &= f_{i,j,k+1} - f_{i,j,k-1} \end{aligned} \tag{3.5}$$

Tri-linear interpolation is used to determine grey-level gradients between voxel centres. The components of the normals to a surface are normalised as:

$$N_u = \frac{F_u}{\sqrt{(F_x)^2 + (F_x)^2 + (F_x)^2}}, \qquad u = x, y, z. \tag{3.6}$$

Figure 3.6(b) illustrates the rendering of a skull from X-ray CT data with diffuse surface shading using grey-level gradients. Superior detail is apparent in the bone surface which appears featureless in the depth-only shaded image in Figure 3.6(a). In general, the grey-level gradient highlights surface detail and provides a more realistic depiction. Small details not apparent with depth shading alone are revealed.

Other object-based gradient techniques, such as Chen's contextural shading method [Chen *et al.* 1985], have been specifically developed for the the cuberille (binary voxel) surface representation. These methods tend to result in inferior shading since the dynamic range of surface normals is somewhat limited by the finite number of topological possibilities in a binary voxel neighbourhood.

The grey-level gradient approach can fail in the case of structures which are small or thin relative to the size of a voxel. Here the grey level gradient is no longer representative of the surface inclination and may result in artifacts. The clinician must bear this in mind. Erroneous gradients are likely in CT data around regions such as the orbits. Data which is noisy or contains speckle, such as ultrasound data, is also likely to result in poor gradient estimation using the grey-level gradient.

3.3.2 Z-buffer gradient estimation

Gradients can also be estimated directly from a depth-map (Z-buffer) without reference to the 3D data. A depth-map can be viewed as a function of two variables $z = z(x, y)$ and the surface normal may be obtained from the gradient vector $\bigtriangledown z$. The partial derivatives $\frac{\partial z}{\partial x}$ and $\frac{\partial z}{\partial y}$, form the components of the normal in image-space.

$$\mathbf{n} = (-\frac{\partial z}{\partial x}, -\frac{\partial z}{\partial y}, 1). \tag{3.7}$$

This involves determining the derivatives $\frac{\partial z}{\partial x}$ and $\frac{\partial z}{\partial y}$ numerically from the Z-buffer distances $z_{i,j}$, where i, j are integer indices specifying the location of a pixel in the Z-buffer. Consider evaluating $\frac{\partial z}{\partial x}$ at the pixel (i, j). $\frac{\partial z}{\partial x}$ may be approximated by the *forward difference*

$$\delta_f = z_{i+1,j} - z_{i,j}, \tag{3.8}$$

or by the *backward difference*

$$\delta_b = z_{i,j} - z_{i-1,j}, \tag{3.9}$$

or by the average of the two,

$$\delta_c = \frac{1}{2}(z_{i+1,j} - z_{i-1,j}), \tag{3.10}$$

The last is known as the *central difference*. The central difference will generally provide the better approximation if z is a continuous function of x and is adequately sampled by Nyquist's theorem. Figure 3.6(c) uses the central difference to estimate both $\frac{\partial z}{\partial x}$ and $\frac{\partial z}{\partial y}$ and to shade the image.

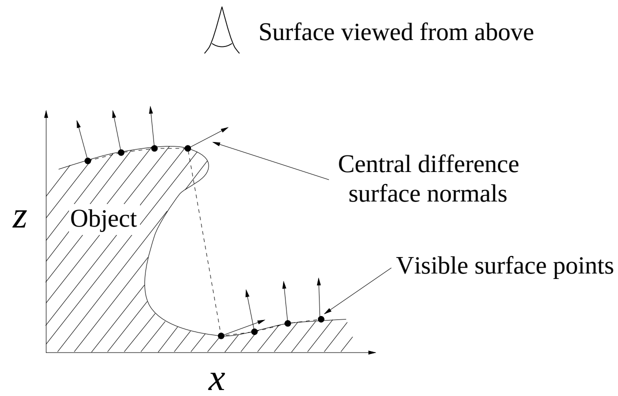

Figure 3.5 Failure of the central difference method to correctly estimate surface normals when a discontinuity occurs in the depth-map due to a hidden surface.

The central difference estimate can fail when surface and tangent discontinuities occur in the depth-map. Consider Figure 3.5 where the viewing direction and primary light source are in the direction of the negative z axis, i.e. the surface is viewed from above. The hidden object surface results in a discontinuity in the depth-map which is treated as a steep gradient by the central difference method. Consequently, the surface normals estimated either side of the discontinuity are incorrect and the discontinuity is rendered as a dark band in the image since little light is reflected in the direction of the viewer. Such artifacts are apparent in Figure 3.6(c) in the region of the zygomatic arch. This 'shadow' artifact can sometimes appear pleasing to the eye since it emphasises the presence of a discontinuity in the surface by enhancing edge contrast — in effect behaving like an edge detection operator.

3.3.2.1 Correcting for discontinuities

Gordon [1985] has proposed a method which takes a weighted average of the forward and backward differences to provide a more reliable estimate of surface gradient from parallel projections,

$$\frac{\partial z}{\partial x} = \frac{W_b(|\delta_b|) + W_f(|\delta_f|)}{W_b + W_f} \tag{3.11}$$

where

$$W(\delta) = \begin{cases} 1 & \text{if } \delta \leq a \\ \epsilon & \text{if } \delta \geq b \\ \frac{1+\epsilon}{2} + \frac{1-\epsilon}{2}\cos(\frac{\delta - a}{b-a}\pi) & \text{otherwise.} \end{cases} \tag{3.12}$$

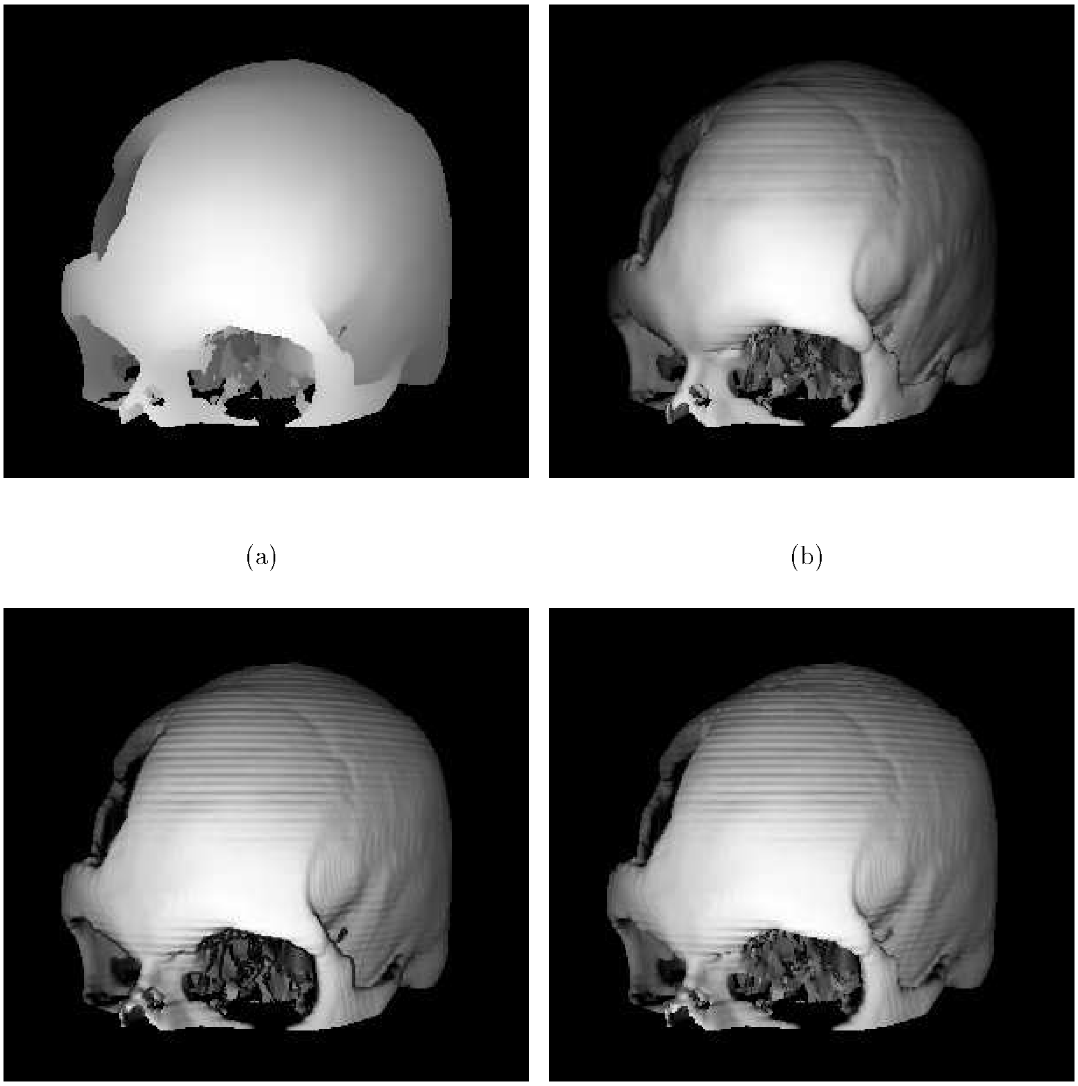

(a) (b)

(c) (d)

Figure 3.6 Surface rendering of a CT data set. (a) Depth rendered view, (b) voxel gradient shading, (c) Z-buffer gradient shading, (d) Corrected Z-buffer gradient shading.

Gordon recommends the values $a = 2, b = 5$ and $\epsilon = 10^{-5}$. The parameter a specifies the relative depth below which the maximum weight is assigned, b is the relative depth above which the minimum weight is assigned, in between the weight assigned varies smoothly and continuously between the two extremes according to the cosine function.

If the forward and backward weights are equal then the central difference is obtained. The value $a = 2$ corresponds to a gradient of 63° . Thus, gradients steeper than this are given little weighting by (3.12). When a difference is steeper than 78° (corresponding to $b = 5$) the minimum weight is assigned. If both the forward and backward differences are large, then the central difference will be obtained since both will be assigned small, but equal weights. The evaluation of $\frac{\partial z}{\partial y}$ is similar.

Gordon's method is widely used; however, it was originally developed to estimate surface gradients from discrete voxel surfaces. The main concern was to distinguish between actual and *artificial* edges in order to minimise significant changes in shading due to small changes in viewing orientation. Artificial edges are those corresponding to discrete jumps in depth due to discretization of a single continuous surface rather than abrupt changes from one surface to another.

3.3.2.2 A new method for determining surface gradients

Gordon's method can fail in places of rapidly changing slope, resulting in incorrect normals. This is because the method does not explicitly identify the presence of a discontinuity and the weights in the model are a function of the absolute value of the forward and backward differences — no account is taken of the possibility of differing sign. In this thesis a predicate-based approach was developed to assign normals using the second derivative (the change in gradient) to explicitly identify the existence of a discontinuity and act accordingly. The algorithm considers the general case of a perspective projection, and calculates gradients in object-space.

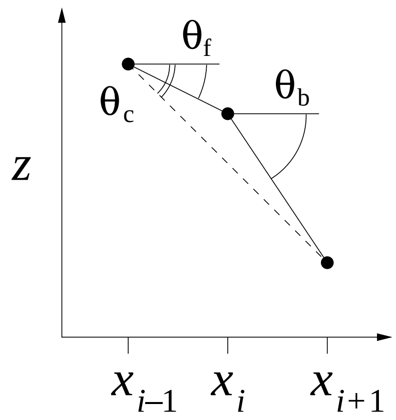

Figure 3.7 Geometry and gradients for estimating $\frac{\partial z}{\partial x}$ at x_i.

Consider determining $\frac{\partial z}{\partial x}$ at the point x_i (Figure 3.7) where the forward, backward and central differences are expressed by the angles, θ_f, θ_b and θ_c, respectively. Note that the x_i's need not necessarily be equally spaced. Figure 3.8 summarises the algorithm used to estimate $\frac{\partial z}{\partial x}$ at the point x_i. Nine cases are sequentially identified and each is illustrated with a typical example. The first three cases deal with the possibility of one or both of the adjacent Z-values being *undefined*, which occurs when the corresponding ray does not intersect a surface. The following six cases consider the prospect of a discontinuity occurring in the intervals x_{i-1}, x_i and x_i, x_{i+1}.

A large change in depth between two adjacent points in the depth-map alone is not indicative of a discontinuity in the surface profile. A smooth continuous surface segment oriented obliquely to the viewing direction may also exhibit a steep profile. The change in gradient, must also be considered when determining whether a discontinuity occurs. Gordon's weight allocation indirectly achieves this in some cases.

Two parameters, θ_{max} and $\Delta\theta_{max}$, are used to detect discontinuities. θ_{max} relates to the magnitude of the forward and backward differences while $\Delta\theta_{max}$ relates to the magnitude of the change in gradient between the forward and backward differences. $\Delta\theta_{max}$ determines the maximum permissible change in surface gradient. If the change in gradient, $|\theta_b - \theta_f|$, is less than $\Delta\theta_{max}$, then no discontinuity is deemed to have occurred in the vicinity of x_i, regardless of whether the gradients over the intervals x_{i-1}, x_i, and x_i, x_{i+1} exceed the threshold θ_{max} (case 4).

If a significant change in gradient does occur at x_i, then the magnitude of the forward and backward differences are considered. In case 5 the magnitudes of θ_f and θ_b are compared to the threshold θ_{max}. If both exceed this threshold and are of opposite sign, then a 'spike' (or an inverted spike) is deemed to have occurred. Both intervals x_{i-1}, x_i and x_i, x_{i+1} therefore contain a discontinuity. If both gradients exceed the threshold θ_{max}, but have the same sign, then no discontinuity is deemed to have occurred even though the change in slope exceeds $\Delta\theta_{max}$. Cases 7 and 8 consider the more usual occurrence where a single discontinuity occurs, either in the interval x_{i-1}, x_i (case 7) or x_i, x_{i+1} (case 8). Case 9, the default case, corresponds to the situation where a significant change in slope has occurred, $|\theta_b - \theta_f| > \Delta\theta_{max}$, but neither θ_f nor θ_b exceeds the threshold θ_{max}. In this case no discontinuity is deemed to have occurred. This provides some immunity to the problem of ripple which may occur when the sampled data contain a noise component (Figure 3.9). It also compensates for a low choice for $\Delta\theta_{max}$.

From experience, the values of 65° for θ_{max} and 20° for $\Delta\theta_{max}$ have been found to produce good results and are used throughout this thesis. Figure 3.6(d) has been shaded using the algorithm described in Figure 3.8. The absence of 'shadowing' on the zygomatic arch is apparent when compared to Figure 3.6(c). Figure 3.10 compares this subregion in detail. The corresponding grey-level gradient image is regarded as a better

Case	**Example**	**Conditions**	**Return value** $(\frac{\partial z}{\partial x})_i$
1	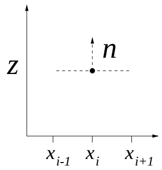	$z_{i-1} = \mathit{undefined},\ z_{i+1} = \mathit{undefined}$	0
2	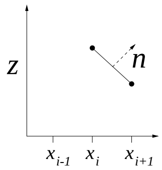	$z_{i-1} = \mathit{undefined}$	$\tan\theta_b$
3	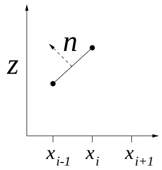	$z_{i+1} = \mathit{undefined}$	$\tan\theta_f$
4	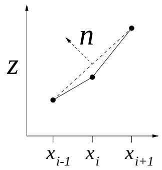	$\lvert\theta_b - \theta_f\rvert < \Delta\theta_{max}$	$\tan\theta_c$
5	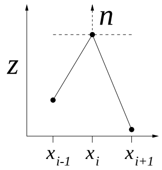	$\lvert\theta_f\rvert > \theta_{max},\ \lvert\theta_b\rvert > \theta_{max},\ \theta_f\theta_b < 0$	0
6	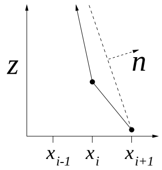	$\lvert\theta_f\rvert > \theta_{max},\ \lvert\theta_b\rvert > \theta_{max},\ \theta_f\theta_b \geq 0$	$\tan\theta_c$
7	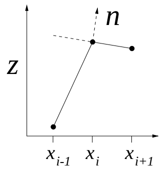	$\lvert\theta_f\rvert > \theta_{max}$	$\tan\theta_b$,
8	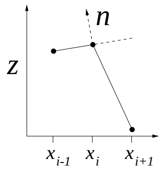	$\lvert\theta_b\rvert > \theta_{max}$	$\tan\theta_f$,
9	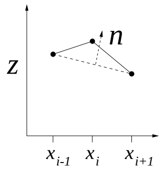	*default case*	$\tan\theta_c$

Figure 3.8 Algorithm for estimating $(\frac{\partial z}{\partial x})_i$ using parameters from Figure 3.7, where the constants $\Delta\theta_{max}$ and $\theta_{max} > 0$.

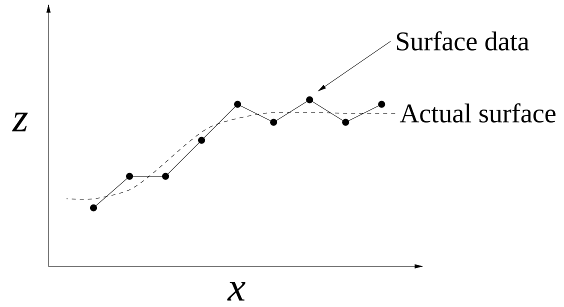

Figure 3.9 Ripple resulting from noisy surface data.

indicator of the actual surface gradient in 3D and is also shown for comparison. The new algorithm has correctly assigned normals in most cases, eliminating the 'shadow' artifact apparent in Figure 3.10(a).

(a) (b) (c)

Figure 3.10 Detailed comparison between Z-buffer shading methods and grey-level gradient shading. (a) Z-buffer shading based on a central difference estimate without discontinuity compensation, (b) Z-buffer shading with discontinuity compensation, (c) grey-level gradient shading.

All Z-buffer shading methods are subject to adequate sampling of the surface to correctly estimate surface gradients. Discontinuities may be falsely assigned or not detected using the method in Figure 3.8 if a surface is under-sampled. The presence of noise in the depth data may also result in a ripple artifact of varying degree, depending on the magnitude of the noise and the spatial sampling frequency. The latter dependency is due to the fixed 3x3 neighbourhood used to determine surface normals. A variable sized neighbourhood, or equivalently, smoothing the surface data at an appropriate scale, is complicated by the presence of discontinuities. The possible complexity of discontinuities in the surface data mean that this is not easy nor practical to compute; finding an adequate mathematical model to describe discontinuities is the first difficulty. These problems are discussed in detail by Terzopoulos in *Regularization of Inverse Visual Problems Involving Discontinuities* [Terzopoulos 1986] and also in Terzopoulos [1988]. As yet, no practical and readily implementable algorithms have been developed for fitting semi-smooth surfaces to discontinuous range data.

Summary

Image-space shading methods typically do not produce images of the same quality as object-space methods. This is partly because the dynamic range of possible surface angles is lower (being limited by the spacing between adjacent rays in the Z-buffer) and surfaces tangential to the viewing direction are not assigned accurate normals. In particular, shallow surface features suffer since the spacing between adjacent surface samples is large relative to the change in depth. Shallow features are generally rendered with lower contrast compared to an object-space shading method such as the grey-level gradient. If these small features are due to noise then reducing their prominence may be desirable. The viewer must therefore be wary of distinguishing between shading artifacts and reality.

A further difficulty with image-space shading is that the normal determined at a surface point depends on the viewing direction and may vary slightly from view to view. However, this effect is not usually noticeable in the final image.

Despite these difficulties, image-space shading is applicable when grey-level gradients are not available or cannot be determined. Such a situation occurs when rendering ultrasound data directly from non parallel slices, since a voxel reconstruction is not available (Section 4.3). Z-buffer shading more accurately portrays the depth-map extracted from the volume data than the interpolated grey-level gradient described in Section 3.3.1.1. For example, the visual ripples apparent in Figure 3.10(b) are due to ripples in the depth-map, yet the same depth-map appears smooth when a grey-level gradient is used to render the surface in Figure 3.10(c). Grey-level volume shading may depict the anatomy better by producing a more pleasing image, but it does not depict the actual depth-map. This is acceptable when the goal is purely visualisation of anatomical structures, but accurate depiction of the depth-map is important when it is to be used for quantitative purposes such as designing an implant.

The next section returns to the viewing algorithm and considers the analytical determination of surface normals by fitting a mathematical function to the volume data. In this way an improved depth-map can be extracted from the data and the *ad hoc* methods currently used for determining surface normals are avoided.

3.3.3 B-Spline Interpolation

Tri-linear interpolation provides continuity of the scalar parameter being visualised, but not of the first and higher derivatives. As a result, the gradient of an iso-value surface rendered from the volume data is discontinuous, i.e. whenever the surface crosses a voxel boundary a crease may occur. The shading of the rendered image is strongly influenced by the surface gradient and therefore small creases in the surface are highlighted. This is particularly apparent when gradients are estimated from the Z-buffer and the surface is semi-parallel to the viewing plane. By using higher order interpolants in the ray-caster it is possible to provide higher degrees of continuity. In this section piecewise cubic B-splines are fitted to the 3D data and used to quasi-interpolate between the sampled data points in order to render a surface which is continuous in the first and second derivatives. Because the volume data lie on a regular grid, the B-spline is formulated in 3D as a tensor product of 1D splines. Consequently, the difficulties associated with a parametric definition are avoided. The aim of this work is to render an improved estimate of the actual surface of an object which is known *a priori* to be smooth and continuous. This is generally true for the problem of rendering the surface of the skull from CT data.

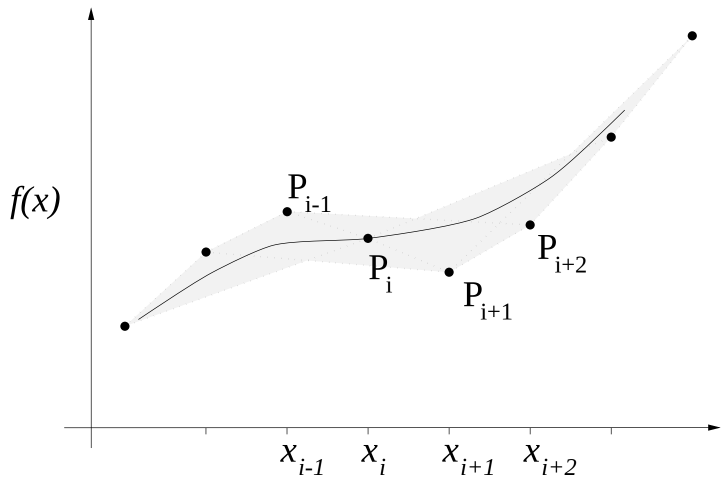

Figure 3.11 1D example of a B-spline cubic through eight points with the convex hull illustrated.

The 3D formulation is developed by first considering the analogous 1D and 2D B-splines in the special case where the control points lie on a regular grid. Consider Figure 3.11. A B-spline $f(x)$ is specified over the interval $x_i \leq x \leq x_{i+1}$ as the weighted average of the four points $P_{i-1}, P_i, P_{i+1}, P_{i+2}$, which are equally spaced along the x axis. Without loss of generality, if the origin is translated to the point x_i and the interval $x_i \leq x \leq x_{i+1}$ scaled such that x varies over the range $0 \leq x \leq 1$, then a B-spline segment which provides continuity of the first and second derivatives with adjacent segments at its endpoints, is given by,

$$f(x) = TM \begin{bmatrix} P_{i-1} \\ P_i \\ P_{i+1} \\ P_{i+2} \end{bmatrix}, \tag{3.13}$$

where

$$T = \begin{bmatrix} x^3 & x^2 & x & 1 \end{bmatrix}, \tag{3.14}$$

$$M = \frac{1}{6} \begin{bmatrix} -1 & 3 & -3 & 1 \\ 3 & -6 & 3 & 0 \\ -3 & 0 & 3 & 0 \\ 1 & 4 & 1 & 0 \end{bmatrix}. \tag{3.15}$$

In general, the spline does not pass through the control points $P_{i-1}, P_i, P_{i+1}, P_{i+2}$. Rather, when the product in (3.13) is evaluated it can be seen that the spline is constrained to lie within the *convex hull* of the four points since the coefficients of the four control points range in value between 0 and 1 and sum to 1 for $0 \leq x \leq 1$. This property is illustrated in Figure 3.11 where the convex hulls for five adjacent segments are shown.

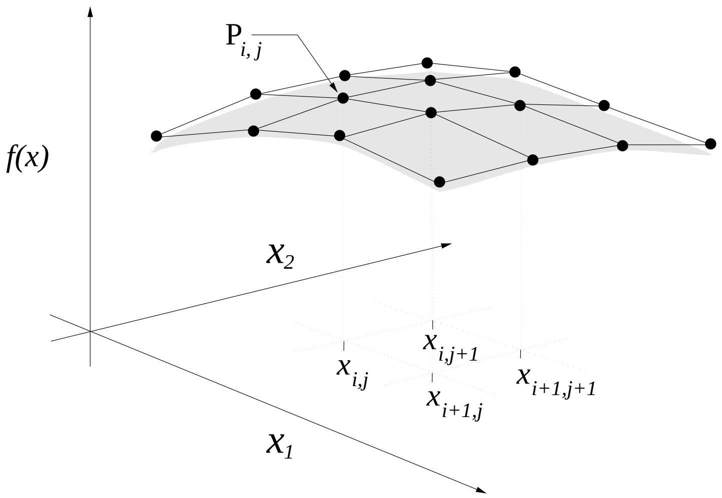

Figure 3.12 Sixteen control points for a 2D B-spline cubic patch.

The 1D formulation (3.13) generalises straightforwardly to 2D [Foley and Van Dam 1982]. In 2D the control points lie on a regular 2D grid in $\mathbb{R}^2$ and the data points are denoted by $P_{i,j}$ where i, j are integer indices indicating the grid coordinates. Consider the patch illustrated in Figure 3.12 bound by the points $\mathbf{x}_{i,j}, \mathbf{x}_{i+1,j}, \mathbf{x}_{i+1,j+1}, \mathbf{x}_{i,j+1}$ where $\mathbf{x} = (x_1, x_2)$. If the nodes defining the patch are first scaled and shifted so that they lie in the unit square, then $f(\mathbf{x})$ is defined over the patch by the 2D B-spline (3.16) for $0 \leq x_1 \leq 1$ and $0 \leq x_2 \leq 1$.

$$f(\mathbf{x}) = SMQM^TT, \tag{3.16}$$

where

$$T = \begin{bmatrix} x_1^3 & x_1^2 & x_1 & 1 \end{bmatrix}, \tag{3.17}$$

$$S = \begin{bmatrix} x_2^3 & x_2^2 & x_2 & 1 \end{bmatrix}, \tag{3.18}$$

and

$$Q = \begin{bmatrix} P_{i-1,j-1} & P_{i,j-1} & P_{i+1,j-1} & P_{i+2,j-1} \\ P_{i-1,j} & P_{i,j} & P_{i+1,j} & P_{i+2,j} \\ P_{i-1,j+1} & P_{i,j+1} & P_{i+1,j+1} & P_{i+2,j+1} \\ P_{i-1,j+2} & P_{i,j+2} & P_{i+1,j+2} & P_{i+2,j+2} \end{bmatrix}. \tag{3.19}$$

The matrix Q consists of the control points from the 4x4 neighbourhood centred on the patch. Continuity of the second derivative is achieved across the patch boundaries and the surface is constrained to lie within the convex hull of the 16 control points.In 3D, where $\mathbf{x} = (x_1, x_2, x_3)$, the control points are located on a regular grid in $\mathbb{R}^3$ and are denoted by $P_{i,j,k}$ where i, j, k are integer coordinates referring to grid locations. The 3D B-spline is a weighted average of 64 control points in a 4x4x4 neighbourhood describing a spatial distribution $f(\mathbf{x})$ within the volume bound by the eight central nodes (Figure 3.13). If these are scaled and shifted such that they lie within the unit cube, then the spline can be formulated as,

$$f(\mathbf{x}) = RM \begin{bmatrix} SM & Q_{k-1} & M^TT \\ SM & Q_k & M^TT \\ SM & Q_{k+1} & M^TT \\ SM & Q_{k+2} & M^TT \end{bmatrix}, \qquad 0 \leq x_1, x_2, x_3 \leq 1 \tag{3.20}$$

where

$$R = \begin{bmatrix} x_3^3 & x_3^2 & x_3 & 1 \end{bmatrix}. \tag{3.21}$$

Here the 4x4x4 support for the spline has been viewed as four planes of 16 points stacked one upon the other. The 3D spline is thought of as a 1D interpolation between 2D surfaces fitted to each plane.

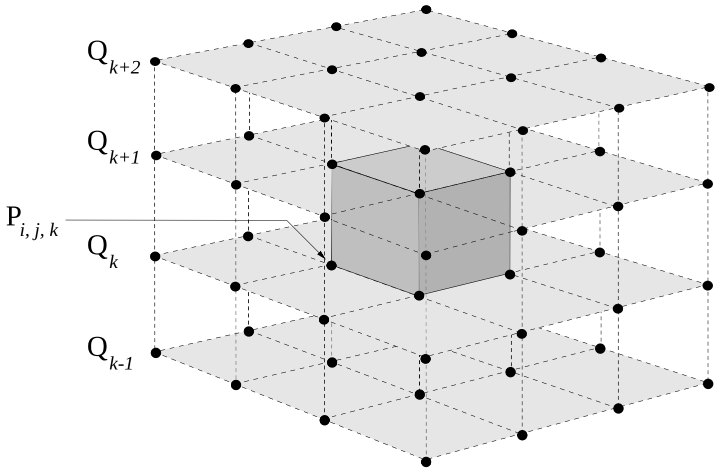

Figure 3.13 The sixty-four control points which define the cubic B-spline in three dimensions.

In the ray-casting context, voxels values are treated as control points and used to fit piecewise cubic splines which describe the scalar parameter being visualised. The aggregate of B-spline segments determines the scalar distribution over the entire data volume. The resulting spatial distribution is smooth and continuous. The convex hull property of the B-spline means that the the fitted function is constrained to have values within the range of the voxel values forming the support. Consequently, the B-spline is a *quasi-interpolant* of the voxel data since it approximates the voxel values at the interpolation nodes. It is not unreasonable to approximate the sampled data at voxel centres if the sampled data do not represent the actual value of the parameter at the grid coordinates but are in fact averages of the true parameter over the volume of the voxel. This is often the case with tomographic data such as X-ray CT. In such a case, the data can be viewed as having a noise component arising from the partial voluming effect.

In practice, the B-spline interpolant closely approximates the data except where sudden changes occur at a scale which is small relative to the size of the B-spline support. The piecewise nature of the B-spline interpolant avoids the propagation of spurious ripples due to sudden changes or discontinuities in the data. Furthermore, the gradient at a point $\mathbf{x}$ is precisely known since it can be analytically determined by differentiating (3.20).

Figure 3.14 illustrates part of a typical cranial surface extracted from CT data by ray-casting using tri-linear interpolation and B-spline interpolation. The CT data consists of rectangular voxels of dimension 1x1x2.3mm. Figure 3.14(a) was obtained

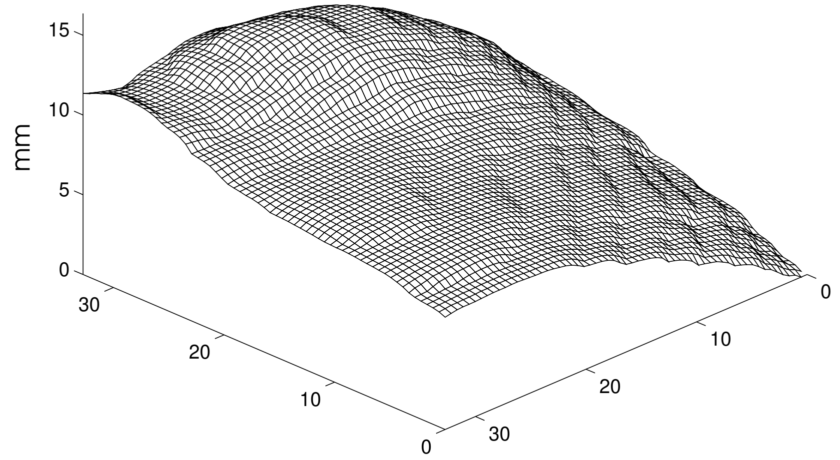

(a)

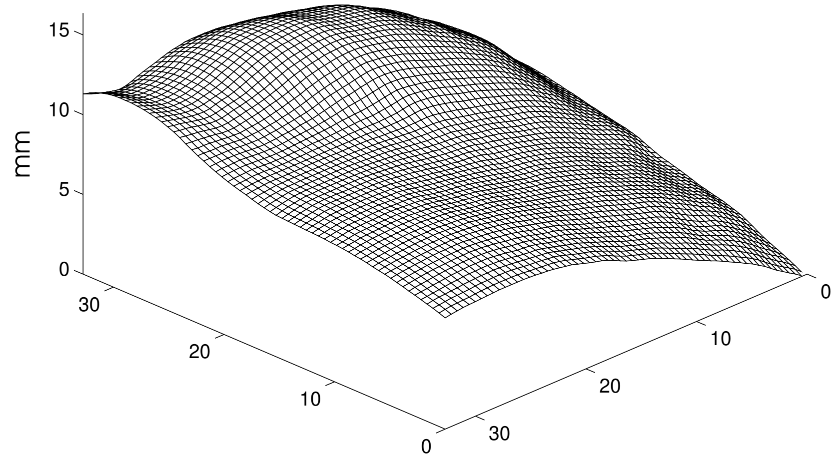

(b)

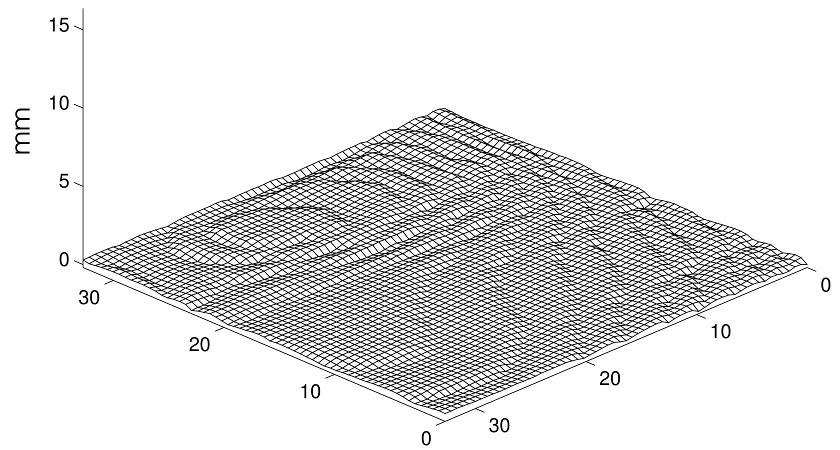

(c)

Figure 3.14 Comparison of ray-casting interpolation artifacts. (a) Tri-linear voxel interpolation. (b) B-spline voxel interpolation. (c) Difference between (a) and (b).

using conventional tri-linear interpolation while Figure 3.14(b) is the corresponding surface obtained using B-spline interpolation. The scales are in millimetres. Apparent in Figure 3.14(a) is a step artifact occurring at voxel boundaries. This is largely absent from the appreciably smoother surface depicted in Figure 3.14(b). In Figure 3.14(c) the difference between the two surfaces is displayed. Surface differences range between -0.19mm and 0.45mm, the mean being 0.21mm with a standard deviation of 0.016mm. Significantly, the difference between the two surfaces is very small. Furthermore, Figure 3.14(b) differs from Figure 3.14(a) primarily in the absence of ripples.

In Figure 3.14 it is not possible to directly determine which representation is the more accurate since the actual skull surface is not known, nor can it be easily measured. However, given that the difference between the two surfaces is very small (a fraction of a voxel), and that this is primarily due the ripples in the tri-linear surface, then it can be argued that the B-spline is the more accurate representation since it is known that skulls do not contain such ripples.

Figure 3.15 compares skull surfaces rendered using tri-linear and B-spline interpolation. The same threshold choice has been applied in both cases. Figure 3.15(a) is the B-spline surface rendered using the B-spline gradient. Figure 3.15(b) is the corresponding tri-linear surface rendered using a grey-level volume gradient (this figure is reproduced from Figure 3.6(b)). The B-spline surface gradient results in shading that is very similar to the grey-level volume gradient. In Figure 3.15(c) the B-spline depth-map is shaded using Z-buffer gradients determined by the algorithm described in Section 3.3.2. Figure 3.15(d) is the corresponding Z-buffer shaded image for the tri-linear ray-cast surface, reproduced from Figure 3.6(d). The smoother B-spline surface is particularly apparent in the Z-buffer surface renderings. The observation that the B-spline gradient produces surface shading very similar to that of the grey-level volume gradient is typical of the B-spline interpolant.

Figure 3.6(c) also shows that a high quality image, similar to that achieved using object-space shading, can be obtained using Z-buffer shading. This suggests that a poor Z-buffer shading may actually reflect the quality of the interpolation strategy used in the ray-caster, while the popular grey-level shading method masks the inadequacy of the actual surface extracted.

B-spline interpolation results in smoother surfaces with reduced discretization artifacts like bumps and ripples which otherwise become prominent when the spacing between slices is large relative to the slice resolution. B-spline interpolation is particularly suited to situations where the underlying surface is naturally smooth and continuous, but it requires significantly more computation than tri-linear interpolation (138 floating point multiplications are required to evaluate the B-spline at a point and a further 130 operations are required to evaluate the gradient). The B-spline is not so well suited to visualising structures which are small or thin relative to the size of

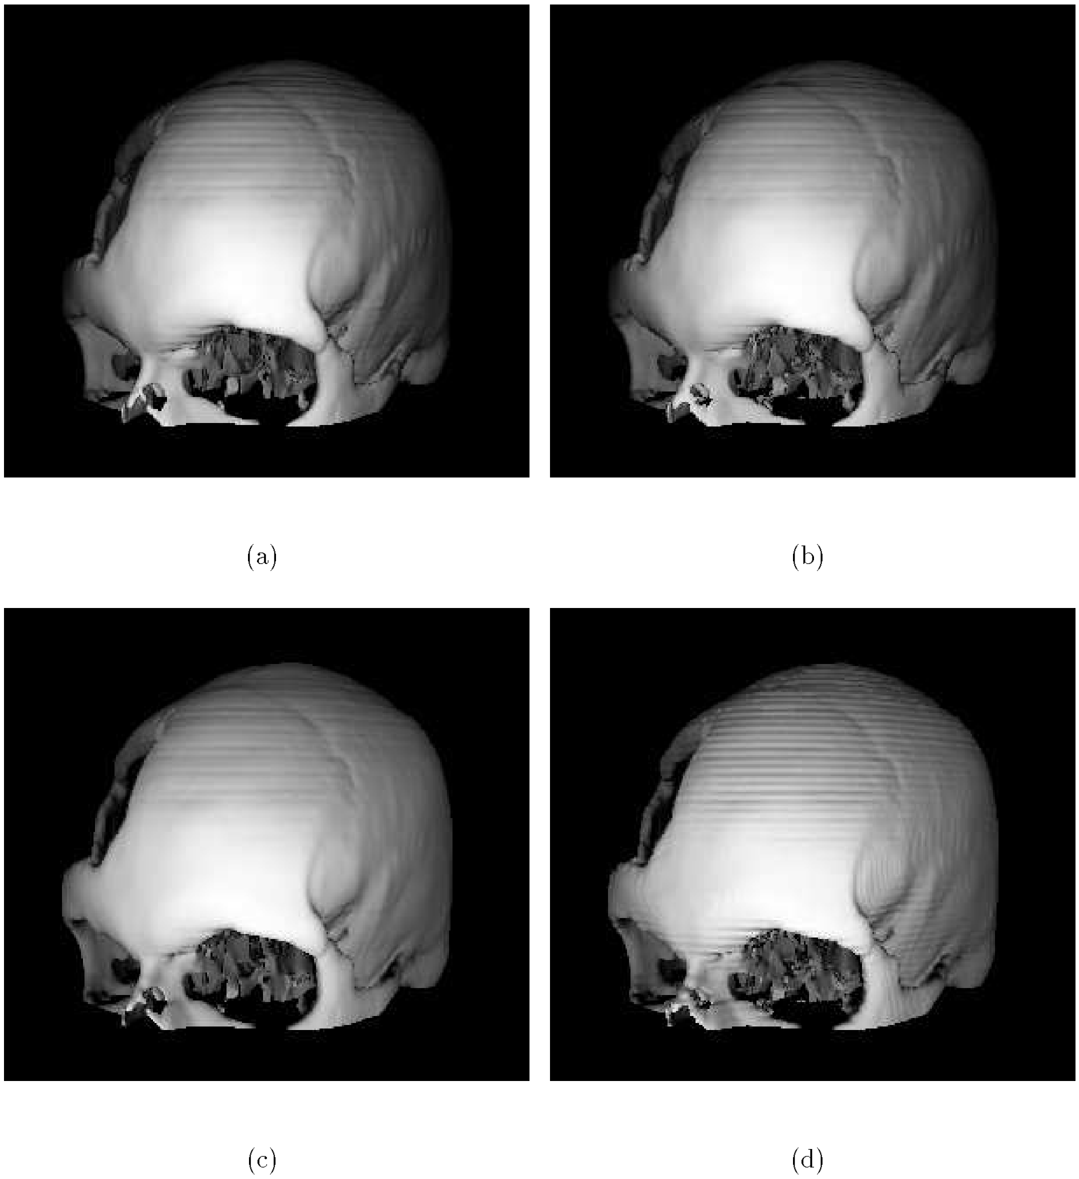

Figure 3.15 Comparison between tri-linear and B-spline ray-casting algorithms. (a) B-spline depth-map shaded using the B-spline gradient, (b) Tri-linear depth-map shaded using a grey-level volume gradient, (c) B-spline depth-map shaded using the Z-buffer gradient, (d) Tri-linear depth-map shaded using the Z-buffer gradient.

the B-spline support. In these instances the tri-linear method may produce better results. The only significant differences between the B-spline and tri-linear renderings in Figure 3.15 occurs in the nasal region where the bone is thin and the apparent 'hole' varies in shape. Such thin structures cause problems irrespective of the interpolation method since the sampled data no longer exhibit values typical of the structure due to the partial voluming effect.

The B-spline interpolation technique is applicable when the data lie on a regular cubic grid or can be mapped on to a regular cubic grid by an appropriate linear transformation. If such a transformation cannot be found, then a parametric representation can be formulated where f, x_1, x_2, x_3 are each written as a function of three variables, say r, s, t, which vary between 0 and 1, i.e. $f(r,s,t), x_1(r,s,t), x_2(r,s,t), x_3(r,s,t)$. The difficulty with a parametric representation is that evaluating f at a point (x_1, x_2, x_3) requires finding the corresponding point in the parametric space. This involves solving third order polynomials in three variables. Consequently, multiple solutions are possible. Although $f(r,s,t)$ is a single valued function, $f(x_1, x_2, x_3)$ may be multi-valued and this has no physical meaning in the medical imaging context.

The B-spline interpolant is used in Chapter 6 to render cranial surfaces from CT data and fit prostheses to the extracted depth-maps. This application is ideal since surfaces are only fitted to smooth continuous parts of the skull. Surface data near a discontinuity is usually associated with a defect to be repaired and is therefore discarded.

Webber [1990] has also tried to fit mathematical functions to volume data to improve the ray-casting algorithm. His method involves fitting a local bi-quadric surface as a function of two variables, $w = f(u,v)$. A 3x3x3 voxel neighbourhood determines coefficients of the polynomial f. In order to do this, a coordinate system $u - v - w$ is chosen such that the rendered surface is locally a function of u and v. A surface which is vertical relative to the $u - v$ plane cannot be represented. Choosing an appropriate $u - v$ plane for each voxel is difficult and no such plane may exist in the general case. To simplify the problem Webber restricts the $u - v - w$ coordinate system to be some translation and 90° rotation of the data $x - y - z$ coordinate system. The transformation is then chosen using a series of heuristics. The method results in a series of local quadric surface patches which form a discontinuous surface. Consequently, there is no continuity of the first and higher derivatives. Webber demonstrates the approach by locally approximating the surface of a sphere and justifies the use of a discontinuous surface on the basis of the degree of precision commonly present in voxel data and the density of rays cast. Coordinate changes generate the greatest patch discontinuities and are to be avoided. However, the complexity of anatomical structures is such that many coordinate changes are likely. This method suffers from many difficulties because it is essentially a 2D approach to a 3D problem.

3.4 CONCLUSION

3D imaging techniques generate large amounts of data and a key factor in their clinical usefulness is how this data can be presented so that a clinician can extract the maximum amount of useful information. This chapter has presented a brief overview of medical graphics and described the rendering of opaque surfaces using the ray-casting algorithm. Clinically, opaque surfaces are preferable to semi-transparent surfaces which are not suitable when quantitative measurements are required.

The B-spline interpolant has been introduced for ray-casting surfaces known *a priori* to be locally smooth and continuous. Besides being less susceptible to partial voluming effects, and therefore producing smoother surfaces, gradients can be determined analytically. Gradient estimation is important because realistic images require determining surface normals. An improved method for estimating surface normals from depth-map data has also been described.

This chapter has focused on rendering bone surfaces from CT data where the volume data is in the form of a regular array of voxels. The next chapter considers the more difficult problem of visualising ultrasound data from non-parallel slices.

Chapter 4

VISUALISATION OF ULTRASOUND DATA

This chapter discusses 3D graphics from pulse-echo ultrasound data. A number of difficulties peculiar to ultrasound make the visualisation of anatomic structures a challenging problem compared to that from 3D X-ray CT and MRI data. The complex nature of ultrasound-tissue interactions and the subsequent difficulty distinguishing tissue boundaries have resulted in much variation in the rendering techniques used. These are discussed in Section 4.1. Due to the presence of significant noise in the data, many authors have favoured semi-transparent renderings since these avoid explicitly defining tissue boundaries. However, these methods provide information of limited clinical value and, consequently, this chapter focuses on the rendering of opaque surfaces.

A prerequisite for the application of conventional ray-casting is that the data form a regular grid. In Section 4.2 an algorithm is presented for reformatting non-parallel slice data on to a regular voxel array. A direct method for producing graphics from non-parallel slices, as they are acquired, is presented in Section 4.3. The intention is to provide instant feedback to the sonographer at the time of acquisition.

Filtering is an essential, albeit controversial, aspect of rendering surfaces from 3D ultrasound data. Speckle reduction is a major field of research in itself, and not the primary focus of this thesis; Section 4.4 briefly discusses this aspect of ultrasound imaging. Section 4.5 explains the interactive approach taken to filtering in the current system and describes the filters used to process 2D and 3D B-scan data.

Section 4.6 concludes the chapter with the results from 3D *in vivo* studies.

4.1 REVIEW

Severe noise sources make volume rendering of ultrasonic data a challenge compared to rendering MRI and X-ray CT data. High quality imaging of X-ray CT and MRI data relies on an initial classification of the data. However, there is a lack of general and reliable methods for tissue characterisation in ultrasonic imaging. Ideally, high intensities recorded from pulse-echo ultrasound correspond to regions with high variation in acoustic impedance. However, the nature of specular reflection means that similar structures may result in strong or weak received echoes, depending on the direction of the incident sound. Furthermore, high density material may totally absorb transmitted ultrasound waves, leaving structures further from the transducer in shadow. Reverberation between a strong reflector and the transducer face may also lead to the imaging of false *ghost* interfaces.

Specular reflections occur between layers of tissues with different acoustic impedances when the boundary is smooth. High frequency spatial variations in acoustic impedance result in additional diffuse scattering. This is characterised by the backscatter coefficient per unit volume [Wear *et al.* 1994]. Backscatter varies between different kinds of tissue, and therefore causes grey level variations in the displayed image. Constructive and destructive interference between neighbouring scatterers, due to coherent detection, gives rise to multiplicative noise and speckle which modulates the grey-level image.

For these reasons the grey-level intensity is not always representative of the presence of an interface or a particular tissue type. Consequently, techniques described in Chapter 3, which render iso-value surfaces and have been successfully used with X-ray CT data, are likely to be less reliable in ultrasound imaging. Surface rendering with geometric facets will depend on the robustness of the surface detection method, while semi-transparent volume rendering will require more sophisticated tissue classification schemes. Despite these concerns, reasonable results have been achieved for specific applications, particularly vascular and foetal imaging.

4.1.1 Surface rendering

Surface rendering, with triangular facets and polygon meshes, has been successfully used in 3D intravascular imaging systems [Pretorius *et al.* 1992, Kitney *et al.* 1991, Burrell *et al.* 1989, Selzer *et al.* 1989, Pretorius *et al.* 1992]. Due to the nature of intravascular imaging, the artery walls are usually depicted with high contrast and artifacts, such as shadowing, reverberation and speckle, rarely affect the region of interest. Consequently, the extraction of contours is a simpler task than in other forms of ultrasound imaging. Restricted radial edge detectors and *a priori* knowledge of the circular shape of arteries and veins is taken advantage of in many vascular imaging systems. For example, Nixon *et al.* [1992] uses a Hough transform to find circular features which guide the extraction of artery walls from B-scans taken with an external transducer.

Verlande *et al.* [1991] and Martin *et al.* [1989] use geometric rendering to reveal the shape of the left ventricle of the heart and the mitral valve. Contours are extracted semi-automatically from constituent ultrasound slices and are connected with triangular facets. Thune and Olstad [1991] also describe a surface visualisation method for the heart wall. A supervised edge detection algorithm, which allows the user to specify an initial estimate, extracts edges from the inner ventricle wall from 2D slices in the volume data. A 3-D spline is used to connect the segments together to form closed contours. Eldemellawy and Martin [1989] take an alternative approach and render a geometric model fitted to 3D data. In their case an ellipsoid is fitted to the left ventricle.

Nelson and Elvins [1993] have studied a variety of techniques including surface and volume rendering methods. They use a marching cubes algorithm to render iso-valued surfaces from ultrasound data, but report that the results are sensitive to the selection of an appropriate threshold and that false positive and negative surface pieces can arise. Incorrect handling of small features and branching structures often occurs. However, good results have been obtained for imaging the foetus.

4.1.2 Semi-transparent rendering

Semi-transparent rendering, using models described by Levoy [Levoy 1988] and Dreblin [Dreblin *et al.* 1988], have been used by Mills and Fuchs [1990], Ohbuchi and Fuchs [1990] and Nelson and Elvins [1993] to render ultrasound data. Nelson and Elvins use a splatting algorithm (cf. 3.1.1.2) rather than ray-casting. Manual inspection of a representative data slice is usually used to assign attenuation coefficients (opacities) to grey-level intensity values.

Thune and Olstad [1991] describe a semi-transparent shading method which involves classifying the volume data into three substance categories according to intensity. At a transition between two substances a surface normal is calculated from the central difference. Colour and opacity are assigned according to the number of voxels of the substance which the ray traverses. The shading of an image pixel is determined by summing all contributions along the ray. Opacity tables, colour tables and classification tables for different substances are all user-specified.

Steen and Olstad [1994] have also used Sabella's [Sabella 1988] volume rendering method to make semi-transparent renderings of ultrasound data. This method models the scalar field as a varying density emitter field and was originally developed for rendering naturally "cloud-like" objects; for example, it has been used to render the appearance of the planet Saturn. Sabella's model does not require the simplifying assumptions made by the more commonly used transparency methods of Levoy [Levoy 1988] and Dreblin [Dreblin *et al.* 1988], but the physical phenomenon it approximates is less intuitive. Steen reports that the method produces results very similar to another, simpler method, known as the *maximum intensity projection method.*

The maximum intensity projection method displays the highest intensity value encountered along a ray and has been used by a number of researchers specifically for rendering ultrasound data [Steen and Olstad 1994, Thune and Olstad 1991, Herrington *et al.* 1991, Nelson and Pretorius 1995, Nelson and Elvins 1993]. The method has been shown to give clinically interesting renderings, but usually for a limited number of viewpoints. These depend on the anatomy and the presence of intervening structures which otherwise confuse the image. A single voxel is projected per ray, therefore noisy voxels with high intensities may be projected instead of voxels corresponding to regions with maximum variation in acoustic impedance.

Another method described in the literature is the *sum of values projection* method. This method accumulates a weighted average of the intensity values encountered along each viewing ray. Although less sensitive to noise than the maximum intensity projection, this method has been found to give low contrast renderings and its ability to display 3D structure is poor in both static and temporal sequences. Nelson and Elvins [1993] and Steen and Olstad [1994] have found that clinically useful renderings are only possible if averaging is restricted to a limited depth range.

A common finding with transparent methods, and methods such as the maximum intensity projection, is that visual ambiguities often arise since no geometric model is associated with the data. These ambiguities can be partially removed by computing an animated sequence or by interactive manipulation of the viewing parameters [Steen and Olstad 1994] or through stereo viewing [Nelson and Elvins 1993]. Depth shading is commonly used to produce images from these methods. This is due to the difficulty in determining gradients and the inadequacy of weighted summations of gradients, calculated along viewing rays, to depict useful anatomical information. Z-buffer shading is only useful for rendering opaque, smooth surfaces and cannot be used with methods like the maximum intensity projection which do not define a continuous surface.

Despite these shortcomings, Herrington *et al.* [1991] favour the maximum intensity projection combined with multi-planar reformatting. Nelson and Elvins [1993] have also found that maximum intensity projection can provide clear views, although it is best applied to restricted regions. Steen and Olstad [1994] differ slightly by preferring to project the magnitude of the grey-level gradient. First they smooth the data with an edge preserving filter before projecting the magnitude of the grey-level gradients estimated using the central difference method (Equation (3.5)). Because the ultrasound intensities represent a change in acoustic impedance, they are effectively visualising the second derivative of the acoustic impedance distribution. Appropriate filtering is critical in the determination of noise-free gradients which correspond to changes in tissue backscatter.

4.1.3 Opaque surface rendering

Pini *et al.* [1991], Balen *et al.* [1993], Mintz *et al.* [1993] and Cheng *et al.* [1993] have used the ray-casting technique described in Section 3.2 to render opaque iso-value surfaces of the heart and foetus. Cheng *et al.* [1993] perform a classification of ultrasound voxels into three classes, liquid, boundary, and soft tissue in their studies of the foetus. They use fuzzy logic and relaxation techniques to achieve this. They then render a particular substance surface using a conventional ray-casting algorithm.

McCann *et al.* [1988] have successfully used both surface and volume rendering to depict the heart by defining regions of interest (ROI). These are derived either by thresholding grey-levels or, more often, by manual segmentation. McCann comments that volume rendering preserves valuable backscatter information in a form that the cardiologist is familiar with. However, they find that semi-transparent volume rendering shows little anatomical detail because of "the confusing cloud-like appearance (of the heart)". The clinically crucial details, such as the valvula dimensions and morphology are not apparent. Although they acknowledge that the backscattered data is rich in anatomic information, the identification of morphological detail and appreciation of surfaces is difficult. Surface displays, where possible, were found to show clinically useful information including the shape and size of cavities. Nelson and Elvins [1993] similarly comment that surface rendered images are more easily interpreted by the viewer and overall spatial relationships more easily comprehended.

In the following sections two ray-casting approaches are described. Both use a simple threshold to render iso-value surfaces. The first involves forming a regular voxel array so that a conventional ray-casting algorithm, as described in Section 3.2, can be applied. Although the data volume is reconstructed incrementally, the computational overhead of the reconstruction, combined with that of the ray-caster, prevents a real-time implementation. Nevertheless, the reconstruction of a regular volume is useful for multi-planar reformatting and making quantitative 3D measurements. The second method described is an incremental method for directly rendering iso-value surfaces from non-parallel slices, as they are acquired. This method is significantly faster than the first but is more sensitive to noise and the rendered surface varies depending on the viewing direction.

4.2 REFORMATTING NON-PARALLEL SLICE DATA

A prerequisite for the use of conventional ray-casting is that the ultrasound data lie on a regular grid. If the pixel data, which constitute a slice, are simply mapped into an appropriate voxel space, where a contribution is made to each voxel through which the slice passes, then, in general, there will be some voxels to which no data is mapped. These 'holes' in the array will result in distracting artifacts when an iso-valued surface is rendered. Authors who take this approach apply *ad hoc* schemes to identify holes in the array and then average non-zero neighbours to fill them in [McCann *et al.* 1988, Steen and Olstad 1994].

The ray-casting methods developed for irregular, distorted meshes by Giertsen [1992] and Shirley and Neeman [1989] require the identification of a mesh which passes through the data. Due to the arbitrary nature of data acquired by a six-degrees-of-freedom hand-held probe, the formation of a mesh is a non-trivial problem, particularly if new data is continually being acquired.

The developers of 3D probes use reconstruction methods which take advantage of the fixed geometry of the acquired 2D scans. For example, Thune and Olstad [1991] acquire 3D images by tilting a sector scanner through a 45° arc and acquiring 8 equally spaced slices. They use linear interpolation along the 45° arc spans passing through the slice pixels to construct a regular voxel array. Nixon and Hames [1993] use linear interpolation along one axis of the reconstructed voxel array, restricting their slices to be semi-parallel. Such methods are not possible when the acquired slices have an arbitrary geometry.

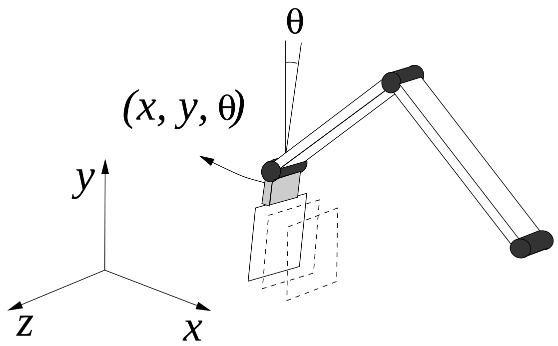

Figure 4.1 Coordinate system for the scanning system of Ohbuchi *et al.* The mechanical arm restricts the movements of the probe to translation in the $x - y$ plane and rotation θ about the z axis. This simplifies the volume reconstruction algorithm.

Ohbuchi and Fuchs [1990] propose an incremental volume renderer based on Levoy's rendering algorithm with the intention of rendering semi-transparent surfaces in real-time, as 2D slices are acquired by a scanner. They assume that the scalar field being sampled is a continuous function of 3-space and that temporally consecutive slices are geometrically next to each other. A combination of bi-linear interpolation and splatting is used to re-sample the slices on a regular grid and fill in gaps between slices. The slices are acquired with a three-degrees-of-freedom mechanical arm which restricts the probe to movements in the x and y directions with rotation (θ) about the z axis (Figure 4.1). Values are first linearly interpolated at regular y intervals within each slice and then inter-slice values are linearly interpolated from these at regular intervals along the x axis to fill in gaps between the slices. Each value is then distributed or splatted among the neighbouring voxels with weightings determined by the coefficients of a low-pass filter kernel. A triangular shaped filter kernel is used and the weights are accumulated in a separate buffer for later normalisation. Like Nixon and Hames's method, this method is biased towards interpolating in a particular direction. Good results require that the slices are semi-parallel and ideally perpendicular to the x axis of the voxel array over the entire sweep of the probe. This is not an unreasonable assumption in their system due to the restrictions imposed by the mechanical arm. Significantly, both the volume reconstruction and the rendering processes are incremental in Ohbuchi's method; only each new sub-volume added to the data is sampled by the ray-caster. Ohbuchi does not describe how his system would cope if consecutive scans intersected each other, or if a previously scanned sub-volume was re-sampled. Multiple evaluation of a voxel centre would require some mechanism to update and possibly undo the shading calculated from previous slices. Presumably, this consideration places a further restriction on the acquisition of slices in the system.

4.2.1 Swept-volume reconstruction

Volume reconstructions from arbitrary ultrasound slices pose several problems. If the same anatomy is insonified from many different directions, then differing amounts of attenuation, specular reflection, shadow and reverberation experienced along each path may result in the imaged distribution being multi-valued. Some method of compounding a potentially large range of ultrasound numbers recorded at the same point in space is therefore required.

In the following reconstruction method, restrictions are placed on the acquisition of slices in order to simplify the reconstruction problem and minimise the possibility of multi-valued data. Slices are assumed to be acquired during a continuous sweep of the probe such that consecutive slices are semi-parallel. This assumption is less restrictive than the previous methods, since over a number of scans the path of a sweep can be quite curved. In practice, the sweep scanning motion corresponds closely to the natural manner of scanning and a simple reconstruction method results where a volume is incrementally constructed as each new slice is acquired. This method is called the *swept volume* reconstruction.

A volume is incrementally constructed on a regular array by interpolating intensity at voxel centres which lie between the two most recently acquired scans. A perpendicular is dropped from the voxel centre on to each scan plane (Figure 4.2). Depending on the relative coarseness of the voxel grid compared to the resolution of the B-scan slices, either the intensity of the nearest pixel is chosen from each scan or the intensity is bi-linearly interpolated from the four neighbouring pixels. The value at the voxel is then computed as the sum of these two intensities weighted according to their perpendicular distance from the voxel centre. The slices are typically of a higher resolution

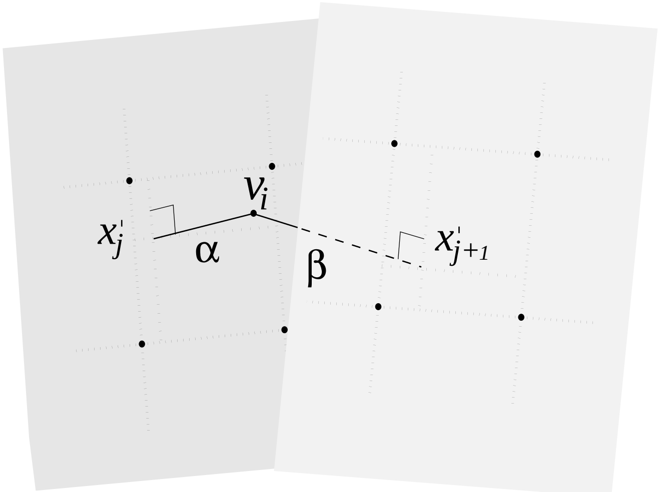

Figure 4.2 Interpolation strategy to evaluate the ultrasound intensity at a voxel centre v_i between two semi-parallel B-scan slices.

than the voxel grid, hence choosing the nearest pixel is usually adequate. The value of the i^{th} voxel, v_i, is given by,

$$v_i = \frac{\alpha}{\alpha+\beta} x'_j + \frac{\beta}{\alpha+\beta} x'_{j+1} \tag{4.1}$$

where v_i lies between the j and $j+1$ scan planes, and x'_j and x'_{j+1} are the intensities evaluated at the intersections of the scan plane normals which pass through v_i with their respective planes. α is the distance from the voxel to the j^{th} plane and β is the distance to the $j+1$ plane. Only voxels lying between the two most recent slices are updated in the volume.

The possibility that two consecutive scans intersect is allowed. Multiple contributions to a voxel within a sweep are either averaged or the maximum value is chosen. Similarly, multiple sweeps of the probe can be compounded to image a larger volume or to improve the signal to noise ratio. Average and maximum intensity compounding are described further in the context of speckle reduction in Section 4.4.2.1.

Figure 4.3 is a depth-shaded image of a foetus rendered from ultrasound data supplied courtesy of John Gardener at University College London. The ultrasound data set consists of 68 carefully acquired semi-parallel B-scans. The slices were reconstructed in a regular volume as if they were parallel and equally spaced, then an iso-value surface was rendered using the ray-casting algorithm described in Section 3.2.

For the purpose of illustrating the reconstruction process and rendering techniques, the foetal data volume was re-sampled to simulate the acquisition of 50 non-parallel slices from two sweeps of a probe. Figure 4.4(a) is a slice through the original data

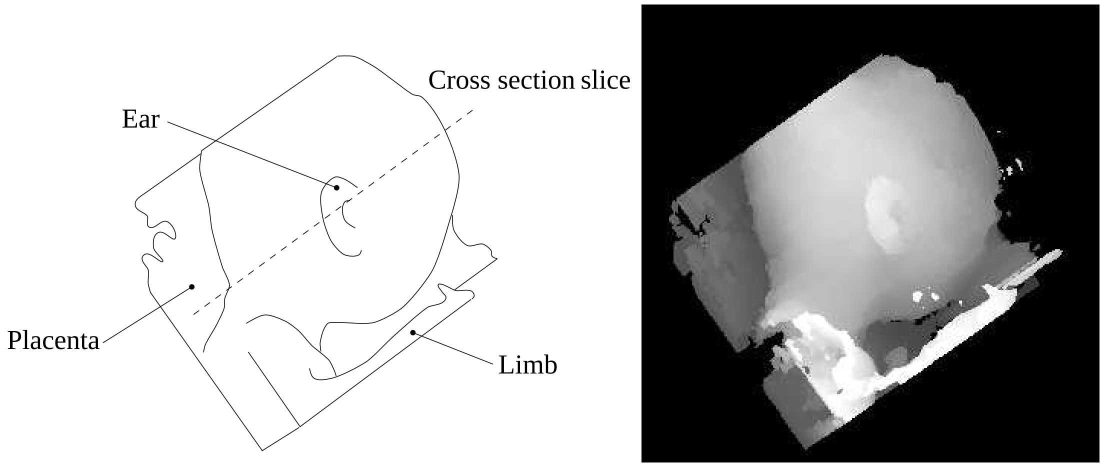

Figure 4.3 Foetal data set. The position of the slice studied in Figure 4.4 is illustrated.

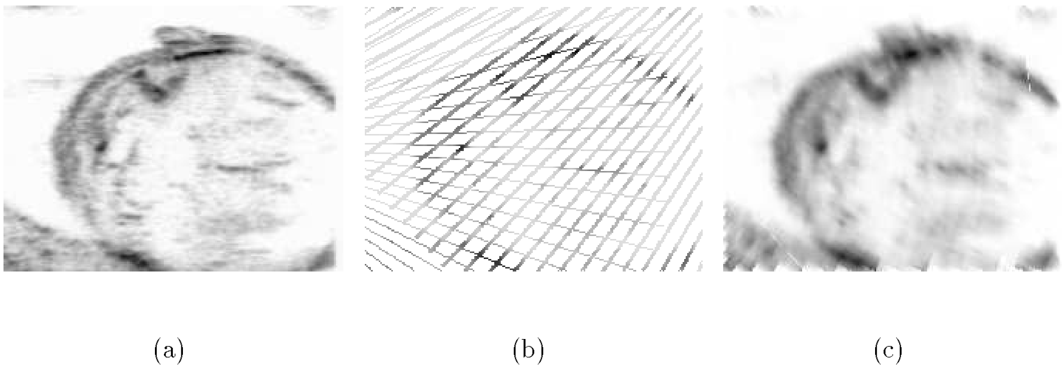

Figure 4.4 Volume reconstruction from non-parallel slices illustrated by re-sampling a regular voxel array. (a) Slice through the original, complete data volume, (b) the slice re-sampled by non-parallel scan planes, (c) the reconstructed slice.

volume. Figure 4.4(b) is the same slice re-sampled by the set of non-parallel planes. Figure 4.4(c) is the corresponding slice from the swept volume reconstruction. Figure 4.5(b) is the depth shaded view of the reconstructed volume with the same viewing parameters as those used to render the orignal complete data set in Figure 4.5(a). Figures 4.5(c) and 4.5(d) compare the original and reconstructed data sets using a volume gradient to shade the two images. Neither the original data slices nor the re-sampled slices have been filtered in these reconstructions.

The re-sampling of the original data apparent in Figure 4.4(b) is coarser than that which would normally be acquired and this accentuates the effects of the reconstruction method. Superficially, the depth-shaded images of the reconstructed and original volumes appear similar because depth shading does not emphasise surface roughness and is more tolerant to the presence of noise. However, small discernible structures which appear below the ear in Figure 4.5(a) are noticeably absent in Figure 4.5(b). Furthermore, some new holes have occurred in the skull in Figure 4.5(b). These effects are a combination of the lower sampling frequency of the simulated scans and the process of averaging multiple contributions to a voxel from the two sweeps.

Apparent in Figure 4.4(c) is a blurring effect due to the interpolation process in the swept volume reconstruction. Similarly, the volume-gradient shaded view in Figure 4.5(d) appears blurred compared to Figure 4.5(c). This is particularly noticeable in the region of the ear. Unlike depth shading, volume gradient shading emphasises surface detail and is particularly susceptible to noise. The poor quality of the gradient shaded images in Figures 4.5(c) and 4.5(d) is typical of unfiltered ultrasound data. Z-buffer gradient shading results in similar image quality due to its sensitivity to surface roughness.

The swept volume reconstruction has been introduced as an incremental method. Alternatively, the scans can be stored and mapped onto a regular voxel array off-line. The holes in the array can then be identified and filled in by determining which pairs of scan planes they lie between and averaging the contribution from each pair. Although not consistent in the treatment of each voxel, this approach will generally be more efficient to compute. Unlike other hole-filling strategies, which impose some arbitrary weighting to voxels within a fixed neighbourhood, this strategy retains knowledge of the geometry of the scan data and uses it to determine the neighbourhood of contributing data values.

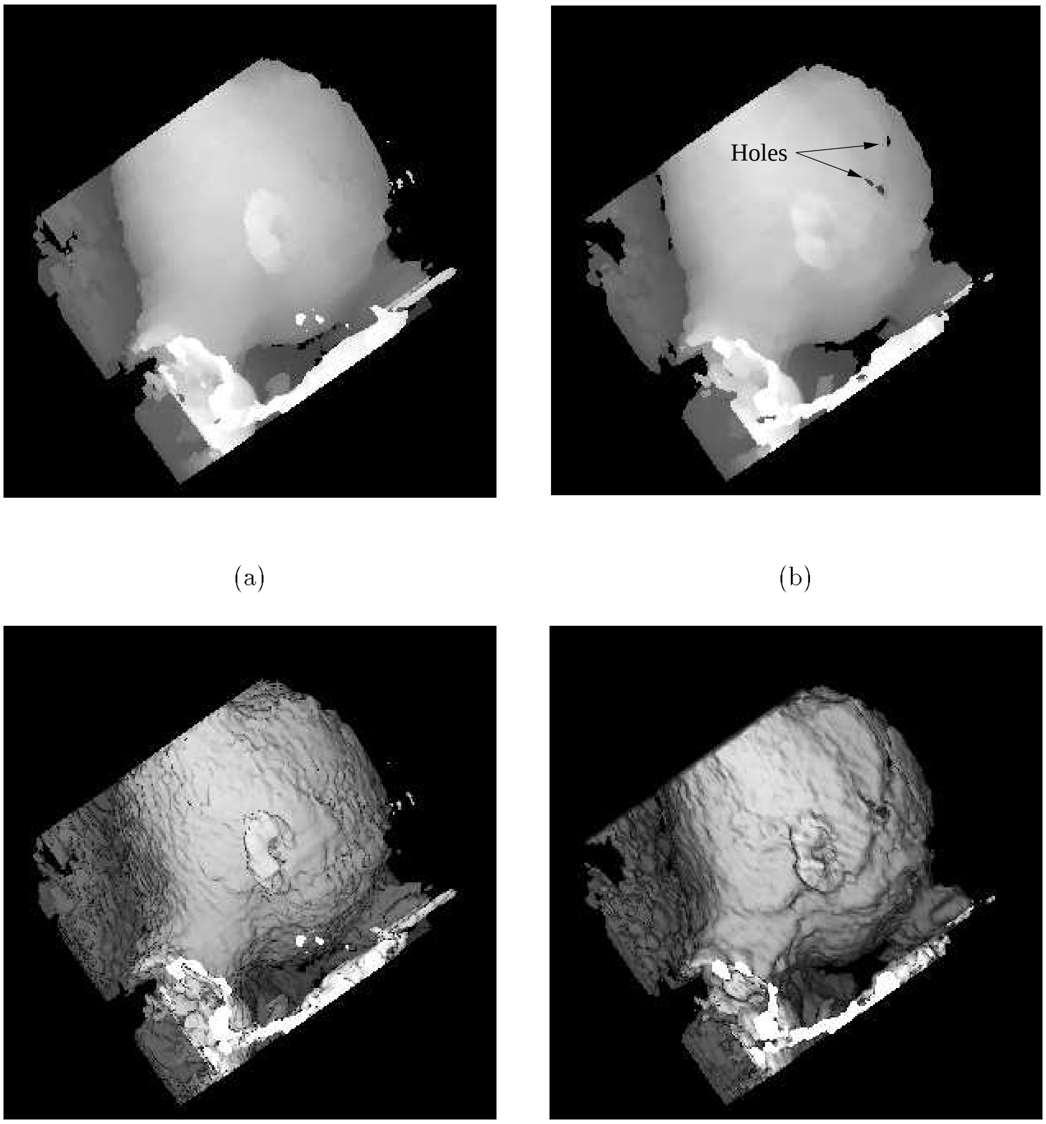

(a) (b)

(c) (d)

Figure 4.5 Comparison of renderings from a complete data set with those constructed from a subset of non-parallel slices. (a) Depth shaded view of original data, (b) depth shaded view of reconstructed data, (c) voxel gradient shading of original data, (d) voxel gradient shading of reconstructed data.

4.3 DIRECT RAY-CASTING

The swept volume reconstruction method is suitable for ray-casting and multi-planar reformatting because it constructs a regular array of voxels. Although it allows incremental reconstruction of the volume, the associated computation prevents a real-time implementation, particularly when combined with the cost of rendering the volume data with a standard ray-casting algorithm. Consequently, a new method for directly producing images from non-parallel slices is now described. This method is comparatively fast but does not construct a regular voxel array. The method is intended to provide some immediate feedback to the clinician during scanning and is not capable of providing images of the same quality as the swept volume reconstruction, nor can it provide a facility for multi-planar reformatting or 3D quantitative measurements.

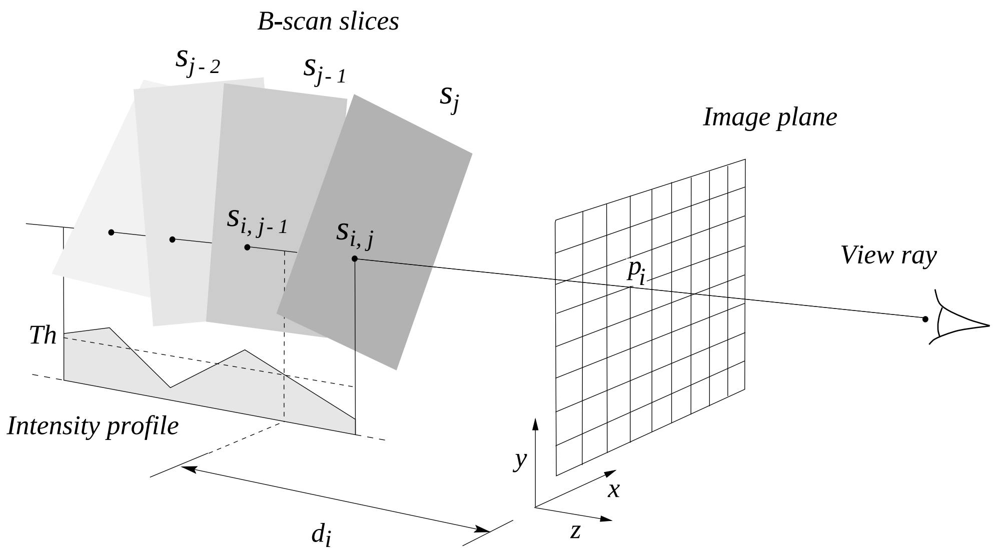

Figure 4.6 Direct rendering of non-parallel slices. Linear interpolation along viewing rays is used to estimate threshold transitions between slices.

Figure 4.6 illustrates the principle of the method. The image plane consists of a 2D depth-map. Each location in the depth-map (grid point) corresponds to a view ray. Two depths are stored in each location. One records the value and depth of the closest B-scan pixel which exceeds the threshold, the other records the value and depth of the B-scan plane immediately in front of it, if one exists. The depth of the actual threshold transition is estimated by interpolating between the two values. If the distance between the two values exceeds a user defined limit then the threshold transition depth is not interpolated but taken as the depth of the value above the threshold. As each scan is acquired, the intersection point with each view ray is calculated, and if the ray lies within the B-scan field of view then the value of the nearest pixel is chosen and the depth-map is suitably updated. Depth-shading or Z-buffer shading is used to render an image since volume gradients are unavailable.

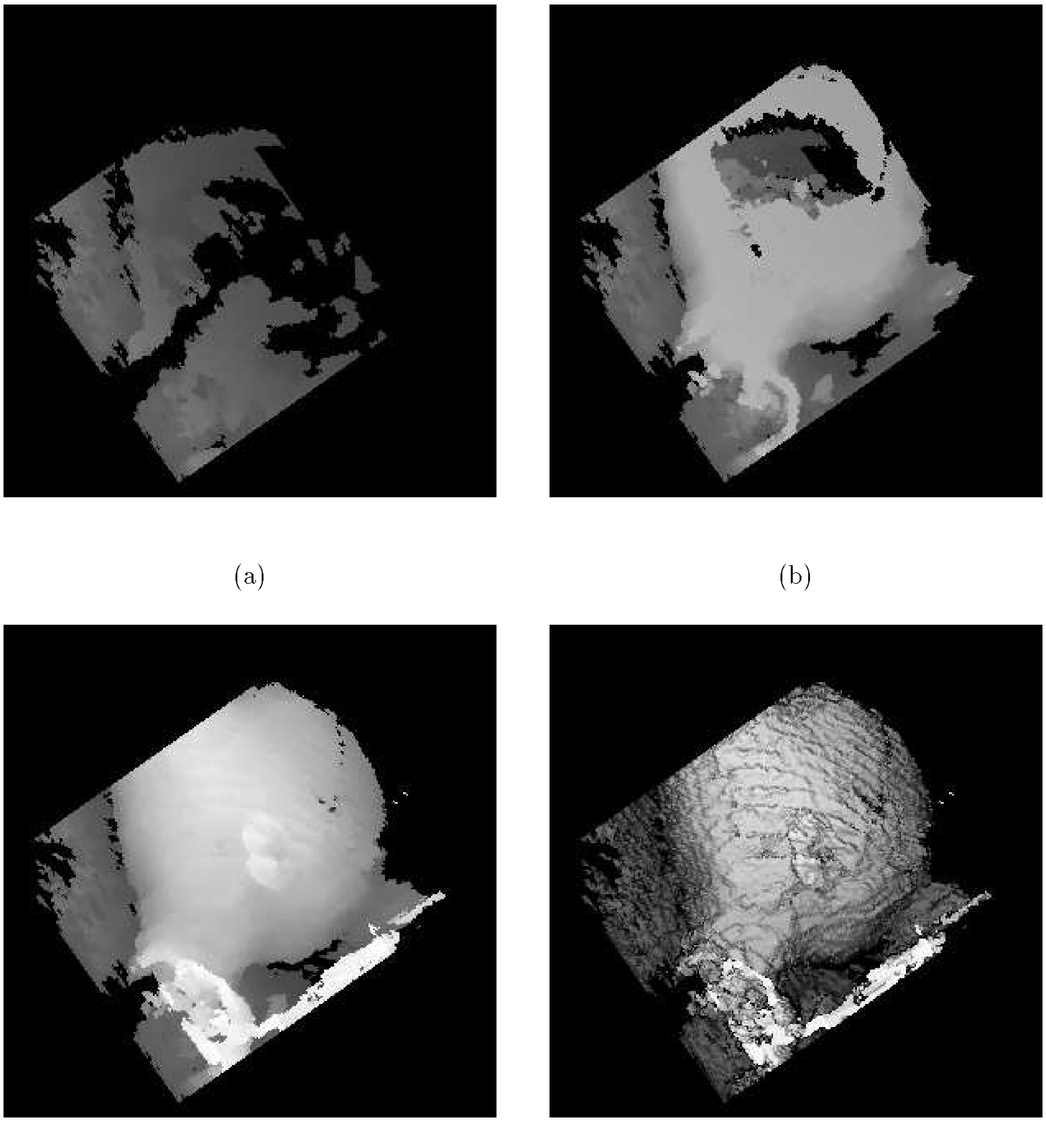

(a) (b)

(c) (d)

Figure 4.7 Direct rendering of non parallel slices without performing an intermediate voxel representation. The data is rendered incrementally as slices are acquired using depth shading and interpolation along the viewing rays. (a) Depth-shaded view after processing 10 slices, (b) depth-shaded view after processing 18 slices, (c) depth-shaded view after processing 50 slices, (d) corresponding Z-buffer gradient shading.

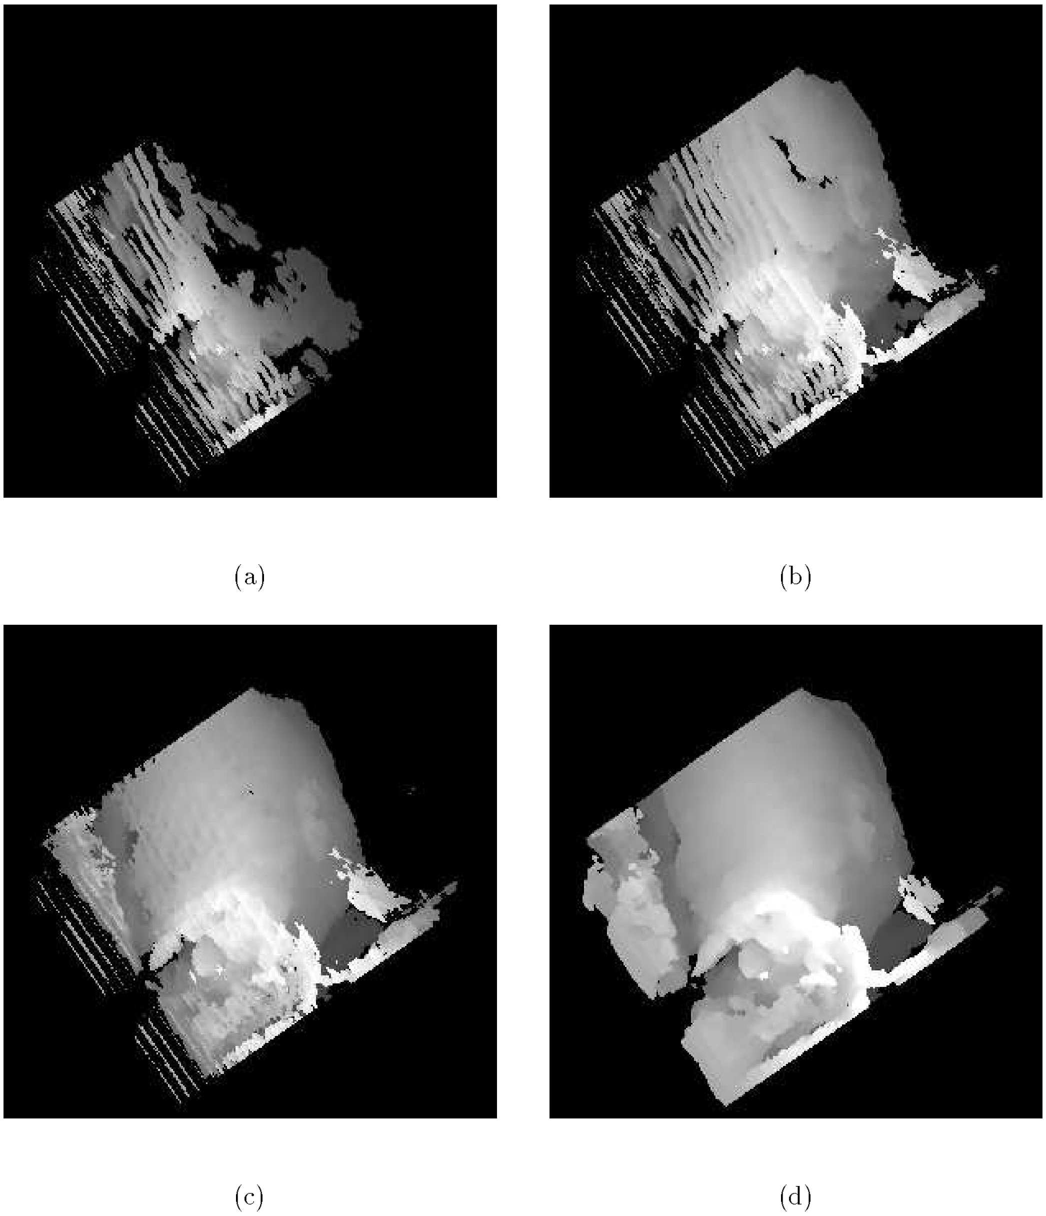

Figure 4.8 Incremental rendering of the slice data from Figure 4.7 when the majority of the scans are oriented obliquely with respect to the viewing direction. (a) Depth-shaded view after processing 14 data slices, (b) depth-shaded view after processing 33 data slices, (c) depth-shaded view after processing 50 data slices, (d) corresponding view of the swept volume reconstruction.

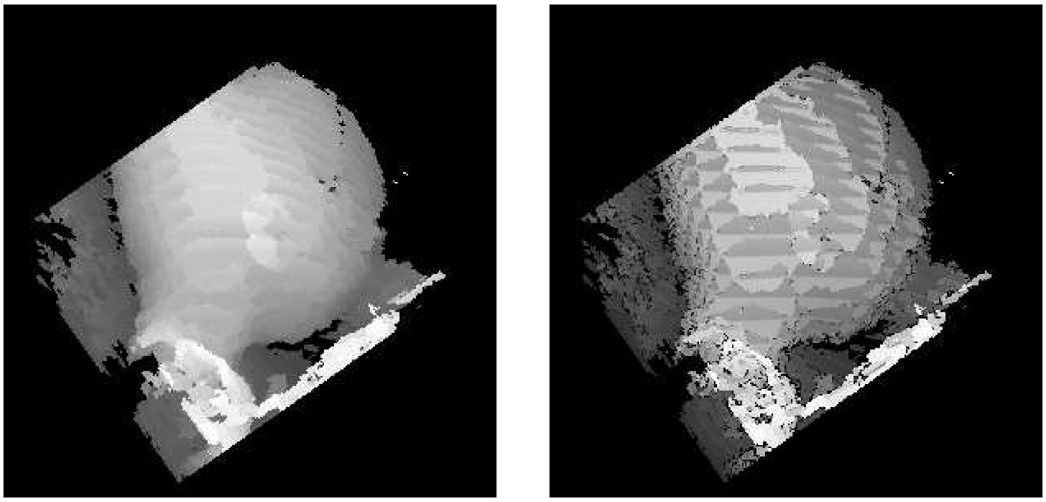

(a) (b)

Figure 4.9 Direct rendering of ultrasound slices without interpolating along the viewing rays. (a) Depth shaded image, (b) Z-buffer shaded image.

The rendered image evolves as scans are acquired. Figure 4.7(a-c) illustrates the incremental rendering of the foetus from the non-parallel slice data used to illustrate the swept volume reconstruction in Figure 4.5. Figure 4.7(d) corresponds to the Z-buffer shaded view for the complete set of slices. The Z-buffer shading may vary erratically as new slices are acquired and the shading emphasises the discontinuity which occurs between slices. Unlike the swept volume reconstruction, consecutive slices need not be geometrically next to each other. The order of the slices has no effect on the final image; the order affects only the evolving image. Because interpolation is performed along the viewing rays, quality may vary significantly with the viewing direction. A viewing direction which results in rays which intersect the majority of slices at an oblique angle will not produce particularly good images. This is illustrated in Figure 4.8(a-c) where the foetal data set has been rendered from an alternate viewing angle. Figure 4.8(d) is the corresponding view of the swept volume reconstruction. The swept volume reconstruction is unaffected by the viewing angle.

The potentially difficult problem of specifying appropriate viewing parameters prior to scanning is simplified by using the probe and Polhemus spatial locator to define the viewing plane by acquiring an initial slice perpendicular to the viewing direction. The probe is then swept from this initial point such that a 3D image is built-up from back to front, with respect to the viewing direction, as slices are acquired. This avoids old data obscuring new slices. This initial pre-sweep slice can also be used to assist the selection of an appropriate threshold.

Multiple threshold transitions can be implemented at little extra cost by providing multiple buffers which represent increasing threshold levels. Each new intersection between a scan plane and a view ray is evaluated and the appropriate threshold buffers updated. The depth-shaded view can then be toggled between thresholds without having to repeat the ray-casting process for all the previous scans. The selection of a new viewpoint, however, requires reconsideration of the entire data set.

The method can be simplified and made faster by eliminating the interpolation between slices. This results in a noticeable slice artifact since the surface depth values stored in the threshold buffer will lie on one or other of the scan planes. Figure 4.9 illustrates the final view of the data rendered in Figure 4.6 without interpolation. Figure 4.9(a) is the depth-shaded view and Figure 4.9(b) the corresponding Z-buffer shading. The latter illustrates the scan plane artifact more clearly.

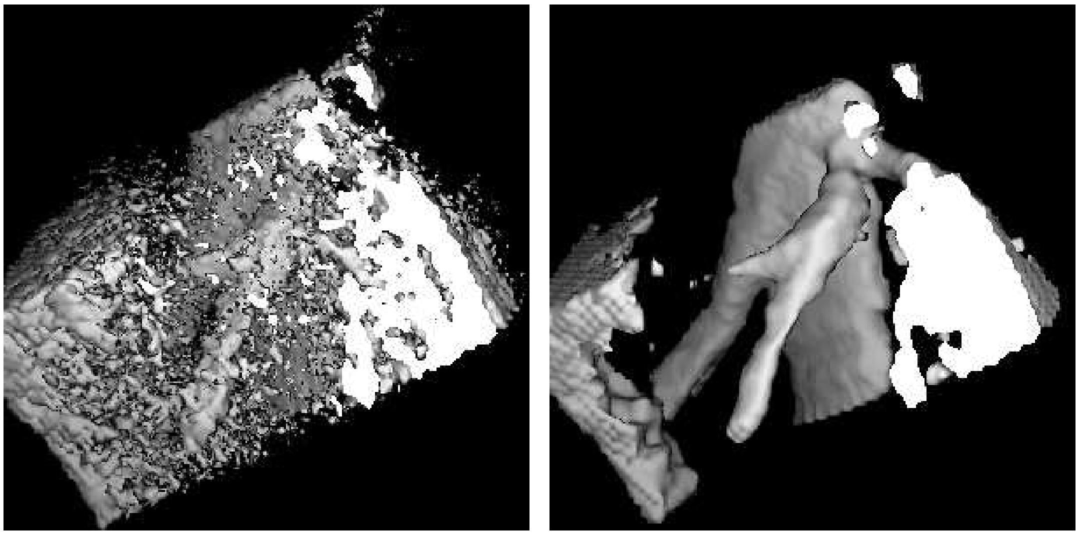

(a) (b)

Figure 4.10 Volume reconstruction of the inferior vena cava and hepatic veins, with and without filtering. Occluding structure in the foreground has been removed with the aid of a cut-away view. Both figures use volume gradient shading. (a) No filtering, (b) reconstruction from filtered data.

4.4 FILTERING

Surface renderings from 3D foetal ultrasound data, often taken during the first six months of pregnancy, are commonly found in the 3D ultrasound literature. This is because the nature of the subject, being surrounded in amiotic fluid, results in high contrast images in which the surface of the foetus is readily extracted by simple thresholding, albeit combined with the facility to cut away occluding structures. Filtering was not required to render the main structures present in the foetal data in the preceding figures, although filtering could have improved the shading of these images and the estimation of volume gradients. Structures embedded in soft tissue, however, are not so easily rendered by thresholding the raw data. Filtering to reduce speckle and enhance contrast becomes essential.

Figure 4.10 illustrates a swept volume reconstruction of the hepatic veins and inferior vena cava, with and without filtering. Volume gradient shading has been used in both images along with identical viewing parameters. Data values below the threshold have been treated as opaque while those above are transparent. This is done to depict the fluid filled veins since they are characterised by an absence of echoes. In Figure 4.10(a) speckle and signal drop-out have obscured the principal structures and resulted in distracting artifacts which also affect the determination of surface gradients. In Figure 4.10(b) 2D filtering has been applied to the constituent B-mode slices prior to reconstruction and 3D filtering to the reconstructed volume. In this figure, speckle reduction and filtering were necessary to visualise the main structures in the data. The purpose of the next section is to discuss speckle reduction and the basic filtering techniques used in this thesis.

4.4.1 Speckle

The characteristic granular, or mottled, appearance of ultrasound images, commonly referred to as 'speckle', is the result of a complex superposition of many factors including interference effects, attenuation and electrical processing. It occurs in virtually all medical ultrasound images and can severely degrade their clinical information content. Consequently, speckle reduction and filtering are important aspects of many ultrasound visualisation and segmentation methods.

Conventional pulse-echo ultrasound involves propagating a short pulse of ultrasound, at a fixed frequency, into the body and coherently detecting the backscattered waves with either a single transducer or an array of transducers. The received radio frequency signal (rf A-line) is envelope-detected, processed and displayed as a single, brightness modulated, line in the image. Thus, the B-mode image is a crude representation of the ultrasonic reflector/scatterer density distribution in a plane through the imaged medium. The phenomena perceived by the observer as speckle is predominantly due to the severe reduction of the received signal envelope by destructive interference effects. Interference effects also modify the spectral properties of the signal and can compromise both Doppler and pulse-echo attenuation.

Burckhardt [1978] presents a theoretical analysis of ultrasound speckle as an interference phenomenon. The analysis is based on an object comprising many point scatterers per resolution cell, with a random phase associated with each scattered echo. The large number of scatterers allows for the possibility of coherent interference between overlapping echoes, sometimes predominantly constructive, sometimes destructive. A point on the backscattered signal, originating from the object and measured at the transducer face, is therefore composed of a number phasors. Burckhardt states that if the number of scatterers per resolution cell is large and their phases are uniformly distributed between 0 and 2π, then the phasor sum becomes a random walk problem resulting in an image with predictable global statistics, such as a Rayleigh probability density for the received envelope amplitude. However, his treatment ignores other mechanisms which may also contribute to speckle, such as multiple scattering, high side or grating lobe levels and other factors which distort the phase of the wave across the receiving aperture.

Although of random appearance, speckle is not random in the sense of electrical noise, and consequently speckle patterns are readily repeatable under identical transducer-scanning conditions. As a result, some authors argue that speckle is important and can be used to segment an image. For example, Akiyama *et al.* [1990] use an estimate of the local fractal dimension to segment unfiltered B-scans while Wear *et al.* [1994] use an autoregressive method to form high resolution images of the backscatter coefficient. However, the majority of researchers view speckle as a distracting artifact which removes important image detail. In particular, its effect on edge positions is of critical importance when determining tissue boundaries.

4.4.2 Speckle reduction

Speckle reduction methods fall into three broad categories: statistical compounding, speckle identification and B-scan image filtering. Each of these is now discussed.

4.4.2.1 Compounding

Methods based on spatial or frequency compounding are essentially statistical in nature and rely on obtaining an ensemble of images whose distortions are statistically independent. The aim is to remove spurious image detail due to speckle while preserving the detail of the pristine image.

Spatial compounding entails combining a number of images of an object taken from different angles of view. These are then compounded by some procedure in the hope that the speckle-like pattern will be averaged to a low level, compared to the more stable image features that correspond to true tissue structure. In practice, spatial compounding has been applicable in only a limited number of regions of interest due to the practical limitations of conventional equipment. However, the advent of 3D imaging, using a spatial locator attached to a conventional sector probe, makes spatial compounding readily applicable as a method of speckle reduction. Average and maximum intensity compounding of B-scan images is inherent in the swept volume reconstruction described in Section 4.2.1. In the case of uncorrelated speckle, averaging on an intensity basis will produce an image with an improved signal to noise ratio (SNR) in proportion to the square root of the number of images averaged. Choosing the maximum intensity, also known as "maximum amplitude writing", is simpler to implement and recognises the fact that destructive interference effects, which cause drop-out in the signal envelope, are the main concern. Burckhardt's analysis of speckle shows that expected improvements in the SNR are almost as good for maximum intensity compounding as for averaging.

Frequency compounding involves averaging pulse-echo images produced using differing carrier frequencies [Trahey *et al.* 1987]. Alternatively, an ensemble of images can be derived by spectral decomposition of the wide band diffraction data. Minimally overlapping, narrow bandpass filters are applied to split the received wide-band signal into separate channels [Robinson 1982]. Each channel is usually envelope detected, and all the channels compounded by some method. This approach is also known as the *frequency diversity* method [Leeman *et al.* 1989].

Robinson [1982] has applied speckle processing techniques used successfully in astronomical imaging, such as speckle interferometry, shift and add, and Wiener filtering, to ultrasound data. However, these techniques have not been realised in practical medical imaging systems with soft tissues.

4.4.2.2 Speckle identification

This approach involves the explicit recognition and correction of speckle effects in the rf amplitude signals received by the transducer. It is based on a theoretical understanding of speckle as purely an interference phenomenon [Leeman *et al.* 1989, Healey *et al.* 1991, Healey and Leeman 1993].

Leeman *et al.* advocate processing the rf amplitude signal, which they call the A-line signal, rather than the 2D analysis of image data because the A-line signal represents the format in which pulse-echo systems acquire their raw data. Ideally, the A-line signals contain phase information which carries a deterministic imprint of the interference phenomena which lie at the root of the speckle artifact. Leeman proposes an interesting scheme called the Zero Adjustment Procedure (ZAP) which describes and corrects for speckle artifact in the complex z-domain. In this domain the signal amplitude and phase characteristics can be manipulated simultaneously. Large excursions in the instantaneous frequency from the carrier frequency of the imaging pulse are used to demarcate local speckle-corrupted segments in the A-line signal. Those A-line segments identified as being affected by destructive interference are completely described in the z-domain by the location of their zeros. The correction procedure attempts to modify the relative phase between local backscattered pulses to reduce destructive interference by adjusting the position of those complex zeros which are effectively responsible for the appearance of severe interference. Leeman argues that only destructive interference effects, which cause drop-outs in the signal envelope, need to be corrected.

The advantage of this system is that correction for speckle is localised and no specific statistical distribution is assumed. Furthermore, A-line processing is more likely to be achieved in real-time than 2D processing methods.

4.4.2.3 B-scan filtering

The most popular speckle reduction technique is 2D filtering of the B-scan image data. Most methods involve the application of a local filter operator whose output at a pixel is a function of the input grey-level values within a small neighbourhood of that pixel. The neighbourhood is thought of as a *window* since for each output pixel, only the pixel values within the window are used. The window is scanned across the input image, each position contributing to one pixel in the output image. The window can be of any size and shape, although it is almost always symmetrical about a centre point and usually square for implementation reasons. Herrington *et al.* [1991] differ by using anisotropic filters for radial intra-vascular catheter scans. They use a weighted Gaussian average to enhance tissue interfaces in a radial direction together with maximum intensity compounding of scans acquired from a stationary catheter position over time. Isotropic filters, however, are generally used in external B-scan systems.

Given that the speckle artifact has its origin in the individual A-line scans from which the image is computed, 2D processing of A-line signal envelopes appears a less satisfactory way to proceed. However, the raw rf A-line data is rarely available from modern scanners. Furthermore, practical ultrasound imaging systems deviate from theoretical models of speckle effects due to nonlinearities in the signal processing chain, envelope smoothing, dynamic range compensation (the mapping of echo amplitudes to grey-levels) and noise characteristics that are specific to individual machines. The grey-level image is also affected by user interaction. For example, the time-gain-compensation characteristic (TGC), which directly affects the grey-level intensities as a function of distance from the transducer, can be adjusted by the sonographer and is often varied during an examination. This important variable is not easily monitored, particularly on the Acuson XP10 scanner. All these effects hinder the development of general processing methods for B-scan images based on a theoretical understanding of speckle.

In the system described in this thesis, the Acuson XP10 scanner is effectively a 'black box' where only the B-scan image data is accessible. Filtering is necessarily restricted to processing B-scan grey-levels.

4.5 2D B-SCAN FILTERING TECHNIQUES

In practice, B-scan filtering is a compromise between noise suppression and preservation of image detail. Maximum noise suppression and preservation of edges are conflicting requirements which make development of a robust edge detection operator difficult. A plethora of filtering methods exist, partly because the presence of speckle within a B-scan image is not well defined, nor agreed upon. There is some debate over what constitutes the successful removal of speckle. For example, the compounding methods described above which use averaging imply that a signal is speckle-free if some average interference value is achieved, while the maximum compounding methods imply that an image is speckle-free if all the interference is totally constructive in nature.

Speckle reduction is an active field of research itself, and is not a major focus of this thesis. The philosophy adopted in the 3D imaging system has been to provide the user with a set of tools which are interactively applied, rather than attempting to formulate a general-purpose technique suitable for all imaging situations. The following describes the basic 2D filtering techniques used in ultrasound imaging which have a 3D analogue.

4.5.1 Gaussian and LoG Filters

A number of researchers [Coppini *et al.* 1995, Herlin and Giraudon 1993, Nixon and Hames 1993, Herrington *et al.* 1991] use Gaussian weighted averaging to smooth images and the Laplacian-of-Gaussian (LoG) [Marr and Hildreth 1980, Torre and Poggio 1986] or related Canny operator [Canny 1983] to detect tissue boundaries. The latter are rotationally invariant differential operators whose response does not depend on the edge orientation. The support for these filters is determined by the standard deviation of the Gaussian σ which controls the degree of spatial smoothing. An important property of the Gaussian filter for edge detection is that it does not introduce additional zero-crossings as the scale (σ) increases.

Initial inspection of these filters reveals that a relatively large support, corresponding to three to four standard deviations, is required (ideally infinite). Fortunately, the LoG kernel is separable in two variables and a fast convolution, consisting of a series of 1D row and column convolutions, is possible in 2D [Huertas and Medioni 1986, Chen *et al.* 1987]. Further computation is reduced by taking advantage of the symmetry within these 1D convolutions. The Canny-Deriche operator [Deriche 1990] is a recursive implementation which becomes increasingly efficient with larger supports (> 12 pixels) since the number of multiplications per pixel is fixed. These fast implementations generalise straightforwardly to 3D.

The main criticism of these filters is their inability to cope with multiple boundaries properly, and that the localisation of boundaries is accurate only for small window sizes, where the smoothing capabilities of the filter are correspondingly reduced.

The use of the zero-crossings of the Laplacian to detect edges and segment an image limits the topology of the segmented image. The zero-crossings form closed contours and the image is divided into regions where the Laplacian is positive and negative. A consequence of this binary classification is that no 'T' junctions may occur between edges detected in the image. Despite being inspired by similarities with the human lower visual cortex [Marr and Hildreth 1980], LoG-detected edges in images contaminated by noise often do not agree with those perceived by a human viewer. Bomans *et al.* [1990] and Hohne *et al.* [1990] have applied LoG segmentation to the rendering of 3D X-ray CT and MRI data but have found it necessary to apply binary morphological filters, erosion and dilation, to LoG zero-crossings prior to rendering, or alternatively, to manually edit zero-crossings.

In general, linear low-pass filters have the disadvantage that edges are blurred. The use of a linear high-pass filter to sharpen or detect edges has the disadvantage that high-spatial frequency noise is amplified and spurious oscillations are produced near edges. Non-linear filters overcome some of these disadvantages inherent in the use of linear filters.

4.5.2 The Median Filter

The most popular speckle filter, which is the basis of most comparisons, is the median filter. The median filter is a member of the family of rank filters. Median filtering, typically with a square window, achieves less blurring and better edge preservation than simple averaging and linear filtering techniques [Hodgson *et al.* 1985]. Whereas outlying values have a larger effect on the determination of the mean than those values close to the mean, the median gives equal weighting to pixel values. Consequently, the median is closer to the bulk of the intensity values than the mean.

The effect of applying a median filter to a region of constant intensity I and variance V is to reduce the variance. This is observed as a narrowing of the intensity histogram. The high spatial frequencies associated with noise are attenuated; in particular, oscillations in intensity with a period less than the window width are smoothed. Smoothing occurs because pixels of approximately the same intensity are selected from one window position to the next. However, in the special case of a binary oscillation of fixed frequency, whose period is smaller than the filter width, the oscillation will persist in the output.

Median filters are especially useful for eliminating impulsive noise, which speckle often manifests itself as. Provided that only a small proportion of pixels within the window are likely to be noise pixels, and these occupy the extreme rank positions, then the noise pixels will not be selected by the median. While rank filters preserve the general shape of edges, the median also preserves the position of an edge if the window is symmetric about its centre. Irregularities in the edge, which are less than the width of the window, are smoothed. Ringing at edges is also avoided since no new intensity values are generated by the application of a rank filter.

A drawback of the median is that it only preserves the mean intensity when applied to a region where the noise distribution is symmetrical. When the noise distribution is not symmetric, then the general effect of applying a median filter is to shift the mean intensity in the direction of the skew. Thus, if speckle results in predominantly low intensities, corresponding to destructive interference, then the median will reduce the mean intensity.

4.5.3 The Mode Filter

Evans and Nixon [1993] have proposed using a variation of the median filter to estimate the local mode within ultrasound images. It is difficult to determine the mode within a small population, typically 9x9. Consequently, they use a truncated median filter to estimate the mode of the local distribution. The truncated median filter truncates the frequency distribution of the data until the median bisects the range of the distribution, so that a new median value is found that is closer to the mode. In the original technique proposed by Davies [Davies 1988] this is an iterative process, where with each pass of the filter, the truncated median moves closer to the mode of the distribution. However, Evans has found that a single iteration is sufficient.

Evans and Nixon argue that near edges the mode will suppress noise where little is present and enhance an edge in the case of high noise. However, when directly on an edge, the mode can be either one of two peaks since data distribution is bimodal at an edge. This may lead to rather coarse decisions being taken. The mode may therefore offer an improvement over the median filter near edges but is likely to give erratic results on an edge.

4.5.4 Adaptive Filters

For fully developed speckle, when the number of scatterers present within the imaging system resolution cell is large, the image statistics are constant but as the number of scatterers decreases, and at boundaries between features, the local statistics vary. The non-stationary nature of ultrasonic images has motivated the development of adaptive filtering schemes. These adjust their smoothing properties according to the local image statistics. Many variations exist [Schistad and Taxt 1991, Herlin and Giraudon 1993, Evans and Nixon 1993]. Filters usually vary their output between the original grey-level value and the mean or median of the local region. The proportion of each is controlled by the ratio of the local mean to the local variance with the output tending towards the original value when this ratio is low (i.e. at an edge/discontinuity).

Evans and Nixon [1993] have investigated an adaptive median filter which performs a weighted ranking where the weights are a function of the distance from a pixel in the neighbourhood to the central pixel, as well as the local variance and mean. Adaptive filters rely on the statistical identification of speckle in the image and require a model to quantify the amount of speckle present. These are essentially subjectively chosen [Steen and Olstad 1994, Leeman *et al.* 1989]. Often some degree of user interaction is required to select appropriate parameters, although these may be stable for similar images.

4.5.5 Morphological Filters

Mathematical morphology is a set theoretic approach to image analysis which provides a framework within which multidimensional signal processing, in particular non-linear filtering, can be interpreted [Maragos and Schafer 1990].

A number of researchers have used morphological filters and geometric filters to remove speckle. Klingler *et al.* [1988] uses a combination of grey-level and binary opening and closing operations to remove speckle and segment 2D cardiographic images. Thomas *et al.* [1991] have used morphological techniques to automate foetal femur measurement in 2D ultrasound images, while Belohlavek *et al.* [1991] have used morphological filters to segment 3D ultrasound data.

Although many authors do not elaborate, the basic grey-level dilations and erosions applied to ultrasound data are frequently realised by min and max rank filters. Choosing a rank other than the median preserves the shape of edges in an image, but unlike the median, shifts the position of an edge in the direction of the region of lower intensity for ranks lower than the median, or higher intensity for ranks greater than the median. In other words, rank positions greater or lower than the median will expand or shrink regions of high intensity, respectively.

An opening operation (realised by an erosion followed by a dilation) can therefore be realised by a min filter followed by a max filter. Conversely a closing operation is realised by a max filter followed by a min filter. Opening and closing represent a type of low-pass filter, where the cut-off frequency depends on the size of the window (the larger the size, the larger the features removed and the lower the cut-off frequency). However, unlike true low pass filters they preserve edges.

The main disadvantage of these filters is that the mean intensity level, in general, is not maintained if the input images are noisy. This can be improved by sequentially applying opening and closing filters. An opening followed by a closing operation effectively removes spikes from the image. The opening filter removes positive going features, while the closing filter removes negative going features. The size of the features eliminated depends on the size of the window.

In Klingler's system multiple cardiac images are acquired at the same stage in the cardiac cycle and compounded using averaging to improve the signal to noise ratio. Klingler then removes the background variation in the image by applying a large hexagonal grey-level opening filter to the image and subtracting this from the original data. Isolated regions and narrow peninsulas smaller than the structuring element are completely removed by an opening operation. Subsequent closing operations bridge gaps in structures containing holes smaller than the size of the closing element.

When applied to ultrasound images, morphological filters are capable of filling in drop-out regions, i.e. bridging gaps in structures where no data is present, or alternatively, completely removing small structures present in the data. These filters will alter well-defined edges which are free from speckle, as readily as those corrupted by speckle. They are therefore best applied with caution and suited to applications where a binary segmentation is appropriate, such as defining the ventricular cavities in Klingler's system, or the foetal femur in Thomas's case. Interestingly, Belohlavek *et al.* [1991] extracts 'objects' from morphological filtered ultrasound data and uses these as binary masks to segment the original data volume, thus preserving the original texture information in subsequent rendered views. Belohlavek initially filters the data with small 3D grey-level structuring elements, then applies a median filter to improve edge definition and further reduce image noise. The data is then thresholded and iterative binary dilation is applied until distinct segments in the image are joined (an appropriate halting criterion is not specified). Morphological 'thinning' is then applied to extract closed surfaces which are used to segment the original data.

An introduction to morphological techniques in image processing can be found in Gonzalez and Woods [1992] and Pratt [1991], while Dougherty [1993] provides a more advanced treatment. Bailey [1985] presents a detailed summary of the properties of rank filters. In practice, morphological techniques are applied somewhat arbitrarily. Parameters such as threshold levels, the number of iterations and the sizes and shapes of structuring elements are usually derived empirically or determined by user interaction.

4.5.6 Filtering comparison

Figure 4.11 compares the application of averaging, median filtering and Gaussian weighted averaging to a 2D slice from the foetal data set. Figure 4.11(a) is a representative slice from the data set, Figure 4.11(b) is the same slice with the contours superimposed which correspond to the threshold used in the previous 3D renderings (Figures 4.5-4.9). 63 separate regions were identified by the threshold operation. Figure 4.11(c) illustrates the contours for the same threshold choice superimposed upon the image after convolving with a 5x5 averaging filter. The image was segmented into 17 separate regions. Figure 4.11(d) is the median filtered image with a 5x5 window. 20 regions have been found. Figure 4.11(e) is the corresponding Gaussian weighted average. A standard deviation of $\sigma = \sqrt{2}$ was used to achieve a similar degree of

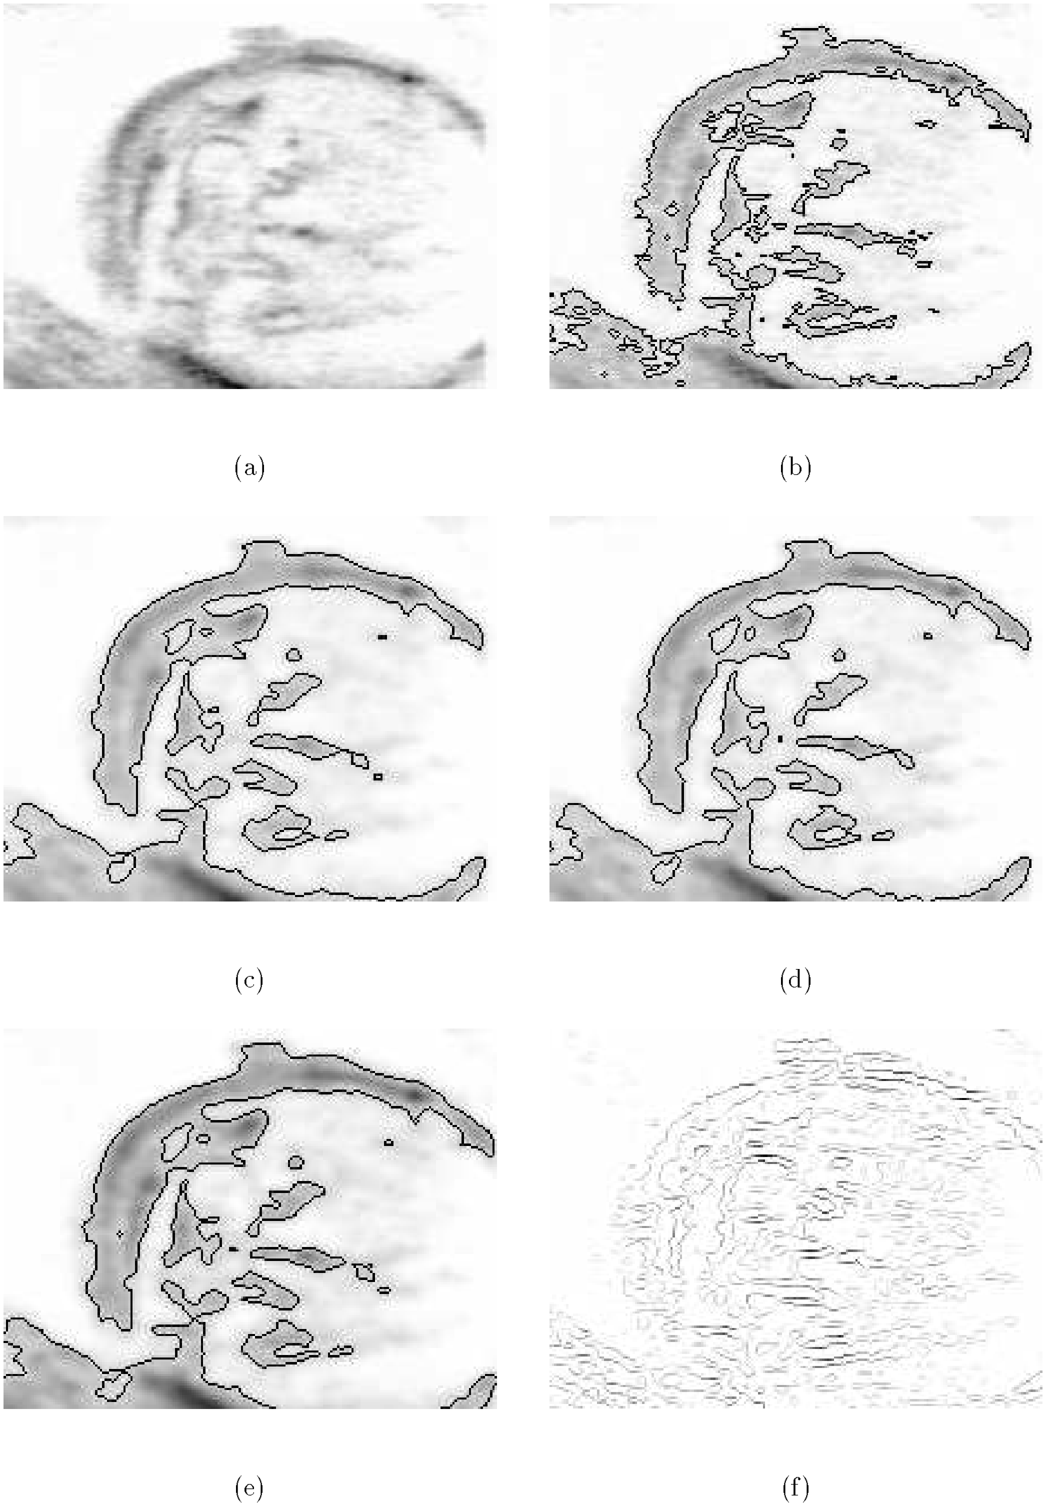

Figure 4.11 (a) B-scan slice, (b) iso-value contours corresponding to the ray-casting surface threshold, (c) 5x5 averaging filter, (d) 5x5 median filter, (e) Gaussian averaging filter, $\sigma\sqrt{2}$, (f) corresponding zero crossings of the Laplacian, the pixel intensities are proportional to the detected edge strength.

smoothing as for the 5x5 averaging and median filters. 17 regions have been identified by the threshold operation. Figure 4.11(f) is the corresponding zero-crossings of the Laplacian, where the intensity of the displayed contours is proportional to the edge strength.

In Figure 4.12 2D morphological opening and closing operations have been applied to the data slice depicted in Figure 4.11(a). Figure 4.12(a) is the result of grey-level morphological closing with a circular structuring element 5 pixels in diameter. 18 regions have been found. Figure 4.12(b) is the result of applying a grey-level opening operation to Figure 4.12(a), resulting in 14 regions. In Figures 4.12(c) and 4.12(d) the order of the operations has been reversed. Figures 4.12(c) is the result of first applying the opening operator, with 26 regions identified. The closing operator was then applied in Figure 4.12(d), resulting in 10 regions.

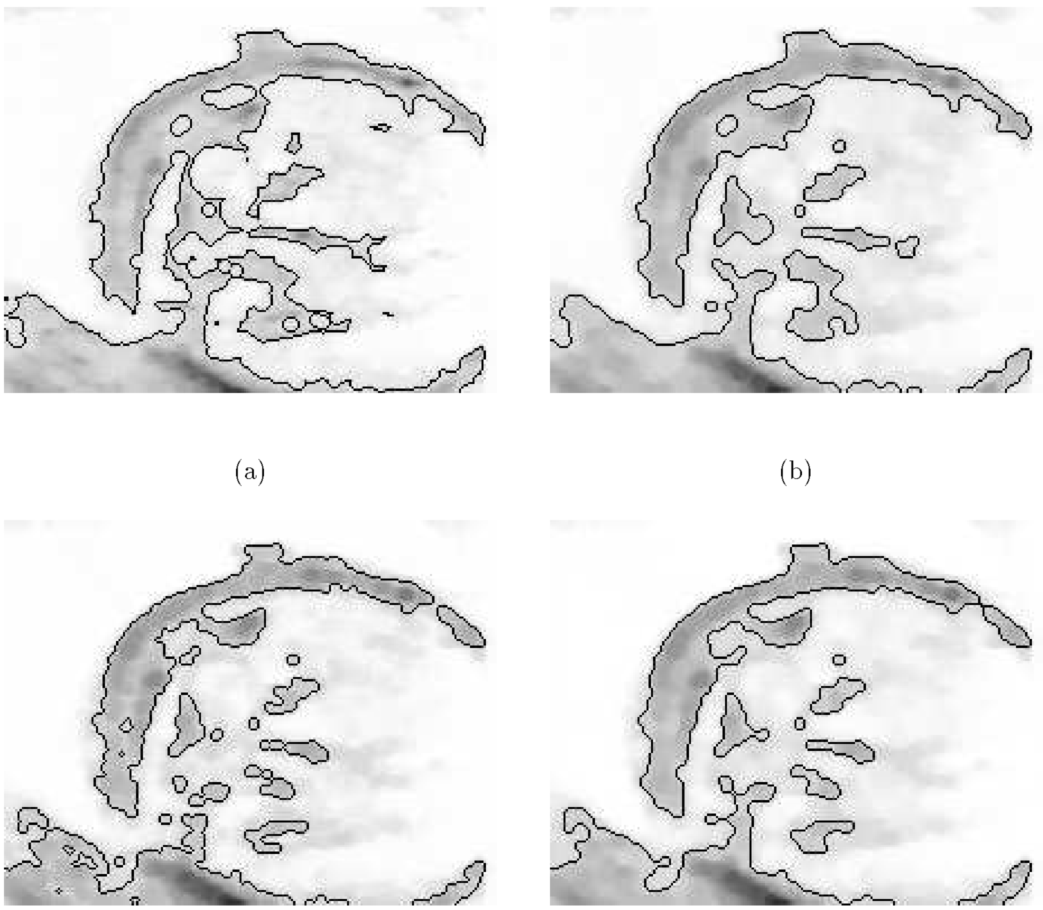

(a) (b) (c) (d)

Figure 4.12 (a) Grey-level closing with a circular structuring element 5 pixels in diameter, (b) closing followed by opening, (c) grey-level opening with a circular structuring element 5 pixels in diameter, (d) opening followed by closing.

Little difference is apparent between the average, Gaussian weighted average and median filters in Figure 4.11; the thresholded images all have similar topologies. The morphological filtered images, however, differ significantly from the former; edges and features present in the median and average filtered images have been removed or altered. The order of the opening and closing operations results in the identification of different regions in the thresholded image. This will be true in general. Morphological operators must therefore be applied cautiously and the viewer must be aware of the potential alteration of the data, either by the removal of backscatter structures or through the bridging of backscatter structures.

4.5.7 Discussion

It is difficult to draw conclusions from studies of specific images. Many authors present numerous comparisons of filtering techniques applied to *in vivo* images, but the correct position of edges and details is not known and it is therefore difficult to quantify speckle reduction or whether real edges have been modified. Evans and Nixon [1993] are among the few who have used a phantom to compare methods. They have also proposed a quantitative measure of speckle. The *Contrast to Speckle Ratio* (CSR) measures the observer's ability to observe a void or dark area against a background of speckle and is defined as,

$$CSR = \frac{\tilde{x_o} - \tilde{x_i}}{(\sigma_o^2 - \sigma_i^2)^2} \tag{4.2}$$

where $\tilde{x_i}$ and ${\sigma_i}^2$ are the average signal value and variance, respectively, inside the void and $\tilde{x_o}$ and ${\sigma_o}^2$ are those outside the region. They found that the median and adaptive median filters offer the greatest improvement in SNR and CSR away from edges, while the strongest edge was displayed by the truncated median which approximates the mode. Overall, ignoring the significant computational cost over a standard median filter, they conclude that an adaptive median filter is the best choice.

The arguments for and against particular filtering techniques are often empirically based and assessment is subjective in the sense that a filter is judged according to how it affects perceived edges and structures in a representative set of images. However, perception is influenced by knowledge of the underlying anatomy and of the nature of ultrasound artifacts. Consequently, a general filtering process, optimal for all applications, is unlikely to be developed from consideration of image grey-level values alone. Automatic analysis of ultrasound images is outside the scope of this thesis; nevertheless, segmentation is closely related to the process of visualisation. Rather than attempting to define a general-purpose speckle reduction technique, an interactive approach has been taken to filtering. A small set of basic filters, including the Gaussian weighted average, the median and morphological opening and closing filters, have been found to provide adequate renderings of high contrast structures in 3D data. Surveys by Schistad and Taxt [1991] and Thune and Olstad [1991] support the general applicability of median and rank filters in preference to linear low pass filters to ultrasound data.

2D filtering of B-scan slices prior to volume reconstruction is performed for two reasons. Firstly, if the 3D reconstruction is at a lower resolution than the constituent 2D slices then an effective down-sampling occurs in the reconstruction process; hence low pass filtering of the slice data is appropriate. Secondly, interpolation between slices in the reconstruction process spreads speckle artifacts over a finite volume. Furthermore, 2D filtering is more efficient to implement than 3D filtering. Experience has shown that reducing speckle prior to reconstruction minimises the need for 3D filtering. A median filter has been initially used on the 2D data in the reconstructions following in Section 4.6. In the case of severe speckle artifact, such as occurs with imaging the veins in Figure 4.10, morphological opening and closing operators are useful for removing large speckle structures.

3D median and rank filtering has been implemented with a cubic window rather than a spherical window. The cubic window is convenient and efficient to implement, but has a directional frequency response. This directional dependence does not have a noticeable effect on the performance of the median filter, but a blocking artifact is introduced when large dilation and erosion operators are applied. However, efficient methods do exist for implementing min and max filters in 2D which use rank positions other than the extremes to reduce the directional dependence [Hodgson *et al.* 1985].

Figure 4.13 demonstrates the application of a 9x9x9 median filter (3mm^3) to the foetal data set. The data set consists of 220x172x204 voxels. Because the voxel resolution corresponds to the pixel resolution in the original 2D data slices, 2D filtering of the slices was not applied. Figure 4.13(a) is the unfiltered data set rendered using a grey-level volume gradient. In the filtered image, Figure 4.13(b), small speckle regions have been removed and the shaded surface is noticeably smoother.

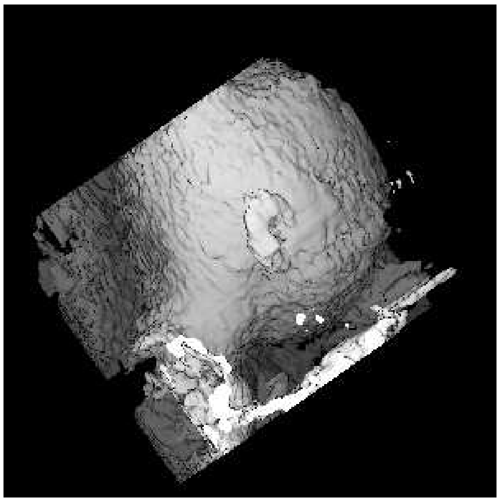

(a)

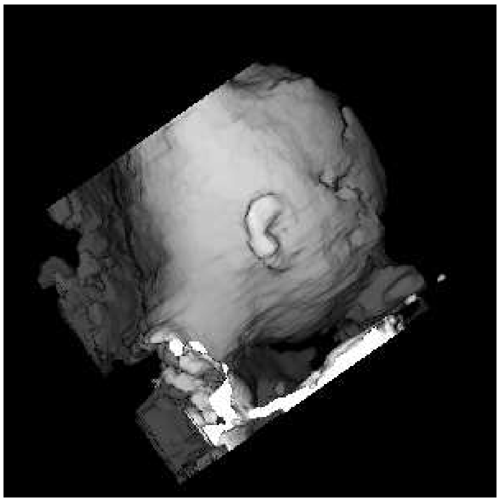

(b)

Figure 4.13 Comparison of unfiltered and median filtered foetal data rendered with volume gradient shading. (a) Unfiltered data, (b) Median filtered data with a 9x9x9 window. The data volume consists of 220x172x204 voxels.

4.6 3D RECONSTRUCTIONS FROM *IN VIVO* DATA

The following sections present the results of *in vivo* imaging of human tissue. The filtering applied to the data and visualisation parameters are detailed. Practical scanning considerations and artifacts are also discussed when they arise.

4.6.1 Hepatic veins

In Figure 4.10 33 B-scan slices of the liver were acquired with a 7MHz sector probe. The slices were initially filtered with a 7x7 median filter ($1.75 mm^2$) followed by a 5 pixel diameter opening filter to remove the persistent speckle. A volume consisting of $1 mm^3$ voxels was then constructed and filtered with a 3x3x3 ($3 mm^3$) median filter to smooth the interpolated volume data.

Two factors are important when interpreting Figure 4.10. The first is that the filtering process has removed fine vascular structures from the image at a scale smaller than 1mm, as well as speckle. The second factor is that no cardiac gating (synchronisation to the heart beat) was used during the acquisition of the data. The acquisition took 3 seconds during which the subject held their breath in order to minimise tissue movement due to respiration. The pulsing and translation movement due to circulation was very pronounced in the smaller vessels. Movements of up to 5mm were observed in the periphery of the constituent B-scan slices. Consequently, a sub-volume of the

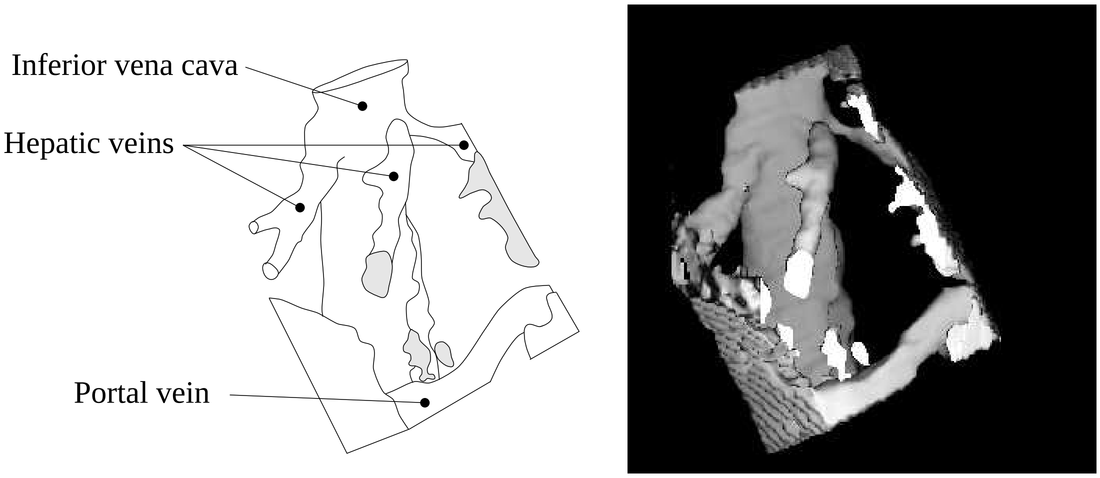

(a)

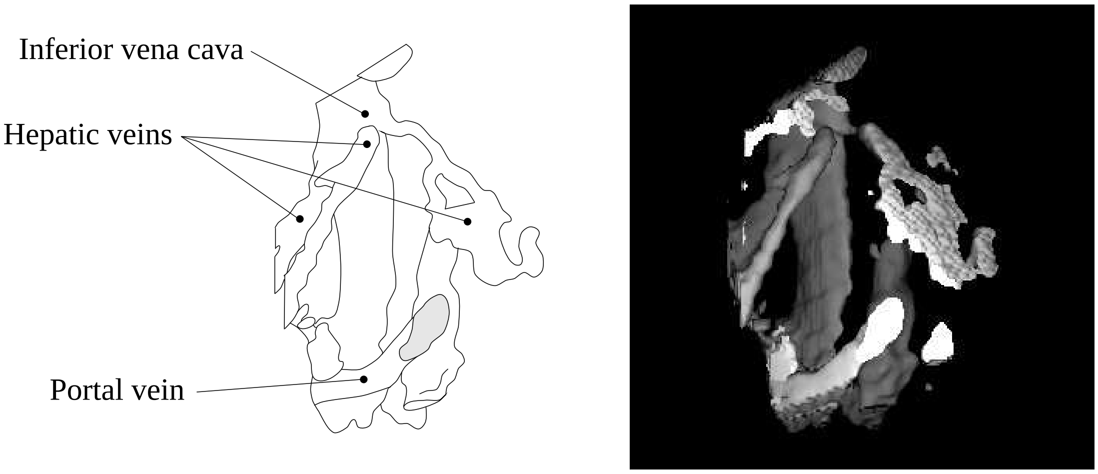

(b)

Figure 4.14 Two views of a volume reconstruction of the hepatic veins and inferior vena cava using volume gradient shading. (a) Cut-away view, (b) Cut-away view of the same data set, rotated by 45° about the vena cava.

entire volume imaged by the probe is depicted here. This contains the larger vessels which were less affected by movement.

A fundamental assumption in the reconstruction technique is that no tissue movement relative to the Polhemus transmitter occurs during acquisition. A number of authors synchronise the acquisition of data to the ECG waveform and/or the respiratory cycle. Klingler *et al.* [1988] use digital subtraction to detect significant movement between acquired frames and then reject slices with too much variation.

Figure 4.14 illustrates two views of the hepatic veins and inferior vena cava acquired during another examination. 57 data slices were acquired with a 7MHz sector probe and filtered with an 11x11 pixel (2.7mm^2) median filter followed by a 5 pixel diameter opening filter. The reconstructed volume was median filtered with a 3x3x3 window. The threshold selected to render Figure 4.14 differs from that for Figure 4.10 due to operator variation of the time-gain characteristic and grey-level mapping characteristic on the scanner. The thresholds were selected interactively for each data set. The three hepatic veins branching off from the inferior vena cava in Figure 4.10 are apparent in Figure 4.14(a). Figure 4.14(b) is a view of the same data set rotated 45° about the vena cava. This view corresponds to the view of the data set depicted in Figure 4.10. A larger volume, extending 100mm along the length of the inferior vena cava to the portal vein, is shown here. The structures depicted are sufficiently complex to illustrate the benefit of a 3D reconstruction. The relatively small difference in the viewing direction between Figures 4.14(a) and 4.14(b) reveals the need to interactively manipulate the viewing direction in order to perceive the structures present. Animated sequences of rotations about a fixed axis are also useful for understanding images.

Cut-away views have been used in both Figures 4.14 and 4.10 to reveal the structures of interest. Bone surfaces rendered from X-ray CT data in Chapter 4 depicted CT data above the applied threshold as solid structure. Because echoes are absent from the fluid-filled vessels of interest in Figures 4.10 and 4.14, regions below the applied threshold have been depicted as solid. Consequently, regions of signal drop-out or shadow may become confused with genuine fluid-filled cavities. Such regions result in the small irregular structures observed in the periphery of the volumes rendered in Figures 4.10 and 4.14. The false depiction of signal dropout as anatomic structure is more apparent in studies of the gall bladder.

4.6.2 Gall bladder

The gall bladder is situated under the right lobe of the liver and attached by connective tissue. It is essentially a bag which fills with bile produced by the liver and empties in to the duodenum via the cystic bile duct. Its shape is determined by the surrounding tissue and deforms depending on the physical position of the body and how full the bladder is. It is not a rigid structure and its shape may change depending on whether the patient is standing, lying flat on their back or over to one side.

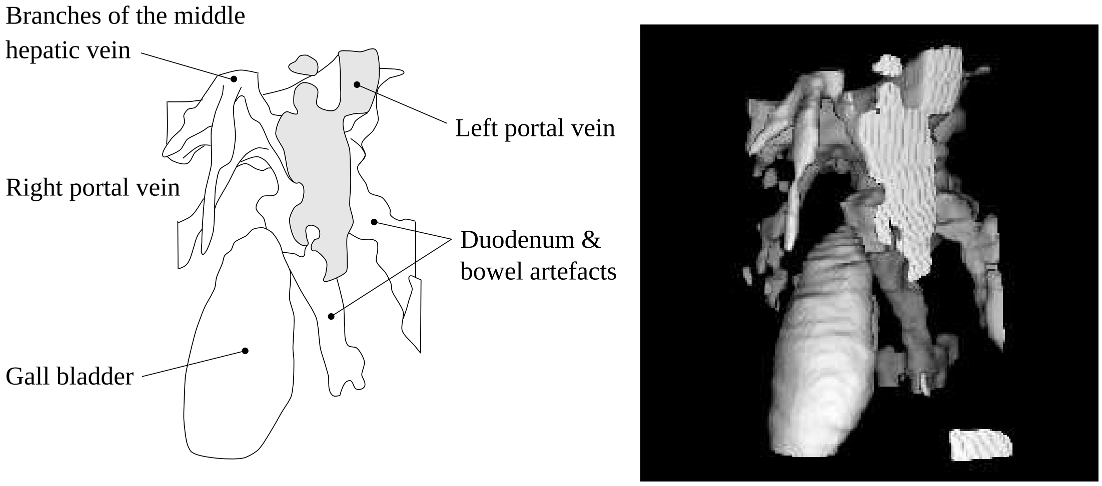

Figure 4.15 Reconstruction of the gall bladder.

Figure 4.15 is a study of the gall bladder reconstructed on a regular grid of 0.5mm^3 voxels from 57 transverse sector scans. The constituent 2D B-scan slices were median filtered with a 9x9 filter (2.25mm^2). The 3D reconstruction was filtered with a 2.5mm^3 median filter. The gall bladder is about 41mm in length in this view.

Figure 4.16 is one of the constituent 2D B-scans from which Figure 4.15 was constructed. For the purpose of clarity, the grey-scale in Figure 4.15 has been inverted with respect to how images are commonly displayed on a scanner. Strong reflections are depicted in black, while regions free of echoes are shown in white.

The duodenum lies immediately behind the gall bladder and is a rich source of artifacts. Gas in the bowel results in strong reflections which are often characterised by reverberation artifact, also known as 'comet tail' artifact. When the wall of the gall bladder aligns parallel with the direction of sound propagation, high attenuation along the wall of the bladder may result in acoustic shadows being cast behind the bladder. Sonographers expect to observe these artifacts in 2D B-scans of the gall bladder and assign limited meaning to echoes arising from structures behind the liver. When a simple threshold is applied to the reconstructed volume data, regions of signal drop-out become indistinguishable from those due to a fluid-filled cavity.

In Figure 4.15 irregular shaped solid structures appear to lie behind the gall bladder. It is tempting to assign meaning to these structures. However, consultation with sonographers from Christchurch Hospital revealed that these structures are a mixture of genuine tissue interfaces within the bowel and shadow artifact arising from attenuation and refraction. Consequently, these structures have little anatomic meaning.

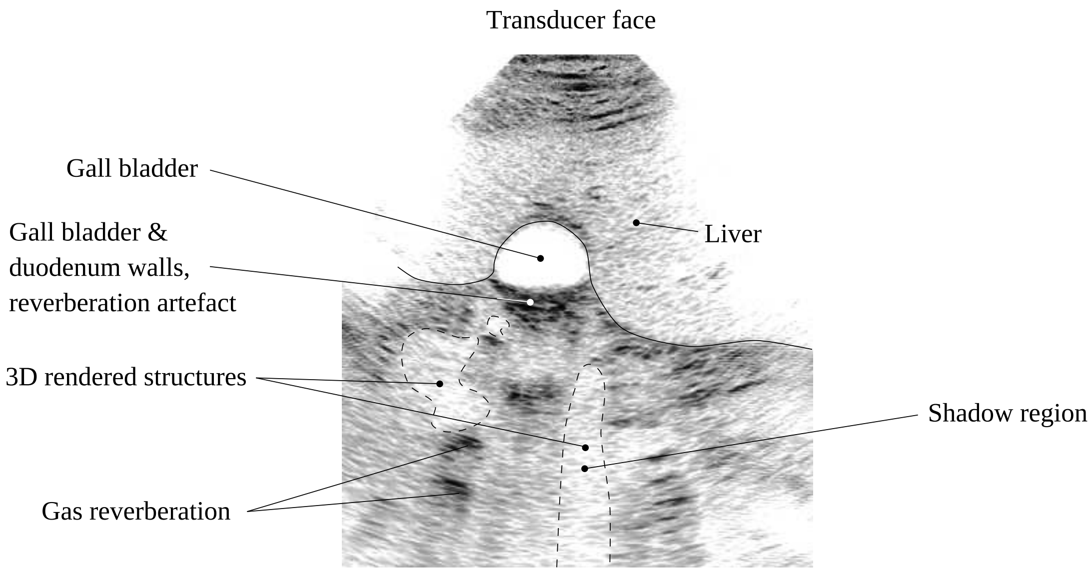

Figure 4.16 Transverse B-scan image of the gall bladder.

Importantly, the sonographers found it easier to recognise artifacts in the 2D B-mode images than in the 3D surface display. This is partly due to extensive experience analysing 2D B-scans, but also because valuable backscatter information, visible in the 2D slices, is lost in the surface rendered view. Reformatting the reconstructed volume to produce an animated sequence of 2D slices through the data was helpful for interpreting the artifacts in Figure 4.15.

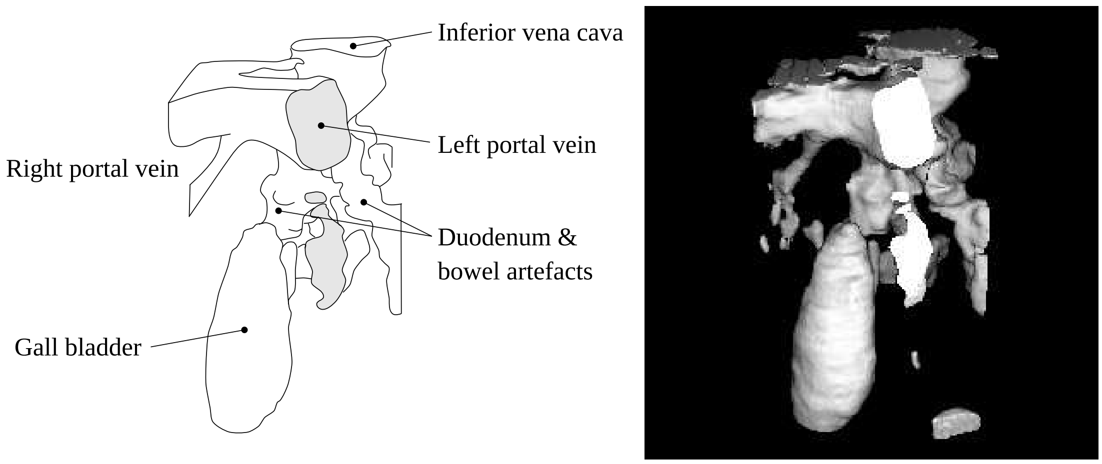

Figure 4.17 Cut-away view of the gall bladder.

Figure 4.17 is another view of the reconstructed volume in Figure 4.15. The foreground hepatic veins have been removed and the volume tilted forward to reveal the inferior vena cava and portal veins more clearly.

Figure 4.18 Reconstruction of the gall bladder depicting part of the cystic bile duct.

Figure 4.19 Radiograph of gall bladder from Wier and Abrahams [1986].

Figure 4.18 is a separate study of the gall bladder. Much of the surrounding structure has been removed by truncating the data volume with planar slices. In this figure the neck and beginning of the cystic bile duct can be seen. The volume is constructed of 0.5mm^3 voxels and has been filtered with a 2.5mm^3 median filter.

Figure 4.19 is an X-ray of a gall bladder from Wier and Abrahams [1986]. To form this image a contrast agent, calcium ipodate, was given orally then absorbed in the intestine and concentrated in the gall bladder. The X-ray reveals a similar shape to that observed in the reconstructed ultrasound images.

When performing a 2D examination, the ability to resolve certain anatomic features is better in certain orientations of the scan plane than others. The same is true for 3D graphics constructed from a sweep of 2D scans. Figures 4.15, 4.17 and 4.18 were formed from transverse sweeps of the probe. Figure 4.20 is a further study of the gall bladder constructed from a longitudinal sweep. This reconstruction reveals the top of the gall bladder folded over on itself (indicated by the presence of the crease). A similar shape is revealed in the X-ray in Figure 4.22 from McMinn *et al.* [1993]. This image was achieved by introducing a contrast agent directly into the billary system via an endoscope passed through the mouth.

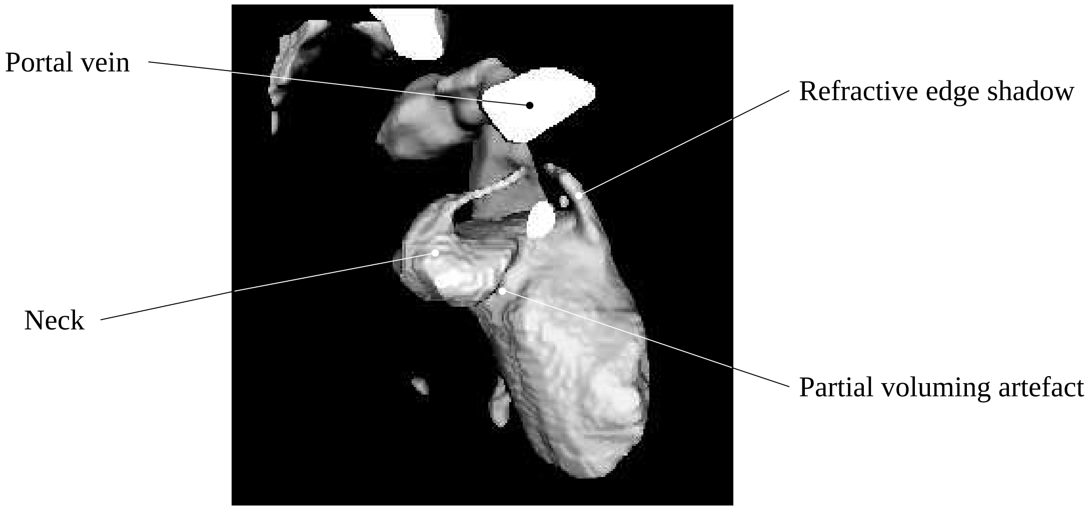

Figure 4.20 Gall bladder reconstructed from longitudinal B-scans.

Unlike the transverse scans, the longitudinal scans of the gall bladder proved to be difficult to segment by simple thresholding. In Figure 4.20 manual assistance was used to guide the determination of the tissue boundary in some of the slices through the reconstructed data volume. Figure 4.21 is a constituent slice from the data set.

In Figure 4.21 the wall of the gall bladder produces a specular reflection which results in reverberation artifact inside the gall bladder when the wall is perpendicular to the direction of sound propagation. When the wall is obliquely oriented to the direction of sound propagation, the received echo comes mainly from backscattered waves and is comparatively weak; little sound is reflected directly back to the transducer face. Strong reflections from gas in the adjacent duodenum result in reflections being erroneously displayed inside the bladder. This is because the scanner displays echoes from strong interfaces in the periphery of the beam as weak reflections along the beam axis — a partial voluming effect. The slice thickness also results in the partial voluming of structures normal to the plane of the scan. In longitudinal scans, the shape of the gall bladder changes rapidly in this direction. Consequently, the wall of the gall bladder

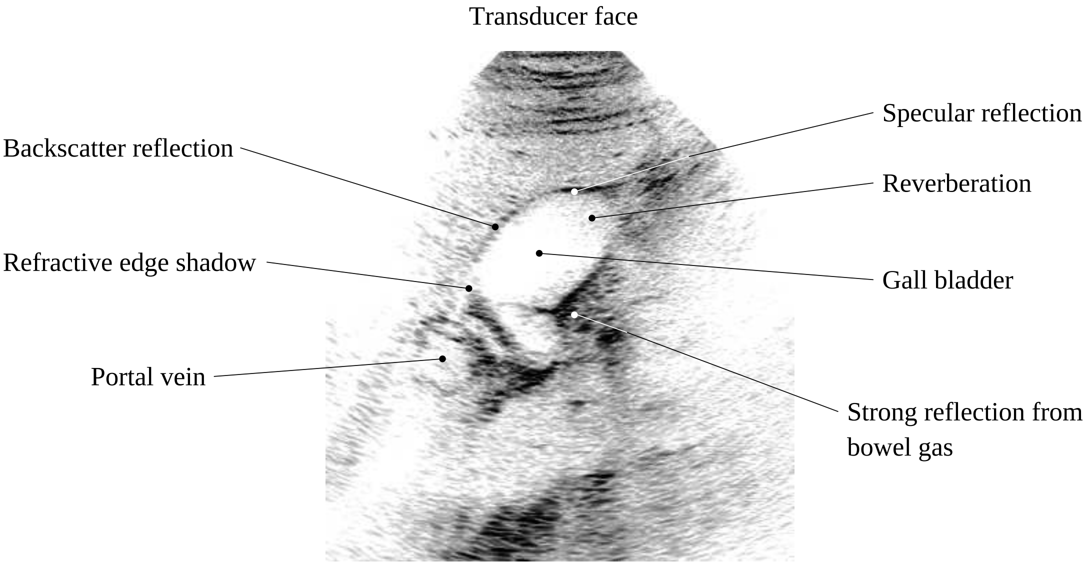

Figure 4.21 Artifacts in a longitudinal B-scan of the gall bladder.

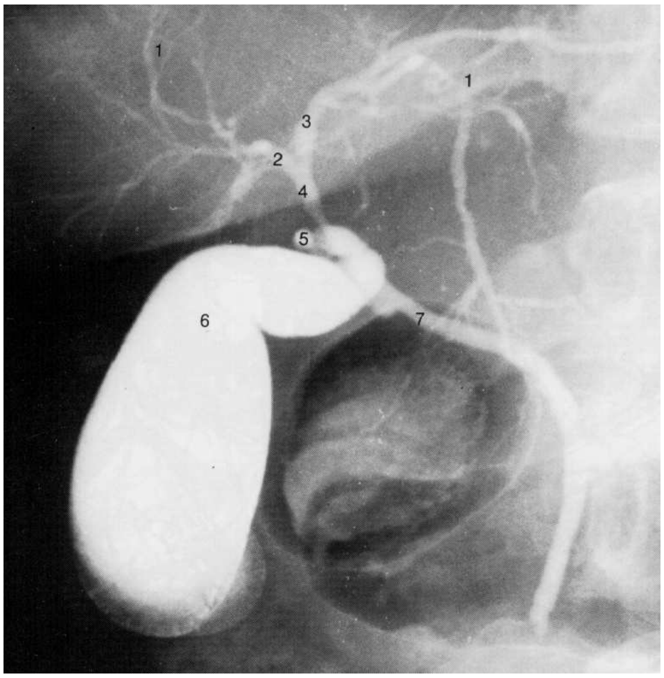

Figure 4.22 X-ray of Gall bladder from McMinn *et al.* [1993].

appears blurred. All these effects result in the gall bladder wall being depicted with lower contrast compared to the transverse scans. This makes it difficult to choose a global threshold which adequately segments the gall bladder, particularly in the region of the neck and base.

A further artifact apparent in Figure 4.21 is refractive edge shadow. This occurs when the critical angle of radiation incident on an interface is achieved, and the incident wave is refracted along the interface boundary. High attenuation along the plane of the gall bladder wall results in signal drop-out (i.e a shadow) in the corresponding line of the B-scan image. Such a shadow is indicated in Figure 4.21 and is readily recognised, and even expected, by trained sonographers in 2D scans. In Figure 4.20 this artifact manifests itself as a well-defined protrusion from the right side of the bladder. If unaware of the artifact, it might be tempting to erroneously assign meaning to this structure. It is therefore necessary to know where the probe is in relation to the rendered view in order to be able to interpret a 3D image reliably. In Figure 4.20 the gall bladder has been insonified from below, in line with the refraction artifact.

4.6.3 The urinary bladder

Figure 4.23 is a study of the urinary bladder. This structure is also a non-rigid fluid-filled bag. Quantitative measurement of the volume of the bladder, as a function of time, is of clinical interest. Currently, the volume can be estimated from a series of 1D measurements taken from 2D sector scans [Kruczkowski *et al.* 1988] or determined accurately by catheterisation. The latter, being an invasive method, is prone to introducing infection, hence a non-invasive method is favoured. In the view shown in Figure 4.23 the bladder is 70mm high, 75mm wide and 50mm deep containing approximately 260ml. A posterior view is shown and the interface furthest from the transducer face is depicted in this view.

Figure 4.24 is a second bladder study showing posterior and anterior views of the bladder. This set of data suffers from poor depiction of the interface closest to the transducer face. This is apparent in the constituent 2D slices and is illustrated by the representative slice shown in Figure 4.24. The poor contrast, due to reverberation, in the 2D scans causes the threshold process to fail and the reverberation artifact manifests itself as a series of concavities in the front surface of the bladder.

The signal drop-out regions at the periphery of the 2D scans have been depicted in 3D as large peripheral solid structures and are not to be confused with genuine anatomy.

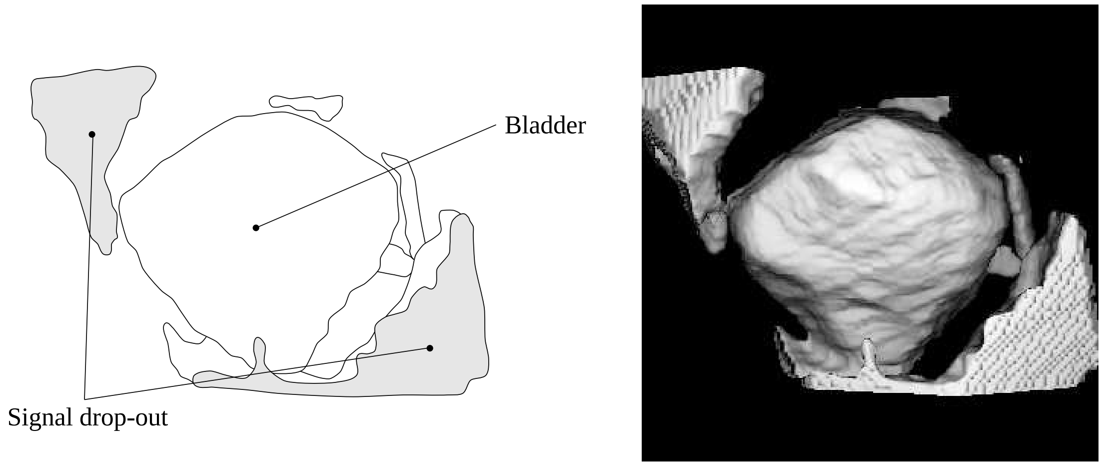

Figure 4.23 Reconstruction of the bladder, posterior view.

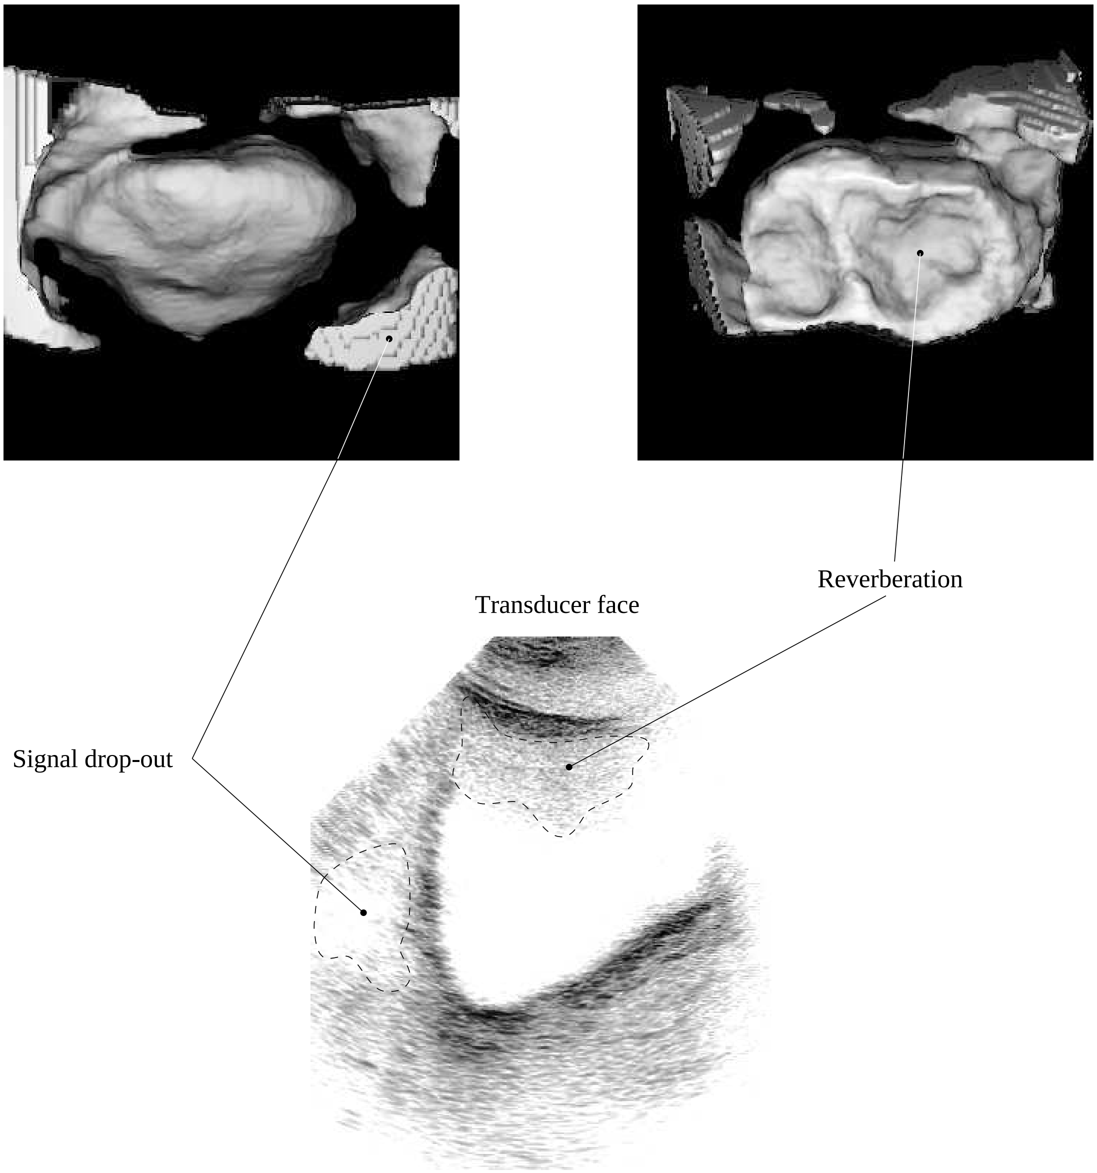

Figure 4.24 Bladder reconstruction affected by reverberation artifact. Left hand image is the posterior view, right hand image is the anterior view.

4.7 CONCLUSIONS

This chapter has discussed the visualisation of 3D ultrasound data and presented two methods for rendering surfaces from a series of registered B-scans. These methods have been demonstrated with simulated B-scan slices and through *in vivo* imaging. How 2D artifacts manifest themselves in 3D reconstructions has also been illustrated.

The clinical usefulness of rendering opaque surfaces, versus semi-transparent techniques, is offset by the difficulty of defining surfaces in ultrasound data. Specular reflections, speckle interference effects, reverberation, resonances and shadowing effects mean that simple thresholding approaches are not robust segmentation techniques without prior filtering of the data. Consequently, the quality of images is degraded compared to those obtained from X-ray CT data. Volume shading is particularly susceptible to noise effects. Depth shading alone produces comparatively good results since it does not emphasise surface roughness, but consequently surface detail is also lost.

In Section 4.2.1 a simple reconstruction method which forms a regular array of voxels from non-parallel slice data has come about by restricting the acquisition of slices to smooth sweeps of the probe. Data in this format is then suitable for conventional ray-casting and multi-planar reformatting. In Section 4.3 a method is described which produces graphics directly from non-parallel slices by projecting them into a depth buffer as they are acquired. A threshold surface is rendered from the data by linearly interpolating the data between slices in the direction of the viewer. Because of this directional dependence, the reconstruction varies with the view-point and the image quality degrades when the viewing rays intersect slices at oblique angles.

In practice, the usefulness of a real-time reconstruction is limited by the difficulty associated with choosing an appropriate viewing direction, specifying rendering parameters, filtering parameters, and cut-away planes, prior to scanning. The latter are frequently required to remove occluding structures. The use of the probe to define the viewing plane and viewing direction simplifies the specification of these parameters, but the best results restrict scanning to be in the direction of the 'viewer'. Nevertheless, this method is fast and useful for quick, depth-shaded renderings immediately after, if not during, acquisition.

Spatial compounding, inherent in the swept-volume reconstruction, improves the signal to noise ratio of the data, but the amount of compounding is data-dependent and varies from voxel to voxel. Speckle reduction and some smoothing is required to render meaningful images of structures embedded in soft tissue. The B-mode image is a crude representation of scattering interference effects, being an envelope-detected signal modified by dynamic range conversions and time-gain-compensation. In Section 4.6 minimal filtering, using median, average and morphological filters has mostly been sufficient to render high contrast structures such as the foetus and fluid filled cavities like the bladder and blood vessels.

3D surface rendering has been useful for revealing high contrast structures, but multi-planar reformatting of the reconstructed grey-level volume is useful for distinguishing artifact from genuine structure and is preferable in regions which are difficult to segment using a threshold. Interpretation of 3D surfaces rendered from ultrasound data is subject to knowledge of the filtering processes applied to the data, tissue movement during acquisition and false structures resulting from artifacts. These artifacts are inherent in the underlying 2D scans from which 3D images are constructed.

Visualisation is a process which involves the selective reduction and presentation of information within a data volume. Studies of the bladder and gall bladder from longitudinal scans have demonstrated the need for more sophisticated segmentation techniques, for example, adaptive thresholding or edge detection. Developing semi-automatic segmentation tools is therefore an important area of future work necessary for extracting clinically useful information from 3D ultrasound data. However, these are difficult skills which human experts learn through experience and are unlikely to be easily automated.

The possibility of combining multiple sweeps, discussed in Section 4.2.1, increases the likelihood of temporal variation in the imaged distribution. An intriguing possibility is the use of a second or third spatial locator, suitably attached to the patient, to monitor and compensate for tissue movement. However, such a system would require an accurate model of the deformation of soft tissues, and is likely to be limited in its applicability.

The repeatable nature of the reconstructions performed with the current system provides some confidence in the 3D reconstructions. However, quantitative assessment is required. There is the need to develop a 3D tissue equivalent phantom for developing and quantifying speckle reduction techniques, visualisation techniques and compounding methods for eliminating artifacts. Such a phantom requires careful design and must remain accurate over a reasonable lifetime. An anticipated first step towards verifying the accuracy of *in vivo* imaging is the verification of bladder volumes estimated from 3D reconstructions. Due to refraction, which is inherent in ultrasound imaging, measurements from 3D reconstructions are not expected to be 100% accurate. However, the degree of accuracy needs to be quantified over a range of volumes. Verification requires accurate measurement of bladder volumes for comparison. This is best achieved by catheterisation — an unpleasant procedure.

Chapter 5

SHAPE RECONSTRUCTION FROM INCOMPLETE SURFACE DATA

Iso-value surfaces were rendered from filtered ultrasound data in Chapter 4 by simply thresholding the data. This was sufficient to visualise regions in which echoes were absent, i.e. fluid-filled cavities embedded in soft tissue. However, ultrasound primarily provides information concerning the location of tissue interfaces within the body. When rendering structures other than fluid-filled cavities, thresholding reveals incomplete noisy surfaces between tissues rather than continuous regions [Klingler *et al.* 1988, Steen and Olstad 1994]. In this chapter the problem of extracting closed surfaces from incomplete surface data is studied. The explicit identification of a closed surface is necessary to extract quantitative information such as organ volume.

Many authors have commented on the difficulty of constructing reliable 3D edge-detectors for ultrasound data or reliable methods for tissue classification from ultrasonic backscatter [McCann *et al.* 1988, Herlin and Giraudon 1993, Thune and Olstad 1991]. Due to poor intensity contrast, speckle, and shadow regions, tissue boundary detection is difficult. Many researchers use semi-transparent rendering to visualise structures within the data thus avoiding the issue. The novel approach to the problem of extracting surfaces described in this chapter was inspired by the observation that structures which appear ambiguous when depicted by semi-transparent methods become clearer when viewed from multiple viewpoints in an animated sequence. In this method multiple projections from many viewpoints are used to reconstruct incomplete object surfaces. A binary classification of the data results which is suitable for making quantitative measurements and surface rendering.

In the swept volume reconstruction described in Chapter 4 (section 4.2.1) ultrasound slices were restricted to smooth sweeps of the probe where temporally adjacent slices were also spatially adjacent and semi-parallel. In this chapter ultrasound data from completely arbitrary slices is considered. No temporal or spatial ordering of the data is assumed. The clinical problem, however, is simplified by considering the reconstruction of a single isolated object. Specific aspects associated with fitting closed surfaces to medical ultrasound data are addressed by the proposed method. These include the non-uniform distribution of the data (due to the arbitrary nature of manual scanning), the noisy nature of tissue interfaces and the fact that the data may not be associated exclusively with the surface of the object of interest.

The chapter is organised as follows. Section 5.1 discusses related work and Section 5.2 introduces the principle of the reconstruction method. The constraints placed on a surface formed by this method are not obvious and Section 5.3 attempts to shed light on these by analysing the reconstruction of simple geometric objects. In Section 5.4 an algorithm is developed which is applicable to binary, 2D pixel data. In Section 5.5 the algorithm is extended to continuous-valued data and in Section 5.6 an iterative version which overcomes the major difficulties of the initial algorithm is introduced. Techniques for handling the problem of reconstructing cavities within a shape are discussed in Section 5.7. The reconstruction method is then applied to 3D voxel data in Section 5.8 and results are presented for a 3D ultrasound phantom.

5.1 RELATED WORK

The problem of constructing surfaces from ultrasound data may be viewed as an example of the general problem of fitting surfaces to scattered, unorganised data. The data is unorganised in the sense that the adjacency relation between surface points is not known. Two points may be in close proximity in 3-space but not necessarily adjacent on a surface passing through them. In contrast, where surface data is obtained from a laser scanner, the relationship between surface points is known by virtue of the methodical way in which an object is scanned. Consequently, many techniques developed for recovering surfaces from scattered surface data applicable to laser surface scanners, stereo vision and ultrasonic ranging rely on the adjacency relationship being known. Hoppe *et al.* [1992] are among the few who have considered the difficult problem of unorganised range data. Their algorithm reconstructs surfaces of arbitrary topology without additional information. However, the input data are assumed to be exclusively surface data and associated with a single surface. A uniform sampling density is also assumed in the data. These requirements are rarely satisfied by typical ultrasound data which is noisy and detects multiple tissue interfaces.

Snyder *et al.* [1992] have developed a method for closing gaps in two and three dimensional edge data, a problem with similar characteristics. The method assumes binary data in the form of a 3D voxel array where voxels are identified as either being edges or not edges. It involves forming a *chamfer map* which records the minimum distance from each voxel in the image to an edge voxel. Voxels which exceed a specified distance threshold are assigned a maximum value. The image is segmented by grouping these voxels into contiguous regions. The resulting regions are disjoint and the intervening voxels are classified as *questionable*. Questionable voxels are then iteratively assigned to neighbouring regions according to a set of rules. The method has some desirable qualities: it can be implemented in parallel and it requires a minimum amount of *a priori* knowledge to be specified. The selection of the distance threshold, however, determines the maximum gap size which can be interpolated across and the heuristics for the assignment of questionable voxels strongly influence the final segmentation. In the presence of noise (voxels falsely identified as edge data), many regions are likely to be found which do not correspond to physical objects and a large set of questionable voxels will arise which will not be easily resolved. Snyder's method is applicable to binary data, where the presence or absence of an edge is specified, but does not generalise to continuous-valued data in an obvious manner other than increasing the dimensionality of the problem and defining a 4D distance measure. In the analysis of ultrasound data it is advantageous to associate a confidence with the presence of an interface boundary. This results in a continuous-valued distribution.

An alternative approach to this type of problem involves fitting deformable models such as superquadrics [Terzopoulos and Metaxas 1991] to the data, spherical meshes [Hebert *et al.* 1995] or inflatable balloons Cohen and Cohen [1990]. These methods require prior knowledge of the location, shape and size of objects to be reconstructed. Performance depends on the specification of parameters which control the surface and determine the relative influence of data points. Determining these parameters from actual data to achieve optimum fit is difficult, especially in the presence of noise [Raja and Jain 1992]. Model-fitting procedures constrain the topology of objects which can be reconstructed. Superquadric surfaces, for example, are parameterised in terms of latitude and longitude, implying a spherical topology. In this thesis a general approach is sought that will impose few restrictions on reconstructible shapes.

In this chapter the surface of an isolated object is reconstructed indirectly by identifying the likely interior of the object. Computed tomography (CT) style projections are taken about the object data, as if a source of 'radiation' were being moved around the data. The resulting cumulative distribution of radiation seen at each point in space is interpreted as indicating how likely a point is to be inside the object. A surface is then defined by applying a threshold to the distribution. In Section 5.6 an iterative scheme is developed to overcome the limited ability of the basic method to reconstruct large gaps in the surface of an object. In this scheme a new set of object data is estimated from the 'radiation' distribution and employed as the object data for a following iteration; as this process is repeated the object data grow to form continuous regions. User interaction resolves ambiguities in how a surface may be fitted to the data by varying the threshold applied at each iteration.

The method is unique in the way projections are used to reconstruct 3D objects. A related problem in model-based vision is the matching of 2D projections to rigid 3D objects in order to determine the position and orientation of an object from sensor data [Lavallee and Szeliski 1995]. The projections yield 2D contours for various viewpoints. These are assumed to be complete and uncontaminated by noise except for the possibility of a small fraction of the object being occluded by another structure, thus affecting the associated contour. The problem is of limited interest since the aim is not to reconstruct the shape of unknown 3D objects, but rather to determine whether known shapes could have produced a given set of projections.

Recent work by Laurentini [Laurentini 1993, Laurentini 1995] has addressed the problem of inferring 3D shape from 2D silhouettes. The silhouette can be thought of as a binary-valued projection. A set of silhouettes specifies a bounding volume obtained by back-projecting each silhouette and finding the associated intersection volume. This volume approximates the object to greater or lesser extent depending on the silhouette viewpoints and the object itself. Such methods are also known as *volume intersection* methods. Yuan [1995] describes a method for automatically choosing new viewing directions in order to efficiently constrain the reconstructed volume. However, the algorithm is known to fail and require operator intervention in certain circumstances. Laurentini considers the question as to how much of the shape of the unknown object can be inferred from the volume formed by the intersection of back-projected silhouettes. He introduces the concept of *hard* points which belong to any possible object originating the set of silhouettes and *soft* points which may or may not belong to the object. Laurentini calls the closest approximation to the object which can be obtained by a set of silhouettes the *visual hull*. The object $\mathbf{O}$, its visual hull $\mathbf{VH(O)}$ and its convex hull $\mathbf{CV(O)}$ are related as follows, $\mathbf{O} \leq \mathbf{VH}(\mathbf{O}) \leq \mathbf{CV}(\mathbf{O})$. Simple geometric objects, consisting of opaque planar polygons are considered by Laurentini and algorithms for computing the visual hull of polygonal sets, polyhedra and solids of revolution are given. While this work is of theoretical importance, it does not tackle the practical problem of reconstructing real objects from continuous-valued projections obtained from noisy and incomplete data.

Murch [1990] has investigated the reconstruction of 2D binary images from a few (1D) projections. The work focuses on convex objects and an algorithm for 2D pixel data is described. Extensions to non-convex objects are suggested. However, noise in the object data, from which projections are calculated, is not considered. Murch extends the work to 3D objects, but there the primary aim is to find the orientation of an object from 1D projections taken at equally spaced angles through 3 or 4 planes which intersect the object.

In this chapter projections of incomplete object data are used not only to reconstruct the given object data but also to interpolate across missing data. Both binary and continuous-valued object data are considered. Algorithms for 2D pixel and 3D voxel data are developed and applied to ultrasound data. Extensions to the method allow the reconstruction of cavities. Both inexact projections and the presence of noise in the object data are considered.

5.2 THE SHAPE RECONSTRUCTION PROBLEM

Consider a single object isolated from its surrounds and sampled in an irregular manner to give a scattered distribution of points deemed to belong to the object. In three dimensions the aim is to reconstruct the surface of the object which encloses the data. The approach developed here is also applicable to the analogous two dimensional problem of reconstructing a closed curve from partial edge data in a plane.

Before continuing further it is necessary to state the shape reconstruction problem formally. Consider a vector space $\mathbf{V}$, in which $\mathbf{x}$ defines a point. Let Ω be a subspace of $\mathbf{V}$ bounded by a simple closed surface S. Ω is termed the reconstruction space and must contain the object to be reconstructed. $O(\mathbf{x})$ defines an attribute of the subspace: in this case a binary value is initially considered which defines whether or not a point belongs to the object of interest. Ω_{obj} denotes the region of the space occupied by the object of interest. $\hat{\Omega}_{obj}$ denotes a subset of Ω_{obj}, corresponding to the incomplete object data, $\hat{\Omega}_{obj} \subset \Omega_{obj} \subset \Omega \subset \mathbf{V}$. This is the ideal case. In practice $\hat{\Omega}_{obj}$ may contain noise; hence $\hat{\Omega}_{obj} \not\subset \Omega_{obj}$. The problem is to reconstruct Ω_{obj} given the partial object data, $\hat{\Omega}_{obj}$.

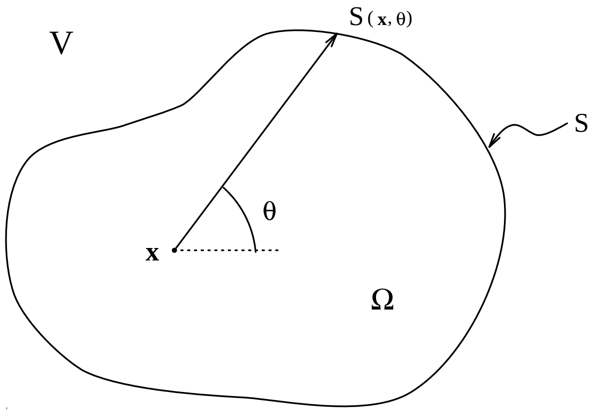

Figure 5.1 Definition of the reconstruction space, Ω, in two dimensions.

Because the method proposed for shape interpolation generalises straightforwardly from two to three dimensions, the two dimensional case is now described and analysed in detail.

The method involves determining an *occluded angle*, denoted by Θ_{occ}, at each point in Ω. A ray emanating from a point $\mathbf{x}$, where $\mathbf{x} \in \Omega$, is occluded if it intersects object data prior to reaching the reconstruction space boundary, S. In two dimensions (Figure 5.1) a ray emanating from $\mathbf{x}$ can be defined by the point $\mathbf{x}$ and the ray's orientation, θ, where $\theta \in [0, 2\pi]$, relative to an arbitrary axis. $S(\mathbf{x}, \theta)$ is the point where the ray *first* intersects the reconstruction space boundary. Note that since the definition of S is arbitrary, a ray may intersect S more than once. The occluded angle at $\mathbf{x}$, $\Theta_{occ}(\mathbf{x})$, may be found by integration over all rays emanating from $\mathbf{x}$ and is expressed in equation (5.1).

$$\Theta_{occ}(\mathbf{x}) = \int_0^{2\pi} F(\mathbf{x}, \theta)d\theta \tag{5.1}$$

where

$$F(\mathbf{x}, \theta) = \begin{cases} 1 & \text{if } S \text{ is occluded along the ray } \theta \text{ emanating from } \mathbf{x}, \\ 0 & \text{otherwise.} \end{cases}$$

From equation (5.1) it can be seen that Θ_{occ} may range from 0 to 2π, where 0 corresponds to no occlusion and 2π implies that $\mathbf{x}$ is totally enclosed by object data.

The occluded angle distribution is formed by calculating $\Theta_{occ}(\mathbf{x})$ at each point in Ω. Points inside the object may be expected to have a higher occluded angle than those outside the object. The occluded angle at a point is therefore regarded as a measure of the likelihood that it is contained within the object. Points in the occluded angle distribution above a particular threshold can be viewed as candidate solution regions. The choice of an occluded angle threshold determines the shape of the contours bounding the solution regions.

Object shapes which can be reconstructed are necessarily restricted to those where every point along the perimeter of the shape can see some part of the reconstruction space boundary, S. In other words, no point on the perimeter is occluded from all points on S. Note that this limitation on shapes which can be reconstructed is dependent on S. Figure 5.2(a) illustrates a shape for which the above condition is not met. The region of the curve which can not be seen along S occurs within what appears as a cavity in the shape, with respect to S. For the same shape this is no longer the case in Figure 5.2(b) due a different choice of S.

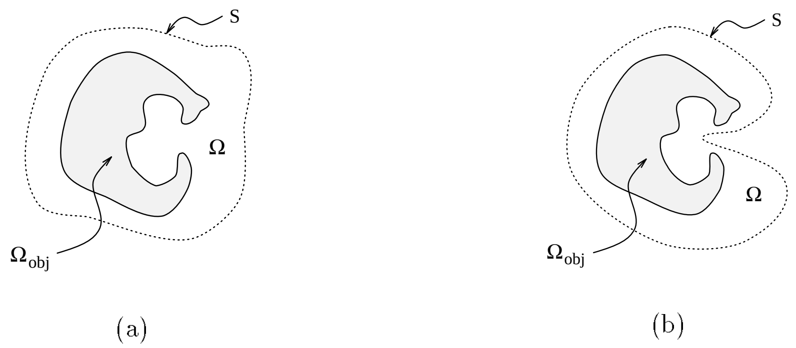

Figure 5.2 (a) Shape with a cavity which can not be reconstructed, (b) alternative choice for S allowing the cavity to be reconstructed.

5.3 ANALYSIS OF OCCLUDED ANGLE DISTRIBUTION

The occluded angle distribution, $\Theta_{occ}(\mathbf{x})$, depends on S and $\hat{\Omega}_{obj}$. In general, $\Theta_{occ}(\mathbf{x})$, evaluated at some point $\mathbf{x}$, $\mathbf{x} \in \Omega$, is a function of all the object data and not just that within the local neighbourhood of $\mathbf{x}$. This, and the observation that the $\Theta_{occ}(\mathbf{x})$ distribution may contain discontinuities, and is hence only piecewise differentiable, makes an analysis difficult. However, some important general characteristics can be inferred by considering a few specific objects which might otherwise be considered of little practical significance. A circular reconstruction space of infinite radius centred on the object data is chosen in the following examples so that the reconstruction dependency on S may be ignored.

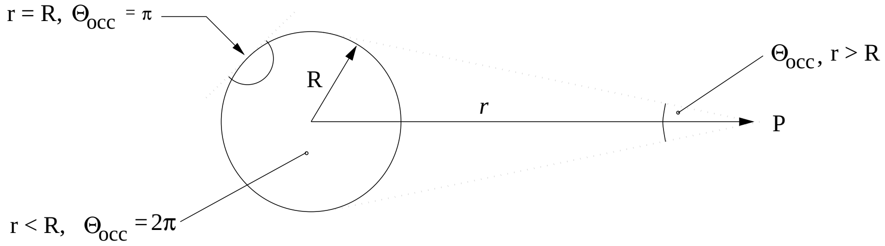

Figure 5.3 Analysis of Θ_{occ} distribution for a closed circle at a point P.

5.3.1 Analysis of a closed circle

Consider the case where $\hat{\Omega}_{obj}$ is a closed circle with radius R (Figure 5.3). Inside the circle, Θ_{occ} is 2π everywhere since the boundary of the reconstruction space, S, is occluded by the circle in all directions,

$$\Theta_{occ}(r) = 2\pi \qquad 0 \leq r < R \tag{5.2}$$

For the same reason, Θ_{occ} is 2π inside any closed curve. At the boundary of the circle a discontinuity occurs and $\Theta_{occ} = \pi$. Along the perimeter of any convex curve, Θ_{occ} will be π, except when a discontinuity in curvature (i.e. a sharp corner) occurs, in which case Θ_{occ} will be less than or equal to π. Taking advantage of the angular symmetry of a circle, Θ_{occ} at a point P outside the circle (Figure 5.3) can be described as a function of radial distance from the circle centre, independent of the angular coordinate:

$$\Theta_{occ}(r) = 2\sin^{-1}(R/r) \qquad r \geq R \tag{5.3}$$

A cross-sectional profile, computed from equations (5.2) and (5.3), of Θ_{occ} along a diameter through the circle is illustrated in Figure 5.4.

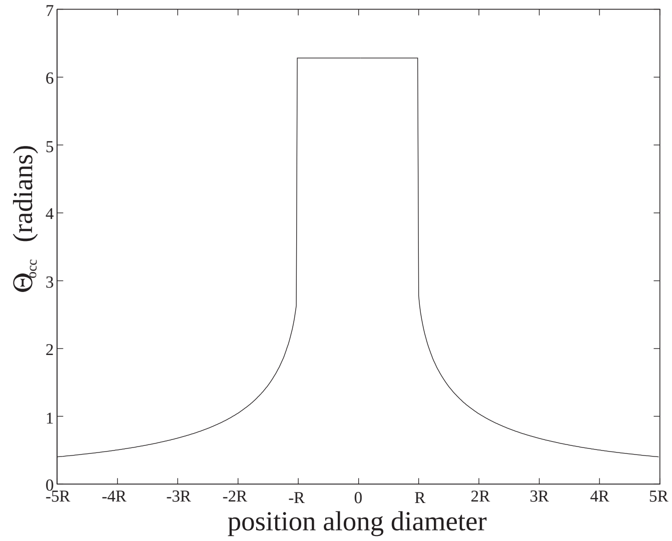

Figure 5.4 Cross-sectional profile of Θ_{occ} through a closed circle.

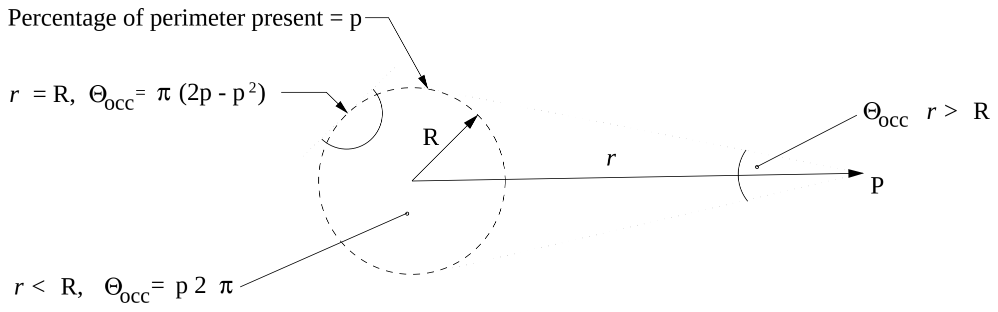

Figure 5.5 Analysis of Θ_{occ} distribution for a partial circle.

5.3.2 Analysis of a partial circle

If the circle now contains gaps such that only a fraction, p, of the curve is present (Figure 5.5), and the gaps are infinitely narrow and evenly distributed about the circumference, then in the limit, Θ_{occ} "inside" the circle will form a uniform distribution with the value $p2\pi$. In this context "inside" means within the convex hull of the data points. The fraction of holes present in the circle is $1-p$, therefore outside the circle, $r \geq R$, Θ_{occ} is similar to that in the closed case except that it is reduced by the coefficient $1-(1-p)(1-p) = (2p-p^2)$.

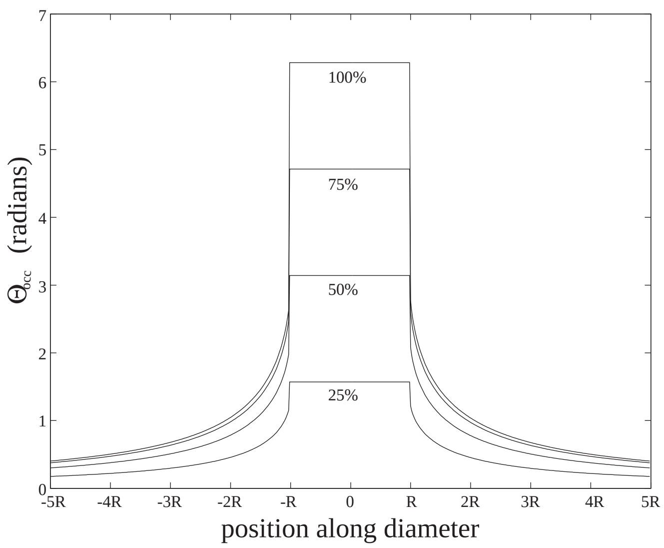

Figure 5.6 Cross-sectional profiles of Θ_{occ} for a circle with 100%, 75%, 50% and 25% of the perimeter present.

$$\Theta_{occ}(r) = 2(2p - p^2)\sin^{-1}(R/r) \qquad r \geq R \tag{5.4}$$

Figure 5.6 illustrates the Θ_{occ} profile along an axis passing through the circle centre for varying percentages of perimeter present, namely 100%, 75%, 50% and 25%. Apparent is the rapid reduction in contrast between Θ_{occ} inside and outside the circle as the percentage of perimeter present decreases. This means that when most of the circle is present, a wide variation in the choice of the Θ_{occ} threshold is tolerated for the same reconstruction accuracy. The range, Th_{occ}, over which a threshold applied to the Θ_{occ} distribution may be varied for a reconstruction to be within a specified tolerance, δ, of the actual circle radius is,

$$Th_{occ}(p, \delta) = [2(2p - p^2)\sin^{-1}(1/1+\delta), 2\pi p]. \qquad 0 < p \leq 1 \tag{5.5}$$

where δ is normalised with respect to the circle radius (Figure 5.7). The range over which the Θ_{occ} threshold can be varied for a small reconstruction tolerance, $\delta \ll 1$, decreases in an approximately squared relationship with decreasing p. We might therefore expect that halving the number of points on a convex curve will lead to an approximate quartering of the range over which the Θ_{occ} threshold may vary for a specified reconstruction tolerance.

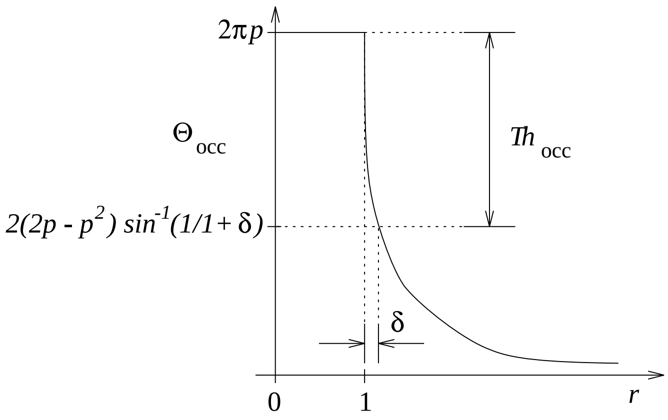

Figure 5.7 The range through which Θ_{occ} may vary to reconstruct a partial circle within the tolerance, δ, of the actual circle radius.

5.3.3 Analysis of a gap in an infinite circle

Now consider the case of a single gap in an infinite circle. In the region of the gap the circle perimeter appears as a straight line (Figure 5.8). Through trigonometric analysis of Figure 5.8, Θ_{occ} at a point **P** inside the circle is determined by equation (5.6).

$$\Theta_{occ}(r_p, \theta_p) = 2\pi - \cos^{-1} \frac{r_p^2 - {\delta_{gap}}^2}{\sqrt[+]{r_p^4 - 2r_p^2\delta_{gap}^2 \cos 2\theta_p + \delta_{gap}^4}} \qquad 0 < \theta_p < \pi \tag{5.6}$$

where **P** is specified in polar coordinates relative to the centre of the gap, (r_p, θ_p). Along the circle perimeter Θ_{occ} is π. It can be shown from equation (5.6) that inside the circle constant values of Θ_{occ} define a family of circles which pass through the gap end points, i.e. the gap defines a chord on the iso-value circles. Along the straight line between the gap end points Θ_{occ} is also π, independent of the gap width (this iso-surface corresponds to a circle of infinite radius). Outside the circle Θ_{occ} is π everywhere since the circle is infinite. As the width of the gap, $2\delta_{gap}$, tends to zero, equation (5.6) (Θ_{occ} inside the circle) tends to 2π as expected. Similarly as r_p, the radial distance from the gap centre becomes large, Θ_{occ} inside the circle also tends to 2π.

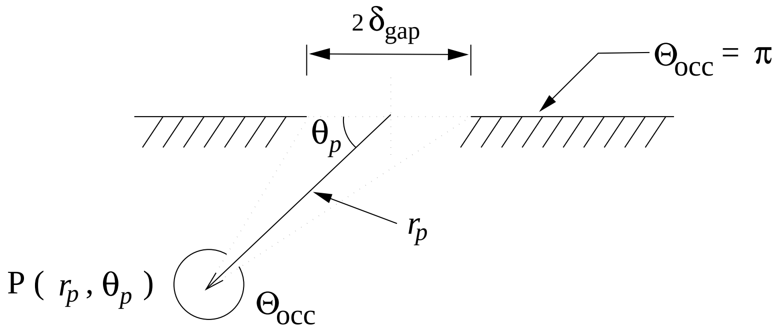

Figure 5.8 Analysis of Θ_{occ} distribution for a gap in an infinite circle.

5.3.4 Analysis of a gap in a finite circle

The case of a gap in a finite circle is now considered. The gap is defined by the angle which it subtends at the circle centre, ψ_{gap} (ref. Figure 5.9). Θ_{occ} inside the circle is determined by equation (5.6). Equation (5.6) applies inside any completely convex curve containing a single gap. In the case of an arbitrary curve containing a gap, Θ_{occ} "inside" the curve will be greater than or equal to the value determined from equation (5.6). Thus equation (5.6) places a lower limit on the occluded angle at a point inside a closed, but otherwise arbitrary, curve containing a gap. For a convex curve containing a single gap, a threshold corresponding to π applied to the occluded angle distribution will therefore interpolate across the gap with a straight line while higher threshold values will lie on circles passing through the gap end points, appearing as 'bites' in the shape of the reconstructed curve.

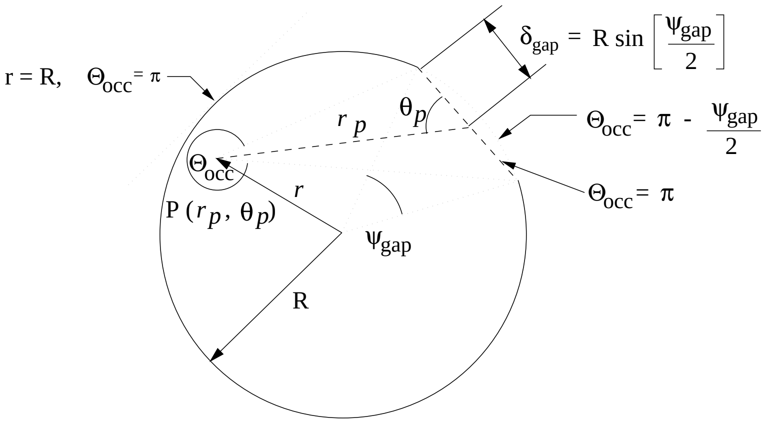

Figure 5.9 Analysis of the Θ_{occ} distribution for a finite circle with a gap.

Along the perimeter of the circle in Figure 5.9, Θ_{occ} is π everywhere except at the end points either side of the gap where a discontinuity occurs and Θ_{occ} is given by

$$\Theta_{occ} = \pi - \frac{\psi_{gap}}{2} \tag{5.7}$$

For the case of a general closed curve, Θ_{occ} will be greater than or equal to π along any smooth segment of the perimeter. This excludes points on the perimeter where a discontinuity in curvature occurs, i.e. at sharp corners. Along the straight line connecting the gap end points Θ_{occ} is π, independent of the size of the gap, ψ_{gap}.

In the special case of a circle containing a gap, Θ_{occ} remains constant with the value $\pi - \frac{\psi_{gap}}{2}$ along the arc defined by the gap end points with $r = R$ (indicated in Figure 5.9). A threshold choice of $\pi - \frac{\psi_{gap}}{2}$ will therefore define a contour which reconstructs the missing section of the circle across the gap. The maximum variation in the reconstructed boundary from the circle for this threshold choice occurs in the region where the gap has no effect, i.e. where the gap is occluded by the remainder of the circle perimeter. Here Θ_{occ} is determined by equation (5.3) which was derived for the case of a closed circle. By rearranging equation (5.3) this maximum variation is,

$$\delta = 1/\sin(\pi/2 - \frac{\psi_{gap}}{4}) - 1 \tag{5.8}$$

where δ is normalised as a fraction of the actual circle radius, R. A gap which subtends $30°$ can therefore be reconstructed with a 0.8% maximum variation in radius from that of the original circle. This rapidly degrades with increasing gap size such that for a $60°$ gap there is a 3.5% maximum variation and at $180°$ (a half circle) there is a 41.4% variation.

A profile of Θ_{occ} along the axis of symmetry through the circle illustrated in Figure 5.10 can be formulated. The analysis has been simplified by restricting ψ_{gap} to the range $[0, \pi]$. There are several cases to be considered. In the regions outside the circle, $r \leq -R$ and $r \geq d2$, the gap has no influence and Θ_{occ} is determined as for the case of a closed circle.

$$\Theta_{occ}(r) = 2\sin^{-1}|R/r| \qquad r \leq -R, r \geq d2 \tag{5.9}$$

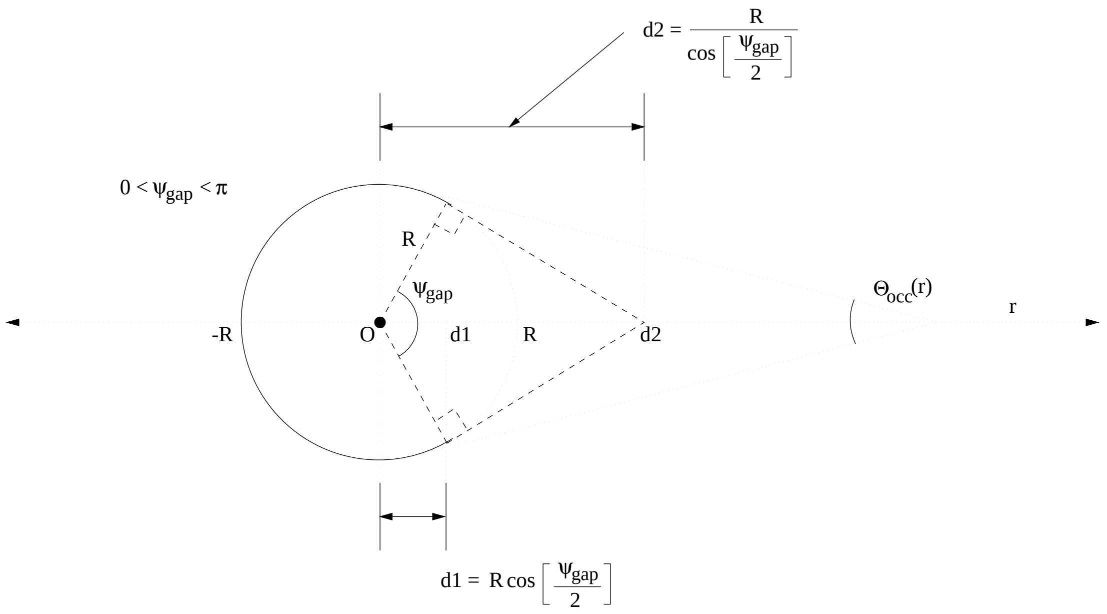

Figure 5.10 Analysis of Θ_{occ} along the axis of symmetry of a circle with a gap, $\psi_{gap} \in [0, \pi]$.

Inside the circle Θ_{occ} is given by equation (5.6) which can been rewritten as a function of r along the axis shown in Figure 5.10 to give equation (5.10).

$$\Theta_{occ}(r) = \pi + 2\tan^{-1}\frac{d1 - r}{R\sin(\frac{\psi_{gap}}{2})} \qquad -R < r < d2 \tag{5.10}$$

where

$$d1 = R\cos(\frac{\psi_{gap}}{2}) \tag{5.11}$$

Equation (5.10) is valid provided arctan is restricted to the range $[-\pi, \pi]$. $d2$ is the point where the tangents to the circle at the gap end points intersect. Figure 5.11(a) illustrates the occluded angle profile, $\Theta_{occ}(r)$, along the axis of symmetry for gaps which subtend angles of 0° , 45° , 90° , 135° and 180° .

Figure 5.11(b) is the corresponding plot of $\frac{d\Theta_{occ}(r)}{dr}$. At the discontinuity in Θ_{occ}, which occurs at the circle boundary, $\frac{d\Theta_{occ}(r)}{dr}$ is infinite. The local minimum corresponds to the inflexion point in the occluded angle profile which occurs at $r = R\sin\frac{\psi_{gap}}{2}$. This is the point which lies along the straight line connecting the gap end points, along which Θ_{occ} is π. The interpolated perimeter becomes more sensitive to thresholds in the vicinity of π, as the gap becomes larger.

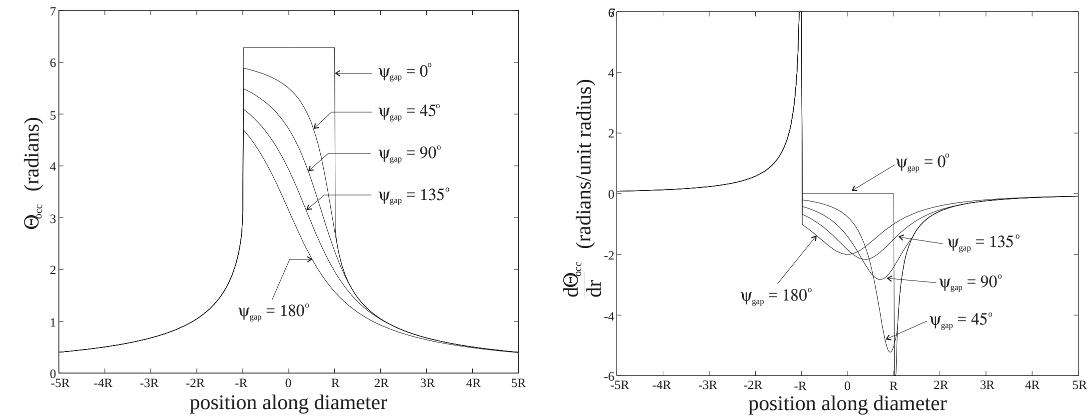

(a) (b)

Figure 5.11 Cross-sectional profiles of Θ_{occ} and $\frac{d\Theta_{occ}}{dr}$ for a circle containing a single gap. Gaps corresponding to $\psi_{gap} = 0°$, $45°$, $90°$, $135°$ and $180°$ are illustrated.

5.3.5 Multiple gaps in a finite circle

In the case of multiple discrete gaps in a circle, the effect of each gap inside the circle can be considered separately. Let Θ_{non_i} denote the non-occluded angle due to the i^{th} gap. Note that $\Theta_{non_i} = 2\pi - \Theta_{occ_i}$. At each point in the circle the non-occluded angle due to each gap is summed. The occluded angle, Θ_{occ}, is the difference between 2π and the net non-occluded angle due to each individual gap,

$$\Theta_{occ} = 2\pi - \sum_{i=1}^{n} \Theta_{non_i} \tag{5.12}$$

where $n =$ the number of distinct gaps. It is important to note that superposition is not generally applicable, and here it applies only inside the circle. This approach may be employed to determine the occluded angle distribution inside any completely convex curve, but not inside an arbitrary curve.

5.4 A DISCRETE ALGORITHM

A discrete 2D data space, represented as an array of regular pixels, is now considered. Following the notation introduced in Section 5.2, a subset of pixels forms the reconstruction space, Ω. $O(j)$ denotes the value of the j^{th} pixel in Ω, where j is an integer index. $O(j)$ is considered constant over the domain of each pixel and has a binary value,

$$O(j) = \begin{cases} 1 & \text{if } j \in \hat{\Omega}_{obj} \\ 0 & \text{otherwise} \end{cases} \tag{5.13}$$

Here $\hat{\Omega}_{obj}$ is the subset of pixels identified as part of the object to be reconstructed. We now want to calculate Θ_{occ} at the centre of each pixel within Ω in order to form the discrete occluded angle distribution, $\Theta_{occ}(j)$.

We return briefly to the continuous domain in order to formulate an approximation for Θ_{occ}. The term *projection* is introduced in the context of the reconstruction space to mean a line integral of the attribute $O(\mathbf{x})$ along a straight line path through Ω. Referring to Figure 5.1, $\mathcal{P}(\mathbf{x}, \theta)$ denotes the projection along the ray emanating from the point $\mathbf{x}$ with orientation θ,

$$\mathcal{P}(\mathbf{x}, \theta) = \int_{\mathbf{x}}^{S(\mathbf{x},\theta)} O(\mathbf{x})dx \tag{5.14}$$

where $S(\mathbf{x}, \theta)$ is the point on S where a ray initiated at $\mathbf{x}$ with orientation θ *first* leaves Ω. Straight line ray paths are assumed here and throughout. The *binary* projection sample, $\mathcal{P}_{bin}(\mathbf{x}, \theta)$, is defined as,

$$\mathcal{P}_{bin}(\mathbf{x}, \theta) = \begin{cases} 0 & \text{if } \mathcal{P}(\mathbf{x}, \theta) = 0 \\ 1 & \text{otherwise} \end{cases} \tag{5.15}$$

Equation (5.1), which defines Θ_{occ} at a point $\mathbf{x}$, may now be interpreted as the integral of binary projections along all rays emanating from $\mathbf{x}$,

$$\Theta_{occ}(\mathbf{x}) = \int_{0}^{2\pi} \mathcal{P}_{bin}(\mathbf{x}, \theta)d\theta \tag{5.16}$$

Equation (5.16) can be approximated by a finite sum of N projections taken at equally spaced orientations,

$$\Theta_{occ}(\mathbf{x}) \approx \frac{2\pi}{N} \sum_{i=1}^{N} \mathcal{P}_{bin}(\mathbf{x}, \theta_i) \tag{5.17}$$

where θ_i is the angle specifying the orientation of the i^{th} ray, $\theta_i = i\frac{2\pi}{N}$. This equation says that the proportion of rays along which object data are "seen" to the total number of rays cast approximates the ratio of the occluded angle to 2π. In general, the larger N is, the better the approximation will be.

In the discrete space, a binary projection may be efficiently calculated with a pixel-based ray tracing algorithm. Here the path of a ray is approximated by traversing Ω in integer pixel increments using Bresenham's line-drawing algorithm [Bresenham 1965]. As each pixel is traversed along the ray it is tested to see if it is an object data pixel; if so, then the projection has unity value and ray tracing may cease; otherwise ray tracing continues until an object data pixel is encountered or the object space is left, in which case the value of the projection is zero.

The direct implementation of equation (5.17) suggests evaluating N projections at every pixel in Ω. However, the contribution to θ_{occ} at each pixel for a particular projection angle may be calculated simultaneously by casting a set of parallel rays through Ω, where the density of the rays cast is sufficient to guarantee that every pixel in Ω is considered. An *accumulator* array is formed where each cell corresponds to a pixel in Ω. The contributions from each projection angle at each pixel in Ω are summed in the accumulator. Each ray is now traced right through Ω, rather than terminating when an object data pixel is encountered. As each pixel along the path of the ray is considered, the content of the corresponding cell in the accumulator is set to one if object data has been encountered along the path of the ray, and set to zero otherwise.

This process is best illustrated by an example. Figure 5.12(a) is a test data set of object pixels. The reconstruction space is the entire 150x150 pixel image. Figure 5.12(b) depicts the contents of the corresponding accumulator array after a single set of parallel rays has been cast through Ω. The resulting image depicts the value of the corresponding binary projection evaluated at each pixel. The dark pixels represent pixels occluded by object data, i.e. the value of the projection is 1 at these pixels. As further sets of rays are traced through Ω over a range of orientations, the resulting binary projections for each orientation are summed together in the accumulator array. The value in each accumulator cell indicates the number of projections occluded at the corresponding pixel in Ω. This is approximately proportional to Θ_{occ}. Figure 5.12(c) illustrates the summation of the projection in 5.12(b) with a second projection taken with a different orientation. The result is a tri-valued accumulator where pixels fall into three categories, those occluded by object data in both directions (the dark-valued pixels), those occluded by object data in only one of the two directions (mid-grey-valued pixels), and those not occluded by object data in either direction (white pixels). Figure 5.12(d) is the accumulator corresponding to Figure 5.12(a) after 25 projections have been taken at equally spaced angular intervals through Ω.

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

Figure 5.12 Estimation of the discrete Θ_{occ} distribution. (a) Object data, (b) binary projection due to a single set of parallel rays cast through Ω, (c) tri-valued accumulator resulting from the summation of two projections, (d) summation of 25 projections taken at regular angular intervals through Ω.

In this implementation, N projections are taken at equally spaced angular intervals of θ over the range $[0, 2\pi]$. Due to the finite number of projections, the occluded angle estimated at each pixel is discretised to integer increments of $\frac{2\pi}{N}$; this is the threshold resolution. Although this suggests that a larger N gives a better approximation to θ_{occ}, there is a limit on N due to the pixel-based ray tracing. Assuming square pixels, if R is the radius of the circle which circumscribes Ω in terms of pixel units, then the maximum possible resolution achievable in Θ_{occ} is approximately $\tan^{-1}(1/R)$ radians. This corresponds approximately to $N = 2\pi R$. The minimum number of projections required to reconstruct a given shape can be very small, but this number is shape dependent and prior knowledge of the shape and its orientation are required to take advantage of it.

No restrictions have yet been applied to the definition of the reconstruction space. Since Ω is a general subset of pixels, ray tracing is modified such that only object data pixels within Ω are considered. Evaluation of the projection integral commences when a ray enters Ω and halts when it leaves Ω. If a ray, when continued along its course, re-enters Ω further on then evaluation of the projection is recommenced afresh. This process is repeated until a ray leaves the domain of the image. Consequently, an arbitrary reconstruction space may be defined.

5.4.1 Contour extraction

At the conclusion of ray tracing, the contents of the accumulator represent an estimate of the occluded angle distribution. Shapes are formed by thresholding the occluded angle distribution. A contour-following algorithm tracks threshold transitions in the accumulator array. Thus contours surrounding regions in Ω where Θ_{occ} is greater than some specified threshold are extracted.

5.4.2 Artifacts

In practice, object data are contaminated by noise. An artifact pixel is an extraneous pixel in the discrete input data which is not actually part of the object but is falsely identified as such, i.e. $\hat{\Omega}_{obj} \not\subset \Omega_{obj}$. If square pixels are assumed and are approximated as being circular with a radius equal to half the pixel width, then the Θ_{occ} distribution due to a single pixel is the same as that for a closed circle. A single noise pixel produces a Θ_{occ} distribution determined by equation (5.3) and illustrated in Figure 5.3. The value of Θ_{occ} at the centre of an adjacent pixel will be approximately the value indicated at twice R in Figure 5.11, where R is the pixel radius. For the next outer ring of pixels (those pixels a distance of one pixel away) Θ_{occ} will be approximately the value indicated at $4R$, and so on. A pseudo-superposition argument applies in the sense that the value of Θ_{occ} at a point in Ω may be increased due to the presence of an artifact pixel by *at most* the amount determined by equation (5.3). Pixels which are occluded from the artifact pixel by intervening object data are not influenced by it. The effect of artifact regions, i.e. clusters of extraneous object pixels, can be estimated in a similar manner by considering the Θ_{occ} distribution due to a circle which circumscribes each artifact region.

Θ_{occ} within an artifact pixel will have the maximum value 2π since the pixel is modelled as a closed curve. This means that the pixel will be occluded in all projections taken through Ω, resulting in a Θ_{occ} value corresponding to 2π. Consequently, noise data, along with original object data, will always be identified when the accumulator is thresholded since these pixels will have the highest value possible in the accumulator, that is the number of projections taken. This is a consequence of taking binary projections.

5.4.3 Application to 2D data

Figure 5.13(a) is a 200x200 pixel test image of a circle containing a gap. The gap subtends an angle of 40° ($\frac{2\pi}{9}$ radians) at the circle centre. The reconstruction space, Ω, is defined over the domain of the image. Figure 5.13(b) is the result of summing 100 projections taken at equi-angular intervals about Ω. The grey level is proportional

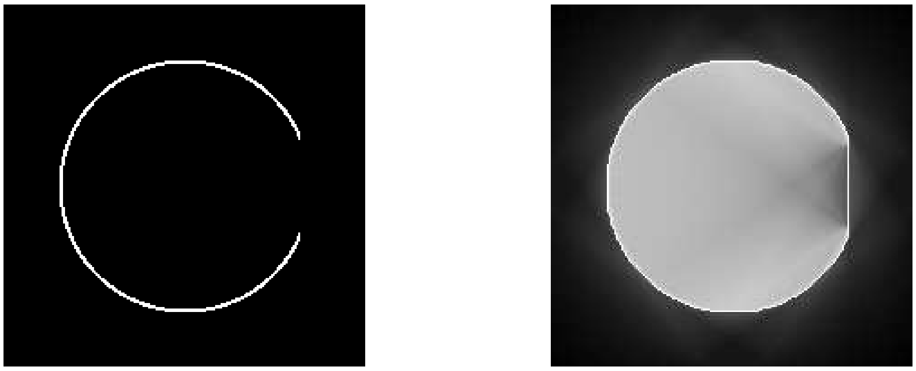

(a) (b)

Figure 5.13 Interpolating a circle with a gap, $\psi_{gap} = \frac{2\pi}{9}$. (a) 200x200 pixel reconstruction space with object data, (b) 100 projection accumulator with a threshold contour corresponding to $\Theta_{occ} \geq \pi$ superimposed.

to the occluded angle, Θ_{occ}. The brighter a pixel in the accumulator is, the higher the proportion of projections occluded at that pixel. A contour is superimposed on the accumulator enclosing the region which is occluded in at least 50 projections, i.e. pixels within the contour are occluded in more than half the projections taken through Ω. This corresponds to an occluded angle threshold of π. As anticipated in Section 5.3, the gap is interpolated by a straight line for this choice of threshold.

The effect of the number of projections used to estimate the Θ_{occ} distribution is illustrated by considering the incomplete circle analysed in Section 5.3. Figure 5.14(a) is an accumulator corresponding to 50 binary projections taken about a partial circle. The original test data are a 200x200 pixel image of a circle with approximately three quarters of the circumference pixels removed, i.e. p, the proportion of the circle present, is 25%. In Figure 5.14(b) the number of projections is trebled from 50 to 150; consequently the Θ_{occ} distribution inside the circle is smoother due to the improved resolution (2.4° *vs.* 7.2°) and accuracy with which Θ_{occ} is determined.

A limited number of projections leads to a criss-cross or mottled pattern in the accumulator. This is particularly apparent in Figure 5.14(a). The occluded angle distribution should be smooth and continuous throughout Ω except at object data, where discontinuities may arise. The criss-cross patterns observed in Figure 5.14 result from the finite number of rays cast and aliasing due to pixel-based ray tracing. These practical defects can be reduced by median filtering the accumulator prior to extracting contours. The median filter removes high frequency speckle regions in the accumulator, the kind produced by a limited number of projections, while preserving edges. Median filtering results in a significant improvement at a cheaper computational cost than taking more projections.

Figure 5.14(c) is the 150 projection accumulator thresholded according to $\Theta_{occ} \geq 0.5\pi$. This is the critical value predicted in Section 5.3 which should reconstruct the circle in the ideal case. Despite the relatively high number of projections taken, thresholding the Θ_{occ} distribution has resulted in holes. Figure 5.14(d) is the result of thresholding the median-filtered accumulator. Filtering with a 5x5 window has smoothed out local variations in Θ_{occ} while retaining the location of the circle boundary. Consequently, the threshold process has been able to reconstruct the interior of the circle uniformly.

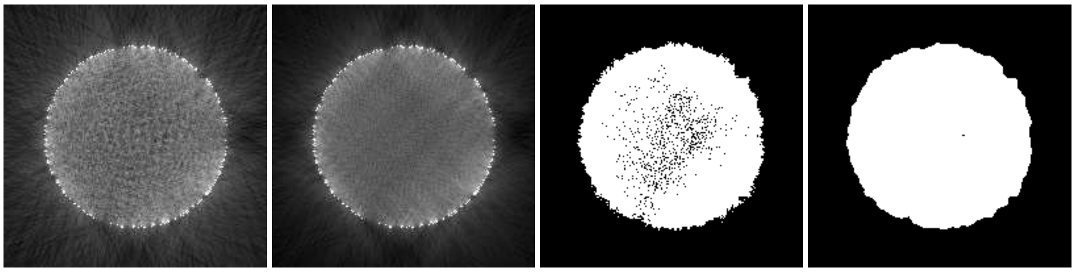

(a) (b) (c) (d)

Figure 5.14 Analysis of a partial circle with 25% of perimeter present. (a) 50 projection accumulator, (b) 150 projection accumulator, (c) 150 projection accumulator thresholded for $\Theta_{occ} \geq 0.5\pi$, (d) 150 projection accumulator filtered with a 5x5 median filter and thresholded for $\Theta_{occ} \geq 0.5\pi$.

5.4.4 Ultrasound data analysis

A reconstruction from 2D ultrasound data of a simple phantom is now considered. Figure 5.15(a) is a B-scan slice taken with a linear probe through a rubber ball phantom suspended in jelly. The strong reflection in the middle of the scan comes from a wire support used to hold the ball in place. The probe was positioned at the left of the image which is why the interface which the ball makes with the jelly is stronger on the left hand side of the scan. The image shown represents a section approximately 50x55mm extracted from a 58x120mm linear scan. This was done to meet the criterion established in section 5.2 that the data be associated with an isolated object. The resolution is approximately 0.25mm^2.

The raw scan in Figure 5.15(a) was thresholded to give Figure 5.15(b), the required binary representation for the algorithm. The choice of this threshold is a compromise between a conservative one which ensures that weak interfaces likely to result from speckle and reverberation are neglected, and a relaxed threshold which risks highlighting extraneous interface pixels.

Figure 5.15(c) is the corresponding filtered accumulator formed by taking 100 binary projections about 5.15(b); the contour $\Theta_{occ} \geq 0.86\pi$ is superimposed upon the

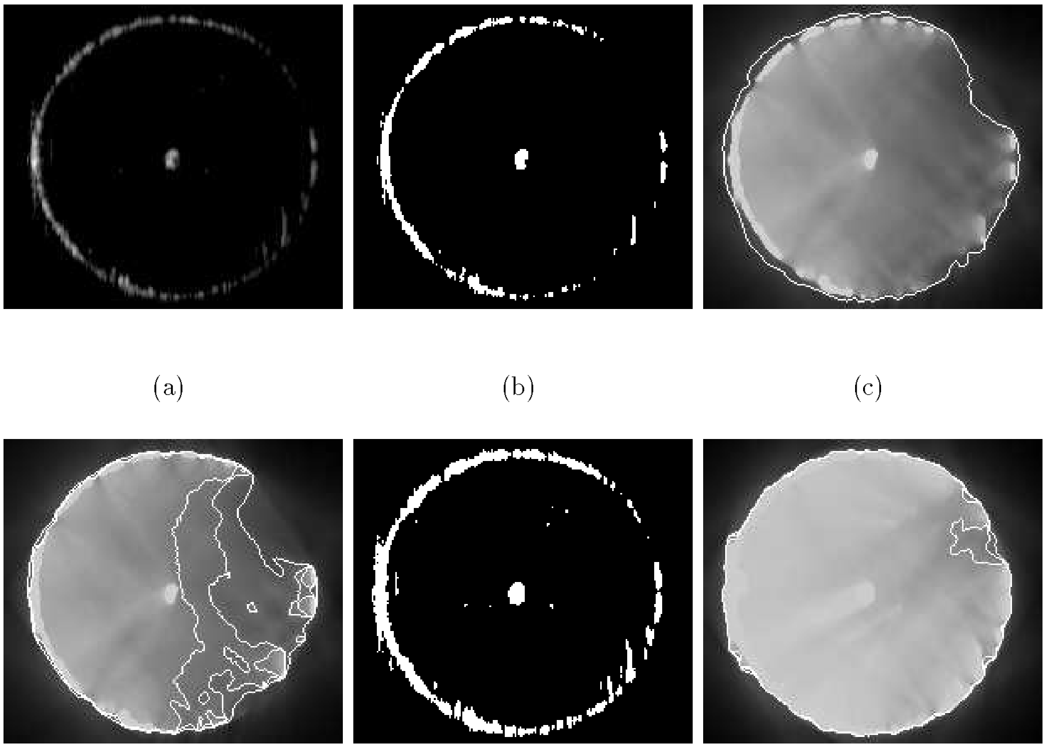

(a) (b) (c)

(d) (e) (f)

Figure 5.15 Reconstruction of an ultrasound phantom. (a) 2D ultrasound image, (b) thresholded ultrasound slice, (c) 100 projection accumulator with $\Theta_{occ} \geq 0.86\pi$ contour shown, (d) accumulator threshold contours for $\Theta_{occ} \geq \pi, 1.20\pi, 1.33\pi$, (e) alternative object data generated by a lower threshold choice, (f) accumulator contours constructed for (e), $\Theta_{occ} \geq 1.20\pi, 1.30\pi, 1.50\pi$

accumulator. This threshold choice resulted in the best reconstruction of the phantom. Figure 5.15(d) illustrates the reconstructions resulting from higher threshold choices corresponding to $\Theta_{occ} \geq \pi, 1.20\pi, 1.33\pi$. The object data in 5.15(a) are relatively sparse on the right hemisphere of the ball and consequently the higher thresholds illustrated in 5.15(d) tend to converge to a half circle reconstruction.

Figure 5.15(e) is the object data resulting from the application of a lower threshold to the raw data in 5.15(a). The 100 projection accumulator for this new set of object data is shown in Figure 5.15(f) with threshold contours superimposed for $\Theta_{occ} \geq 1.20\pi, 1.30\pi$ and 1.50π. The reconstructed curves are considerably improved and reflect the better interface data supplied.

Two shortcomings of the method are apparent; the original binary classification of the ultrasound data is critical to the outcome and the application of a single threshold to the occluded angle distribution is not always able to extrapolate across large gaps while closely adhering to interface data. These problems are addressed in the following two sections.

5.5 CONTINUOUS-VALUED INPUT DATA

A binary reconstruction space, where points are identified as either belonging or not belonging to the object under study, has been discussed. A reconstruction space where $O(\mathbf{x})$ is a continuous-valued distribution which reflects the likelihood that a point is part of the object to be reconstructed, is now considered. This avoids a binary classification of the input data and means that the object data may no longer be automatically retained by the thresholding process. Rays cast through Ω are gradually attenuated by object data rather than simply being occluded. The model for ray attenuation proposed is an exponential one based on monoenergetic X-ray attenuation. A ray undergoes exponential attenuation as it traverses the reconstruction space, where the linear attenuation coefficient is defined by $O(\mathbf{x})$. Referring to the two dimensional reconstruction space illustrated in Figure 5.1, $\Theta_{occ}(\mathbf{x}, \theta)$, the amount of attenuation experienced by a ray passing through Ω along the path defined by θ and $\mathbf{x}$, is given by,

$$\Theta_{occ}(\mathbf{x}, \theta) = 1 - e^{-\int_{\mathbf{x}}^{S(\mathbf{x},\theta)} O(\mathbf{u})du} = 1 - e^{-\mathcal{P}(\mathbf{x},\theta)}, \tag{5.18}$$

where $S(\mathbf{x}, \theta)$ is the point on S where a ray initiated at $\mathbf{x}$ with orientation θ *first* leaves Ω. As before, the total amount of attenuation at $\mathbf{x}$, $\Theta_{occ}(\mathbf{x})$, is found by integrating over all rays emanating from $\mathbf{x}$,

$$\Theta_{occ}(\mathbf{x}) = \int_0^{2\pi} (1 - e^{-\mathcal{P}(\mathbf{x},\theta)})d\theta = 2\pi - \int_0^{2\pi} e^{-\mathcal{P}(\mathbf{x},\theta)}d\theta. \tag{5.19}$$

Raw ultrasound data indicate the strength of pulses reflected by tissue interfaces within the body. Suitable attenuation coefficients therefore need to be assigned to the ultrasound numbers. This assignment should reflect the confidence in the existence of an actual tissue interface versus a spurious response. In the following examples a simple mapping of ultrasound number to linear attenuation coefficient has been employed. However, it is noted that this assignment need not be a function of ultrasound number alone, but could also include the likelihood of reverberation and shadowing artifacts existing within an ultrasound slice.

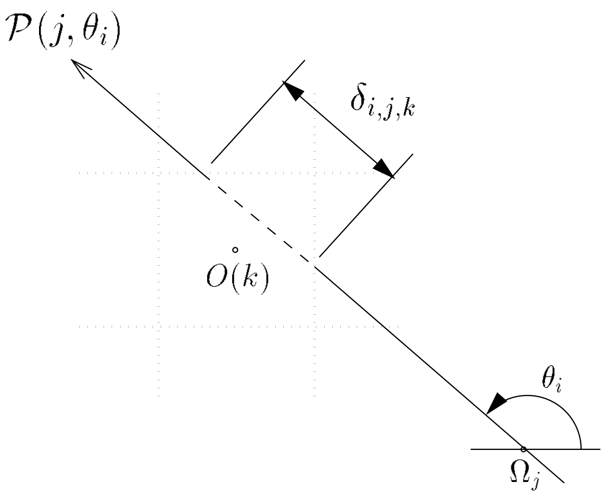

Figure 5.16 The contribution of the k^{th} pixel to the projection integral $\mathcal{P}(j,\theta_i)$.

5.5.1 Discrete implementation

Following the notation of Section 5.4, the discrete 2D attenuation distribution, $\Theta_{occ}(j)$, is calculated at the j^{th} pixel by approximating the integral in equation 5.19 with a summation over a finite number of projections,

$$\Theta_{occ}(j) \approx \frac{2\pi}{N}\sum_{i=1}^{N}(1 - e^{-\mathcal{P}(j,\theta_i)}) \tag{5.20}$$

where N is the number of projections taken and $\mathcal{P}(j,\theta_i)$ denotes the projection sample taken at the j^{th} pixel with the orientation θ_i. Because $\mathcal{P}(j,\theta_i)$ has a continuous value it is appropriate to calculate the contribution from each pixel according to the ray path length through the domain of the pixel. If $O(\mathbf{x})$ is considered constant over the domain of each pixel, then the contribution of the k^{th} pixel, $O(k)$, to the 2D projection sample $\mathcal{P}(j,\theta_i)$ is

$$\Delta_k\mathcal{P}(j,\theta_i) = \int_{\delta_{i,j,k}} O(k)d\delta = O(k)\delta_{i,j,k} \tag{5.21}$$

where $\delta_{i,j,k}$ is the path length through the k^{th} pixel (Figure 5.16). The direct implementation of equation 5.20 suggests taking N projections through Ω at each pixel. Although this guarantees that Θ_{occ} is evaluated at pixel centres, it is computationally intensive. A significant reduction in computation can be achieved if a single set of parallel projections is traced through Ω for each projection angle. As a ray is traced through Ω, the cumulative value of the line integral of the linear attenuation coefficient along the ray is recorded at each pixel boundary which the ray intersects. The amount of attenuation at the centre of the k^{th} pixel can be linearly interpolated from the projection rays either side of the pixel. In Figure 5.17 the attenuation at the k^{th} pixel is linearly interpolated from the adjacent projection samples either side of the pixel, denoted by $\mathcal{P}_\eta$ and $\mathcal{P}_{\eta+1}$. These are in turn calculated from the ray samples $\mathcal{P}_{\eta,k_1}$, $\mathcal{P}_{\eta,k_2}$ and $\mathcal{P}_{\eta+1,k_3}$, $\mathcal{P}_{\eta+1,k_4}$, respectively. The attenuation at a pixel due to each projection angle is then summed in an accumulator array. In this way the computationally intensive task of ray-tracing is greatly reduced.

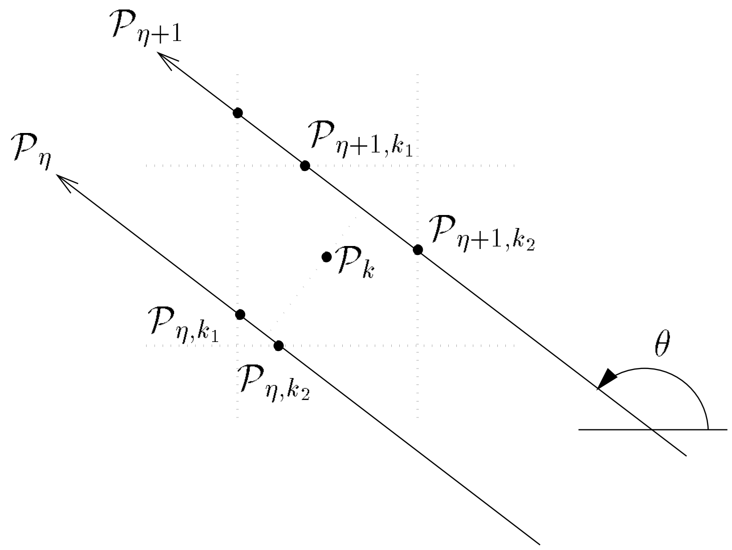

Figure 5.17 Bi-linear interpolation to estimate the attenuation coefficient at the k^{th} pixel for the projection angle θ.

In Figure 5.19 continuous projections are taken through the ultrasound data studied in Figure 5.15. The form of the function which maps an ultrasound number (the strength of a received reflection) to a linear attenuation coefficient is shown in Figure 5.18. Weak reflections are treated with less confidence as being object data compared to strong reflections. The original scan is repeated in Figure 5.19(a). Figure 5.19(b) is the accumulator resulting from 100 continuous projections taken through the ultrasound data. In Figure 5.19(c) the contour corresponding to the threshold $\Theta_{occ} \geq 1.2\pi$ is superimposed upon the accumulator. The spacing between parallel projection samples was the width of a pixel, the pixels in Figure 5.19 being square. Filtering is not required because the interpolation employed between projection samples results in a smooth Θ_{occ} distribution.

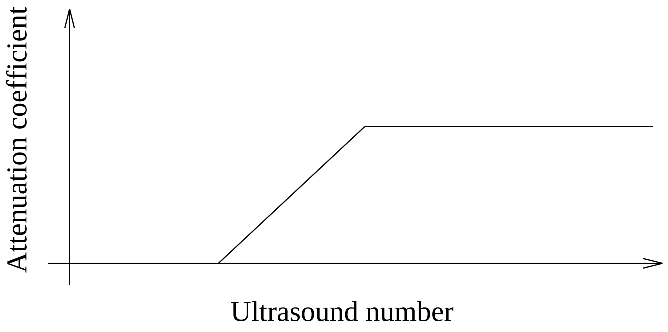

Figure 5.18 Form of the function employed to map ultrasound numbers to attenuation coefficients.

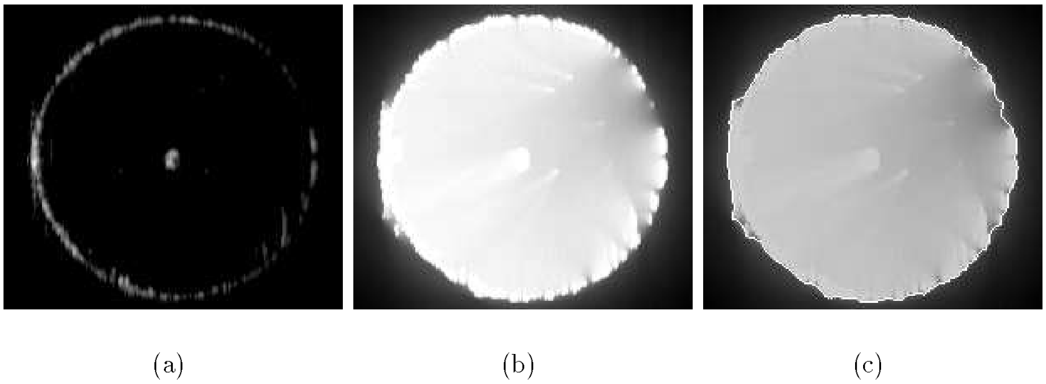

Figure 5.19 Reconstruction of an ultrasound phantom with continuous projections. (a) Raw ultrasound data, (b) 100 projection accumulator, (c) contour corresponding to $\Theta_{occ} \geq 1.2\pi$.

5.6 ITERATIVE ALGORITHM

The choice of an appropriate occluded angle threshold is critical to producing the final reconstructed shape. This choice depends on the object's shape and the object data. A high *conservative* threshold results in surfaces which closely adhere to object data. However, larger gaps are not bridged and separate regions tend to result, as seen in Figure 5.15(d). A lower, *relaxed* threshold is capable of bridging larger gaps in incomplete surface data. However, the resulting surface may no longer adhere closely to object boundary data and consequently the reconstructed surface will tend to be larger than the actual object, as demonstrated in Figure 5.15(c). When the threshold is relaxed further, the surface fitted tends to circumscribe the object data and become more circular. Here a conservative threshold is defined as one which is greater than or equal to π. This follows from the observation in Section 5.3 that along any smooth section of the perimeter of a convex object Θ_{occ} is π. A threshold less than π will define a contour larger than the object.

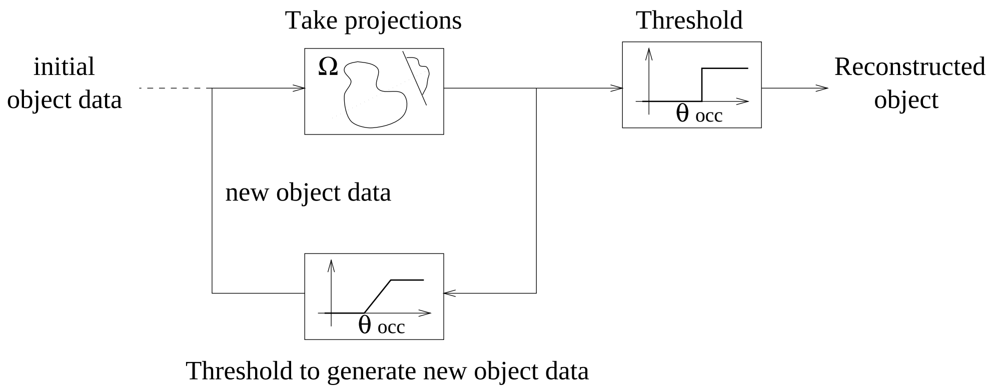

Figure 5.20 An iterative algorithm for object reconstruction.

An iterative algorithm (Figure 5.20) is proposed to allow interpolation across large gaps while adhering closely to likely object boundaries. The attenuation distribution is formed as before and a conservative threshold is chosen which partially extrapolates the data. The regions extracted by thresholding are then employed as the object data for a second iteration of the algorithm. This process is repeated with the result that the object data grow to form a continuous region, interpolating across gaps, yet retaining accurate object boundaries.

Figure 5.21 illustrates the iterative reconstruction of a circle from sparse edge data. The reconstruction space coincides with the domain of the image and contains a small noise region (Figure 5.21(a)). The initial accumulator, after taking 100 projections about the object data, is displayed in 5.21(b). It is not possible by thresholding the accumulator after a single iteration to form a reconstruction which extrapolates across the gaps in the edge data and at the same time adheres closely to the perceived circle. Instead a conservative threshold, $\Theta_{occ} \geq \pi$, is chosen to form the new object data (Figure 5.21(c)) for the next iteration. The growth in the initial data is small, but this quickly increases with further iterations. Figure 5.21(d) and Figure 5.21(e) are accumulators for the fourth and sixth iterations. A contour corresponding to the threshold $\Theta_{occ} \geq \pi$ is superimposed upon the sixth accumulator in 5.21(f); this contour interpolates the original data and represents a possible solution. In this example the conservative threshold was maintained at $\Theta_{occ} \geq \pi$ between iterations; new object data found upon thresholding were universally assigned a high attenuation coefficient such that a ray passing through a full pixel width would be attenuated by 95%, effectively treating the data as opaque.

Figure 5.22 illustrates the iterative reconstruction of an object from sparse data contaminated by noise. The initial object data, Figure 5.22(a), is binary and the reconstruction space coincides with the domain of the image. Four iterations consisting of 100 projections are illustrated in Figure 5.22(b-e). The thresholds applied at each iteration were $1.6\pi, 1.6\pi, 1.3\pi, 1.3\pi$, respectively. As in Figure 5.21, the new object

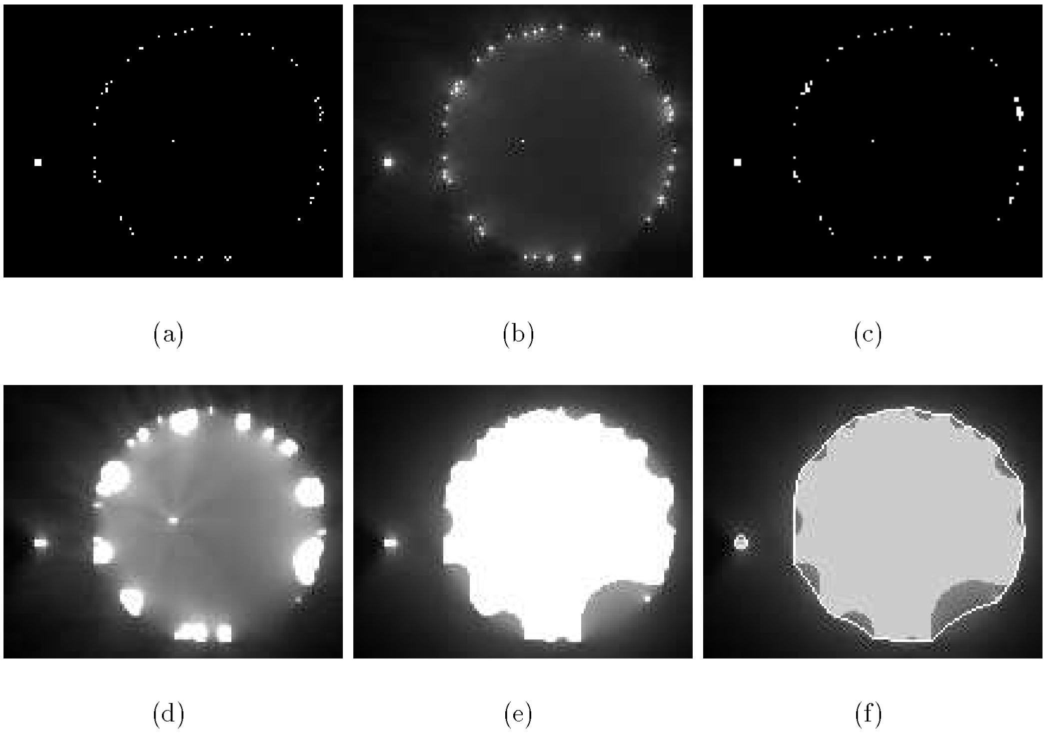

Figure 5.21 Iterative reconstruction of a circle. (a) Object data, (b) 100 projection accumulator, (c) object data for 2nd iteration obtained by thresholding (b), (d) accumulator for 4th iteration (e) accumulator for 6th iteration, (f) 6th iteration accumulator with contour corresponding to $\Theta_{occ} \geq \pi$ superimposed.

data derived by thresholding the previous accumulator were assigned high attenuation coefficients. The binary image depicted in 5.22(f) is the result of thresholding Figure 5.22(e) for $\Theta_{occ} \geq 1.3\pi$.

In these examples a simple threshold, chosen interactively, has been applied at each iteration. However, the continuous-valued data produced by each iteration need not be thresholded and the data above the threshold treated as opaque, but rather attenuation coefficients for the next iteration could be assigned as a function of the current accumulator value. Thresholding is used here because it reduces the number of parameters to be considered and makes it possible to describe the behaviour of the iterative method.

When applying a threshold between iterations, the general requirement is to start with a high threshold and decrease it over subsequent iterations. Ideally, the initial threshold should not be so high that it eliminates genuine object data, or that it fails to augment the initial object data. High initial thresholds are chosen primarily to ensure that only the points most likely to be part of the object are added to the newly derived set of data. If false object points are identified (points outside the actual object) then these may grow in subsequent iterations. In the case of a single iteration, the reconstructed object was determined by the threshold applied to the initial accumulator

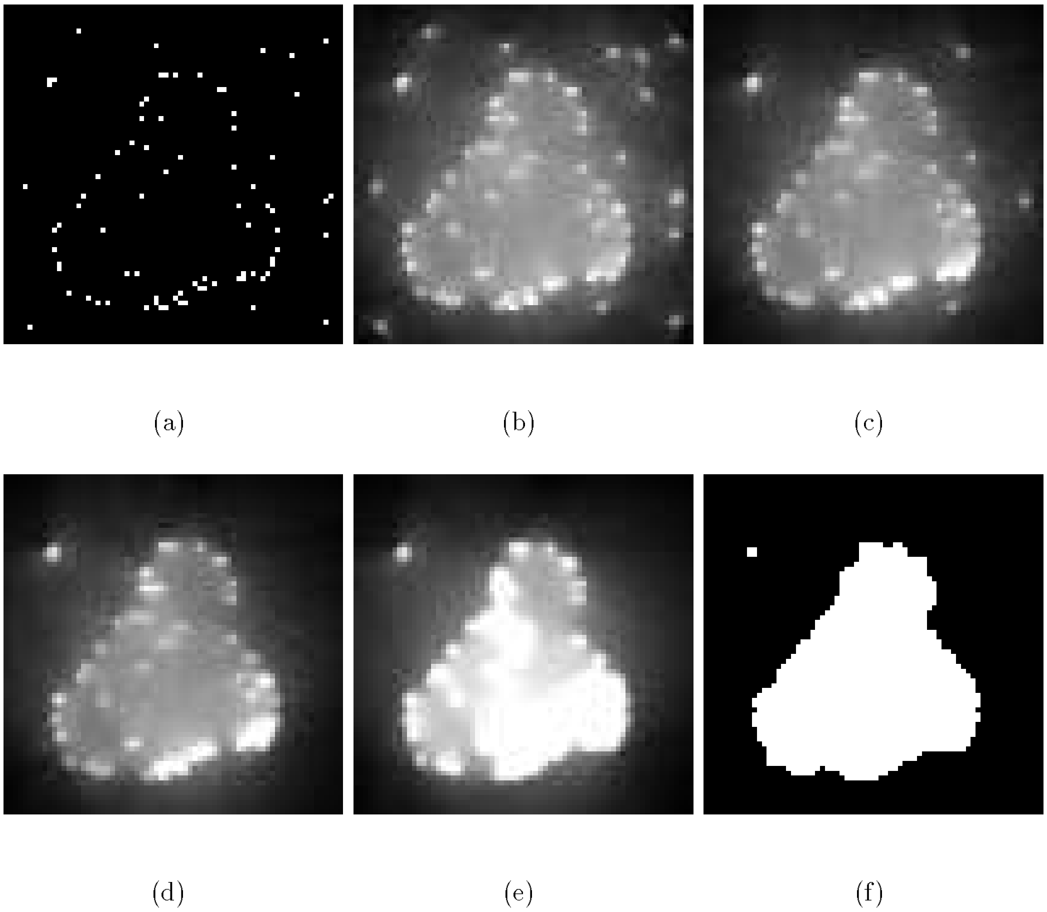

Figure 5.22 Iterative reconstruction of an object in the presence of noise. (a) Object data (b) 100 projection accumulator, (c) accumulator for 2nd iteration, (d) accumulator for 3rd iteration, (e) accumulator for 4th iteration, (f) reconstructed shape obtained by thresholding (e) $\Theta_{occ} \geq 1.3\pi$.

and was very sensitive to this choice. The threshold choice in the iterative scheme is not particularly critical as long as it is initially high, $> \pi$.

The properties of the shape which the iterative scheme converges on, for a fixed or decreasing threshold, are determined by the final threshold. For example, in the case of opaque object data, if the final threshold applied is $\Theta_{occ} \geq \pi$, then the shape constructed will be completely convex. The shape constructed in Figure 5.22(f) is not completely convex; here the final threshold choice of $\Theta_{occ} \geq 1.3\pi$ has preserved the concavities. A lower choice would begin to fill these in. Thus, the final threshold determines the depth of concavities allowed in a shape's perimeter.

5.7 RECONSTRUCTION OF CAVITIES

The limitation which a global threshold places on shapes containing cavities can be overcome by taking advantage of the dependency of the occluded angle on the reconstruction space. Figure 5.23(a) illustrates the problem of reconstructing a shape with a significant cavity. If the domain of the image is employed as the reconstruction space, then the minimum threshold angle required to resolve the cavity is such that the object data will remain as a cluster of discrete regions. Gaps will not be bridged, but treated as cavities. The application of a global threshold cannot distinguish between a gap and an intended cavity of similar size in the perimeter of an object. Figure 5.23(b) is the result of taking 100 projections where the reconstruction space coincides with the domain of the image. The bright pixels inside the cavity indicate a high occluded angle and hence these are seen as having a high probability of being part of the object which is contrary to the desire to preserve the cavity. In Figure 5.23(c) an alternative, 'doughnut' reconstruction space is defined. Figure 5.23(d) is the result of taking 100 projections through this new reconstruction space; it can be contrasted with 5.23(b). Figure 5.23(e) is the result of a second iteration of 100 projections through the data derived by thresholding Figure 5.23(d) for $\Theta_{occ} \geq 1.6\pi$. Figure 5.23(f) is the result of thresholding the third and final iteration of the process such that $\Theta_{occ} \geq 1.6\pi$. The definition of the new reconstruction space has allowed the cavity to be resolved, while still extrapolating across the gaps in the perimeter.

The user can determine whether or not a cavity is resolved by supplying a suitable reconstruction space. Alternatively, by introducing a localised reconstruction space centred on each point $\mathbf{x}$, the ability to reconstruct cavities can be significantly improved without having to supply an initial estimate of the shape. Consider the 2D case of a circular reconstruction space, Ω_x, centred on the point $\mathbf{x}$, where $\mathbf{x} \in V$ (Figure 5.24). In 2D, with continuous-valued data, $\Theta_{occ}(\mathbf{x})$ is now given by,

$$\Theta_{occ}(\mathbf{x}) = 2\pi - \int_0^{2\pi} e^{-\int_0^R O(\mathbf{u},\theta)du} d\theta, \tag{5.22}$$

where R is the radius of the reconstruction space centred on $\mathbf{x}$. The amount of attenuation at each point $\mathbf{x} \in V$ is determined by the object data contained within the circle of radius R surrounding it. Clusters of points circumscribed by the reconstruction space will tend to form regions with successive iterations of the method. The size of R limits the distance over which two data points can be joined, and hence the size of regions formed, at each iteration. Thus, with a small reconstruction space, only data points in close proximity can be grouped together. As R increases, the influence of points further away is felt.

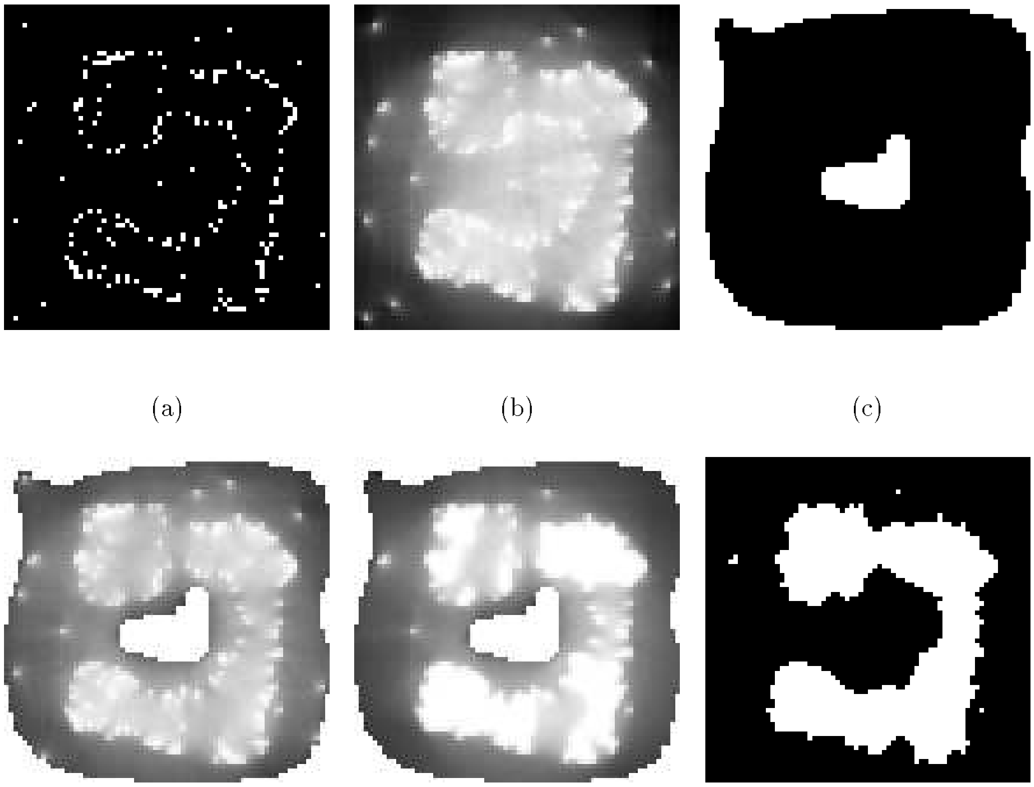

(a) (b) (c)

(d) (e) (f)

Figure 5.23 Iterative reconstruction of a cavity by defining a suitable reconstruction space. (a) Object data, (b) accumulator for 1st iteration where the reconstruction space coincides with the domain of the image (c) alternative reconstruction space, (d) accumulator for 1st iteration, (e) accumulator for 2nd iteration, (f) reconstructed shape after three iterations.

Localising the reconstruction space necessarily compromises the power of the original method to interpolate across large regions of sparse object data. However, in practice starting with a local reconstruction space and increasing R, together with the threshold applied at each iteration, produces better results than employing a reconstruction space that contains all the object data.

Implementing the localised reconstruction space involves a minor alteration to the discrete algorithm for estimating Θ_{occ} proposed in Section 5.5.1. A FIFO buffer (first in, first out) is introduced to calculate the line integral of $O(\mathbf{x})$ along each projection ray. This stores the contribution to the projection integral from the last n pixels which a ray has passed through. The size of this buffer, n, is a function of how many ray-pixel intersections occur over the distance R along the ray path. As the ray is traversed and the i^{th} pixel is encountered, the previous value of the projection, $\mathcal{P}_{i-1}$, is incremented by the contribution from the new pixel, δ_i, and decremented by the contribution from the oldest pixel in the buffer, δ_{i-n}. The new increment is then stored in the buffer in place of the oldest, and ray-tracing continues. Figure 5.25 illustrates the modified procedure. As before, the attenuation at pixel centres is interpolated from the neighbouring projection samples. The additional computation required to realise the localised reconstruction space is therefore not significant and is independent of the size of R.

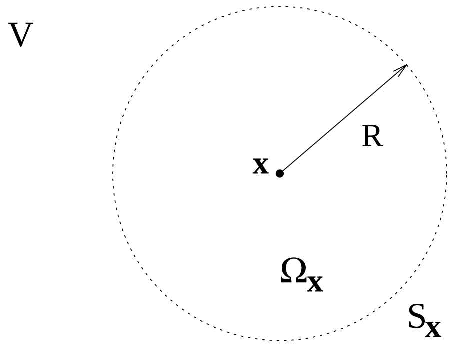

Figure 5.24 Definition of a circular reconstruction space centred on the point $\mathbf{x} \in V$.

In Figure 5.26 the application of a local reconstruction space is illustrated with the cavity data considered in Figure 5.23. Figure 5.26(a) is the accumulator after taking 100 projections. The attenuation at each pixel is calculated from object data within a 10 pixel radius. Figure 5.26(b) is the accumulator after three iterations with a space 20 pixels in diameter and a further iteration with a larger space 30 pixels in diameter. A threshold of 1.4π was applied at each of the first three iterations. A more conservative threshold of 1.5π was then applied to the accumulator in Figure 5.26(b). Figure 5.26(c) is the result of taking a further iteration of projections with a space 30 pixels in diameter and thresholding it for 1.5π. A reconstruction similar to that in Figure 5.23(f) has been achieved without the definition of a shape-specific reconstruction space.

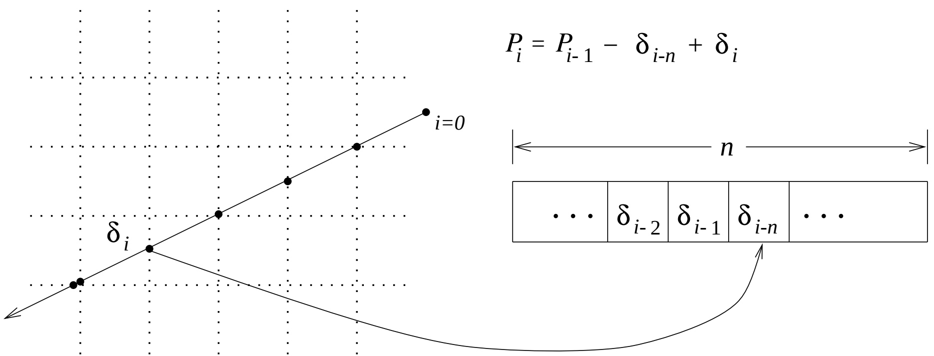

Figure 5.25 Discrete realisation of a localised reconstruction space. As a projection ray is traversed, the contribution to the projection integral from the i^{th} pixel, δ_i, overwrites the oldest element in the buffer. The projection integral at the i^{th} pixel is the sum of the buffer contents.

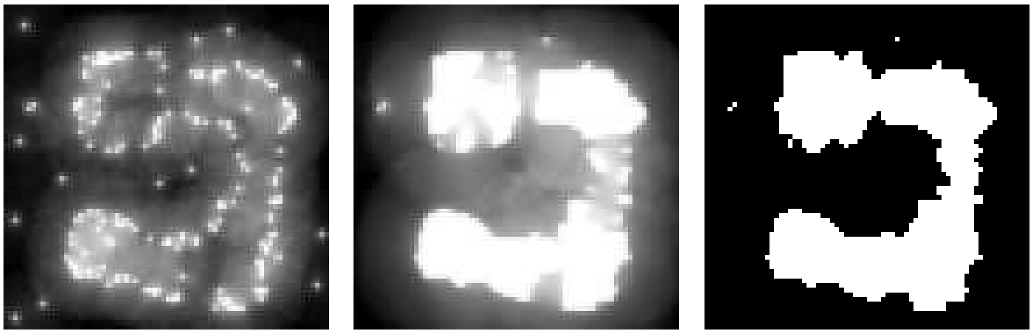

(a) (b) (c)

Figure 5.26 Iterative reconstruction of a cavity by defining a local reconstruction space. (a) Accumulator for first iteration with a circular space 20 pixels in diameter, (b) accumulator for fourth iteration with a circular space 30 pixels in diameter, (c) reconstructed shape after a second iteration with a space 30 pixels in diameter.

5.8 3D SHAPE RECONSTRUCTION

In three dimensions the reconstruction space becomes a volume and Θ_{occ} is interpreted as the occluded *solid angle*. Θ_{occ} is measured in steradians, hence $\Theta_{occ} \in [0, 4\pi]$. The occluded solid angle at the point $\mathbf{x}$, $\Theta_{occ}(\mathbf{x})$, is given by,

$$\begin{aligned}
\Theta_{occ}(\mathbf{x}) &= \int_0^{2\pi} \int_{-\pi/2}^{\pi/2} (1 - e^{-\mathcal{P}(\mathbf{x},\theta,\phi)}) \cos(\phi) d\phi d\theta \\
&= 4\pi - \int_0^{2\pi} \int_{-\pi/2}^{\pi/2} e^{-\mathcal{P}(\mathbf{x},\theta,\phi)} \cos(\phi) d\phi d\theta \qquad (5.23)
\end{aligned}$$

where $\mathcal{P}(\mathbf{x}, \theta, \phi)$ is the projection of $O(\mathbf{x})$ along the ray passing through Ω defined by the point $\mathbf{x}$ and the two angles, $\theta \in [0, 2\pi]$ (azimuth) and $\phi \in [-\frac{\pi}{2}, \frac{\pi}{2}]$ (elevation).

Estimation of the discrete $\Theta_{occ}(\mathbf{x})$ distribution in 3D is analogous to that in 2D. The accumulator becomes a 3D array of voxels. Planes of rays are traced through the reconstruction space and tri-linear interpolation between the four neighbouring projection rays is used to estimate Θ_{occ} at each voxel centre.

In 2D, it is possible to take projections at equally spaced orientation angles for any specified number of projections. In 3D, this is not the case. Orientations are chosen by projecting a uniform random distribution on to the solid angle space defined by θ and ϕ. A cylindrical equal area projection [Maling 1992] is used to do this. An approximately even coverage of orientations is then achieved for sufficiently large N.

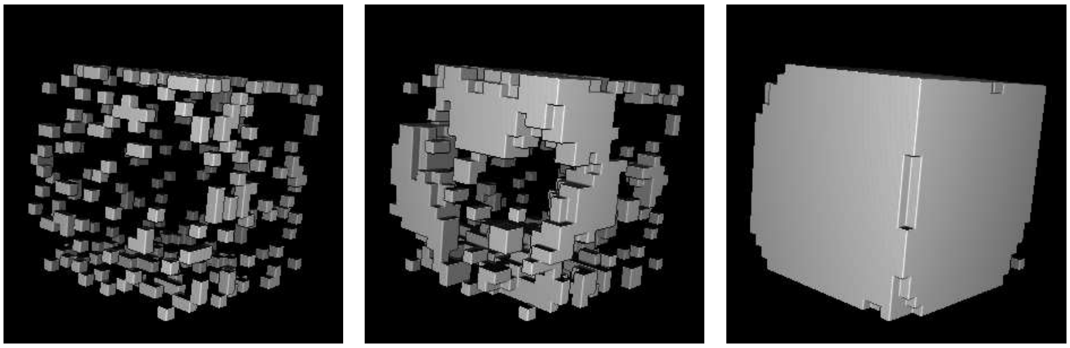

(a) (b) (c)

Figure 5.27 Iterative reconstruction of a cube. (a) Object data consisting of a cube with 90% of its surface voxels missing, (b) reconstructed object data after two iterations of 400 projections, (c) reconstructed object data after four iterations of 400 projections.

5.8.1 3D reconstructions

The iterative 3D reconstruction of a cube from an incomplete set of surface voxels is demonstrated in Figure 5.27. In 5.27(a) 90% of the surface voxels for a complete (20x20x20) cube have been removed. The object data have been assigned a high attenuation coefficient rendering them opaque. Figure 5.27(b) is the reconstructed object data after two iterations of projections with a conservative threshold choice ($\Theta_{occ} \geq 2.5\pi$ steradians). The surface voxels have formed small clusters. Figure 5.27(c) is the reconstructed object data after four iterations of projections. Voxels occluded in at least half the projections taken are shown ($\Theta_{occ} \geq 2\pi$ steradians). 98.8% of the cube has been reconstructed with 0 extraneous voxels.

The iterative scheme was applied to 3D ultrasound data obtained by scanning a phantom which consisted of a kumara (a type of sweet potato) suspended in water. The data was acquired with a 5MHz linear probe with a 58x120mm field of view. Figure 2.2 in Chapter 2 is a photograph taken during the acquisition procedure. 120 ultrasound slices taken from arbitrary positions and orientations around the phantom were mapped on to a 3D array of voxels 190x80x90mm, where each voxel corresponded to 1mm^3. Data were mapped to the nearest voxel with multiple contributions to the same voxel being averaged. The majority of voxels contained no data due to the incomplete coverage of the phantom by the set of ultrasound scans. No filtering was applied to the data to reduce speckle.

Figure 5.28(a) is a ray-traced view of the 3D ultrasound data array. Only those voxels above a minimum threshold have been rendered, these being the input data to the algorithm. Figure 5.28(b) is the corresponding 3D accumulator after two iterations of projections. At each iteration 1000 projections were taken with rays spaced 1mm apart within each projection. The reconstruction space coincided with the domain of the voxel array; a locally defined reconstruction space was not used to reconstruct this object. The raw data was acquired with 8-bit resolution and an attenuation coefficient assigned which varied linearly with the ultrasound number over the range 0.0-5.0. Consequently, strong reflections were effectively treated as opaque and weaker reflections as translucent. The first iteration was thresholded for $\Theta_{occ} \geq \frac{5}{2}\pi$ steradians. The resulting object data was then universally assigned a high attenuation coefficient, resulting in a binary data set. The accumulator formed by taking projections about this data was thresholded for $\Theta_{occ} \geq 2\pi$ steradians. This is the surface displayed in Figure 5.28(b).

For comparison, the phantom was CT scanned with 3mm between slices and a resolution of 0.35mm within each slice. A view similar to Figure 5.28(b) was produced in Figure 5.28(c) via a standard ray-tracing algorithm. The main differences are due to registration errors with the spatial locator for the ultrasound scan planes and incomplete coverage of the object by the probe during scanning.

(a) (b) (c)

Figure 5.28 Reconstruction from 3D ultrasound data. (a) 3D ultrasound data for kumara phantom, (b) 3D accumulator after two iterations of 1000 projections. The surface shown corresponds to a threshold of $\Theta_{occ} \geq 2\pi$ steradians. (c) Rendering of phantom from CT data.

5.9 CONCLUSIONS

A novel approach to the problem of fitting a surface to an incomplete set of object data has been proposed and applied to ultrasound data in two and three dimensions. This approach was determined by the characteristics of 3D ultrasound data. These include the noisy nature of the object data, the non-uniform distribution of the data and the observation that the data are not exclusively associated with the surface of the object. These are significant impediments to the application of surface reconstruction techniques developed for laser scanners, ultrasound ranging and stereo vision.

The iterative scheme, employing continuous-valued projections, reconstructs objects with moderate concavities and is robust to the presence of false object data.

Noisy tissue interfaces are typically detected by ultrasound. The main strength of the approach is its ability to interpolate across regions of sparse data. No attempt is made to explicitly identify false data points in the initial object data nor are points identified which the surface must pass through. The adjacency relationship between data points is not specified either. The user resolves ambiguities by applying the appropriate threshold at each iteration. A complete regular array of voxel data is not required; non-parallel slice data are projected on to a regular grid, but are not interpolated on the grid. The interpolation strategies described in Chapter 4 are therefore avoided.

Only isolated objects in the presence of noise have been discussed. In clinical data this condition is unlikely to be met without the identification of a region of interest within the raw ultrasound data which isolates data pertaining to the object to be reconstructed. This is currently envisioned as an interactive task.

Because reflections invariably arise from the constituent tissue which makes up an object, as well as the interface between the object and its surrounds, the method of surface reconstruction should ideally be insensitive to data specified within the object. However, meeting this requirement is in conflict with resolving complex cavities and it biases the method towards filling-in cavities. This is because of the ambiguity in the interpretation of the data, i.e. whether it is surface data or arises from reflections inside the object. Whether or not a cavity is filled in predominantly depends on the size of a cavity relative to the maximum size of holes to be interpolated across and the amount of surface data given. Two innovations, the introduction of a reconstruction space modified to suit the object and the application of a localised reconstruction space, go some way to addressing this problem. The definition of an appropriate reconstruction space is suitable in the context of an interactive tool. The localised reconstruction space does not require the definition of an appropriate reconstruction space, but it introduces the need for a scheme to determine appropriate combinations of thresholds and sizes for the local reconstruction space. Automating the selection of these parameters is the subject of further work.

Any solution to the problem of constructing surfaces from sparse, unorganised data in the absence of *a priori* information is a compromise between opposing goals. The method proposed here will tend to fit a simpler surface (a lower order approximation) to data associated with complex objects. In the situation where a shape is not reconstructible, this method will at worst result in cavities being filled in. This may restrict the application of the method to objects which do not exhibit the type of cavity which causes a problem, when such structure occurs at the scale of structures which are of interest. Even so, the robust ability to fit a simpler surface to a complex object may provide useful quantitative information.

Chapter 6

SURFACE INTERPOLATION FOR TITANIUM CRANIOPLASTY

In this chapter radial basis functions (RBFs) are presented as a practical solution to the problem of interpolating incomplete surfaces derived from three-dimensional (3D) medical graphics. The specific application considered is the design of cranial implants for the repair of defects, usually holes, in the skull.

A system for designing cranial implants is described. In this system medical graphics are used to reveal bone surfaces in X-ray CT data. A mathematical function is fitted to the rendered surface data in a manner which smoothly interpolates across holes in the surface. From this mathematical description a hard plastic mould can be milled by a computer numerically controlled (CNC) mill. A titanium prosthesis can then be formed by pressing flat titanium plate into the mould under high pressure in a hydraulic press.

The novel contribution of this work is the facility to describe the surface of the skull by a mathematical function which can smoothly interpolate across a defect. Doing so reduces the skilled manual aspects currently associated with making a prosthesis. The desired characteristics of the interpolant are that it is smooth, can interpolate across irregularly shaped holes and that it tends towards a flat plate far from the interpolation centres.

Radial basis functions impose few restrictions on the geometry of the interpolation centres and are suited to problems where the interpolation centres do not form a regular grid. However, their high computational requirements (both speed and storage) have previously limited their use to problems where the number of interpolation centres is small (< 300). Recently developed fast evaluation techniques have overcome these limitations and made radial basis interpolation a practical approach for larger data sets.

The chapter is organised as follows. Section 6.1 backgrounds the cranioplasty application and introduces the shape interpolation aspect of the problem. Section 6.2 discusses various approaches to the interpolation problem. Section 6.3 introduces radial basis functions and section 6.4 illustrates the application of radial basis functions to the problem of interpolating depth-maps of the skull's surface. Section 6.5 presents examples of fitting surfaces to real data and section 6.6 considers computational requirements. A discussion follows in section 6.7.

The notation C^i is used throughout the chapter to denote a function which is continuous in its first i derivatives. The notation $\mathcal{O}(f(n))$ indicates that the computation or storage (depending on the context) required by a process is bounded by a multiple of $f(n)$. For example, the computation in solving an $n \times n$ system of linear equations, with no special structure, is $\mathcal{O}(n^3)$.

6.1 CRANIOPLASTY

Cranioplasty is the procedure of restoring defects, usually holes, in the skull with the aid of cranial implants. Small defects (5cm^2 to 15cm^2) are generally corrected for aesthetic reasons while large defects ($\geq 25\text{cm}^2$) require repair for protection of the brain from external injury and for the prevention of cortical thinning, which may cause unnecessary neurological symptoms. Cosmetic and psychological factors are also of importance, and defects in the frontal or temporal bones are restored to hide disfigurement.

The complexity of cranioplasty depends on the material used to form an implant and upon the size and location of the defect. In general, implants for larger defects ($\geq 50\text{cm}^2$) are not formed intra-operatively. A wide range of materials have been used in the reconstruction of cranial defects. These include hydroxylapatite, alumina ceramic, silicon, polymethyl methacrylate (PMMA), stainless steel and titanium. Issues which influence the choice of a material are tissue compatibility, strength, radiolucency, weight, availability, thermal and electrical conductivity, and the ease with which the material can be manipulated at surgery.

6.1.1 PMMA implants

PMMA is the most popular material currently used in cranioplasty. Autopolymerising PMMA forms a pliable material which can be moulded at the time of surgery and fitted directly on to the skull. PMMA is tissue compatible, has low electrical and thermal conductivity and is radiolucent. The latter property is important as it permits future radiotherapy treatment and X-ray CT scanning after an implant has been fitted. It is the ability to form the implant in surgery which makes PMMA popular. However, this property diminishes in significance if the defect is large or occurs in a region of the skull which has moderate curvature, or where bones are thin (the anterior temporal region for example), or in bony structures that surround vital tissue, such as the orbit of the eye. In these regions the complex shape of an implant is difficult to produce intra-operatively and significantly prolongs surgery. In such cases an accurate implant must be formed presurgically.

Three major physical factors affect the success of an autopolymerising PMMA implant. These are the heat of polymerisation, the residual monomer, and porosity.

The curing process of autopolymerising PMMA resins is exothermic. It has been found that the temperature of autopolymerising PMMA applied directly to the skull can rise to 64° C causing damage to cells and tissue, despite being irrigated with cold saline solution [Van Putten and Yamada 1992]. Consequently, implants formed intra-operatively should be removed from the surgical site prior to final curing. However, the process of removing the implant can in turn result in distortion, affecting the fit and stability (whether the implant stays in place).

The amount of residual methyl methacrylate monomer (MMA) in autopolymerised PMMA resins is also of concern. This has been reported in the range of 1% to 4% immediately after curing and has an adverse effect on surrounding tissue which can lead to infection and complications [Van Putten and Yamada 1992]. MMA has been reported to cause contact dermititis, and allergic reactions on skin and oral mucosa. MMA causes acute local infections in the tissues next to cranial implants, although these are often short lived.

Porosity and the inclusion of blood and tissue fluid can greatly reduce the strength of an autopolymerising PMMA prosthesis [Hinton *et al.* 1992]. Long term results were initially poor for autopolymerising PMMA implants when surgeons casually prepared the cement in the open, often introducing air bubbles. Furthermore, the manual pressing of the cement into the bone increased porosity. Porosity may result in broken implants requiring replacement. For bone implants in other parts of the body the success of autopolymerising PMMA has been improved by preparing the cement in a vacuum and firing it into a joint with a gun. However, this is not applicable when forming cranial implants. Small (0.25 to 0.5 mm) bubbles are found within most autopolymerised PMMA cranial implants. These can cause erroneous evaluation of postsurgical CT scans where the bubbles can mimic cranial abscesses.

These three drawbacks of autopolymerising PMMA can be overcome by using prefabricated heat-processed PMMA resin implants. These have higher tensile and flexural strength due to lower porosity, and have lower levels of residual-free monomer (0.4%). However, the advantage of surgical malleability is lost. The implant must be prefabricated. Consequently, an accurate life-size cast is required to ensure good fit. Combinations of heat treated prefabricated PMMA prosthesis and autopolymerising PMMA filler added to the implant margins have been suggested as a compromise to allow adjustment or augmentation at the time of surgery.

6.1.2 Titanium implants

An alternative to PMMA resins is the use of titanium implants. Titanium is biocompatible, light weight, very strong and, unlike stainless steel, is radiolucent [Chandler *et al.* 1994]. However, unlike autopolymerising PMMA resins, titanium implants must be fabricated prior to surgery.

The biocompatibility of titanium is due to the stable and unreactive layer of oxide that forms on its surface. If the metal is scratched the oxide layer "repairs" itself almost immediately, even when surrounded by bodily fluids. The titanium dioxide layer does not provoke a response from the body's immune system. Thus normal bone repair may proceed after a titanium implant has been fitted. A secure bond between the implant and the bone is formed as bone grows into the titanium screws which attach the implant and into the implant itself. This process is known as osseointegration and was discovered by PerIngvar Branemark while studying the function of blood and marrow [Saul 1994].

In a particular experiment, Branemark wanted to see inside the fibula bone in a rabbit's leg and made a viewing chamber by screwing a cylinder of titanium into the fibula. After completing the study he was frustrated to find that bone had grown into the threads of the titanium cylinder so that it could not be removed. A later study used a similar technique to observe the effects of eating, drinking and smoking on blood cells. A tiny titanium viewing chamber was inserted into the upper arm of medical student volunteers. The chambers were left in place for several months and caused no adverse reaction. The significance that such a metal could form such a strong bond with bone, without triggering a response from the body's immune system, was then realised by Branemark. Since 1965, when Branemark installed the first successful titanium implant, titanium prostheses have become increasingly popular.

At Christchurch Hospital titanium is preferred to autopolymerising PMMA resin for cranial implants and has been used to replace broken or infected PMMA implants [Blake *et al.* 1990]. The fewer complications experienced with titanium and the strength resulting from the process of osseointegration offset the dubious benefit of malleability at the time of surgery, which an autopolymerising PMMA implant offers.

6.1.3 Prefabricated implants

Prefabrication of a cranial implant, whether it be titanium or heat-processed PMMA, requires an accurate model of the defect area to ensure that a good fit is achieved at the time of surgery and the need for alteration is minimised.

Several techniques based on external impressions of the defect are described in the literature. These rely on casts made from the surface of the head, careful mapping and digital palpation of the defect area to determine the boundaries, and sometimes lateral and frontal radiographs with superimposed metal grids. Margins and contours for the implant are extrapolated from the various data sources. This process is tedious and prone to error. Cranial implants made by external methods usually require adjustment or augmentation at the time of surgery to compensate for poor fit. The implants tend to have flatter contours than the skull. Furthermore, the thickness of the bone surrounding the defect is not known. This affects prefabricated PMMA implants which have to approximate the bone thickness at the edges of the defect. In the case of titanium implants, optimal placement of screws to secure the plate is difficult. Finally, external methods are not suited to complex cranial regions such as the orbit of the eye where the bones are thin.

A better model can be formed from X-ray CT data. This practice is becoming common with the wide availability of CT imaging. Early approaches involved manually tracing around the bone in consecutive CT slices and cutting out the resulting bone cross-sections in sheets of perspex or some other suitable material, such that the thickness of the cross-sections matched the spacing between CT slices [Hinton *et al.* 1992]. These cross-sections, once registered, could be stacked upon each other to form a contiguous 3D model. Access to Numerical Control (NC) milling machines has allowed this process to be automated, enabling slices to be cut out directly from the CT data. A step artifact in the model necessarily arises due to the finite spacing between CT slices. Van Putten and Yamada [1992] have found that CT scans taken 3mm apart are adequate for bone edges perpendicular to the slice and reduce this to 1mm for bone edges parallel to the slice where the accuracy of the model would otherwise be reduced.

More recently, 3D models have been milled directly from solid blocks of resin with interpolation between CT slices employed to avoid the step artifact [Joffe *et al.* 1992, Linney *et al.* 1993]. The shapes which can be reproduced in this manner are limited by the size and shape of the cutting tool and the restrictions imposed by the proximity of the milling machine's chuck to the workpiece. There is tradeoff between the length of the cutting tool and the diameter of the tool. Consequently, the ability to machine complex cavities is limited unless the object is partitioned into subsections which can be milled. The mill is ordinarily confined to milling single-valued surfaces for a given orientation of the workpiece. By rotating the workpiece about an axis (or tilting the mill chuck) more complex 3D models can be formed.

The restrictions on cavities which can be milled is not generally a problem when forming a model of a cranial defect since the outward surface of the skull is of interest and this is primarily convex over the regions such as the cranium. Often only a model of the immediate region surrounding a defect is required and the model can be milled from a single orientation of the workpiece.

Nakai *et al.* [1993] describe an alternative optical fabrication method which is not so restricted in the complexity of the shapes which it can form. This involves using lasers and photopolymers. The principle is that a liquid photopolymer is converted to a solid on exposure to UV light by the process of photopolymerisation. A 3D model is built up incrementally from cross-sections, layer upon layer, in a tank of liquid photopolymer. Each layer is cured by exposure to a computer-controlled UV laser and adheres to the previously cured layer beneath it. A platform lowers the model in a tank of liquid photopolymer as it evolves. The laser is controlled so as to map out a binary cross-section of the object of interest. Nakai *et al.* apply this technique to the construction of models of the brain from MRI data. However, the method would be equally applicable to modelling the skull from CT data.

6.1.4 Implant fabrication

There is no standard process for the design and manufacture of prefabricated cranial implants, each treatment facility varying in the details of its method, and even then, each case being approached individually. However, most methods involve forming a model to which the implant, or a template for the implant, is manually fitted.

Until recently, traditional hand forming techniques were used at Christchurch Hospital to manufacture implants from commercially pure 0.7mm titanium plate [Hinton *et al.* 1992]. More than 30 implants ranging in size from 21 cm^2 to 227 cm^2 were made in this way by a skilled technician within the Medical Physics department. In some cases, butt welding of the titanium, using the tungsten inert gas (TIG) welding technique, allowed the modular construction of more complex prosthetic shapes.

The current approach involves forming a female mould in the shape of the implant. The mould is milled out of a block of Lockfast Araldite (K38) epoxy resin by a CNC mill. A custom press designed and built in the Medical Physics department of Christchurch Hospital presses the flat titanium plate into the mould under pressures ranging from 4000psi-6000psi, depending on the complexity of the mould. Intricate moulds require higher pressures. Figure 6.1 is a schematic illustration of the press. A neoprene diaphragm pushes the plate into the mould when inflated. Although the flat plate is initially flexible, it forms a rigid implant when pressed. This new fabrication process removes the manual aspects of working the metal, but the problem of determining the shape of the implant remains.

Linney *et al.* [1993] have attempted to remove the manual aspects of determining the shape of an implant in the case where a defect is unilateral. They reflect the sound side of the anatomy on to the defective side of the body in software and use this as a model for the prosthesis, with the intention of providing symmetry to the shape of the repaired skull. However, this technique is applicable only where a symmetrical undamaged region of the skull exists. Even when this criterion is met, restoring the

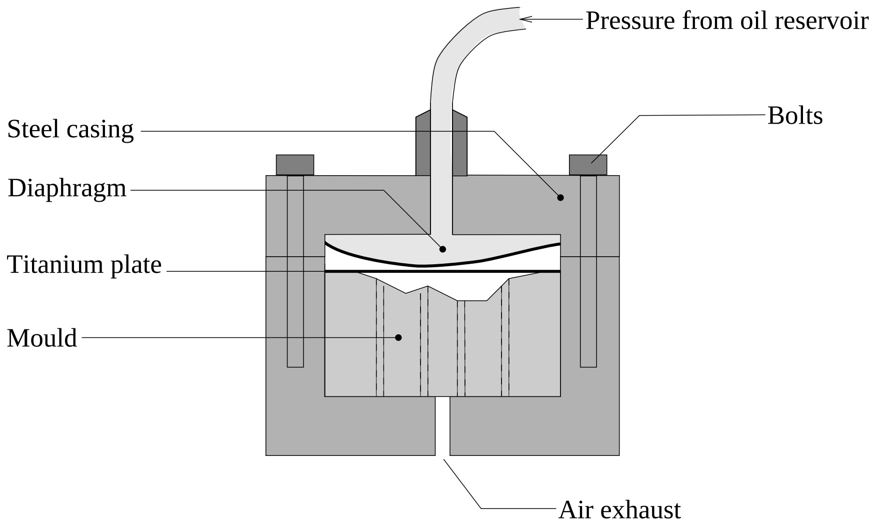

Figure 6.1 Schematic of hydraulic press.

original skull shape is not always desirable if large cavities result between the brain and the implant. Furthermore, the skull is not always naturally symmetric.

In this thesis the problem of designing a cranial implant is viewed as a surface interpolation problem where the surface of the skull is described by a mathematical function. The problem is to fit a function to the bone surface which interpolates across user-defined defect regions. The aim of this work is to create a one-step process for the design of cranial implants by integrating medical graphics and surface interpolation with CNC milling.

6.1.5 The surface interpolation problem

Depth-maps of the skull's surface, obtained from X-ray CT data using ray-casting techniques, have been used to construct models of cranial defects [Van Putten and Yamada 1992, Linney *et al.* 1993]. In this chapter depth-maps are rendered on a graphics workstation and used to fit a surface which will describe the shape of an implant. A user selects a viewing direction to render an appropriate view of the defect region. The workstation mouse is then used to graphically distinguish the defect from the surrounding sound bone by 'painting' the valid surface data. This reflects surgical practice where damaged bone surrounding a defect is often removed. In this way the user determines the extent of the repair. A surface is then fitted to the incomplete depth-map.

A 3D CT data set in the form of a regular array of voxels is obtained by stacking a set of equally spaced CT slices taken through the head and the ray-casting algorithm described in chapter 3 is used to extract a depth-map. No perspective transformation is used in the ray-casting procedure; rather a set of uniformly spaced parallel rays is cast into the CT data. A threshold corresponding to the Hounsfield number for bone (typically 500) is used to identify the bone surface. The resulting depth-map is rendered with a single viewer-centred light source using the Phong shading method described in chapter 3 (section 3.3.1). Additionally, a 3D model of the depth-map can be milled out of a block of epoxy-resin by a CNC mill.

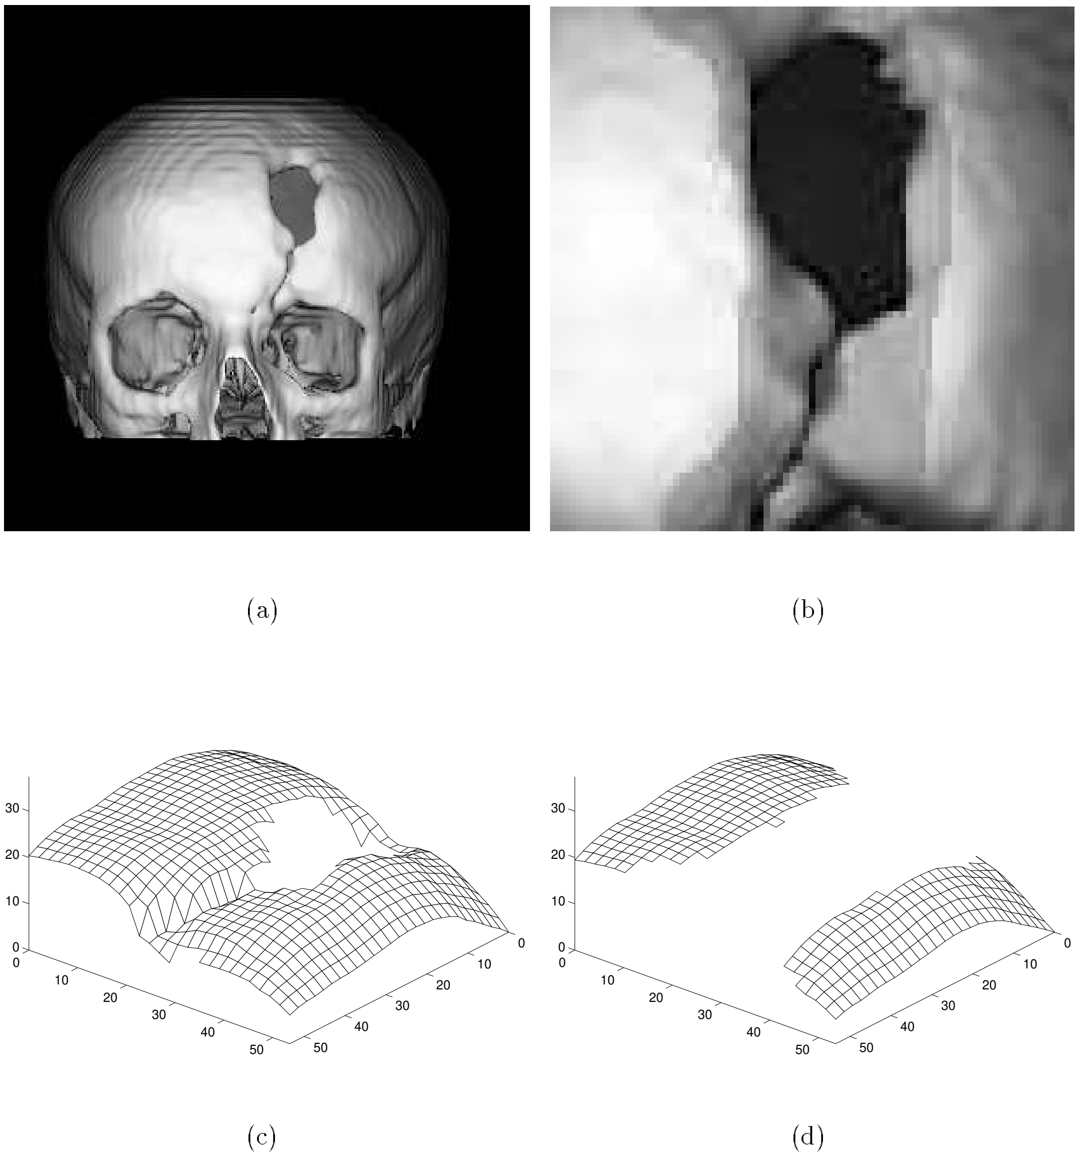

Figure 6.2 Extraction of a depth-map of the skull from CT data. (a) Rendered view of the bone surface from a CT data set, (b) detailed view of the defect with a user-defined region of defect-free bone highlighted, (c) depth-map of the rendered view of the defect in (b), (d) partial depth-map corresponding to the highlighted region in (b).

Depth-maps are rendered at a resolution sufficient for the operator to identify defect regions and then sub-sampled for the purpose of fitting a surface. A depth-map mesh of 2 mm squares provides more than adequate sampling of the surface of the skull for interpolation in most cases. Although the resolution within a CT slice is usually 0.3–0.5 mm, the spacing between slices for cranial examinations is typically 2–3 mm. Consequently, a depth-map which re-samples the data at a higher frequency will contain a ripple which originates from the interpolation procedure employed by the ray-caster. Furthermore, such a depth-map will result in an unnecessarily large number of interpolation nodes, greatly increasing the storage and computational costs of solving the interpolation equations.

Figure 6.2 illustrates the cranial surface interpolation problem. Figure 6.2(a) is an example of a ray-cast view of the surface of the skull rendered from a CT data set. Figure 6.2(b) is a detailed view of the defect in Figure 6.2(a) where the region identified as being free of defect by the user is highlighted. Figure 6.2(c) illustrates the lower resolution depth-map corresponding to the detailed view of the defect while Figure 6.2(d) is the partial depth-map corresponding to the highlighted region in Figure 6.2(b). This is the depth-map to which an interpolating function is to be fitted. In summary, the original ray-cast depth data lies on a regular rectangular grid. The user then interactively identifies a defect region of irregular shape, and the corresponding data points are discarded. This results in a final depth data set to be interpolated which is of the form of a finite grid, with a hole, or holes.

6.2 SURFACE INTERPOLATION

The depth-map is regarded as a single-valued function of two variables (2D), $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, where the interpolation nodes at which the value of the function is known do not generally form a complete regular grid. The problem is to approximate f with a smooth function which interpolates the given data and provides at least C^1 continuity in order to avoid creases in the restored surface.

One approach to this problem is to fit a polynomial to the data. However, an invertible system which uniquely defines the interpolant is not guaranteed for all positions of the interpolation nodes. Moreover, such a polynomial interpolant will typically display spurious bumps and wiggles. An alternative is to fit piecewise polynomials. These require that $\mathbb{R}^2$ be divided into suitable patches, typically rectangles and triangles. One employs different polynomials on the different patches, and the pieces are joined in ways that provide continuity of the prescribed low order derivatives. Examples of such methods are tensor product spline interpolants, parametric spline interpolants, C^1 surfaces made up of Clough-Tocher macro elements over triangles, etc. Surfaces of these types are widely used in Computer Aided Design (CAD) and in Finite Element methods for solving partial differential equations [Astley 1992, Zienkiewicz and Taylor 1989]. However, they are not well suited to scattered data interpolation problems, particularly when there are large data-free regions.

In Computer Aided Design (CAD) piecewise bicubic polynomials, defined parametrically on a rectangular mesh, are a popular choice for representing arbitrary curved surfaces with C^1 continuity. Various forms include the Hermite, Bezier and parametric B-spline [Foley and Van Dam 1982]. Each bicubic patch is specified by a set of 16 coefficients. A simple transformation of the coefficients which define a surface patch in one form can be applied to find those which represent the same surface in another form.

However, the problem in CAD is generally one of representing surfaces in a design context. This is quite different from the interpolation problem where the construction of a surface is data driven rather than directed by a user. The Hermite and Bezier forms are inconvenient for interpolation since they require the determination of tangents and derivatives across patch boundaries. The B-spline form provides C^2 continuity between neighbouring patches and can be conveniently defined by choosing the coefficients for each B-spline as the known function value at the corresponding node of interpolation. The resulting B-spline surface is a quasi-interpolant since it does not usually interpolate to the data. Rather the parametric surface lies within the convex hull of the local data points. Employing a piecewise approach in an interpolation context presents several practical problems. A mesh must be formed which defines the separate patches. The process of automatically forming a mesh is not trivial. Furthermore, to evaluate the surface at a position $\mathbf{x} \in \mathbb{R}^2$ requires determining the appropriate patch in which $\mathbf{x}$ lies and then determining $\mathbf{x}$ in terms of the parametric coordinates which define the surface on the patch. In addition, there is the problem that there are generally not enough constraints to uniquely determine the coefficients, so extra boundary conditions must be introduced at peripheral patches.

The radial basis approximation method offers several advantages over piecewise polynomial interpolants. The geometry of the known points is not restricted to a regular grid and there is no need to define a mesh of patches. Also, the resulting system of linear equations is guaranteed to be invertible under very mild conditions. Finally, polyharmonic radial basis functions have variational characterisations which make them particularly suited to interpolation of scattered data, even with large data-free regions.

6.3 RADIAL BASIS FUNCTION APPROXIMATION

Before considering the specific case of approximating a function of two variables, radial basis function approximation is introduced generally. Consider $f : \mathbb{R}^d \to \mathbb{R}$, a real valued function of d variables that is to be approximated by $s : \mathbb{R}^d \to \mathbb{R}$, given the values $\{f(\mathbf{x}_i) : i = 1, 2, \ldots, n\}$, where $\{\mathbf{x}_i : i = 1, 2, \ldots, n\}$ is a set of distinct points in $\mathbb{R}^d$.

We will consider approximations of the form

$$s(\mathbf{x}) = p_m(\mathbf{x}) + \sum_{i=1}^{n} \lambda_i \phi(\|\mathbf{x} - \mathbf{x}_i\|), \qquad \mathbf{x} \in \mathbb{R}^d, \lambda_i \in \mathbb{R}, \tag{6.1}$$

where p_m is a polynomial of low degree m, or is not present, $||\cdot||$ denotes the Euclidean norm and ϕ is a fixed function from $\mathbb{R}^+$ to $\mathbb{R}$ ($\mathbb{R}^+$ denotes the positive real numbers). Thus, the radial basis function s is a linear combination of translates of the single radially symmetric function $\phi(\|\cdot\|)$, plus a low degree polynomial. If the space of all polynomials of degree at most m in d variables is denoted by π_m^d, then the coefficients, λ_i, of the approximation s are determined by requiring that s satisfy the interpolation conditions

$$s(\mathbf{x}_i) = f(\mathbf{x}_i), \qquad i = 1, 2, \ldots, n, \tag{6.2}$$

together with the side conditions

$$\sum_{i=1}^{n} \lambda_i q(\mathbf{x}_i) = 0, \qquad \text{for all } q \in \pi_m^d. \tag{6.3}$$

Some examples of popular choices of ϕ and the name of the corresponding radial function are given below [Powell 1992].

$$\left.\begin{array}{ll} \phi(r) = r & \text{(linear)} \\ \phi(r) = r^2 \log r & \text{(thin-plate spline)} \\ \phi(r) = e^{-ar^2} & \text{(Gaussian)} \\ \phi(r) = (r^2 + c^2)^{1/2} & \text{(multiquadratic)} \end{array}\right\}, \qquad r \geq 0, \tag{6.4}$$

where a and c are positive constants. None of these functions, except for the Gaussian, has finite support; they grow as $r \longrightarrow \infty$. Consequently, the work involved in evaluating the function at each $\mathbf{x}$ is proportional to n, the number of points where f is known. The choice of ϕ is determined by the dimension of the problem d, the interpolation conditions and the desired properties of the interpolant.

Some typical conditions on the interpolation nodes (the $\mathbf{x_i}$'s) under which the interpolation conditions (6.2) and (6.3) uniquely specify the radial basis function (6.1) are given in Table 6.1. In the case when the polynomial term is absent from the radial basis

ϕ	spatial dimension d	polynomial degree m	Restriction on nodes
linear	any	1	nodes not coplanar
thin-plate	2	1	nodes not colinear
Gaussian	any	absent	none
multiquadric	any	absent	none

Table 6.1 Conditions imposed on nodes for various radial basis interpolants.

function (6.1) there are no corresponding side conditions (6.3). Also "not coplanar" in this context means that the nodes do not all lie in a single hyperplane, or equivalently that no linear polynomial in d-variables vanishes at all of the nodes. Thus, the conditions (6.3) are a formal statement of the intuitive requirement that the $\mathbf{x_i}$'s span the space $\mathbb{R}^d$ in order to constrain the polynomial p_m. The surveys of Powell [1992] and Light [1992] are excellent references for these and other properties of radial basis functions.

In this application the interpolation of 2D (depth-map) data is of concern. In particular, the linear and thin-plate spline interpolants are considered. Now (6.2) and (6.3) imply that the coefficients of the radial basis function, λ, and the polynomial $p_1(\mathbf{x})$, c, can be found by solving the linear system

$$\begin{bmatrix} A & Q \\ Q^T & 0 \end{bmatrix} \begin{bmatrix} \lambda \\ c \end{bmatrix} = \begin{bmatrix} f \\ 0 \end{bmatrix}. \tag{6.5}$$

where

$$A = (a_{ij}) = (\phi(\|\mathbf{x}_i - \mathbf{x}_j\|)), \tag{6.6}$$

$$Q = \begin{bmatrix} 1 & x_1 & y_1 \\ 1 & x_2 & y_2 \\ \vdots & \vdots & \vdots \\ 1 & x_n & y_n \end{bmatrix}, \tag{6.7}$$

$$\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_n)^T, \tag{6.8}$$

$$c = (c_0, c_1, c_2)^T, \tag{6.9}$$

$$p_1(\mathbf{x}) = c_0 + c_1 x + c_2 y \tag{6.10}$$

and

$$f = (f_1, f_2, \ldots, f_n)^T \tag{6.11}$$

The thin-plate, or 2-dimensional biharmonic spline models the deflection of an infinite thin plate [Harder 1972]. While the linear radial basis function will interpolate the data, the thin-plate spline is more attractive since it also provides C^1 continuity and minimises the energy functional

$$E(s) = \int_{\mathbb{R}^2} \left(\frac{\partial^2 s}{\partial x^2}\right)^2 + 2\left(\frac{\partial^2 s}{\partial x \partial y}\right)^2 + \left(\frac{\partial^2 s}{\partial y^2}\right)^2 dxdy \tag{6.12}$$

over all interpolants for which the energy functional (6.12) is well defined. In this sense the thin-plate spline is the *smoothest* interpolator of f. Higher order polyharmonic splines achieve continuity of higher derivatives. For example, Perrin *et al.* [Perrin *et al.* 1987a, Perrin *et al.* 1987b] have used the tri-harmonic thin-plate spline ($\phi(\mathbf{x}) = ||\mathbf{x}||^4 \log ||\mathbf{x}||$) in electroencephalography to interpolate with C^2 continuity distributions of potential and current scalp density measured on the surface of the head.

6.4 RADIAL BASIS INTERPOLATION OF DEPTH-MAPS

This section discusses the application of radial basis functions to the interpolation of cranial depth-maps. The first task is to calculate the interpolant, or rather its coefficients. The solution procedure is as follows. The interpolation nodes (the $\mathbf{x_i}$'s) are first scaled uniformly in x and y, and shifted so that the new nodes lie in the unit square. The variational characterisation of the thin-plate spline interpolant means that it is preserved under this transformation. Then the interpolation problem corresponding to the transformed data is solved using the double precision diagonal pivoting method for symmetric indefinite systems from LAPACK [1992].

Although the interpolation matrix is guaranteed to be invertible, provided the centres $\mathbf{x}_i$ are distinct and not colinear, the condition number of the matrix may be large. The condition number of the interpolation matrix A, $\text{cond}(A) = ||A||||A^{-1}||$, bounds the ratio of the relative change in the solution to the relative change in the data. Thus a large condition number means that it is possible for numerically obtained coefficients, $\mathbf{c}$, to contain large errors. This affects the ability to solve (6.5) accurately on a digital computer. Scaling the interpolation nodes to lie on the unit square helps to improve the conditioning of the interpolation matrix. The accuracy of the computed inverse can be further improved by using iterative refinement (see appendix C).

The current software implementation invokes iterative refinement to improve the accuracy of an initial solution if warranted by the magnitude of the residual. However, experiments have shown that the infinity norm condition number of the system is typically around 6×10^6, the initial relative residual being around 2×10^{-13}, and the initial relative error being around 1.5×10^{-11}. The aim was to obtain a relative error (against the "true" thin-plate interpolant) of 10^{-4}. Overall, more than adequate accuracy in the fitted surface is readily obtained by solving in double precision, even without iterative refinement.

In the case of the thin plate spline basis, care must be taken when computing $\phi(r) = r^2 \log r$, since a singularity occurs at $r = 0$. Although the rate at which r^2 approaches zero nullifies the rate at which $|\log r|$ grows, it is necessary in practice to modify the coding of the method to detect for the $\log r$ singularity and evaluate the expression as zero when r is sufficiently small.

The natural criterion for assessing a restored cranial surface is how closely it matches the original surface prior to the introduction of the defect. However, the original surface is rarely known and does not exist in the case of a congenital defect. Therefore, the performance of the method was simulated by seeing how well it restored artificial holes in complete skulls. The interpolant fitted to the incomplete depth-map was then compared with the original surface.

Figure 6.3 illustrates the test data employed in the following interpolation example in which an artificial hole is introduced into an initially defect-free depth-map of a cranial surface. Figure 6.3(a) is a rendered view of a CT data set which is free from defect. Figure 6.3(b) is an enlarged view of part of Figure 6.3(a) with an artificial hole superimposed. Figure 6.3(c) is the corresponding depth-map. A thin-plate spline radial basis interpolant was fitted to 683 points in Figure 6.3(c). Figure 6.4 compares the fitted surface with the original data. Figure 6.4(a) is the complete depth-map of the original data, while Figure 6.4(b) is the depth-map of the interpolated surface. Figures 6.4(c) and 6.4(d) are the rendered views of the respective depth-maps in 6.4(a) and 6.4(b).

In Figure 6.5(a) the difference between the fitted thin-plate spline surface and the original is displayed. The scales are in millimetres. The maximum error for the thin-plate spline interpolant is 0.60 mm. Figure 6.5(b) is the corresponding comparison between the original surface and a linear radial basis interpolant. The largest error from the original is 0.56 mm. The differences between the two interpolants and the original surface are dependent upon the surface data chosen in Figure 6.3. This example is nevertheless typical of many of the smooth, low curvature surfaces which are to be reconstructed. Although the magnitude of the error in both fitted surfaces is similar, the variational characterisation of the thin-plate spline, and its C^1 continuity, make it the natural choice for cranial implant design.

6.4.1 Surface evaluation

Once the linear system is solved, s can be evaluated at any resolution, anywhere on $\mathbb{R}^2$. Figure 6.6 shows the thin-plate surface from Figure 6.4 evaluated outside the convex

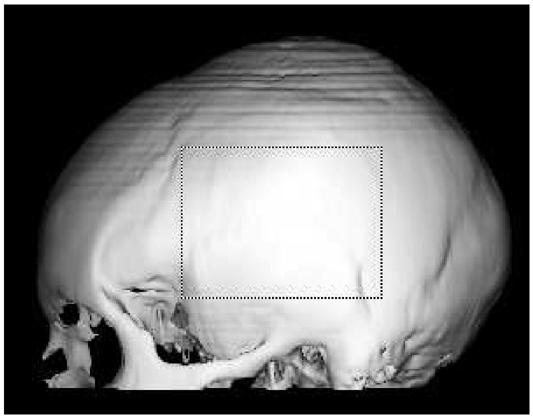

(a)

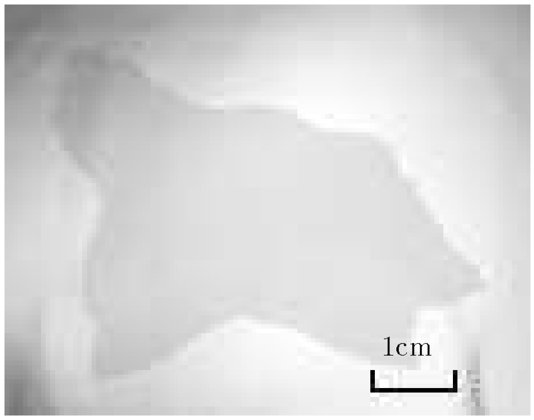

(b)

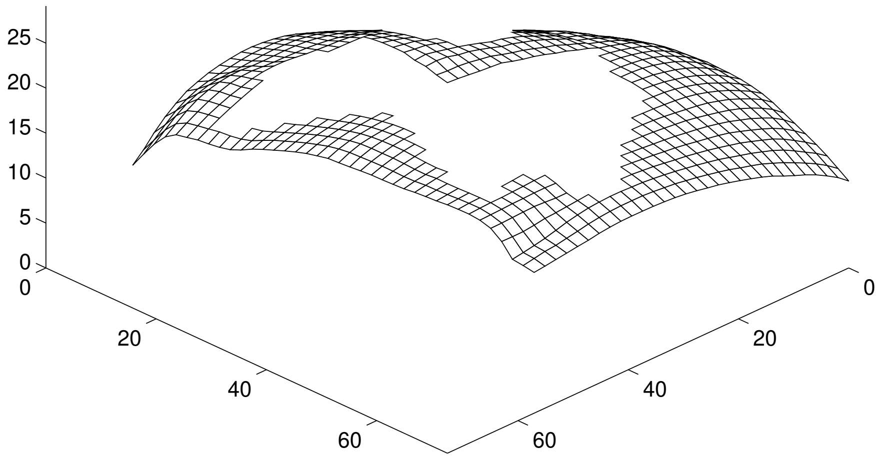

(c)

Figure 6.3 Test data for surface interpolation. (a) Original CT data set. (b) Detail of the test surface with an artificial defect superimposed. (c) Depth-map corresponding to (b).

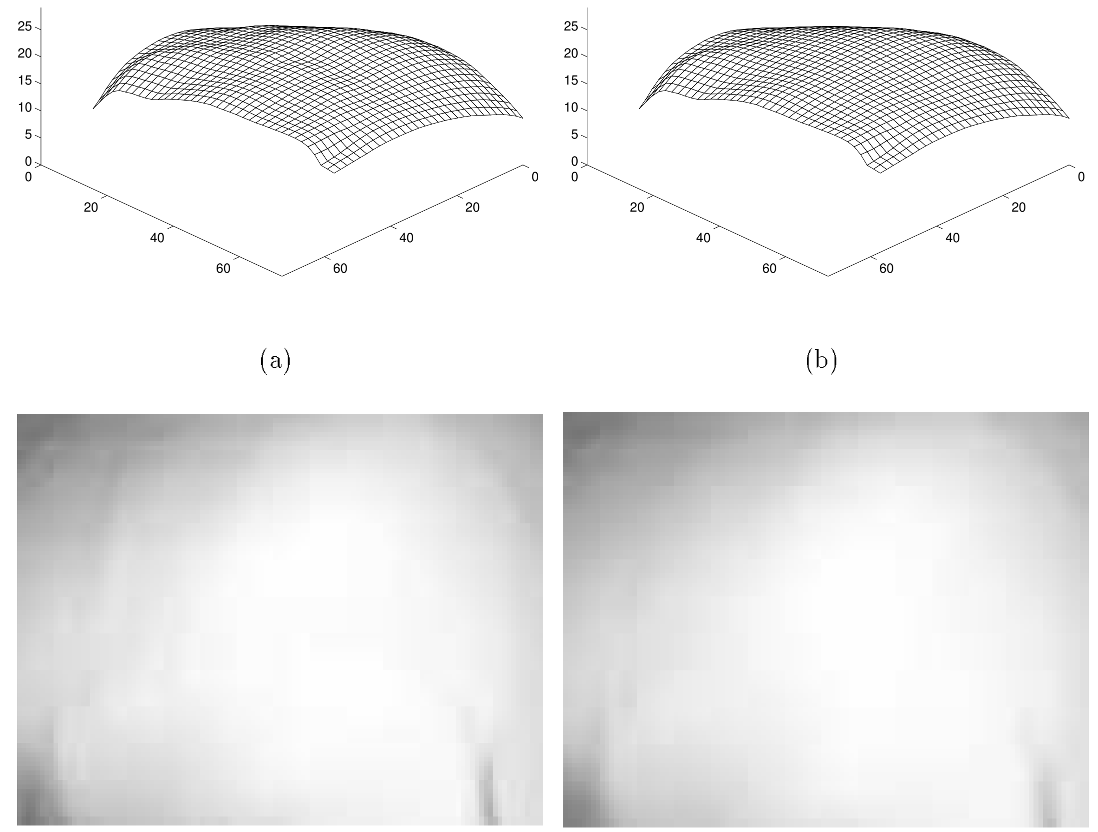

(a) (b)

(c) (d)

Figure 6.4 Comparison of the interpolated surface with the original from Figure 6.3. (a) Depth-map for original data. (b) Depth-map for thin-plate spline surface. (c) Rendered view of original data. (d) Rendered view of thin-plate spline surface.

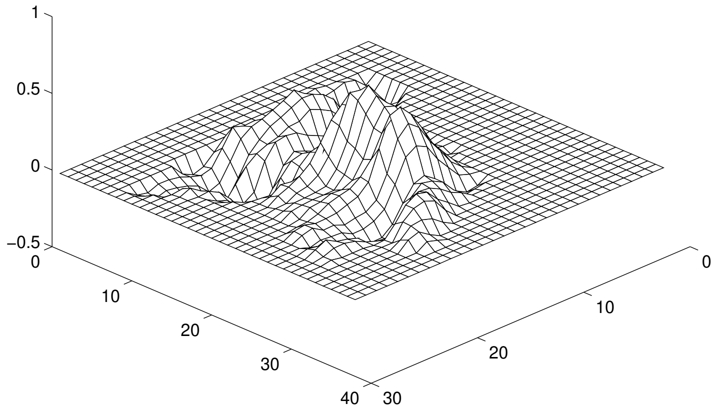

(a)

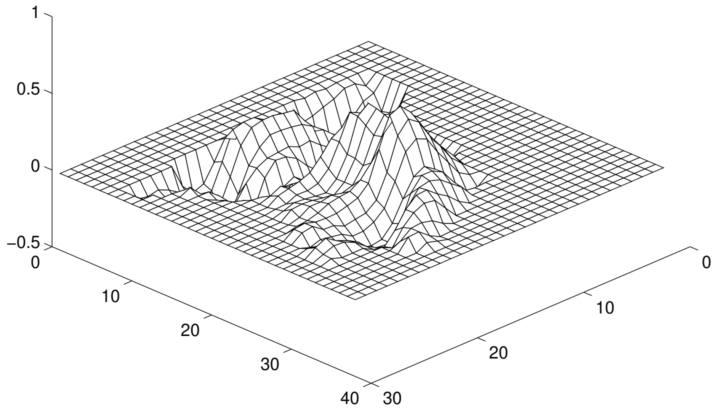

(b)

Figure 6.5 Comparison of surface variation between the thin-plate spline and linear radial bases; the variation is in millimetres. (a) Difference between the thin-plate spline surface and original data. (b) Difference between the linear surface and original data.

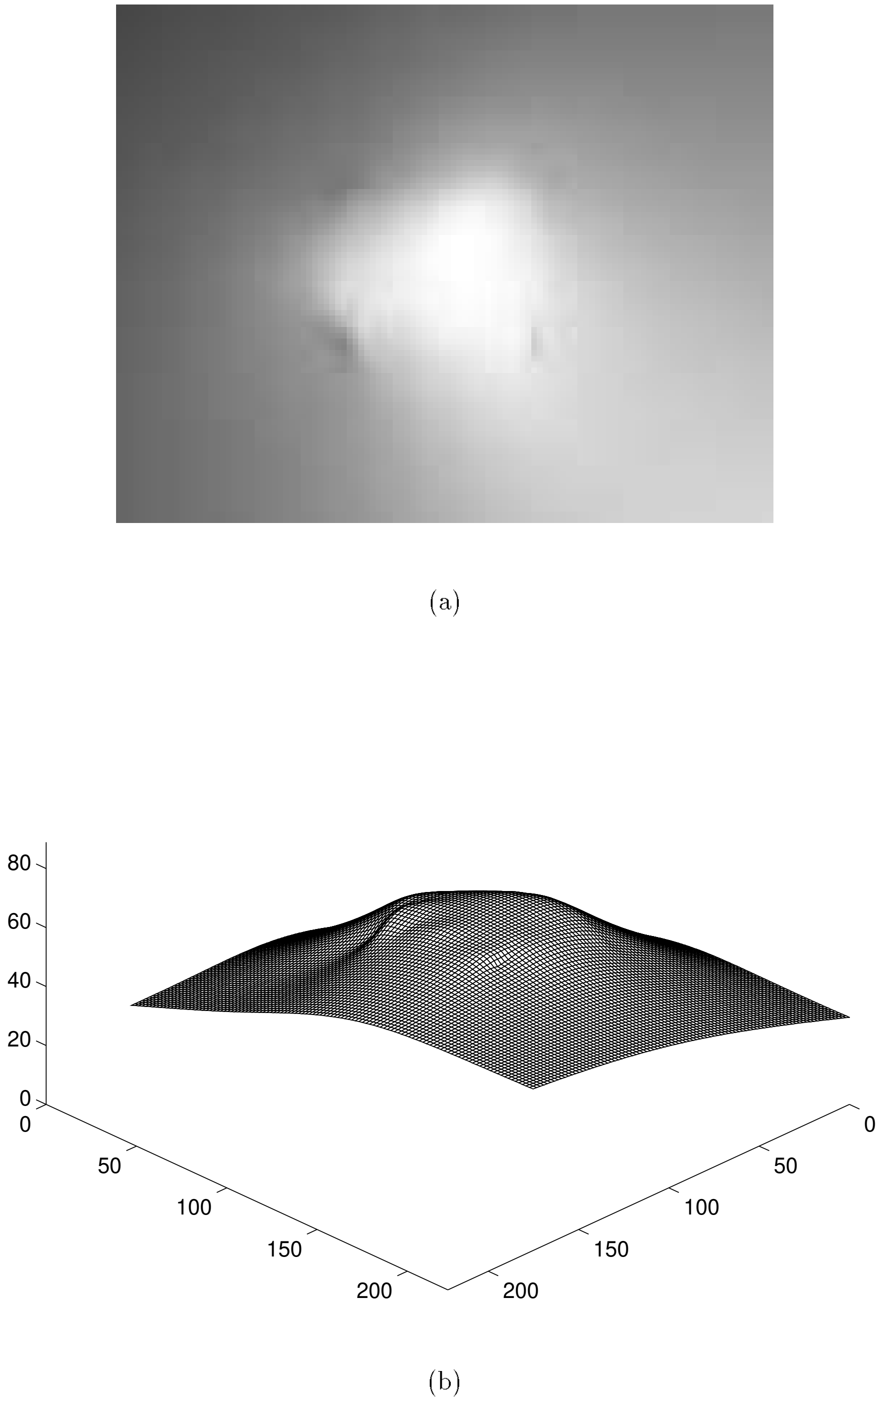

Figure 6.6 Behaviour of interpolant in Figure 6.4(b) far from the data points. (a) Rendered view of fitted thin-plate surface. (b) Mesh plot of thin-plate surface.

hull of the interpolation nodes. Figure 6.6(a) is a rendered view taken from above the surface, while Figure 6.6(b) is the corresponding depth-map. Apparent is the planar component of the solution which is described by the linear polynomial term in (6.1).

The tendency of the interpolant to a flat plane outside the interpolation nodes is an ideal characteristic in this application since cranial implants are made from flat titanium plate which is pressed into a mould in a hydraulic press. The smooth convergence of the mould surface to a flat plane reduces the likelihood of thinning or tearing of the metal at the periphery as it is drawn in to the mould.

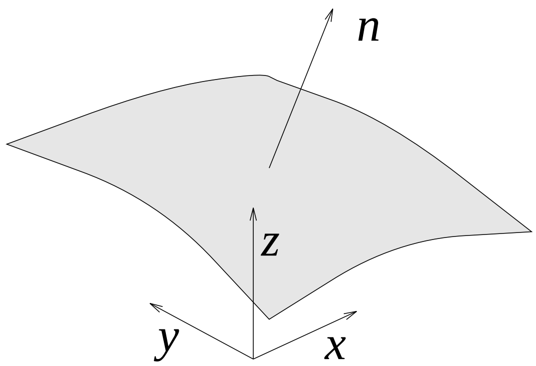

Figure 6.7 Depth-map coordinate system.

It is desirable to orient the mould surface so that the planar component is flat with respect to the plane of the press. The normal of the planar component of the fitted surface (Figure 6.7) is derived directly from the coefficients of the polynomial term $p_1(\mathbf{x})$,

$$\mathbf{n} = (-c_1, -c_2, 1). \tag{6.13}$$

The rotation angles about the x and y axes, denoted by θ and ψ respectively, required to re-orient the interpolated depth-map so that the normal aligns with the z axis (viewing axis) are given by

$$\theta = \tan^{-1}\left(\frac{-c_2}{\sqrt{1+c_1^2}}\right) \qquad \text{and} \qquad \psi = \tan^{-1} c_1 \tag{6.14}$$

A surface can be re-oriented in the CADCAM package which produces the cutting instructions used to mill the mould for the implant. Normally, corrections to the orientation of the surface are small provided an appropriate view of the defect was initially rendered. Should re-orientation of the surface result in a multi-valued surface with respect to the orientation of the mill, then ray-casting can be repeated with an improved viewpoint determined by the orientation angles θ and ψ.

Once the surface has been evaluated over the desired domain and at the desired resolution, the resulting depth-map is exported in IGES format [IGES 1993] to the CADCAM software package *Surfcam* [Surfcam 1993]. IGES, the Initial Graphics Exchange Specification, defines a recognised neutral data format that allows for the digital exchange of information among computer-aided design systems.

The depth-map is imported as a mesh by Surfcam. Surfcam produces a tool path and cutting instructions to drive the CNC mill. The tool type and dimension are specified along with a tolerance for the desired finish of the cut. Surfcam also provides a gouge avoidance facility which prevents the erroneous removal of material. Cavities which cannot be milled due to the size of the cutting tool are left uncut. CNC code is produced from the tool path and passed to the mill controller which drives the mill. The mill used at Christchurch Hospital is a Wintec MV-40 mill equipped with a Fanuc series O-M controller. The cutting tool usually employed for this work has a hemispherical head. Cusps inevitably arise between cuts and the tolerance specified determines the magnitude of the resulting surface ripple. A rough cut is sometimes performed with a 10mm diameter cutting tool followed by a finer cut with an 8mm diameter tool. In theory a ripple less than 0.01mm is possible to achieve. However, in practice 1mm between cutting tool paths is used. This leaves a sufficiently small cusp to be removed by sand paper so that a smooth surface finish is achieved.

6.5 EXAMPLES

Two examples are presented to illustrate the application of radial basis interpolation to real CT data of skull defects requiring repair.

The first example involves the repair of a very large (150 cm^2) hole in a basically convex region of the skull. Figure 6.8(a) is a rendered view of the CT data set which consists of 38 slices at 3 mm spacing with a pixel size of 0.47 mm. The rendered threshold corresponds to a Hounsfield number of 500. Figure 6.8(b) is a detailed view of the defect with the support region identified by the user highlighted. Figure 6.8(c) is a mesh-plot of the support shown in Figure 6.8(b). The mesh resolution is approximately 1 mm. This is the data to which a surface is fitted. Figure 6.9(a) is a rendered view from directly above the thin-plate spline surface fitted to the data in Figure 6.8(c). Figure 6.9(b) is a mesh plot of the fitted surface viewed from the same direction as Figure 6.8(c).

The mould produced from the interpolated surface data is shown in Figure 6.10(a) next to the finished titanium plate. CNC instructions were generated from the surface data to machine the mould from a very hard epoxy-resin. In Figure 6.10(b) the finished plate is shown beside a model of the defect area which has been machined from hard plastic using depth-map data shown in Figure 6.8. The plate has been slotted to allow fine adjustments to be made to the margins at the time of the operation. The peripheral holes in the plate, for screw fixation, were determined taking into account the skull thickness. In the general case, the position of the underlying air and venous sinuses would also be considered. The plate was anodised prior to operative insertion to provide an enriched oxide surface layer.

The second example involves the repair of a hole close to the orbital margin and other regions of high curvature such as the zygomatic arch. Consequently, a smooth concave mould could not be produced from a mirror-image depth-map of the unaffected side. Figure 6.11(a) is a rendered view of the CT data set which consists of 47 slices at 3 mm spacing with a pixel size of 0.89 mm. The rendered threshold corresponds to a Hounsfield number of 500. Figure 6.11(b) is a detailed view of the defect with the support region identified by the user highlighted. Figure 6.11(c) is a mesh-plot of the support shown in Figure 6.11(b). The mesh resolution is approximately 1 mm. Figure 6.12(a) is a rendered view from directly above the thin-plate spline surface fitted to the data in Figure 6.11(c). Figure 6.12(b) is a mesh plot of the fitted surface viewed from the same direction as Figure 6.11(c). The surface continues smoothly outside the support of the defect region, which is essential for pressing titanium into the mould without tearing.

6.6 COMPUTATIONAL CONSIDERATIONS

Radial basis functions have previously been proposed [Nielson 1993] as interpolants for 3D volume data where the interpolation centres do not lie on a regular grid. However, Nielson notes that conventional methods limit the application of radial basis functions to data sets of moderate size (n=300 to 500). Radial basis functions have also been used in electroencephalography to interpolate potentials measured on the surface of the head to produce topographical maps [Perrin *et al.* 1987a, Perrin *et al.* 1987b, Satherley 1994]. However, the number of electrodes (interpolation nodes) used in this application are typically fewer than 40.

Generating a radial basis interpolant involves solving a linear system like (6.5) for $(\lambda|c)^T$. Since the ϕ's used are typically non-zero, and growing away from zero, the matrices of these systems are neither sparse nor do they display any of the special structures such as bandedness or positive definiteness commonly exploited to enable fast solution of linear systems. Thus solution of these systems by direct, or even simple iterative,

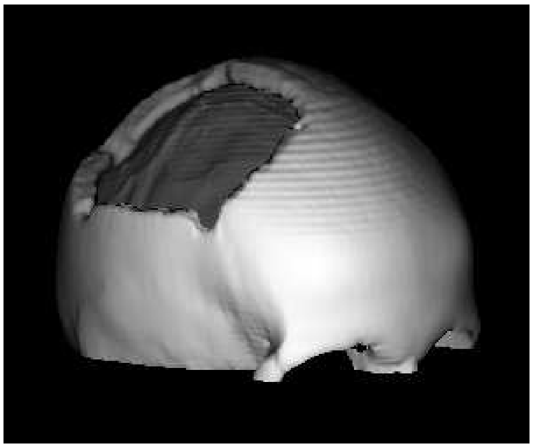

(a)

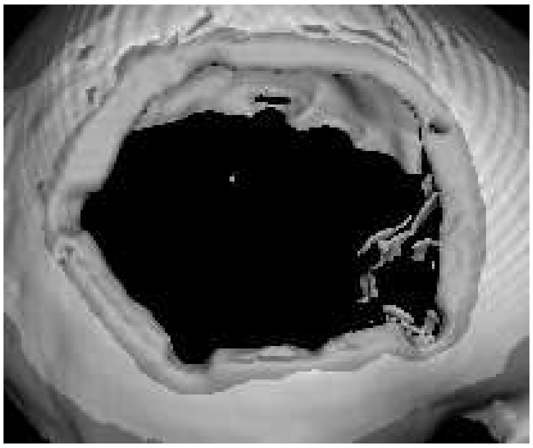

(b)

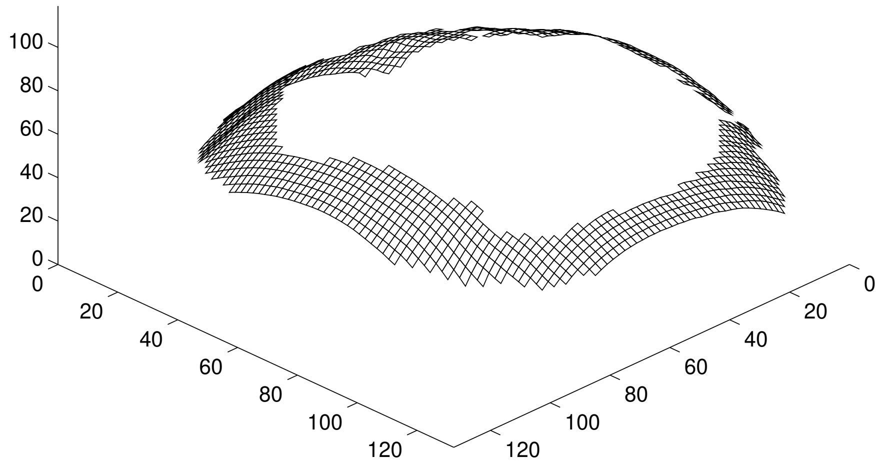

(c)

Figure 6.8 Example of fitting a surface to a cranial defect. (a) Rendered view of full CT data set. (b) Detail of the defect region with support region for surface highlighted. (c) Depth-map corresponding to (b).

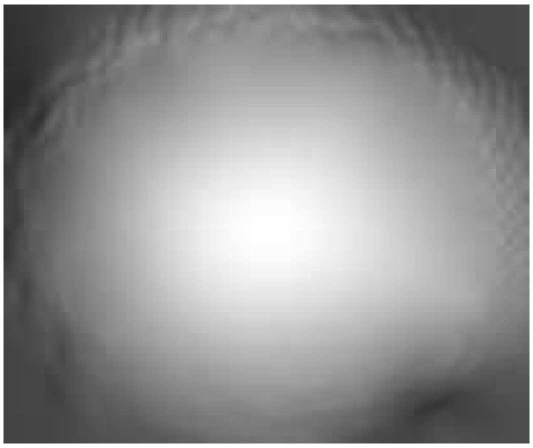

(a)

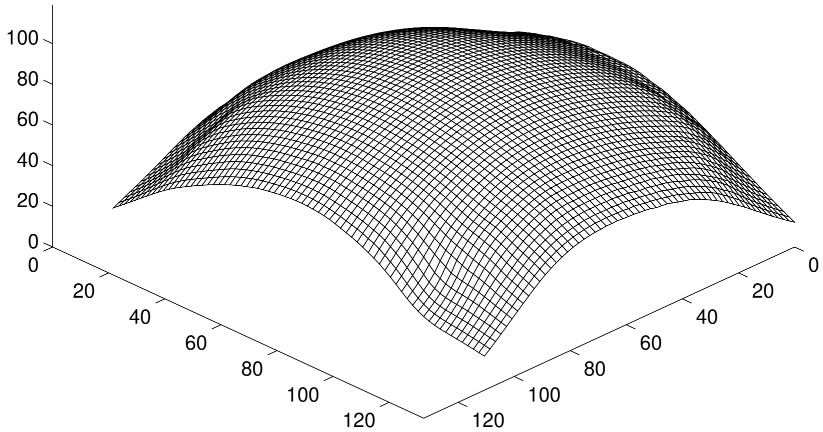

(b)

Figure 6.9 (a) Rendered view of surface fitted to depth-map in Figure 6.8(c). (b) Depth-map of fitted surface.

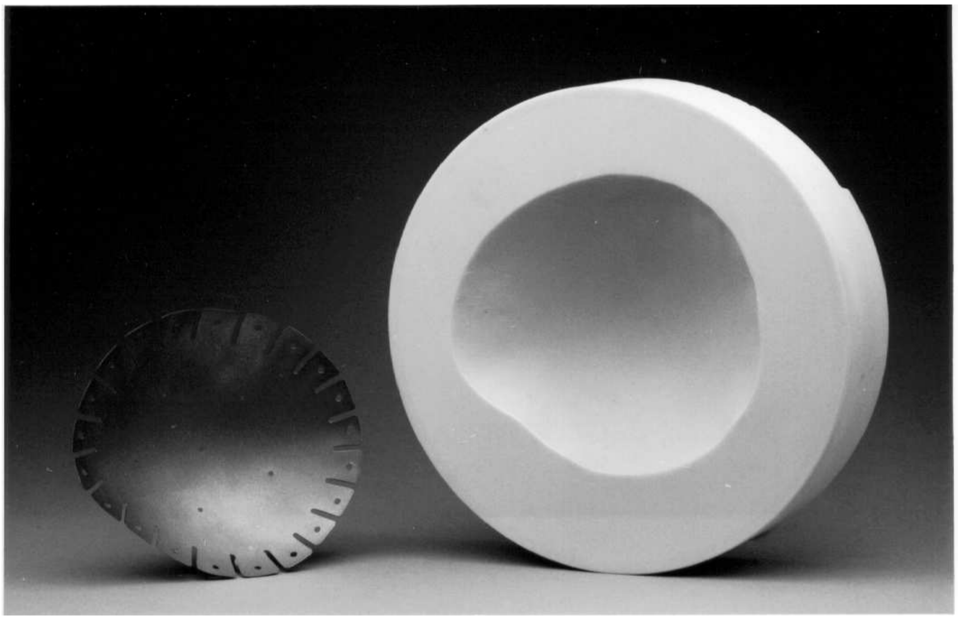

(a)

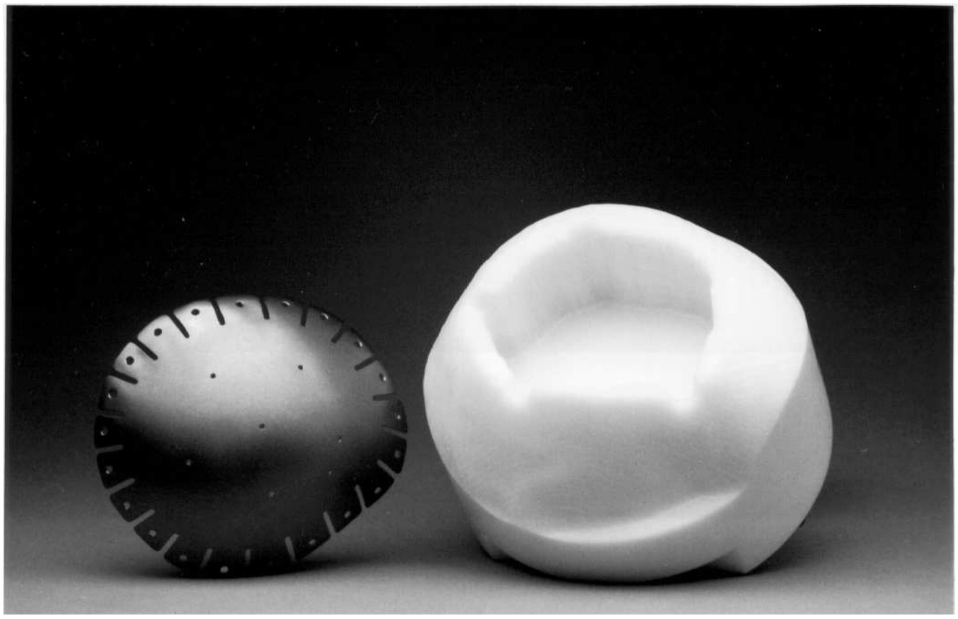

(b)

Figure 6.10 (a) Finished plate with the mould used to press the titanium plate. The mould was machined from epoxy-resin using interpolated surface data shown in Figure 6.9. (b) Finished plate with the plastic model of the defect area machined using depth-map data shown in Figure 6.8.

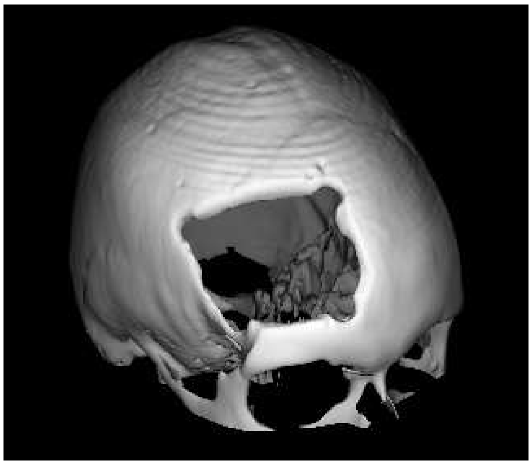

(a)

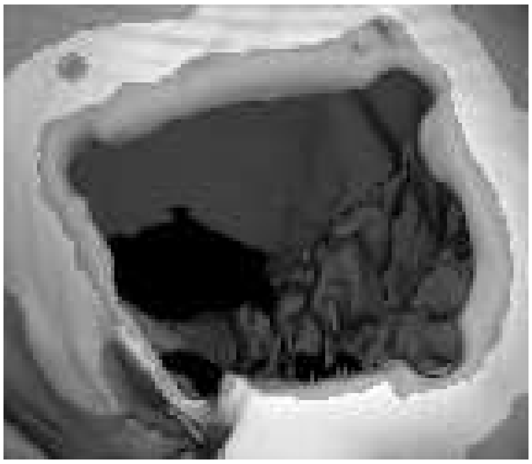

(b)

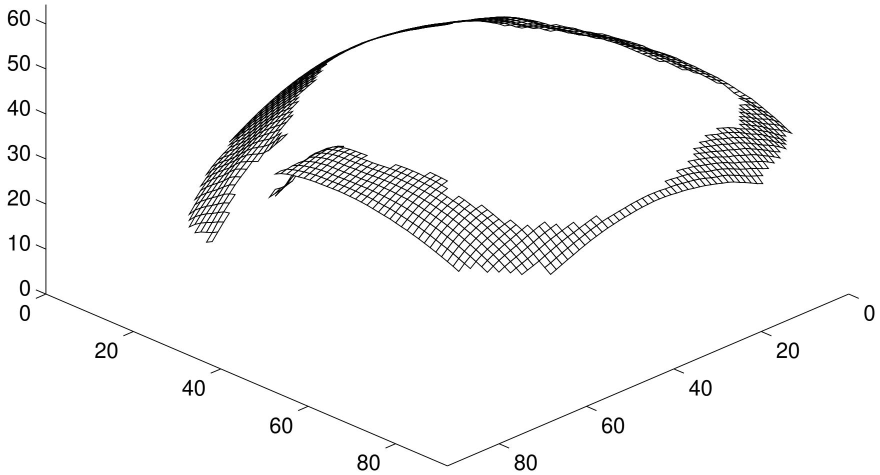

(c)

Figure 6.11 Example of fitting a surface to a cranial defect. (a) Rendered view of full CT data set. (b) Detail of the defect region with support region for surface highlighted. (c) Depth-map corresponding to (b).

(a)

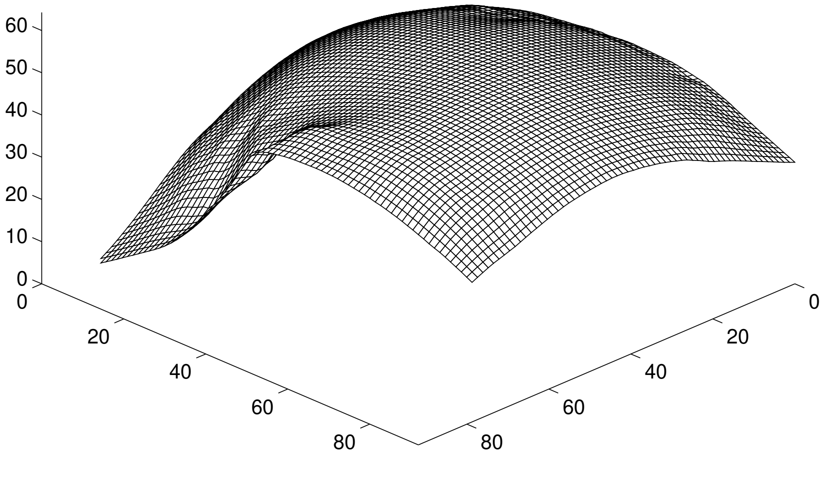

(b)

Figure 6.12 (a) Rendered view of fitted surface in Figure 6.11. (b) Depth-map of fitted surface.

Spline of	205 × 205 evaluation grid		512 × 512 evaluation grid	
	direct	fast	direct	fast
Figure 6.9	176s	4.34s	1103s	9.53s
Figure 6.12	146s	4.35s	910s	9.33s

Table 6.2 Solution times for fast and direct evaluation of thin-plate spline interpolants.

methods requires $\mathcal{O}(n^3)$ operations and $\mathcal{O}(n^2)$ storage. Taking advantage of the symmetry of the system will halve the usual operations count for Gaussian elimination but even symmetric solvers will still be impractical when n is large. These requirements used to be a considerable obstacle to the use of radial basis functions when n is large. However, recently published methods now exist which allow the solution of the thin-plate spline interpolation equations with $\mathcal{O}(n)$ storage and $\mathcal{O}(n^2 \log n)$ operations. These can be found in Beatson and Powell [1994] and the references there. Furthermore, improved methods under development by M.J.D. Powell and others promise solution in $\mathcal{O}(n \log n)$ operations.

The current calculations, however, involve moderately sized data sets and conventional techniques are satisfactory for solving the corresponding linear systems. The 697 node interpolation problem underlying Figure 6.9 required approximately 40.5 seconds of CPU time to solve on a SUN Sparc 2, while the 578 node interpolation problem underlying Figure 6.12 required approximately 23.5 seconds of CPU time to solve on the same machine.

Once a radial basis function has been fitted to the data there is the problem of evaluating it. Due to the behaviour of ϕ, the work required for direct evaluation of s at a single extra point, x, is proportional to the number of nodes n. Thus, the amount of computation required for direct evaluation of s on a fine mesh becomes significant, even for moderate node counts n. However, recently developed fast evaluation methods (see Beatson and Newsam [1992] and Beatson and Light [1994]) exploit hierarchical data structures and series expansions, to reduce the incremental cost of one evaluation of a thin-plate spline to within precision ϵ, to $\mathcal{O}(1 + |\log \epsilon|)$ operations. Of course the setup time before evaluation does depend on the node count, nodal positions, the values of the spline coefficients, and the desired accuracy.

Table 6.2 shows times in seconds for evaluation tasks associated with Figures 6.9 and 6.12. The improvement to be had by using a fast evaluation code is dramatic. All computations were carried out in double precision on a SUN Sparc 2 machine, and the fast evaluator was required to evaluate the fitted surface with an infinity norm relative error, ϵ, of less than 10^{-4}. In these examples, this meant an accuracy better than .01mm for the fast-evaluated surface.

6.7 DISCUSSION

6.7.1 CT data artifacts

Sometimes reconstruction artifacts, for example due to dental fillings, occur within a CT data slice and affect the rendering of bone surfaces from a set of slices. Occasionally a slice may be misaligned due to patient movement during scanning or even missing if the set of scans were acquired in two stages. In this case it is not always possible to repeat the examination, due to the desire to minimise the patient's exposure to ionising radiation. One of the benefits of the approach presented here is that the user may arbitrarily discard any data from the depth map whether it be due to the presence of an actual defect or an artifact.

6.7.2 Interactive surface manipulation

In the current software, depth-map data can be manually edited, thereby allowing the user to further constrain the shape of the surface by adding more interpolation nodes. However, this process is tedious and involves resolving the interpolation matrix. Future development of the software will provide the surgeon with an interactive facility for sculpting the surface via adjustment of the depth at user specified "control points". For example, restoring the natural shape of the skull across a large hole may lead to a large cavity forming between the brain and the prosthesis. The user may want to indent an initial depth-map in the defect region to avoid such cavities. Similarly, the surface data around the periphery of the hole may not be sufficient to restore the appropriate shape to the skull. Prerequisites for the new interactive software are an efficient means for updating an initial matrix factorisation as interpolation nodes are added and deleted, and also a fast means of recalculating the resulting surface. Fortunately, algorithms for both tasks are available.

One approach for the first task is the matrix updating method which is amongst the things discussed in Beatson and Powell [1994]. They first change the interpolation problem into an equivalent problem, with symmetric positive definite matrix, and then consider means for updating the Cholesky factorisation of this matrix as the nodes of interpolation change. In the current context this would enable solving the system with a few extra nodes of interpolation at the cost of $\mathcal{O}(n^2)$ extra operations, rather than the $\mathcal{O}(n^3)$ operations required if one started from scratch. An alternative method is available if one is prepared to specify *a priori* the positions of the possible extra nodes, but not the values of f at them. Essentially it involves deducing the solution of the original unsupplemented system without extra nodes by means of solving a slightly larger system for several right-hand sides. This provides initial values at the extra nodes which can then be fine tuned by the user to achieve the desired shape. This method is restricted to situations where one can nominate the extra nodes *a priori*, making it less flexible than the more sophisticated method of Beatson and Powell [1994]. However, restricting where the user can 'pull' and 'push' the surface may be acceptable and might even encourage a methodical approach to the design of the surface. A suitable method for the second task is the fast evaluation code described in section 6.6.

6.8 CONCLUSION

Surface interpolation has a role to play in computer-assisted surgical planning through the reconstruction of missing surfaces and the prediction of soft tissue movement. In this chapter the application of radial basis functions to cranial implant design has been demonstrated. Radial basis function approximation is suited to this type of problem due to the few constraints placed upon the geometry of the nodes of interpolation. Few alternative techniques are available when the data do not lie on a complete regular grid. The thin-plate spline is a particularly appropriate choice since it fits a smooth C^1 surface to the data and is the smoothest C^1 interpolant in the sense that it minimises the energy functional (6.12). The property that the interpolant tends to a flat plate far from the interpolation centres is ideally suited to the manufacture of titanium prostheses in a hydraulic press. This new approach to the design and manufacture of titanium prostheses reduces the previous reliance on one skilled person and should give the surgeon greater influence over the final design of an implant. Prototyping in software reduces the time and expense of the procedure while ensuring the accurate fit and appearance of the implant. This will in turn reduce the overall period of treatment experienced by the patient.

Radial basis function approximation has wider significance and there are many potential applications. In particular, radial basis functions are well suited to 2D and 3D image warping problems. Image warping is relevant to the registration of data sets which are not related by a simple geometric transformation. In a medical imaging context, this situation may arise when registering scans taken over a period of time, or when registering scans of anatomy taken with different imaging modalities. Similarly, predicting the movement of soft tissue in cranio-facial surgery when underlying bone structures are altered is another example. These problems are characterised by the matching of common landmarks in two data sets. The interpolation problem is to find the function which maps one data set on to the other, given the known movements of the landmarks. The latter do not usually lie on a regular grid. This problem is also of interest in computer graphics and animation. Current 'morphing' techniques, as they are known, involve specifying piecewise interpolants over user-defined meshes. The ability to arbitrarily specify warping centres without having to specify a mesh would be a significant advantage.

Figure 6.13 illustrates the warping of a 2D image. Movements were specified at 12 points in Figure 6.13(a), four on each face and four in the periphery of the image. The vector field was then interpolated using the thin-plate spline basis and Figure 6.13(b) was evaluated by finding where each pixel mapped to in the original image. Both Figures 6.13(a) and 6.13(b) are corresponding portions of the original and warped images.

(a) (b)

Figure 6.13 Example of 2D image warping using a thin-plate spline radial basis interpolant. (a) Subsection of the original image, (b) corresponding subsection of warped image.

Radial basis functions also have relevance to the visualisation of scattered 3D data, provided the data is noise-free and the underlying distribution is smooth and continuous. In 3D interpolation problems the number of data points is typically large (10^3–10^6), consequently the fast solution and evaluation methods described in section 6.6 are essential. In theory, a function could be fitted to the scattered data and evaluated on a regular grid, thus allowing conventional ray-casting methods to be applied.

Chapter 7

CONCLUSIONS

This thesis has addressed two problems in medical imaging, the development of a system for 3D imaging with ultrasound and a system for making titanium prostheses for cranioplasty. Central to both problems is the construction and depiction of surfaces from volume data where the data is not acquired on a regular grid or is incomplete.

7.1 3D IMAGING WITH ULTRASOUND

7.1.1 Data acquisition

The development of systems for 3D imaging with ultrasound is an area of active research in medical imaging. The motivation for developing 3D ultrasound is to assist the visualisation of complicated anatomy and to allow quantitative measurements concerning organ size and volume to be made. Additional benefits are the ability to obtain 2D slices from orientations previously not possible and the potential to improve images through spatial compounding.

Chapter 2 described the construction of a system to acquire 3D ultrasonic data. 3D volume data is acquired as a set of registered 2D B-scans with the aid of an electro-magnetic spatial locator. This approach has been adopted by a number of researchers. It allows conventional ultrasound equipment to be upgraded to provide a 3D imaging capability, thus avoiding the development and purchase of a new scanner. The use of an electro-magnetic spatial locator maintains the freedom and convenience of a hand held probe, but the arbitrary orientation of acquired slices requires the development of new reconstruction and visualisation algorithms.

The resolution of the imaging system is currently limited by the accuracy of the spatial locator rather than the scanner resolution. Probe movement during acquisition and errors in the synchronisation of position measurements with the acquisition of scan planes also affects the registration of scan planes in the reconstructed volume and degrades the resolution. These error sources limit the velocity with which the probe can be moved. However, because the system monitors the time and amount of movement between acquired slices, excessive probe movement can be detected and error-prone scans rejected. In practice, the scanning velocities of the current system are adequate. An asynchronous acquisition mode, which acquires frames when significant probe movement is detected, allows an approximately uniform sampling of the volume to be achieved independent of variations in the operator's scanning motion.

Determining the exact nature of inaccuracies in the position and orientation measurements made by the Polhemus spatial locator requires a precision mechanical positioning system. Experience with the Polhemus locator has indicated that the primary component in position and orientation variations is not random in nature, as assumed by the analysis in Section 2.2.2.2, but is better modelled as a distortion of the true measurement. Consequently, the imaging system has exceeded the performance predicted by the analysis in Section 2.2.2.2. Reliable reconstructions have been achieved when the probe is restricted to smooth sweeps where consecutive acquired slices are near-parallel. This is because reconstructions from probe sweeps rely on accurate relative measurements over small regions whereas reconstructions from completely arbitrary slices requires accurate absolute measurements.

7.1.2 Surfaces from ultrasound data

Although a system for acquiring 3D ultrasound data has been independently developed in the course of this research, the main contribution of this thesis has been to study the problem of constructing surfaces from ultrasound data. Rather than using semi-transparent rendering methods, this thesis has focused on the task of rendering opaque surfaces from ultrasound data because these are of greater clinical interest and necessary for making quantitative volume measurements.

Three methods for forming 3D graphics from arbitrary slices have been developed. The first of these, the swept-volume method, forms a regular voxel array from the slice data and is described in Section 4.2.1. This representation is suitable for conventional ray-tracing algorithms. The method allows incremental construction of the volume as slices are acquired, provided that consecutive slices are acquired from a smooth sweep of the probe. This is compatible with the operating constraints of the current spatial locator and in practice corresponds to the natural manner of scanning. Multiple sweeps can be compounded to image a larger volume. Averaging multiple contributions to voxels improves the signal to noise ratio of the data. 2D scan slices are filtered prior to reconstruction on typically lower resolution voxel grids. 3D filtering of the volume data is then interactively applied in the visualisation process. Non-linear filtering techniques, in particular, median and rank filters, are used. Although more sophisticated methods have been developed, these filters are recognised as standard speckle reduction techniques in the literature. They are simple to implement, computationally efficient and general-purpose. Ray-tracing is then used to visualise iso-value surfaces from the filtered volume reconstructions.

The second visualisation method renders iso-valued surfaces directly from arbitrary oriented slices without forming an intermediate voxel reconstruction. Unlike the previous method, slices need not be in any particular spatial or temporal order. The method is efficient in terms of computation, but the quality of images formed is dependent on the viewing direction because interpolation along viewing rays is used to determine threshold transitions. Slices oriented at steep, oblique angles to the viewer pose particular problems. However, this directional viewing dependence is minimised when the acquired scans are close together. Because an intermediate volume representation is not formed, spatial compounding of scans and 3D filtering is not possible. This means that iso-value surfaces are more susceptible to noise in the scan data. A further consequence is that volume gradient shading is not possible and surfaces are depicted by depth shading and/or Z-buffer gradient shading.

This second visualisation method was developed to provide real-time 3D graphics from slices as they are acquired. However, in practice, real-time rendering has been found to be of limited value due to the difficulty of choosing viewing parameters prior to acquisition. For example, choosing an appropriate threshold is an interactive process, similarly choosing cut-away viewing planes and applying speckle reduction filters, which are essential for visualising structures, requires some user interaction.

The third method is discussed in Section 7.1.2.2.

7.1.2.1 Artifacts

Rendering iso-valued surfaces has been successful for visualising fluid-filled cavities such as the hepatic veins, portal vein, gall bladder and urinary bladder. Such structures are relatively easy to segment using thresholding techniques because they can (ideally) be distinguished from surrounding backscatter by their absence of echoes. However, regions where echoes are absent due to signal attenuation, shadow and refraction are also depicted by the thresholding process and may be erroneously interpreted as anatomic structures. Alternatively, fluid-filled cavities may not result in completely echo-free regions due to reverberation and resonance artifacts. Consequently, the surfaces of genuine structures may be falsely depicted. For example, the surface of the urinary bladder depicted in Figure 4.24 suffers from reverberation artifact due to the bladder wall.

Reverberation, refraction and attenuation artifacts are readily recognised in conventional 2D scans. In 3D surface displays, knowledge of the primary direction of insonification and the underlying anatomy is required to accurately distinguish genuine structures from artifact. Unfamiliarity with 3D images and the unusual viewing directions which are possible with 3D reconstructions can make interpretation more difficult than in 2D scans. If distinguishing artifact from genuine structure proves to be a difficult skill to acquire, then the usefulness of such 3D displays may be restricted to quantitative measurements of known anatomical structures rather than the detection of abnormal structures.

7.1.2.2 Shape reconstruction from ultrasound data

Many tissues of interest are not distinguished by large homogeneous grey-level regions and consequently can not be robustly identified by thresholding. Pulse-echo ultrasound primarily provides information concerning changes in acoustic impedance, and in general thresholding tends to reveal incomplete, noisy boundaries between tissues. The third reconstruction method, developed in Chapter 5, considered the problem of reconstructing incomplete tissue interfaces.

The method addresses specific aspects associated with fitting closed surfaces to medical ultrasound data. These include the non-uniform distribution of the data, the noisy nature of tissue interfaces and the fact that the data may not be exclusively associated with the surface of an organ. The ultrasound problem is viewed as an example of the general problem of reconstructing an object's shape from unorganised surface data. The method was developed with reconstructions from arbitrary, unrelated slices in mind. Consequently, ultrasound data are projected on to a regular grid, but not interpolated as with the swept-volume reconstruction. The grid is therefore incomplete and contains holes. No use is made of the slice nature of the ultrasound data and no *a priori* knowledge concerning the shapes to be reconstructed is assumed.

The algorithms developed in Chapter 5 are most closely related to work in computer vision concerning the determination of the shape of 3D objects from their silhouettes and the recognition of known 3D objects from their silhouettes. This thesis extends this work significantly by considering the reconstruction of incomplete objects contaminated by noise and by developing practical algorithms for pixel and voxel data. Furthermore, the reconstruction of non-convex, real objects is considered rather than simple geometric objects. 2D and 3D ultrasound data derived from phantoms, as well as artificial data, are used to demonstrate reconstructions.

An infinite number of possible objects may be consistent with a given set of incomplete object data. Consequently the shape reconstruction problem, as stated in Chapter 5, is ill-defined. Ambiguities arise when distinguishing between surface and non-surface data and when deciding whether a gap in a surface indicates missing data or the presence of a cavity in the object. Because data within the convex hull of an object may not be surface data, preserving cavities in the object's shape could conflict with the requirement that the reconstruction be immune to non-surface object data.

The bias of the basic method developed in this thesis is towards interpolating across gaps and therefore filling in cavities.

The specification of an appropriate reconstruction space in Section 5.7 overcomes the bias towards filling in cavities. Although this can be seen as supplying *a priori* knowledge concerning the shape to be reconstructed, it may be acceptable in the development of an interactive tool. Furthermore, the manual identification of a reconstruction space will usually be required with clinical data to ensure that the reconstruction space mainly contains data associated with a single object. Alternatively, the reconstruction of cavities can be assisted by defining a local reconstruction space. This avoids the specification of explicit information concerning the final shape and the method acquires characteristics similar to that of a morphological operator. Small clusters of object data, up to the size of the local reconstruction space (analogous to a morphological *structuring element*), are initially formed and then, through subsequent iterations, join together. In general, cavities which are smaller than the reconstruction space will not be preserved, depending on the data and the threshold chosen.

In the absence of any additional information, the problem of preserving cavities is an intrinsic one. The method developed in Chapter 5 can reconstruct objects with moderate concavities and at worst fills in surface cavities. For reliable performance the reconstruction of complex shapes is best avoided.

7.2 TITANIUM CRANIOPLASTY

In Chapter 6 surface interpolation is used to create an automated process for the design and manufacture of cranial implants. The system improves upon current treatment procedures by avoiding the manual aspects of fashioning an implant. It is also suitable when other techniques which use symmetry to reconstruct the skull (by reflecting the shape of sound bone opposite the defect on to the defect area) are inadequate or not possible. The system has been successfully used to treat patients at Christchurch Hospital.

The main contribution of this thesis has been the use of radial basis functions to interpolate incomplete depth-maps of the skull surface. Radial basis functions are ideal for this problem because they impose few restrictions on the geometry of the interpolation centres, therefore allowing irregular holes in the data. In particular, the thin-plate spline basis produces a surface which is suited to the manufacture of a titanium prosthesis in a hydraulic press. The interpolant is the smoothest interpolant of the data in the sense that it minimises the sum of the second derivative over $\mathbb{R}^2$, Equation (6.12), and it converges to a flat plane far from the interpolation centres.

The utilisation of new fast evaluation methods in Chapter 6 has made an interactive system for designing surfaces possible and demonstrated that RBFs may be of wider interest in medical imaging. Previously RBFs have only been considered for small problems such as the interpolation of EEG potentials where the number of interpolation centres (electrodes) is typically less than forty. Fast evaluation methods make the use of RBFs in large systems, e.g. > 1000 points, practical. Such large systems are typical of 3D imaging problems.

7.2.1 B-spline ray-tracer

In Chapter 3 tensor product B-splines were introduced and applied to the problem of rendering bone surfaces from CT data. A smooth piecewise continuous function exhibiting C^2 continuity was fitted to the CT data distribution to be visualised. Consequently, iso-valued surfaces rendered from the data also exhibited C^2 continuity. The B-spline ray-tracer is applicable to any scalar volume data which can be mapped on to a regular grid. Fitting a function to the volume data means that volume gradients, essential for surface shading, can be determined analytically.

When rendering bone surfaces from CT data, the B-spline interpolant reduces ripples due to the relatively large slice spacing between CT data slices, and improves upon current tri-linear and shape-based interpolation techniques [Herman *et al.* 1992]. The B-spline, however, is a quasi-interpolant and does not generally reproduce the CT data values specified at the voxel centres. For this reason the use of B-splines might be seen as controversial. However, the CT densities represent averages over small volumes rather exact values at voxel centres. When rendering cranial bone surfaces, the B-spline surfaces have been found to differ from the tri-linear surfaces primarily in the absence of ripples. Because it is known that the underlying bone is smooth and does not contain these ripples, the B-spline is regarded as a better representation of the true surface. However, the assumption that the CT distribution is smooth and continuous is not always true, therefore errors are expected when the actual distribution is discontinuous. In practice, little significant difference between surfaces rendered by different methods in the vicinity of discontinuities has been observed. The titanium cranioplasty application is ideal in this respect since regions where bone surfaces are subject to error are discarded because they are usually associated with defects (fractures in the skull) and the RBF interpolant is only fitted to the smooth sound bone surrounding a defect.

7.3 FUTURE WORK

7.3.1 3D ultrasound imaging

The ultrasound imaging system described in Chapter 2 is susceptible to tissue movement during data acquisition. Movement results in misregistration of common tissue structures in the reconstructed volume. The main sources of tissue movement are respiration and circulation. Reducing the time to sample a volume is an obvious solution. However, ignoring the limitations of the current hardware, the time to sample a given volume at a particular resolution is physically limited by speed of sound in tissue. This is because the volume is imaged by a single focused beam. For a given penetration depth, there is an associated pulse travel time and minimum wait period before the next pulse can be sent. Thus, the number of lines which can be imaged in a given time is limited [Wells 1977]. Acquisition at the maximum scanning frame-rate will not allow tissues affected by respiration and circulation to be 'frozen' over any reasonable sized volume. Researchers imaging the heart have reduced movement artifact in reconstructions from 3D probes by synchronising data acquisition to the end-expiration of the respiratory cycle and to the cardiac cycle, via the electrocardiogram (ECG) [Martin *et al.* 1989, Belohlavek *et al.* 1991, McCann *et al.* 1988]. In the studies presented in Chapter 4, the subjects held their breath to reduce diaphragm movement due to breathing. Synchronisation to the ECG will further reduce movement artifacts, particularly in reconstructions of structures such as the hepatic veins, which are noticeably affected by movement due to circulation (Figure 4.14).

The nature of ultrasound imaging is such that the data are relatively plentiful and tissue can be readily rescanned. Consequently, unlike many other imaging situations, data suspected of being affected by registration errors or artifacts can be readily discarded. Because certain ultrasound artifacts, such as shadowing and resonance, are dependent on the scanning direction, the ability to image tissue from a number of different orientations and directions suggests that it might be possible to remove some of these artifacts; tissue obscured in one view may be clearly imaged in another. The development of more sophisticated spatial compounding techniques is therefore a promising area for future work which this thesis has not explored. In the current system, the limited video output of the Acuson scanner prevents important scanner settings which affect grey-level images from being monitored or controlled. Access to signals at a lower level of processing in the formation of B-scan images is necessary in order to develop artifact reduction techniques based on a physical understanding of interference phenomena. Current research into speckle reduction methods for ultrasound highlights the need for a consistent quantitative assessment of speckle reduction methods. Consequently, the development of phantoms or other methods to assess and compare compounding techniques will be an important part of this work.

Conventional 2D ultrasound imaging is based upon assumptions which are not always strictly applicable. These include the speed of sound in tissue being constant and refraction being negligible. As a result, the shape of structures in conventional 2D images can be distorted. Consequently, 3D reconstructions from 2D images will also be distorted. Verification of *in vivo* imaging is therefore required to determine how accurate quantitative measurements from the 3D ultrasound system are. Verification of bladder volumes has been proposed in Chapter 4 as a means of achieving this since the volume of the bladder can be independently determined by emptying the contents. However, the accuracy of volume measurements is not a function of the acquisition hardware alone. Because refraction, variations in the speed of sound and other artifacts are tissue dependent, the accuracy of volume measurements will vary with the anatomy being imaged. Each application will need to be assessed separately.

The images produced by the current system are dependent on a number of processes including filtering, segmentation and rendering. Each of these processes and their respective parameters must be considered when interpreting images produced by the system. Developing scanning protocols, simplifying the interface and reducing the number of variables are important ergonomic issues which will affect the clinical usefulness of the system.

The principal limitations of the acquisition system are the accuracy of the Polhemus spatial locator and the lack of an additional electrical interface to the Acuson scanner. These are technology limitations which can be improved. The visualisation of ultrasound data, however, is a more difficult problem. Thresholding is not a robust segmentation technique and is suitable for anatomic structures which can be distinguished from their surrounds by a homogeneous grey-level distribution. The shape reconstruction method developed in Chapter 5 has considered the reconstruction of isolated objects. This has proved to be a difficult problem. However, without prior processing, clinical data will generally consist of many different tissue interfaces. An initial segmentation of the data is required before the method can be applied. The introduction of continuous-valued projections, the iterative scheme and the local reconstruction space complicates the understanding of the method and further analysis of the effects of these modifications is required.

Automatic surface identification is a very difficult task in ultrasound data, and is a skill mastered by human experts over a period of time. Even then, they often use knowledge of the underlying anatomy, artifacts and physics of ultrasound propagation in tissue to interpret grey-level images. Visualisation of ultrasound data shares many of the characteristics which make computer vision a difficult problem. Here the underlying goal also appears to be mimicking the human perceptual system. Consequently, general, fully automatic segmentation tools are unlikely to be achieved. However, the development of a set of interactive tools and methods which use a priori knowledge to segment particular structures are worthwhile pursuing. Despite practical difficulties with reconstructions from ultrasound data, the method developed in Chapter 5 is of academic and practical interest in the context of shape reconstruction from silhouettes in computer vision problems.

This thesis has focused on pulse-echo imaging and has not investigated other imaging modes such as Doppler and power imaging. The use of contrast agents has not been investigated either. In some applications the problem of visualising anatomy could be significantly simplified by using contrast agents [Balen *et al.* 1994]. For example, Balen *et al.* [1993] describes the successful use of contrast media to improve the definition of the uterine cavity in 3D ultrasound studies.

7.3.2 Radial basis function approximation

A feature of the surface interpolation problem in Chapter 6 was that the ray-traced skull data, to which the thin-plate spline interpolant was fitted, was smooth and relatively noise-free. The development of a quasi RBF interpolant which is not strictly constrained to pass through the interpolation nodes would be useful when interpolating noisy data on an irregular grid. Developing such a characterisation, along with improving the fast evaluation and solution methods, would make RBF approximation applicable in a large number of practical engineering problems. In the titanium cranioplasty work, the use of a quasi B-spline interpolant in the ray-tracer provided the smooth surface data for the interpolation problem.

Potential applications for RBF approximation include 2D and 3D image warping and the visualisation of 3D data not acquired on a regular grid. The former is particularly relevant to the problem of registration of medical data sets acquired over extended intervals of time, or by different modalities. The development of fast solution and evaluation methods means that application of RBFs in 3D medical imaging problems can now be investigated.

Appendix A

SPATIAL LOCATOR CALIBRATION

This procedure, referred to in chapter 2, determines the transform which relates points in the scan plane of an ultrasound probe to a Polhemus receiver mechanically mounted on the probe. The scan space (x_s, y_s, z_s) and receiver (x_r, y_r, z_r) coordinate systems are illustrated in Figure A.1. The $x_s - y_s$ scan plane is modelled as a flat slice having no thickness. The origin for the scan reference frame is the centre of the transducer face, which corresponds to the top-centre of the acquired scan image.

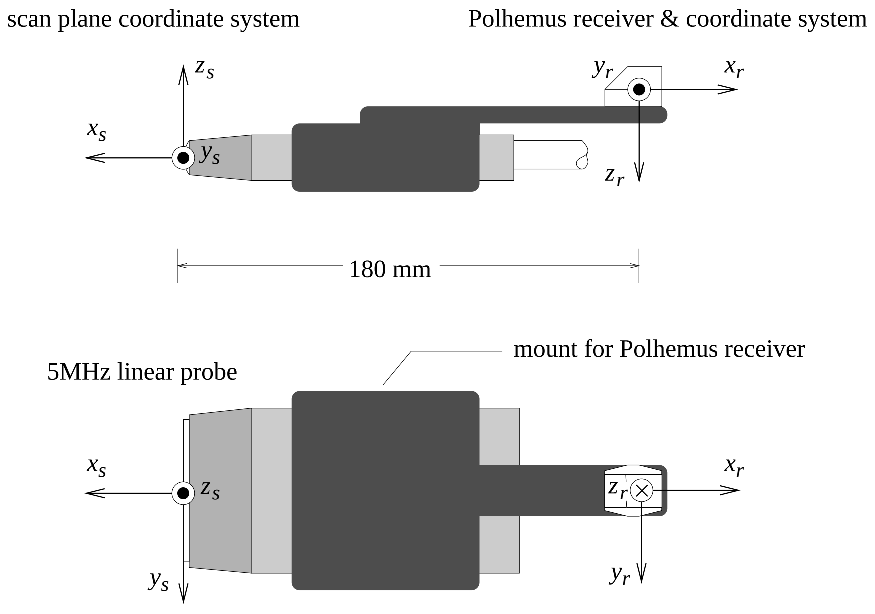

Figure A.1 5MHz probe with Polhemus receiver mounting.

Calibration involves finding the displacement and rotation of the scan's frame of reference from the receiver's frame, in order to form the matrix M_{sr} which transforms points in scan space to receiver space. The matrix parameters are determined in an order consistent with that employed by the 3Space unit, namely, rotation about receiver's z axis followed by rotation about the receiver's new y axis, then by rotation about the receiver's new x axis, and finally translation of the scan origin on to the receiver's origin.

A point $\mathbf{x_s}$ in scan space is transformed to a point $\mathbf{x_r}$ in receiver space by

$$\mathbf{x_r} = M_{sr}\mathbf{x_s} \tag{A.1}$$

where the scan coordinates have zero z component, $\mathbf{x_s} = (x_s, y_s, 0)$.

The complete transformation from $\mathbf{x_s}$ in scan space to $\mathbf{x_t}$ in transmitter (i.e. object) space is described by

$$\mathbf{x_t} = M_{rt}M_{sr}\mathbf{x_s} \tag{A.2}$$

where M_{rt} is the matrix to transform from receiver to transmitter space. This matrix is determined by the receiver's position and orientation parameters, output by the 3Space unit.

A simple phantom, made from perspex, containing three orthogonal nylon wires of 1mm diameter was constructed to determine the transform M_{sr} (Figure A.2). By locating the three wires in a series of scans taken at different locations and orientations about the water-filled phantom, a transformation between the scan plane and the receiver, consistent with the geometry of the wires, can be found.

For convenience, the origin of the phantom's frame of reference is chosen to be the convergence of the three wires, with the wires representing the respective axes, as annotated in Figure A.2. In order to find M_{sr} the transform M_{tp} which transforms points in transmitter space to phantom space must be found.

A point $\mathbf{x_t}$ in transmitter space is transformed to a point $\mathbf{x_p}$ in phantom space by

$$\mathbf{x_p} = M_{tp}\mathbf{x_t} \tag{A.3}$$

The complete transformation from $\mathbf{x_s}$ in scan space to $\mathbf{x_p}$ in phantom space is described by

$$\mathbf{x_p} = M_{tp}M_{rt}M_{sr}\mathbf{x_s} \tag{A.4}$$

where M_{sr} and M_{tp} are functions of six parameters each; three determine the amount of translation along each axis and three determine the angular rotation about each

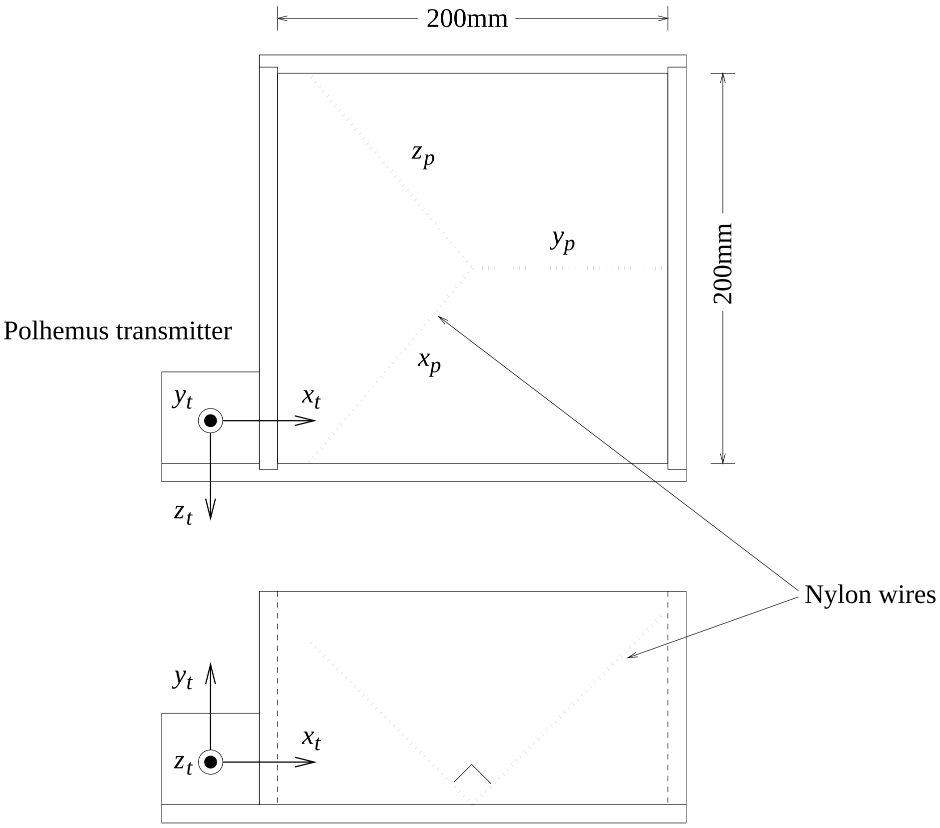

Figure A.2 Perspex phantom for receiver-scan plane offset calibration. The nylon wires, indicated by the dotted lines, correspond to the axes of the phantom coordinate system (x_p, y_p, z_p).

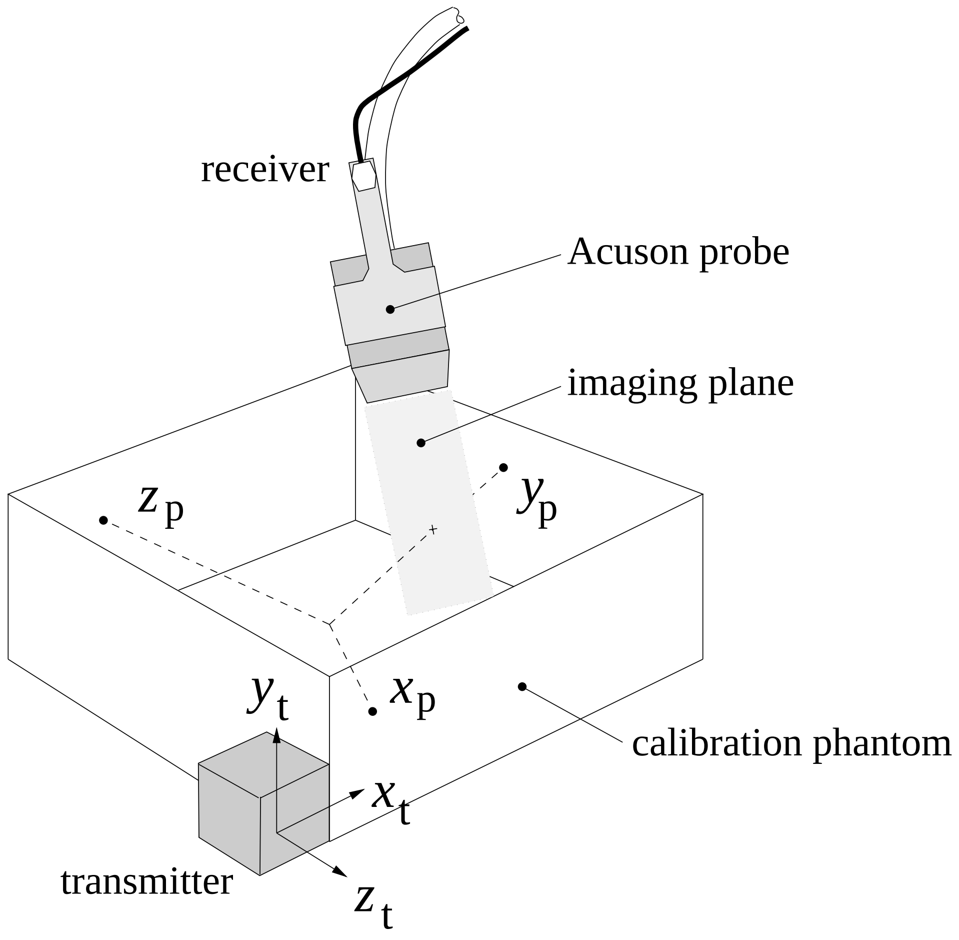

Figure A.3 Calibration procedure to determine the receiver-probe offset. At the intersection of the scan plane with the y_p-wire (axis) of the phantom, the x and z coordinates in phantom space are zero. Thus, each wire imaged constrains two degrees of freedom and two equations are added to the system (A.6) to determine the receiver-probe offset.

axis. It is convenient to write the total transform as a function of these 12 parameters in vector form,

$$M_{sp}(\mathbf{a}) = M_{tp} M_{rt} M_{sr} \tag{A.5}$$

where $\mathbf{a} = (x_{tp}, y_{tp}, z_{tp}, \theta_{tp}, \phi_{tp}, \varphi_{tp}, x_{sr}, y_{sr}, z_{sr}, \theta_{sr}, \phi_{sr}, \varphi_{sr})$. Each scan slice through the phantom which intersects a wire constrains two degrees of freedom. For example, if the Y-wire intersects a scan plane then the X and Z coordinates of the point of intersection in phantom space are zero (Figure A.3). Therefore each scan which intersects a wire in the phantom contributes two equations to form a system of $2n$ equations, where n is the number of scans,

$$Q(\mathbf{a}) = 0 \tag{A.6}$$

where

$$\begin{aligned} Q(\mathbf{a}) &= (q_1(\mathbf{a}) \ldots q_{2n}(\mathbf{a}))^T && \text{(A.7)} \\ q_{2i}(\mathbf{a}) &= (M_{sp}(\mathbf{a})\mathbf{x}^{\mathbf{j}}_{\mathbf{s_i}})_\alpha && \text{(A.8)} \\ q_{2i-1}(\mathbf{a}) &= (M_{sp}(\mathbf{a})\mathbf{x}^{\mathbf{j}}_{\mathbf{s_i}})_\beta, && \text{(A.9)} \end{aligned}$$

and $i = (1, \ldots, n) \quad \alpha, \beta, j \in (X, Y, Z) \quad \alpha, \beta \neq j, \quad \alpha \neq \beta$. The notation $(M_{sp}(\mathbf{a})\mathbf{x}^{\mathbf{j}}_{\mathbf{s_i}})_k$ signifies the selection of the k^{th} component of the point $\mathbf{x}^{\mathbf{j}}_{\mathbf{s_i}}$ in phantom-space. $\mathbf{x}^{\mathbf{j}}_{\mathbf{s_i}}$ is the position in scan-space of the intersection between the i^{th} scan and the j axis of the phantom coordinate system. $M_{sp}(\mathbf{a})$ is the matrix which transforms points in scan space into phantom space. In other words, for each wire imaged only the two zero-valued coordinates of the intersection point in phantom space contribute an equation to the system, hence $k \neq j$. Since there are 12 parameters to be determined, a minimum of six scans are required to determine the two transforms, M_{sr} and M_{tp}. In practice, there is an uncertainty associated with determining the point of intersection between a nylon wire and the scan plane; this is ±1mm for the l558 probe with a 120mm field of view. This is particularly the case when the scan plane intersects a wire at an oblique angle, since the scan plane has a definite thickness. Consequently, many more scans than the six required to determine the system are taken. An iterative least squares technique is used to solve the over-determined system of equations. The technique offers quadratic convergence, provided the initial estimate of the parameters $\mathbf{a}$ which determine M_{sp} are sufficiently close to the solution [Burden and Faires 1989]. The new estimate $\mathbf{a}'$ is calculated from the current estimate $\mathbf{a}$ by equation A.10,

$$\mathbf{a}' = \mathbf{a} + \mu \mathbf{C} \tag{A.10}$$

where

$$\begin{aligned} \mathbf{C} &= (D^t D)^{-1} D^t \mathbf{E} && \text{(A.11)} \\ \mathbf{E} &= Q(\mathbf{a}) && \text{(A.12)} \\ D_{ij} &= \frac{dq_i(\mathbf{a})}{da_j} \qquad i = 1 \ldots 2n, \quad j = 1 \ldots 12 && \text{(A.13)} \end{aligned}$$

and $\mu = (\mu_1 \ldots \mu_{12})$ is the gain factor by which each parameter is altered per iteration. Equations corresponding to outlying erroneous data points can be removed from the system by considering the magnitude of the corresponding error component $|\mathbf{E_i}|$, where $\mathbf{E_i}$ is the error due to the i^{th} equation. Those equations with large errors can be removed from the system and an improved solution recalculated. The need to do this has not yet arisen nor have convergence problems been encountered. Typically 6-10 scans of each wire in the calibration phantom are taken, yielding a system of 30 equations. The net RMS calibration error E_{rms} (equation A.14) has consistently been found to be less than 0.5mm.

$$E_{rms} = \sqrt{\frac{\sum_{i=1}^{2n} E_i^2}{2n}} \tag{A.14}$$

Figure A.4 is a photograph taken during the calibration procedure for a new probe. In this case a Polhemus receiver is mounted on a 7MHz endo-vaginal probe.

Figure A.4 Calibration of a 7MHz endo-vaginal probe at Christchurch Women's Hospital.

Appendix B

ITERATIVE REFINEMENT FOR RBF SOLUTION

Iterative refinement is a method referred to in chapter 6 for improving the accuracy of a numerically computed matrix inverse.Consider solving the linear system $A\mathbf{x} = \mathbf{d}$ for $\mathbf{x}$. If $\mathbf{x}_0$ is the solution found by inverting A, say by standard Gauss-Jordan elimination methods on a digital computer, then an improved estimate $\mathbf{x}_1$ can be found by solving the new system,

$$A\mathbf{x}' = \Delta, \tag{B.1}$$

where

$$\Delta = d - A\mathbf{x}_0, \tag{B.2}$$

and

$$\mathbf{x}_1 = \mathbf{x}_0 + \mathbf{x}'. \tag{B.3}$$

This process can be repeated iteratively, if warranted by the magnitude of the residual, Δ. The i^{th} estimate is given by

$$\mathbf{x}_i = \mathbf{x}_{i-1} + A^{-1}\Delta_{i-1} \qquad i = 1, 2, 3, \ldots \tag{B.4}$$

where

$$\Delta_j = \begin{cases} \Delta_{j-1} - A\mathbf{x}_j & j = 1, 2, 3, 4, \ldots \\ d - A\mathbf{x}_1 & j = 0 \end{cases} \tag{B.5}$$

Note that, in practice, forward elimination on the interpolation matrix A need only be performed once. Each subsequent iteration requires forward elimination on Δ_{i-1} alone followed by back substitution.

REFERENCES

AKIYAMA, I., SAITO, T., NAKAMURA, M., TANIGUCHI, N. and ITOH, K. (1990), 'Tissue characterization by using fractal dimension of B-scan image', In *IEEE 1990 Ultrasonics Symposium Proceedings*, IEEE, pp. 1353–5.

ANDERSON, F. (1991), '3D real time ultrasonic imaging using ellipsoidal backprojection', *Proceedings of SPIE*, Vol. 1443, pp. 62–80.

ASTLEY, R. (1992), *Finite Elements in Solids and Structures - An introduction*, Chapman and Hall, London, 1st ed.

BAILEY, D. (1985), *Hardware and Software Developments for Applied Digital Image Processing*, PhD thesis, Department of Electrical and Electronis Engineering, University of Canterbury, Christchurch, New Zealand.

BALEN, F.G., ALLEN, C.M., GARDENER, J.E., SIDDLE, N.C. and LEES, W.R. (1993), '3-dimensional reconstruction of ultrasound images of the uterine cavity', *British Journal of Radiology*, Vol. 66, No. 787, July, pp. 588–91.

BALEN, F.G., ALLEN, C.M. and LEES, W.R. (1994), 'Ultrasound contrast agents. [review]', *Clinical Radiology*, Vol. 49, No. 2, February, pp. 77–82.

BEATSON, R. and LIGHT, W. (1994), *Fast evaluation of radial basis functions: Methods for 2-dimensional polyharmonic splines*, Technical Report 119, Mathematics department, University of Canterbury, Christchurch, New Zealand, December.

BEATSON, R. and NEWSAM, G. (1992), 'Fast evaluation of radial basis functions: I', *Computers and Mathematics with Applications*, Vol. 24, No. 12, pp. 7–19.

BEATSON, R. and POWELL, M. (1994), 'An iterative method for thin-plate spline interpolation that employs approximations to the Lagrange functions', In GRIFFITHS, D. and WATSON, G. (Eds.), *Numerical Analysis 1993*, Longman Scientific, pp. 17–39.

BELOHLAVEK, M., GREENLEAF, J.F., FOLEY, D.A. and SEWARD, J.B. (1991), 'Utility of image enhancement methods in three-dimensional ultrasound reconstruction', In *IEEE 1991 Ultrasonics Symposium Proceedings*, New York: IEEE, pp. 1219–22.

BLAKE, G.B., MACFARLANE, M.R. and HINTON, J.W. (1990), 'Titanium in reconstructive surgery of the skull and face', *British Journal of Plastic Surgery*, Vol. 43, pp. 528–535.

BOMANS, M., HOHNE, K., TIEDE, U. and RIEMER, M. (1990), '3-D segmentation of mr images of the head for 3-D display', *IEEE Transactions on Medical Imaging*, Vol. 9, No. 2, June, pp. 177–183.

BRATKO, I., MOZETIC, I. and LAVRAC, N. (1988), *KARDIO: A Study in Deep Qualitative Knowledge for Expert Systems*, MIT Press.

BRESENHAM, J.E. (1965), 'Algorithm for computer control of digital plotter', *IBM System Journal*, Vol. 4, No. 1, pp. 25–30.

BRICHIERI-COLOMBI, Y. (1993), *Assessment of the Accuracy of a Magnetic Tracking Device When Used in Conjunction With an Ultrasound Probe*, Technical Report, Dept. Medical Physics, Christchurch Hospital, Christchurch Hospital, Christchuch, New Zealand, April.

BURCKHARDT, C. (1978), 'Speckle in ultrasound B-mode scans', *IEEE Transactions on Sonics and Ultrasonics*, Vol. 25, No. 1, January, pp. 1–6.

BURDEN, R.L. and FAIRES, J.D. (1989), *Numerical Analysis*, PWS-Kent Publishing Company, Boston, Massachusetts, 4th ed.

BURRELL, C.J., MCDONALD, A.H. and ROTHMAN, M.T. (1989), '3D computer visualisation of arteries and blood flow - in vitro and in vivo', *Computers in Cardiology*, September, pp. 41–46.

CANNY, J.F. (1983), *Finding Edges and Lines in Images*, Technical Report 720, MIT Artificial Intelligence Laboratory.

CARR, J.C. (1991), *Reconstruction of the KARDIO Expert System*, Man-Machine Studies, report no UC-DSE/40, (ed) J. H. Andreae, Department of Electrical and Electronic Engineering, University of Canterbury.

CHANDLER, C.L., UTTLEY, D., ARCHER, D.J. and MACVICAR, D. (1994), 'Imaging after titanium cranioplasty', *British Journal of Neurosurgery*, Vol. 8, No. 4, pp. 409–14.

CHEN, G.T.Y., HERMAN, G.T., REYNOLDS, R.A. and UDUPA, J.K. (1985), 'Surface shading in the cuberille environment', *IEEE Computer Graphics and Applications*, Vol. 5, No. 12, December, pp. 33–43.

CHEN, J., HUERTAS, A. and MEDIONI, G. (1987), 'Fast convolution with laplacian-of-gaussian masks', *IEEE Transactions on Pattern Recognition and Machine Intelligence*, Vol. 8, pp. 584–590.

CHENG, X.Y., OHYA, A., NATORI, M. and NAKAJIMA, M. (1993), 'Boundary extraction method for three dimensional ultrasonic echo imaging using fuzzy reasoning and relaxation techniques', In *Proceedings of the Nuclear Sciences Symposium and Medical Imaging Conference, San Francisco.*

CLINE, H.E., LORENSEN, W.E., LUDKE, S., CRAWFORD, C.R. and TEETER, B.C. (1988), 'Two algorithms for the three-dimensional reconstruction of tomograms', *Medical Physics*, Vol. 15, No. 3, May, pp. 320–327.

COHEN, L. and COHEN, I. (1990), 'A finite element method applied to new active contour models and 3D reconstruction from cross sections', In *Proceedings of Third International Conference on Computer Vision*, IEEE Computer Society Press, pp. 587–591.

COPPINI, G., POLI, R. and G., V. (1995), 'Recovery of the 3-D shape of the left ventricle from echocardiographic images', *IEEE Transactions on Medical Imaging*, Vol. 14, No. 2, June, pp. 301–317.

DAVIES, E. (1988), *Pattern Recognition*, No. 7, pp. 87–97.

DERICHE, R. (1990), 'Fast algorithms for low-level vision', *IEEE Transactions on Pattern Recognition and Machine Intelligence*, Vol. 12, No. 1, January.

DOUGHERTY, E. (Ed.) (1993), *Mathematical morphology in image processing*, Marcel Dekker, Inc.

DREBLIN, R.A., CARPENTER, L. and HANRAHAN, P. (1988), 'Volume rendering', *Computer Graphics*, Vol. 22, No. 4, August, pp. 51–58.

DUNN, F. (1991), 'Ultrasound', *IEEE Transactions on Education*, Vol. 34, No. 3, August, pp. 266–268.

ELDEMELLAWY, A.A. and MARTIN, R.W. (1989), 'Cross correlation for aligning 3D ultrasound derived endocardial surface', In *Images of the 21st Century - Proceedings of the Engineering in Medicine and Biology Society 11th Annual International Conference*, June, pp. 405–406.

ENRIGHT, S.A. (1992), *Towards Quantitative Computed Tomography*, PhD thesis, Electrical and Electronic Engineering, University of Canterbury, New Zealand, April.

ENRIGHT, S.A., DALE, S.M., SMITH, V.A., MURCH, R.D. and BATES, R.H.T. (1992a), 'Towards solving the bent-ray tomographic problem', *Inverse Problems*, Vol. 8, No. 1, February, pp. 83–94.

EVANS, J.A. (Ed.) (1986), *Physics in Medical Ultrasound*, The Institute of Physical Sciences in Medicine, 47 Belgrave Square, London SW1X 8QX, England.

EVANS, A. and NIXON, M. (1993), 'Speckle filtering in ultrasound images for feature extraction', In *International Conference on Acoustic Sensing and Imaging*, IEE, March, pp. 44–49.

FOLEY, J. and VAN DAM, A. (1982), *Fundamentals of Interactive Computer Graphics*, Addison-Wesley Publishing Co.

FOLEY, J.D., VAN DAM, A., FEINER, S.K. and HUGHES, J.F. (1990), *Computer Graphics Principles and Practice*, Addison-Wesley, Reading, Mass.

FRIGHT, W.R. and ALLAN, R. (1994), *Performance Evaluation: Diasonics ADA 400*, Technical Report, Dept. Radiology, Christchurch Hospital, Christchurch Hospital, Christchuch, New Zealand, April.

FRISBY, J.P. (1979), *Seeing: Illusion, Brain and Mind*, Oxford University Press.

GARDENER, J.E. (1991), 'Volume imaging of soft tissues with ultrasound', In *IEE Colloquium on 3-D Imaging Techniques for Medicine*, IEE London, UK, pp. 61–63.

GIERTSEN, C. (1992), 'Volume visualization of sparse irregular meshes', *IEEE Computer Graphics and Applications*, March, pp. 40–48.

GONZALEZ, R.C. and WOODS, R.E. (1992), *Digital image processing*, Addison-Wesley.

GORDON, D. (1985), 'Image space shading of 3-Dimensional objects', *Computer Vision, Graphics, and Image Processing*, Vol. 29, pp. 361–376.

GRIMSON, W.E.L. (1981), *From Images to Surfaces, A Computational Study of the Human Early Visual System*, MIT Press.

HANRAHAN, P. and LAUR, D. (1991), 'Hierarchical splatting: A progressive refinement algorithm for volume rendering', *Computer Graphics*, Vol. 25, No. 4, August, pp. 285–288.

HARDER, R.L. (1972), 'Interpolation using surface splines', *Journal of Aircraft*, Vol. 9, No. 2, February, pp. 189–191.

HEALEY, A.J. and LEEMAN, S. (1993), 'Speckle reduction methods in ultrasound pulse-echo imaging', In *Acoustic Sensing and Imaging*, IEE, March, pp. 68–76.

HEALEY, A.J., FORSBERG, F. and LEEMAN, S. (1991), 'Processing techniques for speckle reduction in medical ultrasound images', In *IEE Colloquium on Image Processing in Medicine*, IEE, April, pp. 6/1–4.

HEBERT, M., IKEUCHI, K. and DELINGETTE, H. (1995), 'A spherical representation for recognition of free-form surfaces', *IEEE Transactions on Pattern Recognition and Machine Intelligence*, Vol. 17, No. 7, July, pp. 691–689.

HERLIN, I.L. and GIRAUDON, G. (1993), 'Performing segmentation of ultrasound images using temporal information', In *Proceedings of 1993 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, IEEE Computer Society Press, pp. 373–378.

HERMAN, G.T. and LIU, H.K. (1979), 'Three-dimensional displays of organs from computed tomograms', *Computer Graphics and Image Processing*, Vol. 9, No. 1, January, pp. 1–21.

HERMAN, G.T., ZHENG, J. and BUCHOLTZ, C.A. (1992), 'Shape-based interpolation', *IEEE Computer Graphics and Applications*, May, pp. 69–70.

HERRINGTON, D., SANTAGO, P., JOHNSON, T., DOWNES, T., BRADEN, G. and SNYDER, W. (1991), 'Image processing and display of 3D intra-coronary ultrasound images', In *Proceedings Computers in Cardiology*, IEEE Computer Society Press, September, pp. 349–352.

HINTON, J.W., MACFARLANE, M.R. and BLAKE, G.B. (1992), 'Repairing people - titanium in reconstructive surgery', *Shadows: New Zealand Journal of Medical Radiation Technology*, Vol. 35, No. 1, March, pp. 27–32.

HODGSON, R.M., BAILEY, D., NAYLOR, M.J., NG, A. and MCNEILL, S. (1985), 'Properties, implementations and applications of rank filters', *Image and Vision Computing*, Vol. 3, No. 1, February, pp. 3–14.

HOHNE, K.H. and BERNSTEIN, R. (1986), 'Shading 3D images from ct using gray-level gradients', *IEEE Transactions on Medical Imaging*, Vol. MI-5, No. 1, March, pp. 45–47.

HOHNE, K.H., BOMANS, M., POMMERT, A., RIEMER, M., SCHIERS, C., TIEDE, U. and WIEBECKE, G. (1990), '3D visualization of tomographic volume data using the generalised voxel model', *The Visual Computer*, Vol. 6, No. 1, February, pp. 28–36.

HOPPE, H., DEROSE, T., DUCHAMP, T., MCDONALD, J. and STUETZLE, W. (1992), 'Surface reconstruction from unorganised points', *Computer Graphics*, Vol. 26, No. 2, July, pp. 71–78.

HUERTAS, A. and MEDIONI, G. (1986), 'Detection of intensity changes with subpixel accuracy using laplacian-gaussian masks', *IEEE Transactions on Pattern Recognition and Machine Intelligence*, Vol. 8, No. 5, September, pp. 651–664.

HUSSEY, M. (1975), *Diagnostic Ultrasound: An Introduction to the Interactions Between Ultrasound and Biological Tissues*, Blackie.

IGES (1993), *American National Standard IGES 5.2*, ANS US PRO, November.

JAGO, J.R. and WHITTINGHAM, T.A. (1993), 'The application of computed tomography techniques to imaging with transmitted and reflected ultrasound', In *International Conference on Acoustic Sensing and Imaging*, London: IEE, March, pp. 257–265.

JOFFE, J.M., MCDERMOTT, P.J., LINNEY, A.D., MOSSE, C.A. and HARRIS, M. (1992), 'Computer-generated titanium cranioplasty: report of a new technique for repairing skull defects', *British Journal of Neurosurgery*, Vol. 6, No. 4, pp. 343–50.

KAUFMAN, A. (1991), *Volume Visualisation*, IEEE Computer Society Press, Los Altimos, CA.

KELLY, I.M.G., GARDENER, J.E. and LEES, W.R. (1992), 'Three-dimensional fetal ultrasound', *The Lancet*, Vol. 339, April, pp. 1062–1064.

KELLY, I.M., GARDENER, J.E., BRETT, A.D., RICHARDS, R. and LEES, W.R. (1994), 'Three-dimensional us of the fetus. work in progress', *Radiology*, Vol. 192, No. 1, July, pp. 253–9.

KITNEY, R.I., BURRELL, C.J., MOURA, L., STRAGHAN, K. and ROTHMAN, M.T. (1990), '3-D visualisation of arterial structures - tissue differentiation techniques', In *Optical Fibres in Medicine*, pp. 505–513.

KITNEY, R.I., BURRELL, C.J., STRAGHAN, K., MOURA, L. and ROTHMAN, M.T. (1991), '3D visualisation for the study of arterial disease and tissue characterisation', In *Proceedings Computers in Cardiology*, IEEE Computer Society Press, pp. 3–6.

KLINGLER, J.W., VAUGHAN, C.L., FRAKER, T.D. and ANDREWS, L.T. (1988), 'Segmentation of echocardiographic images using mathematical morphology', *IEEE Transactions on Biomedical Engineering*, Vol. 35, No. 11, November, pp. 925–934.

KRIEG, J.C. and SANDERS, S.P. (1984), 'Stereotaxic localization of an ultrasonic transducer', In *37th ACEMB*, September, p. 89.

KRUCZKOWSKI, P.J.C., MYLREA, K.C., ROEMER, R.R. and DRACH, G.M. (1988), 'A non-invasive ultrasonic system to determine residual bladder volumes', In *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology*, New York: IEEE, November, pp. 1623–1624.

LAPACK (1992), *LAPACK users guide*, Society for Industrial and Applied Mathematics, Philadelphia.

LAURENTINI, A. (1993), 'Inferring the shape of the real object from the object reconstructed by volume intersection', In *1993 IEEE Computer Society Conference on Computer Vision Pattern Recognition and*, IEEE Computer Society Press, June, pp. 280–285.

LAURENTINI, A. (1995), 'How far 3D shapes can be understood from 2D silhouettes', *IEEE Transactions on Pattern Recognition and Machine Intelligence*, Vol. 17, No. 2, February, pp. 188–195.

LAVALLEE, S. and SZELISKI, R. (1995), 'Recovering the position and orientation of free-form objects from image contours using 3D distance maps', *IEEE Transactions on Pattern Recognition and Machine Intelligence*, Vol. 17, No. 4, April, pp. 378–390.

LEEMAN, S., GATENBY, J.C. and FORSBERG, F. HODDINOTT, J.C. (1989), 'Speckle reduction in two easy steps', In *IEEE 1989 Ultrasonics Symposium Proceedings*, IEEE, pp. 927–930.

LEVOY, M. (1988), 'Display of surfaces from volume data', *IEEE Computer Graphics and Applications*, May, pp. 29–37.

LEVOY, M. (1990), 'A hybrid ray tracer for rendering polygon and volume data', *IEEE Computer Graphics and Applications*, March, pp. 33–40.

LEYTON, M. (1992), *Symmetry, Causality, Mind*, MIT Press.

LIGHT, W. (1992), 'Some aspects of radial basis function approximation', In SINGH, S. (Ed.), *Approximation Theory, Spline Functions and Applications*, Kluwer Academic Publishers (Dortrecht), pp. 163–190.

LINNEY, A.D., TAN, A.C., RICHARDS, J., GARDENER, S., GRINDROD, S. and MOSS, J.P. (1993), 'Three-dimensional visualisation of data on human anatomy: Diagnosis and surgical planning', *Journal of Audiovisual Media in Medicine*, Vol. 16, pp. 4–10.

LORENSEN, W.E. and CLINE, H.E. (1987), 'Marching cubes: A high resolution 3D surface construction algorithm', *Computer Graphics*, Vol. 21, No. 4, July, pp. 163–169.

MALING, D.H. (1992), *Coordinate Systems and Map Projections*, Pergamon Press, 2nd ed.

MALZBENDER, T. (1993), 'Fourier volume rendering', *ACM Transactions on Graphics*, Vol. 12, No. 3, July, pp. 233–250.

MARAGOS, P. and SCHAFER, R. (1990), 'Morphological systems for multidimensional signal processing', *Proceedings of the IEEE*, Vol. 78, No. 4, April, pp. 690–710.

MARR, D. (1982), *Vision*, W. H. Freeman and Company.

MARR, D. and HILDRETH, E. (1980), 'Theory of edge detection', *Proceedings of the Royal Society of London*, Vol. 207, pp. 187–217.

MARTIN, R.W., BASHEIN, G. and NESSLY, M.L. (1989), 'Left-ventricular transesophageal ultrasound 3D scanning and reconstruction', In *Images of the 21st Century - Proceedings of the Engineering in Medicine and Biology Society 11th Annual International Conference*, June, pp. 421–422.

MCCANN, H.A., SHARP, J.C., KINTER, T.M., MCEWAN, C.N., BARILLOT, C. and GREENLEAF, J.F. (1988), 'Multidimensional ultrasonic imaging for cardiology', *Proceedings of the IEEE*, Vol. 76, No. 9, September, pp. 1063–1072.

MCDICKEN, W.N. (1991), *Diagnostic Ultrasonics*, Churchill Livingstone, 3rd ed.

MCMINN, R., HUTCHINGS, R., PEGINGTON, J. and ABRAHAMS, P. (1993), *A Colour Atlas of Human Anatomy*, Wolfe, 3rd ed.

MEYER, K., APPLEWHITE, H.L. and BIOCCA, F.A. (1992), 'A survey of position trackers', *Presence*, Vol. 1, No. 2, pp. 173–200.

MILLS, P.H. and FUCHS, H. (1990), '3D ultrasound display using optical tracking', In *Proceedings of the First Conference on Visualisation in Biomedical Computing, Atlanta*, May, pp. 490–497.

MINTZ, G.S., PICHARD, A.D., SALTER, L.F., POPMA, J.J., KENT, K.M. and LEON, M.B. (1993), 'Three dimensional intravascular ultrasonography: Reconstruction of endovascular stents in vitro and in vivo', *Journal of clinical ultrasound*, Vol. 21, December, pp. 609–615.

MORITZ, W.E., PEARLMAN, A.S., MCCABE, D.H., MEDEMA, D.K., AINSWORTH, M.E. and BOLES, M.S. (1983), 'An ultrasonic technique for imaging the ventricle in three dimensions and calculating its volume', *IEEE Transactions on Biomedical Engineering*, Vol. 30, No. 8, August, pp. 482–492.

MURCH, R. (1990), *Inverse Scattering and Shape Reconstruction*, PhD thesis, Dept. Electrical and Electronic Engineering, University of Canterbury, Christchurch, New Zealand, April.

NAKAI, T., YASUDA, M., MARUTANI, Y. and OHTA, T. (1993), 'Plastic organ models fabricated from mr images using lasers and photopolymers', *Medical and Biological Engineering and Computing*, November, pp. 647–650.

NELSON, T.R. and ELVINS, T.T. (1993), 'Visualisation of 3D ultrasound data', *IEEE Computer Graphics and Applications*, November, pp. 50–57.

NELSON, T.R. and PRETORIUS, D.H. (1995), 'Visualization of the fetal thoracic skeleton with three-dimensional sonography: a preliminary report', *American Journal of Roentgenology*, Vol. 164, No. 6, June, pp. 1485–8.

NEY, D., FISHMAN, E., MAGID, D. and DREBIN, R. (1990), 'Volumetric rendering of computed tomography data: Principles and techniques', *IEEE Computer Graphics and Applications*, March, pp. 24–31.

NIELSON, G.M. (1993), 'Scattered data modelling', *IEEE Computer Graphics and Applications*, January, pp. 60–70.

NIKRAVESH, P.E., SKORTON, D.J., CHANDRAN, K.B., ATTARWALA, Y.M., PANDIAN, N. and KERBER, P.E. (1984), 'Computerized three-dimensional finite element reconstruction of the left ventricle from cross-sectional echocardiograms', *Ultrasonic Imaging*, Vol. 6, pp. 48–59.

NIXON, M.S. and HAMES, T.K. (1993), 'New technique for 3D artery modelling by noninvasive ultrasound', *IEE Proceedings I [Communications, Speech and Vision]*, Vol. 140, No. 1, February, pp. 86–94.

NIXON, M.S., HAMES, T.K., MARTIN, P., POWELL, S. and D PAIVA, S. (1992), '3D arterial modelling using feature extraction in ultrasound images', In *International Conference on Image Processing and its Applications*, IEE, April, pp. 373–376.

OHBUCHI, R. and FUCHS, H. (1990), 'Incremental 3D ultrasound imaging from a 2D scanner', In *Proceedings of the First Conference on Visualisation in Biomedical Computing*, Los Alamitos, CA : IEEE Comput. soc. Press, pp. 360–7.

PERRIN, F., PERNIER, J., BERTRAND, O., GIARD, M.H. and ECHALLIER, J.F. (1987a), 'Mapping of scalp potentials by surface spline interpolation', *Electroencephalography and Clinical Neurophysiology*, Vol. 66, pp. 75–81.

PERRIN, F., BERTRAND, O. and PERNIER, J. (1987b), 'Scalp current density mapping: value and estimation from potential data', *IEEE Transactions on Biomedical Engineering*, Vol. BME-34, No. 4, pp. 283–288.

PINI, R., COSTI, M., MENSAH, G.A., MASOTTI, L., NOVINS, K., GREENBERG, D.P., GREPPI, B., CEROFOLINI, M. and DEVEREUX, R.B. (1991), 'Computed tomography of the heart by ultrasound', In *Proceedings Computers in Cardiology*, IEEE Computer Society Press, September, pp. 17–20.

POLHEMUS (1992), *3Space Fastrak Users Manual*, Polhemus Incorporated, 1 Hercules Drive, P. O. Box 560, Colchester, Vermont 05446, October.

POWELL, M. (1992), 'The theory of radial basis function approximation in 1990', In LIGHT, W. (Ed.), *Advances in Numerical Analysis II: Wavelets, subdivision algorithms and radial functions*, Oxford University Press, Oxford, UK,, pp. 105–210.

PRATT, W.K. (1991), *Digital image processing*, John Wiley Inc, 2nd ed.

PRETORIUS, D.H. and NELSON, T.R. (1994), 'Prenatal visualization of cranial sutures and fontanelles with three-dimensional ultrasonography', *Journal of Ultrasound in Medicine*, Vol. 13, No. 11, November, pp. 871–6.

PRETORIUS, D.H., NELSON, T.R. and JAFFE, J.S. (1992), '3-dimensional sonographic analysis based on color flow doppler and gray scale image data: a preliminary report', *Journal of Ultrasound in Medicine*, Vol. 11, No. 5, May, pp. 225–32.

RAAB, F.H., BLOOD, E.B., STEINER, T.O. and JONES, H.R. (1979), 'Magnetic position and orientation tracking system', *IEEE Transactions Aerospace Elector Systems*, pp. 709–718.

RAJA, N.S. and JAIN, A.K. (1992), 'Recognizing geons from superquadrics fitted to range data', *Image and Vision Computing*, Vol. 10, No. 3, April, pp. 179–190.

RANKIN, R.N., FENSTER, A., DOWNEY, D.B., MUNK, P.L., LEVIN, M.F. and VELLET, A.D. (1993), 'Three-dimensional sonographic reconstruction: Techniques and diagnostic applications', *AJR*, Vol. 161, October, pp. 695–702.

RICKEY, D.W., PICOT, D.W., HOLDSWORTH, D.W., DRANGOVA, M., MILLER, D.J.M., RANKIN, R.N. and FENSTER, A. (1991), 'Qantitative three-dimensional true velocity colour doppler imaging', In *IEEE 1991 Ultrasonics Symposium Proceedings*, IEEE, pp. 1277–1280.

ROBINSON, B.S. (1982), *Speckle Processing for Ultrasonic Imaging*, PhD thesis, Electrical and Electronic Engineering, University of Canterbury, New Zealand.

SABELLA, P. (1988), 'A rendering algorithm for visualizing 3D scalar fields', *Computer Graphics*, Vol. 22, No. 4, August, pp. 51–58.

SATHERLEY, B.L. (1994), *Zero-based ensemble deconvolution and EEG spectral topography*, PhD thesis, Department of Electrical and Electronic Engineering, University of Canterbury, Christchurch, New Zealand.

SAUL, H. (1994), 'Hipbone connected to the titanium implant', *New Scientist*, July, pp. 34–38.

SCHISTAD, A. and TAXT, T. (1991), 'Speckle reduction in ultrasound images using temporal and spatial context', In *1991 IEEE Nuclear Science Symposium and Medical Imaging Conference*, IEEE, November, pp. 2210–2214.

SELZER, R.H., BLANKENHORN, D.H., LEE, P.L. and LAI, J.Y. (1989), 'Synthesis of 3D arterial images from multiple ultrasound images', In *Images of the 21st Century - Proceedings of the Engineering in Medicine and Biology Society 11th Annual International Conference*, pp. 419–420.

SHIRLEY, P. and NEEMAN, H. (1989), 'Volume visualisation at the centre for supercomputing research and development', In *Proceedings of the Chapel Hill Workshop on Volume Visualization*, May, pp. 17–20.

SHUEY, D., BAILEY, D. and MORRISSEY, T.P. (1986), 'PHIGS: a standard, dynamic, interactive graphics interface', *IEEE Computer Graphics and Applications*, No. 6, August, pp. 50–57.

SNYDER, W.E., GROSHONG, R., HSIAO, M., BOONE, K.L. and HUDACKO, T. (1992), 'Closing gaps in edges and surfaces', *Image and Vision Computing*, Vol. 10, No. 8, October, pp. 523–531.

STEEN, E. and OLSTAD, B. (1994), 'Volume rendering of 3D medical ultrasound data using direct feature extraction', *IEEE Transactions on Medical Imaging*, Vol. 13, No. 3, September, pp. 517–525.

SURFCAM (1993), *Surfcam, Version 4.0 1993, reference manual*, Surfware Incorporated, 421 Park Avenue, San Fernando, CA91340.

SURMAN, K., CARR, J., BRICHIERI-COLOMBI, Y., FRIGHT, R. and GARDEN, K. (1993), 'A system for three dimensional graphics from ultrasound', In *Proceedings of the First New Zealand Conference on Image and Vision Computing, Auckland*, August, pp. 367–373.

TERZOPOULOS, D. (1986), 'Regularization of inverse visual problems involving discontinuities', *IEEE Transactions on Pattern Recognition and Machine Intelligence*, Vol. 8, No. 4, July.

TERZOPOULOS, D. (1988), 'The computation of visible-surface representations', *IEEE Transactions on Pattern Recognition and Machine Intelligence*, Vol. 10, No. 4, July, pp. 1–5.

TERZOPOULOS, D. and METAXAS, D. (1991), 'Dynamic 3D models with local and global deformations: Deformable superquadrics', *IEEE Transactions on Pattern Recognition and Machine Intelligence*, Vol. 13, No. 7, July, pp. 703–714.

THOMAS, J.G., PETERS II, R.A. and JEANTY, P. (1991), 'Automatic segmentation of ultrasound images using morphological operators', *IEEE Transactions on Medical Imaging*, Vol. 10, No. 2, June, pp. 180–186.

THUNE, N. and OLSTAD, B. (1991), 'Visualising 4-D medical ultrasound data', In *Proceedings Visualization 1991*, IEEE Computer Society Press, October, pp. 210–215.

TIEDE, U., HEINZ, K., BOMANS, M., POMMERT, A., RIEMER, M. and WIEBECKE, G. (1990), 'Investigation of medical 3D-rendering algorithms', *IEEE Computer Graphics and Applications*, March, pp. 41–53.

TORRE, V. and POGGIO, T. (1986), 'On edge detection', *IEEE Transactions on Pattern Recognition and Machine Intelligence*, Vol. 8, No. 2, March, pp. 147–163.

TOTSUKA, T. and LEVOY, M. (1993), 'Frequency domain volume rendering', In *ACM SIGRAPH 93*, July, pp. 271–278.

TRAHEY, G., ALLISON SMITH, S. and VON RAMM, O. (1987), 'Speckle reduction acheivable by spatial compounding and frequency compounding: Experimental results and implications for target detectability', In FERRARI, L.A. (Ed.), *International Symposium on Pattern Recognition and Acoustical Imaging*, SPIE, Bellingham: SPIE, pp. 185–192.

TUY, H.K. and TUY, L.T. (1984), 'Direct 2-D display of 3-D objects', *IEEE Computer Graphics and Applications*, Vol. 4, No. 10, November, pp. 29–33.

UDUPA, J.K. and ODHNER, D. (1993), 'Shell rendering', *IEEE Computer Graphics and Applications*, November, pp. 58–67.

VAN PUTTEN, M.C. and YAMADA, S. (1992), 'Alloplastic cranial implants made from computed tomographic scan-generated casts', *Journal of Prosthetic Dentistry*, Vol. 68, No. 1, July, pp. 103–108.

VERLANDE, M., FLACHSKAMPF, F.A., SCHNEIDER, W., AMELING, W. and HANRATH, P. (1991), '3D reconstruction of the beating left ventricle and mitral valve based on multiplanar TEE', In *Proceedings Computers in Cardiology*, IEEE Computer Society Press, September, pp. 285–288.

WEAR, K., WAGNER, R. and GARRA, B. (1994), 'High resolution ultrasonic backscatter coefficient estimation based on autoregressive spectral estimation using burg's algorithm', *IEEE Transactions on Medical Imaging*, Vol. 13, No. 3, September.

WEBBER, R.E. (1990), 'Ray tracing voxel based data via biquadratic local surface interpolation', *The Visual Computer*, Vol. 6, No. 1, February, pp. 8–15.

WELLS, P.N.T. (1977), *Biomedical Ultrasonics*, Academic Press.

WESTOVER, L. (1989), 'Interactive volume rendering', In *Proceedings of the Chaple Hill Workshop on Volume Visualisation*, University of North Carolina Press, May, pp. 9–16.

WESTOVER, L. (1990), 'Footprint evaluation for volume rendering', *Computer Graphics*, Vol. 24, No. 4, August, pp. 367–376.

WIER, J. and ABRAHAMS, P. (1986), *An Atlas of Radiological Anatomy*, Churchill-Livingstone, 2nd ed.

WILSON, J. (1982), *Models of Human Visual Cognition*, PhD thesis, Dept. Electrical and Electronic Engineering, University of Canterbury, Christchurch, New Zealand.

WILSON, S. (1983), 'On retino-cortical mapping', *International Journal Man-Machine Studies*, Vol. 17, No. 18, pp. 361–389.

WOODCOCK, J.P. (1979), *Ultrasonics*, Adam Hilger Ltd., 3rd ed.

YOO, T.S., NEUMANN, U., FUCHS, H., PIZER, S.M., CULLIP, T., RHOADES, J. and WHITAKER, R. (1992), 'Direct visualization of volume data', *IEEE Computer Graphics and Applications*, July, pp. 63–71.

YUAN, X. (1995), 'A mechanism of automatic 3D object modeling', *IEEE Transactions on Pattern Recognition and Machine Intelligence*, Vol. 17, No. 3, March, pp. 307–311.

ZIENKIEWICZ, O.C. and TAYLOR, R.L. (1989), *The Finite Element Method - Basic Formulation and Linear Problems*, Vol. 1, McGraw-Hill (UK), 4th ed.