Quantification of Carotid Artery Disease using Three-dimensional Ultrasound Imaging

Dean C. Barratt

A thesis submitted for the degree of
Doctor of Philosophy from the University of London



School of Medicine,
Imperial College of Science, Technology & Medicine,
London, UK.

2002

BIBL.
LONDIN.
UNIV.

Dedicated to my parents, Brian and Connie.

Abstract

Ultrasound imaging is widely used for the clinical assessment of carotid artery disease. However, current ultrasound scanners only enable quantitative measurements to be made from planar images. Three-dimensional ultrasound (3DUS) is a new development which offers a number of attractions over conventional methods. These include the potential to quantify the severity and distribution of atherosclerotic plaque accurately and comprehensively by using 3D measures, such as volume.

This thesis describes the development and validation of a flexible, low-cost 3DUS imaging system, along with associated computer algorithms to enable the 3D quantification of carotid disease. Incorporated into the system is an electromagnetic tracking device. This allows the spatial 3D position and orientation of a series of planar ultrasound images, acquired using an ordinary commercial ultrasound scanner, to be calculated. The system was designed to take advantage of digital image capture, transfer, and post-processing features, and with minor adaptation can support any ultrasound imaging mode.

Algorithms are described for processing acquired ultrasound images in order to reconstruct a 3D surface model of the diseased carotid bifurcation. The surface model is generated by fitting longitudinal curves (splines) to cross-sectional wall and lumen contours defined for each image. The model has the advantage that, once formed, complex quantitative operations can be performed rapidly and automatically.

A number of quantitative measures determined from 3D carotid reconstructions are considered. These include established planar measures in addition to new 3D measures. Ultrasound phantoms were used to estimate the accuracy of 3D reconstruction and quantification with excellent results. The clinical feasibility of 3DUS was assessed by imaging a series of diseased carotid arteries. Preliminary findings demonstrate that 3DUS may be successfully applied to determine a detailed, quantitative description of carotid plaque in a subgroup of patients. The clinical significance of some of the measures obtainable using 3DUS analysis has yet to be fully established, however.

Declaration of originality

The material presented in this thesis is the original work of the author, and has not been submitted for a higher degree at any other institution. The following journal publications and conference presentations are derived from the research presented here:

Journal publications

Optimisation and evaluation of an electromagnetic tracking device for high accuracy three-dimensional ultrasound imaging of the carotid arteries.
(Barratt DC, Davies AH, Hughes AD, Thom SA, Humphries KN)
Ultrasound Med Biol 2001;27:957-968.

Accuracy of an electromagnetic three-dimensional ultrasound system for carotid artery imaging.
(Barratt DC, Davies AH, Hughes AD, Thom SA, Humphries KN)
Ultrasound Med Biol 2001;27:1421-1425.

Anatomically realistic reconstruction of the carotid artery bifurcation from three-dimensional ultrasound data.
(In preparation)

Conference presentations

9th Congress of the World Federation for Ultrasound in Medicine and Biology, Florence, Italy, May 2000:

Quantitative measures of carotid disease from three-dimensional ultrasound images.
(Barratt DC, Ariff BB, Davies AH, Hughes AD, Thom SA, Humphries KN)
Abstract in: *Ultrasound Med Biol* 2000;26 Suppl.2:A86

Accuracy of three-dimensional ultrasound reconstruction in a clinical environment.
(Barratt DC, Ariff BB, Davies AH, Hughes AD, Thom SA, Humphries KN)
Abstract in: *Ultrasound Med Biol* 2000;26 Suppl.2:A66

The Physiology Society, London, UK, 2000:

A system for 3D vascular ultrasound imaging.
(Barratt DC, Ariff BB, Davies AH, Hughes AD, Thom SA, Humphries KN)
Abstract in: *J Physiol* 2000;525P:528

The following aspects of the thesis are considered to be original contributions:

- The development of an electromagnetic 3DUS acquisition system that captures digital images using a standard, commercial ultrasound scanner (Chapter 4).
- An independent evaluation of the accuracy and precision of the Ascension pcBIRD electromagnetic tracking device in a clinical ultrasound environment (Chapter 5).
- An investigation of the relationship between scan depth and 3D reconstruction precision for a 3DUS system (Chapter 5).
- A method for 3D wall/lumen surface reconstruction of the diseased carotid bifurcation using splines (Chapter 6).
- A method for calculating 'true' transverse plane measures of carotid stenosis from a 3DUS reconstruction (Chapters 6 and 7).
- The quantification of the volumetric stenosis index and stenosis profiles for the carotid arteries (Chapters 6 and 7).

Acknowledgements

Firstly, I would like to thank my supervisors – Mr. Alun Davies, Mr. Keith Humphries, and Dr. Simon Thom – for their help and support. I am also very grateful to Dr. Anil Bharath, Prof. Alun Hughes, and Prof. Kim Parker for their helpful advice, and to Ms. Mary Curle and Mr. John Wolfe for assistance in finding patient volunteers for my study. Also, the help of friends and colleagues at Charing Cross, Hammersmith, and St. Mary's hospitals, and in the Departments of Bioengineering and Chemical Engineering at Imperial College, cannot go without a mention. Worthy of special mention are Tim Pitt, Mary Ellis, Ben Ariff, and Sheila Byrd. Lastly, I am indebted to my parents, Emma and close friends for their unwavering support and encouragement. Thank you.

Contents

Chapter 1 Introduction

1.1 Carotid artery disease

1.2 X-ray angiography

1.3 Ultrasound Imaging

1.3.1 Fundamental physics of ultrasound imaging

1.3.2 A-Mode

1.3.3 M-Mode

1.3.4 B-Mode

1.3.5 Spectral Doppler

1.3.6 Colour and power modes

1.3.7 Other ultrasound modes

1.3.7.1 Contrast imaging

1.3.7.2 Harmonic imaging

1.3.7.3 Tissue Doppler imaging

1.3.7.4 Compound B-mode

1.3.7.5 B-flow

1.3.8 Artefacts

1.3.8.1 Shadowing

1.3.8.2 Enhancement

1.3.8.3 Reverberation

1.3.8.4 Side- and grating-lobe artefacts

1.3.8.5 Mirror artefact

1.3.8.6 In-plane misregistration

1.3.8.7 Artefacts due to finite beam thickness

1.3.8.8 Noise

1.3.8.9 Doppler artefacts

1.4 Magnetic resonance imaging

1.5 3D imaging

Chapter 2 Quantification of carotid artery disease

2.1 Degree of stenosis

2.2 Influential clinical trials

2.3 Angiography-based measurement of degree of stenosis

2.4 Ultrasound-based measurement of degree of stenosis
2.4.1 Direct methods
2.4.2 Indirect methods
2.5 3D quantification of carotid disease
2.6 Motivations for the thesis
2.7 Aims of the thesis

Chapter 3 3D ultrasound imaging

3.1 Applications of 3D ultrasound
3.2 Generation of a 3D ultrasound image
3.3 3D acquisition systems
3.3.1 Constrained systems
3.3.1.1 Mechanical carriage systems
3.3.1.2 Electromechanically-driven transducers
3.3.2 Freehand systems
3.3.2.1 Mechanical arms
3.3.2.2 Acoustic spark-gap devices
3.3.2.3 Optical methods
3.3.2.4 Electromagnetic devices
3.3.2.5 Image correlation methods
3.3.3 Matrix-array transducers
3.4 Segmentation
3.5 3D reconstruction
3.5.1 Surface reconstruction
3.5.2 Voxel-based reconstruction
3.5.3 3D compound scanning
3.6 3D visualisation
3.6.1 Surface rendering
3.6.2 Volume rendering
3.6.3 Interactive visualisation
3.7 3D quantification
3.8 3D image artefacts
3.9 4D ultrasound imaging

Chapter 4 An electromagnetic 3D ultrasound acquisition system

4.1 System architecture

4.1.1 Ultrasound scanner

4.1.2 Electromagnetic tracking device

4.1.3 Personal computer

4.1.4 ECG monitor

4.1.5 Start-stop switch

4.1.6 Software

4.2 System operation

4.3 Calibration

4.3.1 Problem definition

4.3.2 Review of calibration methods

4.3.3 Calibration using a pin phantom

4.3.4 Calibration results

4.4 Discussion

Chapter 5 Accuracy of 3D ultrasound acquisition

5.1 Sources of pixel misregistration

5.1.1 Position and orientation measurement device errors

5.1.2 Temporal misregistration

5.1.3 Finite beam thickness

5.1.4 Imperfect calibration

5.1.5 Motion artefacts

5.1.6 Within-plane image effects

5.2 Error propagation

5.3 Experimental evaluation of the Ascension pcBIRD

5.3.1 Interference in pcBIRD measurements

5.3.1.1 Reorganisation of the scanning environment

5.3.1.2 Optimisation of pcBIRD parameters

5.3.2 Accuracy of position measurements

5.3.3 Accuracy of rotation measurements

5.3.4 Optimal protocol for carotid 3D ultrasound scans

5.4 Experimental validation of reconstruction accuracy

5.4.1 Wedge phantom experiments

5.4.2 Wedge phantom results

5.5 Discussion
5.5.1 Accuracy of electromagnetic POM devices
5.5.2 3D reconstruction accuracy

Chapter 6 Reconstruction and quantification of the diseased carotid bifurcation

6.1 Introduction
6.2 Curves and contours in 2 and 3 dimensions
6.2.1 Spline curves
6.2.2 Contours
6.3 Segmentation of acquired ultrasound images
6.4 Surface reconstruction
6.4.1 Data pre-processing
6.4.1.1 Transformation to a new 3D co-ordinate system
6.4.1.2 Removal of overlapping slices
6.4.1.3 Centroid registration of wall contours
6.4.1.4 Contour resampling
6.4.2 Reconstruction of the vessel wall
6.4.3 Reconstruction of the vessel lumen
6.5 Planar reslicing
6.6 Determination of the vessel central axis
6.7 3D quantification of carotid disease
6.7.1 Plaque volume and volumetric stenosis
6.7.2 True transverse measurements
6.7.2.1 Diameter-based measures
6.7.2.2 Area-based measures
6.7.3 Other measures

Chapter 7 Studies on phantom and patient carotid arteries

7.1 Reconstruction accuracy
7.1.1 Anthropomorphic carotid phantoms
7.1.2 Carotid phantom experiments
7.1.3 Results
7.2 Patient studies
7.2.1 Methods

7.2.2 Results
7.2.2.1 Patient 1 (PA-L)
7.2.2.2 Patient 2 (PG-R)
7.2.2.3 Patient 3 (SA-R)
7.2.2.4 Patient 4 (EK-L)
7.3 Discussion

Chapter 8 Conclusions and further work
8.1 Conclusions
8.1.1 3D ultrasound acquisition system
8.1.2 Image segmentation
8.1.3 Surface reconstruction
8.1.4 Quantification of carotid artery disease
8.2 Further work
8.2.1 Further validation of 3D ultrasound analysis
8.2.2 Clinical studies
8.2.3 Image segmentation
8.2.4 Other applications

Appendix 1 3D matrix transformations
Appendix 2 Miscellaneous algorithms
Appendix 3 A software tool for segmenting ultrasound images
References

List of figures

Chapter 1

1.1 Structure of a normal human artery

1.2 Development of a complex plaque

1.3 Anatomy of the human carotid artery bifurcation

1.4 An image of a diseased internal carotid artery obtained using digital subtraction angiography

1.5 Dependency of a digital subtraction angiogram on the projection plane

1.6 Interactions of ultrasound and soft tissue

1.7 A simple ultrasound transducer assembly

1.8 An example of a B-mode and an M-mode ultrasound image

1.9 The configuration of a Doppler ultrasound instrument

1.10 A spectral Doppler trace for the common carotid artery

1.11 A Duplex ultrasound image of the carotid bifurcation

1.12 Percentage errors in velocity estimates due to errors in Doppler angle

1.13 Examples of colour and power mode images

1.14 Comparison of normal and compound B-mode images

Chapter 2

2.1 Measurement of degrees of stenosis in the internal carotid artery

Chapter 3

3.1 Block diagram of a generalised 3D ultrasound acquisition system

3.2 Stages in the generation and analysis of a 3D ultrasound image

3.3 Three different types of constrained 3D ultrasound acquisition system

3.4 An example of a triangulated surface mesh

3.5 Mapping a 2D ultrasound image into a 3D voxel array

3.6 Weighted nearest-neighbour interpolation for a freely-acquired swept volume

3.7 Ray-tracing for 3D surface rendering

3.8 An example of a surface rendered mesh

3.9 Ray-tracing 3D volume rendering

3.10 3D ultrasound images of the carotid bifurcation displayed rendered using Voxblast

Chapter 4

4.1 The electromagnetic 3D ultrasound acquisition system developed for this thesis

4.2 An illustration of how the pcBIRD receiver is mounted on the scan-probe using a moulded plastic support and nylon screws

4.3 The configuration of the pcBIRD position tracking device for bilateral 3D carotid scanning

4.4 Graphical user interface for the 3D ultrasound acquisition software

4.5 Graphical user interface for the HDILab software

4.6 3D ultrasound data acquisition using the electromagnetic acquisition system

4.7 Sequence of events following the detection of an ECG R-wave during a 3D ultrasound acquisition

4.8 The effect on colour filling of changing the timing of image capture relative to the peak of the ECG R-wave

4.9 The 3D co-ordinate systems used for freehand 3D ultrasound imaging

4.10 Two different co-ordinate systems for acquired images

4.11 Examples of calibration phantoms for freehand 3D ultrasound acquisition systems

4.12 Calibration of the 3D ultrasound system using a pin phantom

4.13 A B-mode image of the pin phantom

Chapter 5

5.1 The lever-arm effect

5.2 Relative timings of the 3D ultrasound acquisition cycle

5.3 The effect of ultrasound beam thickness on locating a small object lying within the beam

5.4 Beam thickness as a function of penetration depth for the L10-5 and L12-5 scan-probes with a single focal zone placed at a depth of 20 mm

5.5 The effect of errors in pcBIRD orientation measurements on image pixel registration

5.6 The predicted dependency of K on the perpendicular distance between the surface of the ultrasound transducer and the origin of the EPOM receiver

5.7 The floor plan for the vascular ultrasound laboratory used for 3D ultrasound imaging of the carotid arteries

5.8 Four orientations of the scan probe and receiver combination used during measurement noise tests
5.9 Mean measurement deviations (ΔD and $\Delta\phi$) plotted against measurement rate for the DC+ACW and DC+ACN filter settings
5.10 Mean ± SD deviations for pcBIRD position and orientation measurements
5.11 The digital micrometer rig used to estimate the accuracy of pcBIRD position measurements
5.12 The apparatus used to estimate the accuracy of pcBIRD position measurements
5.13 The absolute errors in pcBIRD position measurements for translations of 0 to 90 mm
5.14 The modified goniometer used to estimate the accuracy of the accuracy pcBIRD rotation measurements
5.15 The absolute errors in pcBIRD rotation measurements for rotations up to 90°
5.16 The 'wedge phantom' designed for assessing the reconstruction accuracy of the electromagnetic 3D ultrasound system
5.17 Imaging the wedge phantom at varying scan depths
5.18 B-mode image of the wedge phantom
5.19 Reconstructed pinhead locations
5.20 3D co-ordinate system using in the calculation of reconstruction precision
5.21 Reconstruction precision measures for the pin matrix of the wedge phantom
5.22 RMS absolute errors in 3D pin location for 4 different scan depth ranges

Chapter 6

6.1 The generalised cylinder representation
6.2 The longitudinal curve representation
6.3 An example of a univariate spline function
6.4 Conventions for the definition of wall and plaque contours
6.5 The local 3D co-ordinate system for the 3D reconstruction of the carotid bifurcation
6.6 Centroid registration of cross-sectional wall contours
6.7 The application of the iterative weighted smoothing spline algorithm to noisy data with outliers
6.8 A schematic of the reconstruction of the 3D surface of a straight cylinder from cross-sectional contours

6.9 Matching vertices in consecutive contours
6.10 Vertex matching scheme at the carotid bifurcation
6.11 Longitudinal smoothing splines fitted to cross-sectional contours of a normal carotid bifurcation following the vertex matching procedure
6.12 Registration of plaque contours
6.13 Finding the intersection of a spline with a cut plane
6.14 Determination of the vessel central axis from a 3D surface reconstruction using the longitudinal spline representation
6.15 Calculation of lumen diameters from a transverse image
6.16 Co-ordinate system, I', aligned with a diameter line at angle θ to the fixed co-ordinate system, I.

Chapter 7

7.1 Lumen geometries of 3 anthropomorphic carotid ultrasound phantoms
7.2 Example B-mode images of an anthropomorphic carotid phantom
7.3 Sample cross-sectional area and circumference profiles
7.4 Absolute errors in lumen diameter and area for the three carotid artery phantoms
7.5 The steps involved in analysing patient 3D ultrasound sets
7.6 Power Doppler images of the ICA for patient 1
7.7 3D surface renderings of the reconstructed wall and lumen of the left carotid bifurcation for patient 1
7.8 Results of 3D quantitative analysis for patient 1
7.9 Power Doppler images of the ICA for patient 2
7.10 3D surface renderings of the reconstructed wall and lumen of the right carotid bifurcation for patient 2
7.11 Results of 3D quantitative analysis for patient 2
7.12 Power Doppler images of the ICA for patient 3
7.13 3D surface renderings of the reconstructed wall and lumen of the right carotid bifurcation for patient 3
7.14 Results of 3D quantitative analysis for patient 3
7.15 Power Doppler images of the ICA for patient 4
7.16 3D surface renderings of the reconstructed wall and lumen of the left carotid bifurcation for patient 4
7.17 Results of 3D quantitative analysis for patient 4

Appendix 1

A1.1 A left-handed 3D co-ordinate system with a point specified by the co-ordinates (x, y, z)

Appendix 3

A3.1 The GUI of the segmentation software used to define boundary contours on a series of transverse ultrasound images

List of tables

Chapter 2

2.1 Velocity (frequency) thresholds for discriminating > 70% stenosis using the NASCET stenosis index

Chapter 4

4.1 Manufacturer's specifications for the Ascension pcBIRD

4.2 Calibration parameter estimates for the L10-5 scan probe

4.3 Calibration parameter estimates for the L12-5 scan probe

4.4 Standard measures of calibration performance calculated for the L10-5 and L12-5 scan probes

Chapter 5

5.1 Root-mean-square absolute errors in pcBIRD position measurements for translations of up to 90 mm.

5.2 Root-mean-square absolute errors in pcBIRD rotation measurements for rotations of up to 90°

5.3 Reconstruction precision measures for the pin matrix of the wedge phantom

Chapter 7

7.1 Errors in anthropomorphic phantom reconstruction measurements

7.2 Absolute and relative errors for stenosis indices calculated from 3D phantom reconstructions

7.3 Results of quantitative analyses for 4 diseased carotid bifurcations

List of abbreviations

1D	One-dimensional
2D	Two-dimensional
3D	Three-dimensional
3DUS	Three-dimensional ultrasound
AC	Alternating current
ACAS	Asymptomatic Carotid Atherosclerosis Trial
CCA	Common carotid artery
CEA	Carotid endarterectomy
CT	Computerised tomography
CW	Continuous wave
DC	Direct current
DSA	Digital subtraction angiography
ECA	External carotid artery
ECG	Electrocardiogram
ECST	European Carotid Surgery Trial
EDV	End-diastolic velocity
EPOM	Electromagnetic position and orientation measurement
GUI	Graphical user interface
ICA	Internal carotid artery
ISI	Ideal stenosis index
IWSS	Iterative weighted smoothing spline
MIP	Maximum/minimum intensity projection
MRA	Magnetic resonance angiography
MRI	Magnetic resonance imaging
NASCET	North American Symptomatic Carotid Endarterectomy Trial
PC	Personal computer
POM	Position and orientation measurement
PP	Piecewise-polynomial
PRF	Pulse repetition frequency
PSV	Peak systolic velocity
PW	Pulse wave
RMS	Root-mean-square

SD	Standard deviation
SNR	Signal-to-noise ratio
TDI	Tissue Doppler imaging
TGC	Time-gain compensation
TIA	Transient ischaemic attack
TOF	Time-of-flight
VCR	Video recorder

'I often say that when you can measure what you are speaking about and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind.'

Lord Kelvin, 1867

Chapter 1

Introduction

This chapter begins with an introduction to carotid artery disease. The characteristics and consequences of the disease are stated, together with a brief discussion of current theories on its pathogenesis. The principal techniques for imaging the vasculature are then introduced, with particular emphasis on ultrasound imaging. The fundamental principles of ultrasound imaging are considered, and the most common ultrasound imaging modes are described.

1.1 Carotid artery disease

Atherosclerosis is characterised by the formation of fibrous, fatty lesions, called plaques, in the inner layers of the walls of larger arteries. The development of the disease is a gradual process which takes place in all humans, to a greater or lesser extent, with ageing. In the advanced stages, atherosclerosis may become symptomatic, and is a major cause of mortality and morbidity in the Western World. Heart attack, stroke, claudication, gangrene, and aortic aneurysms are all directly connected with atherosclerosis.

Atherosclerosis-related sequelae may occur via one of two mechanisms: the first is local ischaemia, whereby an artery becomes blocked to such an extent that blood flow in that portion of the vessel is severely impeded. The second mechanism is ischaemia in a downstream artery caused by a process known as embolisation. In this case, disruption of the plaque gives rise to small particles, or emboli, which are carried downstream in the blood, and eventually lodge in and occlude a smaller artery. Emboli may comprise either particulate material produced when a plaque ruptures (e.g. cholesterol), or thrombus which forms as a response to plaque rupture [van der Wal 1999; Arroyo 1999].

Figure 1.1 shows the structure of a normal human artery. The wall of the artery comprises three concentric layers – the intima, media and adventitia. The interior volume of the artery which carries blood is called the lumen, and the layer of cells in contact with blood is called the endothelium.

Atherosclerosis occurs in most large arteries, and a general predilection to arterial bifurcations has been observed [DeBakey 1985; Frangos 1999]. The disease is found commonly in the aorta, coronary, carotid, and femoral arteries. The aetiology of the disease is highly complex, and a multitude of biochemical and mechanical factors have been implicated. Among these, haemodynamic factors are thought to be critical [Nerem 1993]. These include parameters which describe mechanical interactions between flowing blood and the endothelium. In particular, there has been much interest in wall shear, and mounting evidence suggests that atherosclerosis tends to occur in areas of low shear stress [Ku 1985; Zarins 1983].

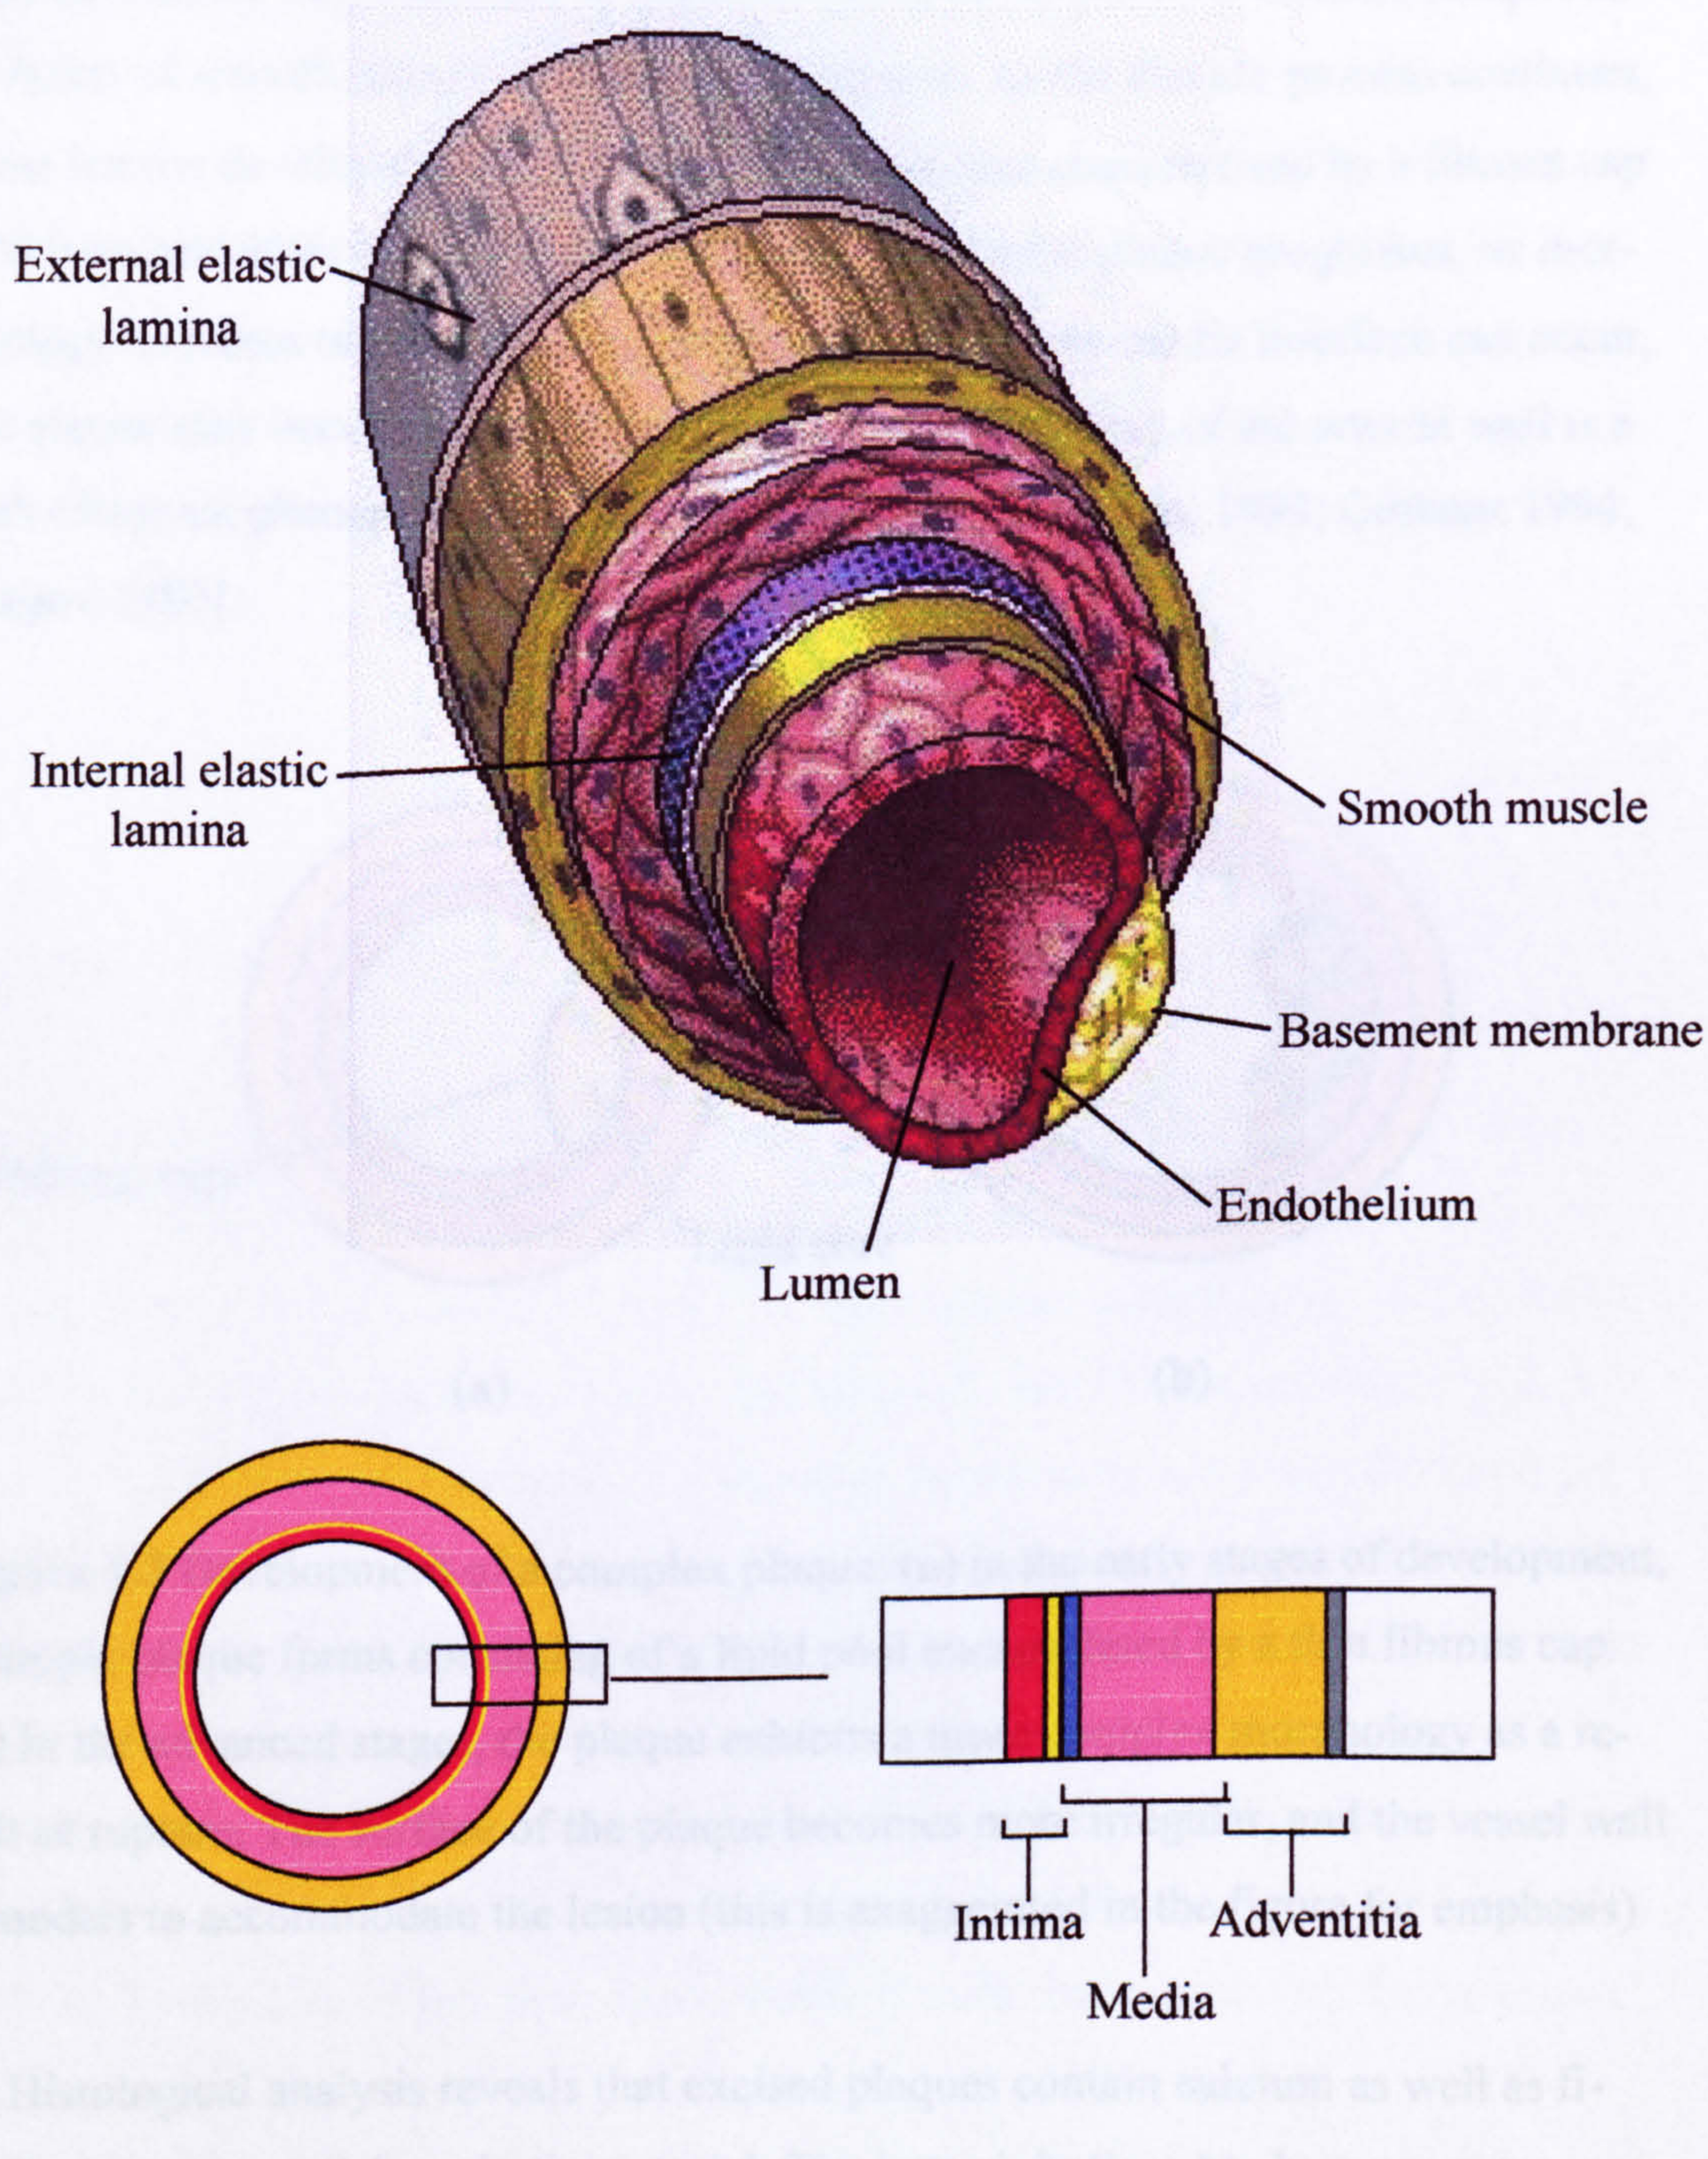


Figure 1.1 Structure of a normal human artery. The artery may be thought of as a tube made up from a number of concentric layers. These are conventionally described as three bands: the intima, the media, and adventitia. The interior volume of the artery, along which blood flows, is known as the lumen.

The earliest evidence of atherosclerosis is the formation of so-called 'fatty streaks', which are found in arteries in early life. The streaks result from an accumulation of lipid-rich macrophages and T-lymphocytes. Animal studies have demonstrated that fatty streaks are the precursors of more advanced atherosclerotic lesions, composed of layers of smooth muscle cells and macrophages. As the disease process continues, these lesions develop further to form 'simple' plaques characterised by a fibrous cap which encapsulates a lipid core (see Figure 1.2a). As the plaque progresses, its morphology becomes more complex: disruption of the intima-media interface can occur, the plaque may become more fibrous, and gross remodelling of the arterial wall is a well-observed phenomenon, as indicated in Figure 1.2b [Dzau 1993; Gibbons 1994; Glagov 1986].

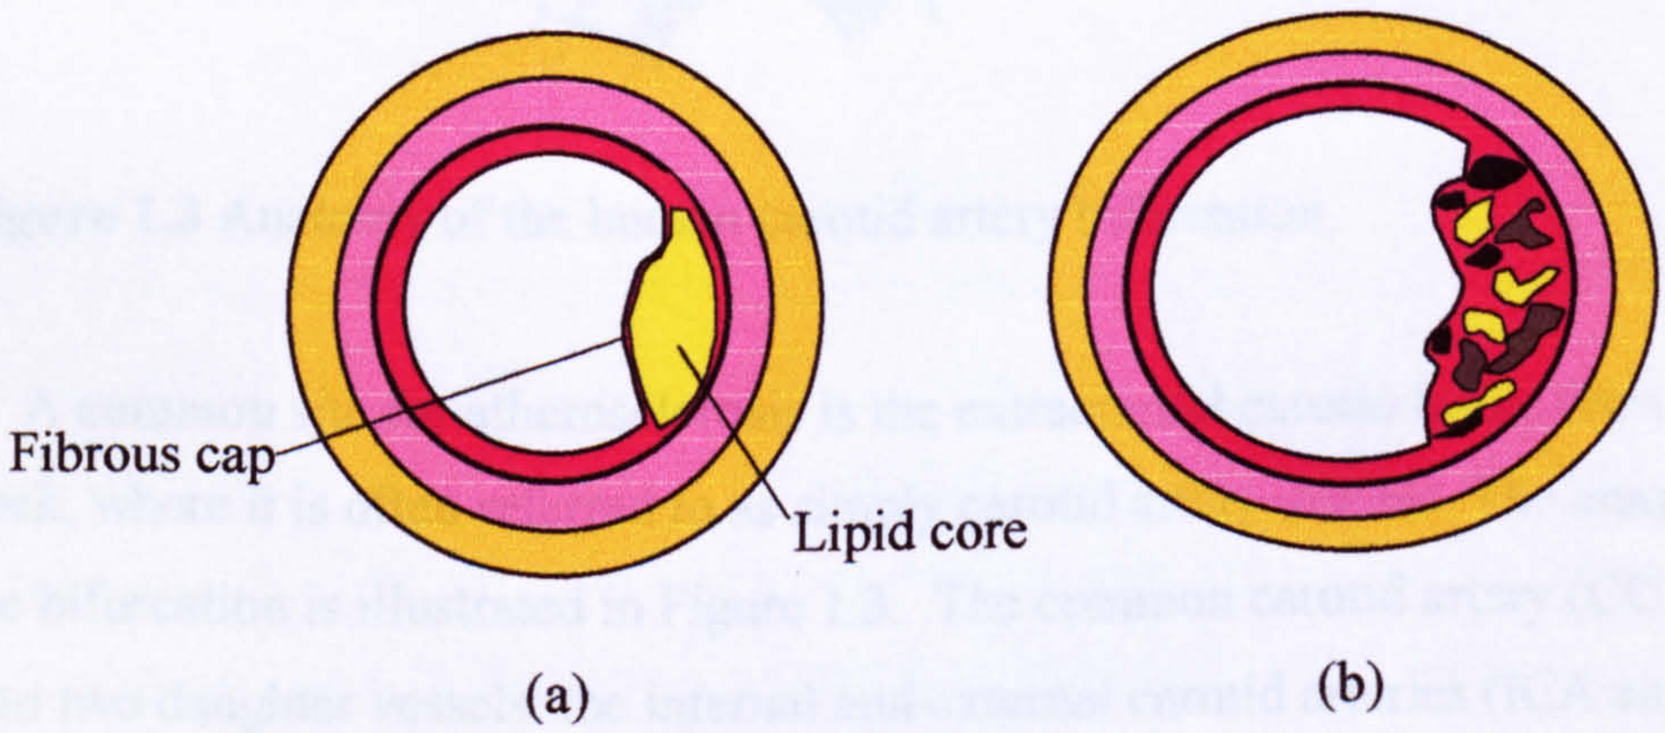


Figure 1.2 Development of a complex plaque. (a) in the early stages of development, a simple plaque forms consisting of a lipid pool encapsulated by a thin fibrous cap. (b) In the advanced stages, the plaque exhibits a more complex morphology as a result of rupture. The surface of the plaque becomes more irregular, and the vessel wall remodels to accommodate the lesion (this is exaggerated in the figure for emphasis).

Histological analysis reveals that excised plaques contain calcium as well as fibrous, necrotic, and thrombotic material. The latter is believed to be a consequence of plaque rupture. There is a large body of evidence which supports the current notion that atherosclerosis is a response to insults on the artery wall due to, for example, smoking, hypercholesterolaemia, diabetes, hypertension, and local haemodynamics. Central to this conjecture is the Response To Injury Hypothesis, which has been refined over many years to explain observed dysfunctions of endothelial and vascular smooth muscle cells when subjected to various insults [Ross 1993].

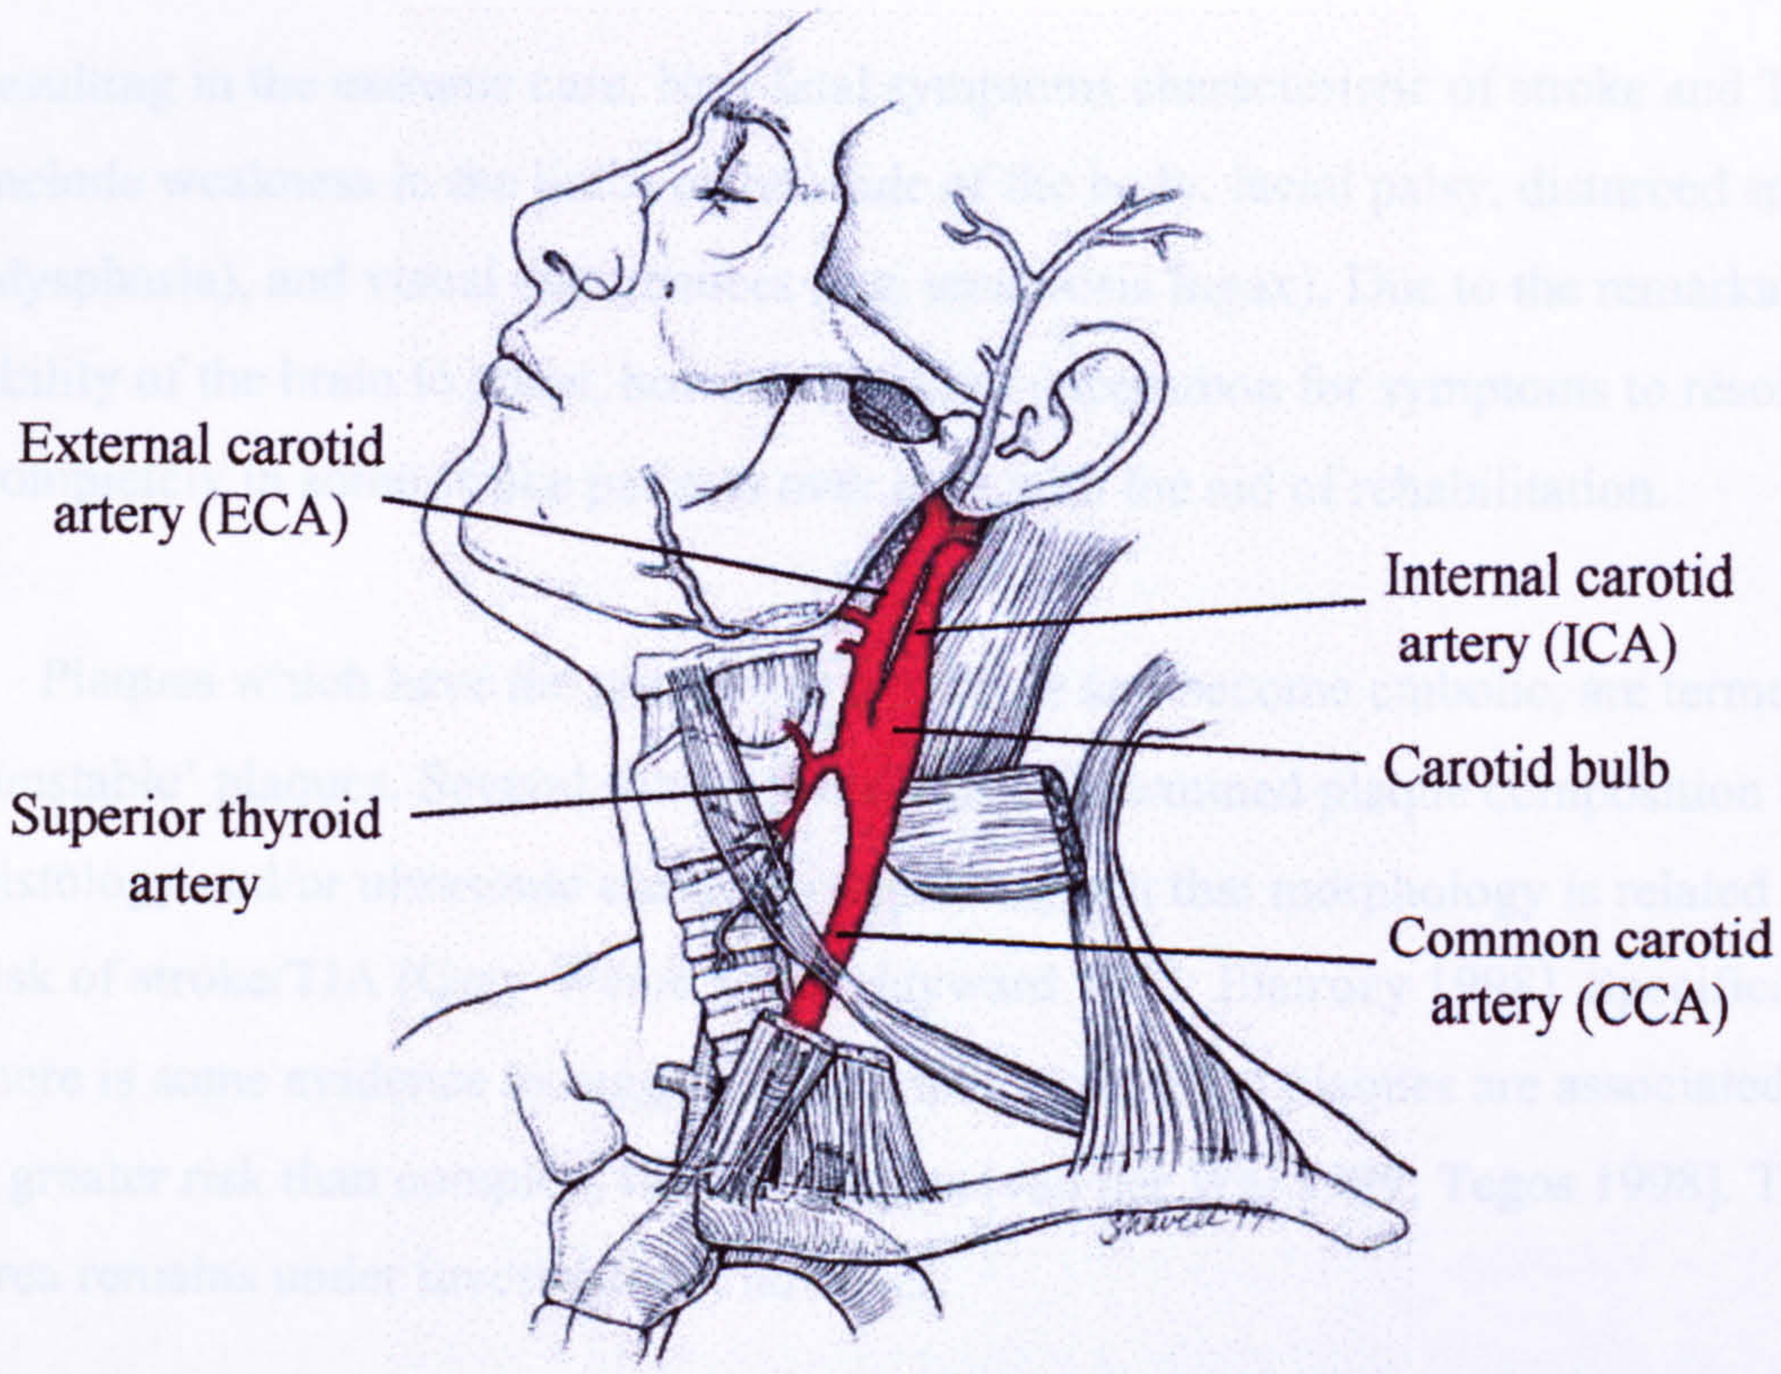


Figure 1.3 Anatomy of the human carotid artery bifurcation.

A common site for atherosclerosis is the extracranial carotid bifurcation in the neck, where it is often referred to as simply carotid artery disease. The anatomy of the bifurcation is illustrated in Figure 1.3. The common carotid artery (CCA) divides into two daughter vessels: the internal and external carotid arteries (ICA and ECA, respectively). Branches, including the superior thyroid artery, are usually found on the ECA close to the bifurcation. The proximal portion of the ICA dilates for the first few centimetres, forming a region known as the carotid bulb. There is one bifurcation on each side of the neck, with each ICA extending into the brain and connecting to a circular arrangement of cerebral arteries called the Circle of Willis. The ICAs are an important route to the brain for oxygenated blood, in addition to the two vertebral arteries. The ECAs, on the other hand, supply muscle, bone, and other facial tissue.

Carotid artery disease is the most common cause of embolisation to the brain, which is, in turn, responsible for a large number of strokes and transient ischaemic attacks (TIA) [Caplan 1993; Solberg 1971]. Symptoms arising from carotid embolisation may either by transient or long-term: when symptoms resolve within 24 hours, a TIA is said to occur; otherwise a stroke is diagnosed. Embolisation causes permanent damage to brain tissue, resulting in a localised region of atrophy called an infarct. Symptoms are dependent on both the location and size of an infarct, with death

resulting in the extreme case. Non-fatal symptoms characteristic of stroke and TIA include weakness in the limbs on one side of the body, facial palsy, disturbed speech (dysphasia), and visual disturbances (e.g. amaurosis fugax). Due to the remarkable ability of the brain to adapt, however, it is not uncommon for symptoms to resolve completely in some stroke patients over time with the aid of rehabilitation.

Plaques which have the propensity to rupture and become embolic, are termed 'unstable' plaques. Several studies which have determined plaque composition using histology and/or ultrasonic characterisation, suggest that morphology is related to risk of stroke/TIA [Gray-Weale 1988; Hayward 1995; Elatrozy 1998]. Specifically, there is some evidence to suggest that simple, lipid-filled plaques are associated with a greater risk than complex, fibrous plaques [van der Wal 1999; Tegos 1998]. This area remains under investigation, however.

As explained above, local haemodynamic forces are also believed to play a large part in the embolisation process. The interaction between plaque morphology and external forces is the most obvious mechanism for plaque rupture and subsequent embolism. Most of the studies in this area are based on histological plaque characterisation, but a few attempts have been made to gain insight into the issue of how and where plaques rupture using computer modelling [Cheng 1993; Richardson 1989].

Stroke is a common cause of mortality and morbidity, particularly among the elderly, with enormous social and economic implications. The cost of treatment and rehabilitation for stroke sufferers is huge, and, therefore, even modest improvements in the detection and management of high risk patients are likely to have a significant impact in reducing the overall burden of stroke.

The ability to image arteries, veins and blood flow is essential for the diagnosis and treatment of many diseases. Fortunately, a number of imaging techniques are now available for investigating the vascular system. Vascular imaging can be conveniently divided into techniques for imaging blood, and techniques for imaging the soft tissue of vessel walls. Blood imaging – historically known as angiography – has become invaluable in the assessment of the circulation of blood within large organs, as well as in the detection of atherosclerotic and aneurysmal disease. Soft tissue im-

aging, on the other hand, provides more detailed information about the vessel wall and associated disease. The principal vascular imaging methods available today are x-ray angiography, magnetic resonance imaging (MRI), and ultrasound imaging. X-ray angiography is the most well established method for imaging blood. MRI encompasses a range of techniques which may be used to image both blood and soft tissue. The latter is rarely used in clinical practice, however, because of the requirement for non-standard receiver coils and spin-echo sequences. Ultrasound imaging is unique in its ability to image both blood and soft tissue simultaneously. Each imaging technique is considered in more detail in the following sections.

1.2 X-ray angiography

X-ray angiography is currently the most widely used vascular imaging technique. An intra-arterial injection of a bolus of radio-opaque contrast agent (typically, an iodine-based compound which attenuates x-rays) is administered. X-rays are emitted and pass through the patient to a detector on the opposite side. An image is then produced from the intensity distribution of detected x-ray energy. Blood vessels are visible in the x-ray image because the attenuation of the contrast agent carried in the blood is higher than that of the surrounding tissue, leading to regions with markedly different intensities. An image obtained in this way is termed an (x-ray) angiogram, and represents a two-dimensional (2D) projection onto a plane of contrast agent carried in the blood. Conventionally, x-ray film is used to detect, display, and store images, but some modern scanners use electronic x-ray detectors to generate digital angiograms.

A variation on the basic technique described above is digital subtraction angiography (DSA). In this method, a sequence of digitised images is acquired immediately before and immediately after the contrast injection. Two images, one selected from each sequence, are then subtracted. This has the effect of greatly increasing the signal-to-noise ratio (SNR) of the blood, since most of the 'clutter' in the image, arising from surrounding structures, is removed. The result is an image where only the blood is visible, as shown in Figure 1.4, and, therefore, the vessel lumen can be easily distinguished.

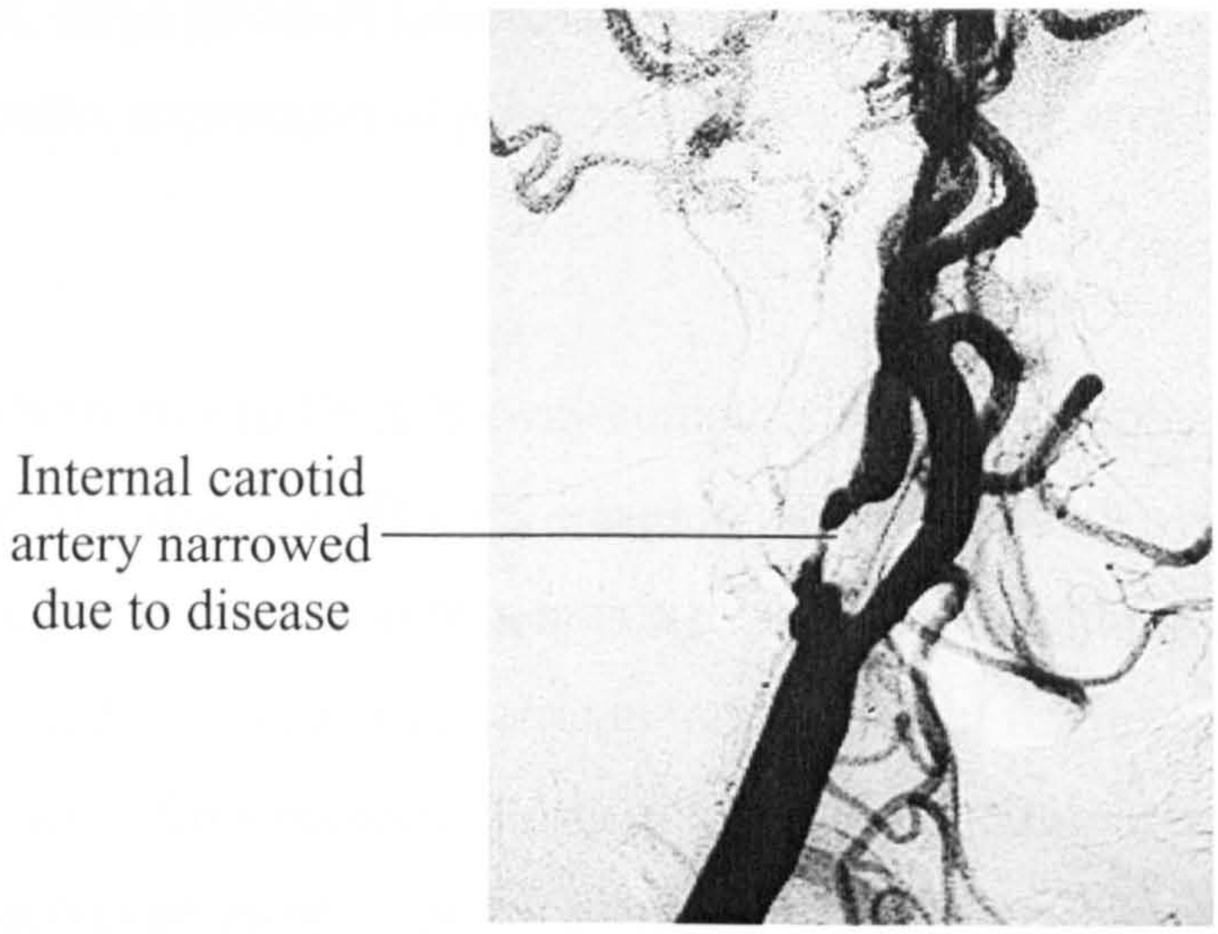


Figure 1.4 An image of a diseased internal carotid artery obtained using digital subtraction angiography. The atherosclerotic portion of the artery is evident from the narrowing of the lumen.

A fundamental limitation of images produced by x-ray projection is that it is impossible to distinguish two structures both lying on the same straight line path between the x-ray source and the projection plane. This situation is illustrated in Figure 1.5.

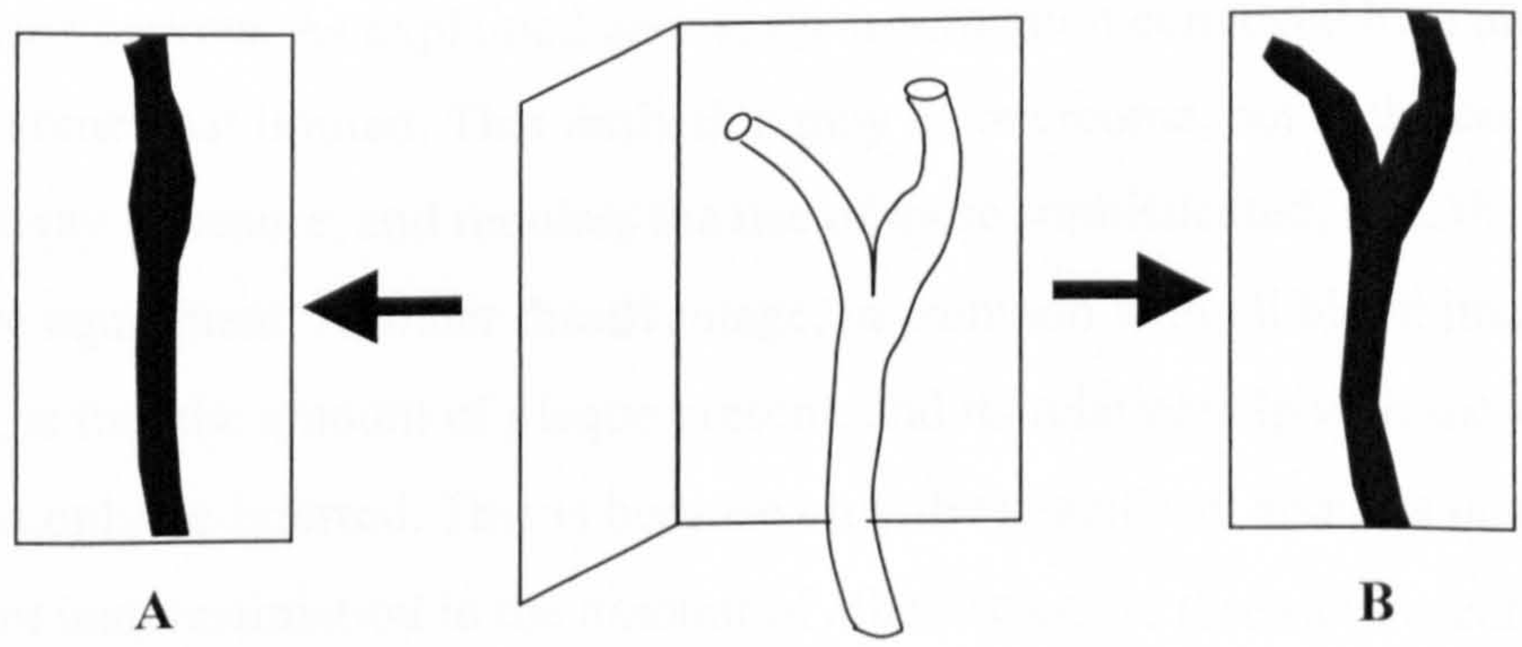


Figure 1.5 The dependency of the DSA image on the orientation of the projection plane. **A** and **B** are images obtained using two orthogonal projection planes. It can be seen from image **A** that the daughter vessels (the ICA and ECA in the this example) cannot be differentiated as in image **B**.

The problem is especially troublesome when the aim of imaging is to visualise or quantify suspected atherosclerotic disease. Where the geometry of the disease is

complex, large errors in lumenal diameter measurements may occur. Therefore, to gain a better impression of plaque geometry, multiplanar or 3D angiography is desirable.

An alternative to DSA is x-ray computerised tomography (CT) [Dillon 1992; Leclerc 1995]. This is a 3D x-ray imaging technique which overcomes the limitations of single plane projection by acquiring a number of contiguous 2D images, or slices, which together form a volumetric representation of the imaged region. Each slice is reconstructed from projection data obtained from multiple paths through the patient. This is achieved using a variety of configurations; for instance, a stationary ring of detectors and a rotating x-ray source. The latest generation of CT scanners use fast helical acquisition techniques, known as spiral CT [Link 1996; Simeone 1997].

X-ray angiography provides high resolution, high contrast images of the vasculature using well-established technology. However, the necessity for a contrast injection means that the technique is invasive, and, consequently, there is an associated risk of complications [Hankley 1990; Davies 1993; Waugh 1992; Pryor 1996; Warnock 1993]. Further, inherent to all x-ray techniques is the use of ionising radiation. Hence, there is a limit on the number of images which may be obtained to ensure that exposure is limited to safe levels. X-ray angiography is, therefore, inappropriate for serial investigations. As explained above, the information conveyed by a uniplanar image is somewhat limited. This limitation may be overcome, but at the cost of increased x-ray exposure, and requires the use of more sophisticated, and, therefore, expensive equipment. Another disadvantage, in common with all blood imaging methods, is that the amount of plaque present, and its relationship with the vessel walls, can only be inferred. This is because only the visualised, and this may lead to a significant underestimation in the amount of atherosclerotic disease present, especially when local dilatation of the vessel was has taken place due to arterial remodelling.

1.3 Ultrasound imaging

Ultrasound has also become a very widely used vascular imaging technique. Due to technical limitations, the technique is currently confined to imaging arteries more than one millimetre in diameter, and there are constraints on which anatomical re-

gions can be investigated, due to their depth or the presence of dense obstructing tissue, such as bone. The method is, however, becoming established as the first-line method of choice for diagnosing peripheral vascular and cerebrovascular disease. From this perspective, ultrasound imaging is attractive for four main reasons:

- Most medical ultrasound is non-invasive – the majority of ultrasound scans are obtained transcutaneously, using a hand-held scan probe placed on the skin.

- Diagnostic ultrasound does not use ionising radiation, and is safe in the sense that there is no evidence of significant biological effects following routine exposure.

- Both blood and vessel walls can be imaged simultaneously.

- Ultrasound is a relatively inexpensive technique, both in terms of the initial cost of a scanner, and the environment required to house it; scanners are moveable, and do not require and special precautions, such as radiation screening, before use. The moveability of ultrasound scanners makes them ideal for use in situations where it is difficult or impossible to move a patient to another area of a hospital. The operational costs of ultrasound are also low, since ultrasound consumables are inexpensive unlike those, for instance, required in nuclear medicine.

1.3.1 Fundamental physics of ultrasound imaging

In general, ultrasound imaging relies on the use of a vibrating transducer to transmit and receive short pulses of high frequency sound waves. Typical frequencies are in the range 2 to 15 MHz, depending on the anatomical region under investigation. When the transmitted pulses propagate through tissue, echo pulses are produced by the interaction between tissue and the ultrasonic waves. These are detected by the transducer and the signals processed to form an image. This type of imaging is frequently referred to as pulse-echo imaging, and is used in nearly all of the ultrasound imaging modes.

(continued on page 31)

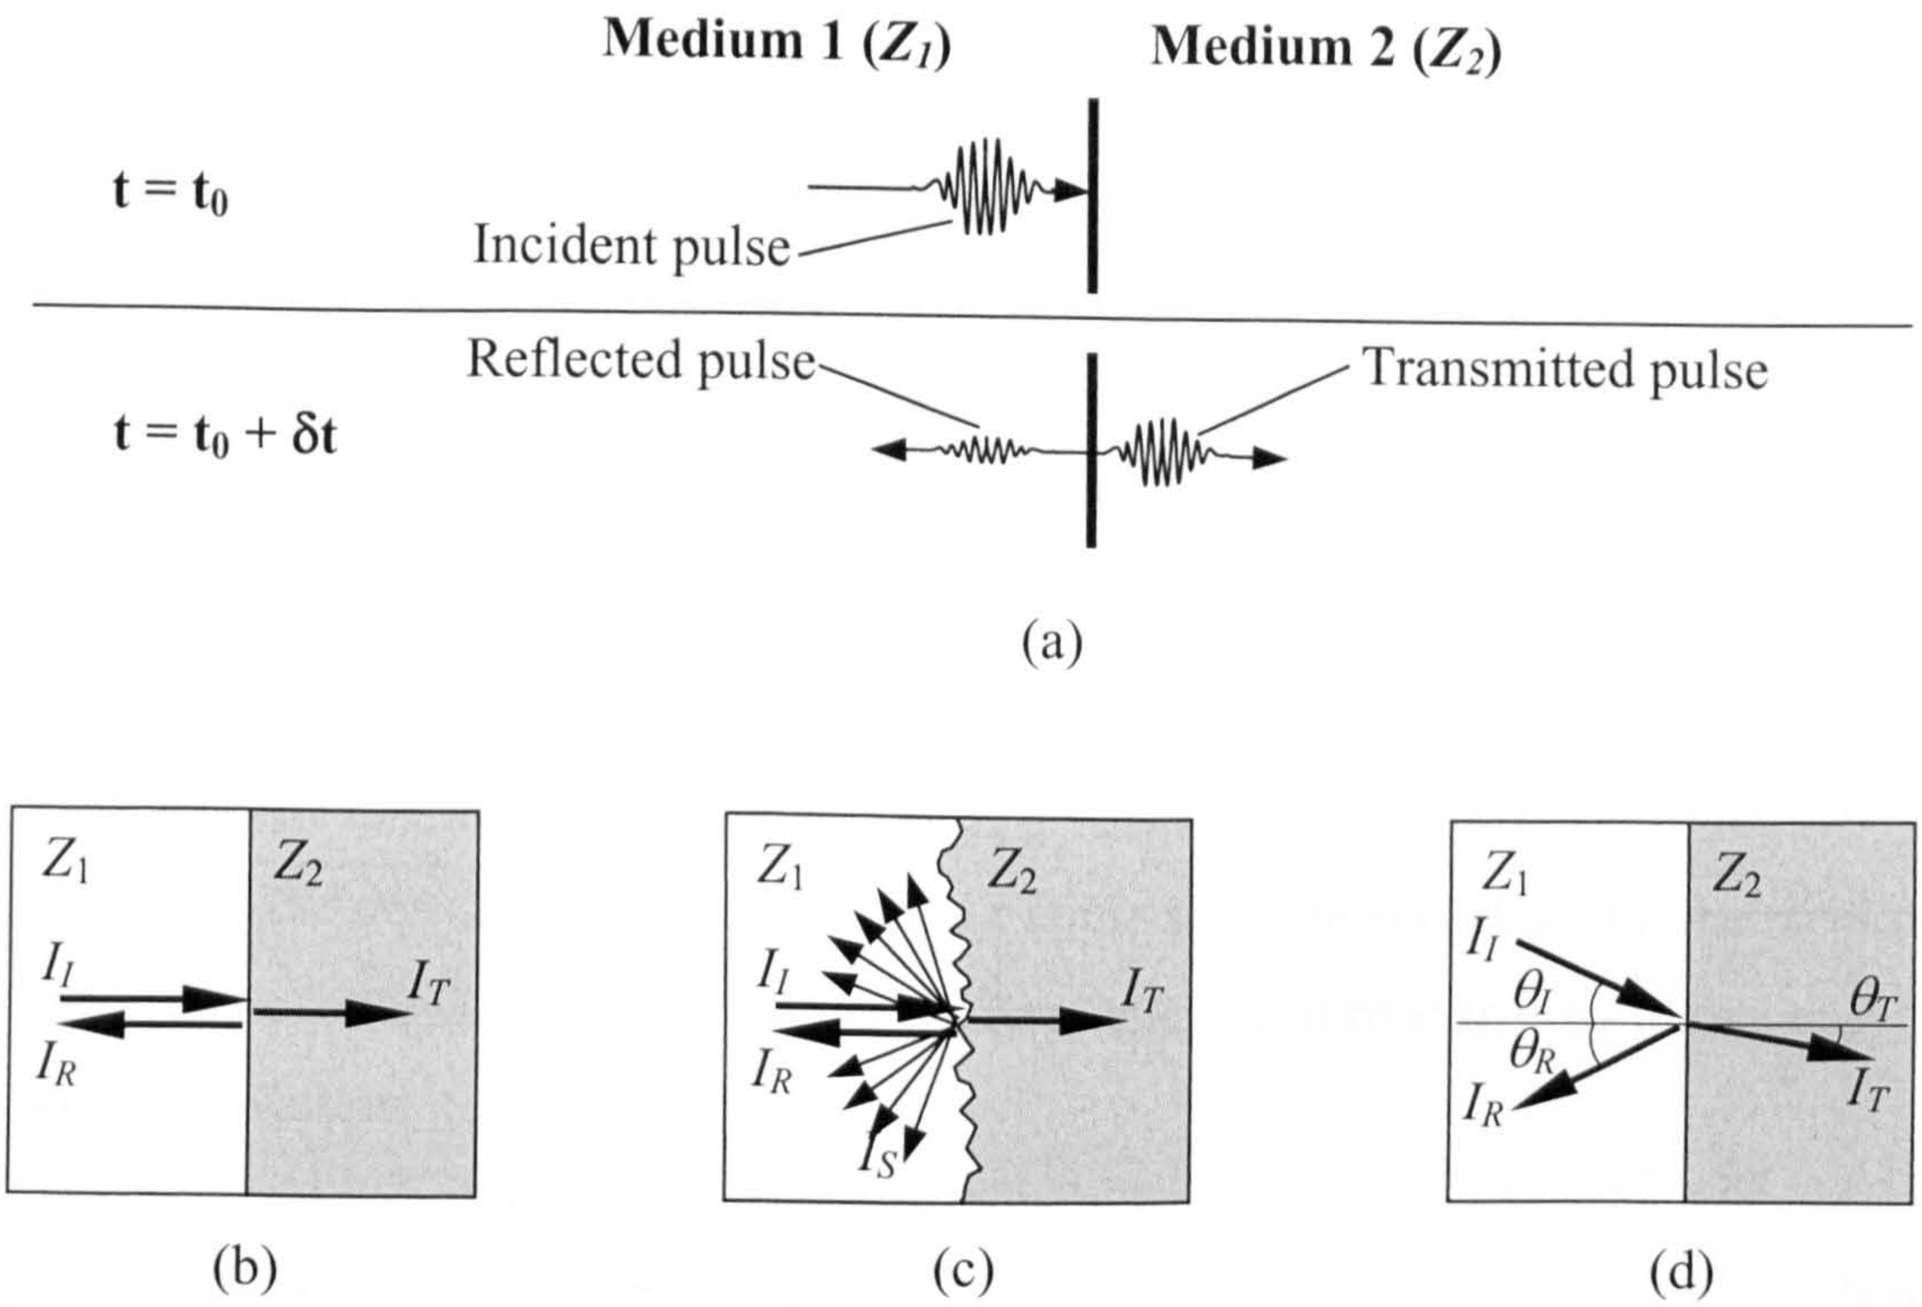


Figure 1.6 Interactions of ultrasound and soft tissue. (a) The behaviour of an ultrasound pulse reaching an interface which separates media with different acoustic impedances, Z_1 and Z_2: at time t_0 the incident ultrasound pulse is propagating through Medium 1. A short time later ($t_0 + \delta t$), the pulse has travelled across the interface into Medium 2 where its energy has diminished, and it is referred to as the transmitted pulse. Between these times, when the incident pulse reaches the interface, a reflected pulse, or echo, is produced, which travels in the opposite direction. (b) Specular reflection occurs when the interface is smooth relative to the wavelength of the incident ultrasound. Here, I_I, I_R and I_T denote the intensities of the incident, reflected and transmitted ultrasound waves, respectively. (c) Reflection from a rough interface produces scattered waves in multiple directions with intensity I_S. (d) When the direction of propagation of the incident ultrasound is not perpendicular to the interface, energy is reflected in the direction indicated, with $\theta_I = \theta_R$. In this case, energy is transmitted at an angle θ_T to the perpendicular. When $\theta_T > 0$, this phenomenon is known as ultrasonic refraction.

Classical wave theory predicts the behaviour of ultrasound as it propagates. The predominant interactions of ultrasound and soft tissue are illustrated in Figure 1.6. In particular, echoes arise when propagating ultrasonic waves travel between tissue types with different acoustic impedances. This is shown in Figure 1.6(a), and, using this effect, it is possible to distinguish tissue interfaces by detecting successive echoes. Moreover, their relative depths can be calculated by measuring the time difference between the arrival of echoes at the transducer surface (assuming a constant speed of sound). The acoustic impedance, Z, of a propagation medium is given by:

$$Z = \rho c, \tag{1.1}$$

where ρ is the density of the medium, and c is the speed of sound within the medium. The speed of sound, c, is a function of the frequency, f, of propagating ultrasound with the relation:

$$c = f\lambda, \tag{1.2}$$

where λ is the wavelength of the transmitted ultrasonic waves. The value of c is usually assumed to be 1540 m/s for soft tissue.

When the wavelength of ultrasound incident to an interface between two media is large relative to the microscopic 'roughness' of that interface, so-called specular reflection takes place, whereby energy is reflected in a single direction. Figure 1.6(b) shows a simple example of specular reflection where the directions of propagation of both the incident and reflected waves are perpendicular to the interface, but antiparallel to each other. The more general case is depicted in Figure 1.6(d) where the incident ultrasound strikes the interface at angle θ_I to the perpendicular, and this is equal to the angle of reflection, θ_R.

Using an idealised model of specular reflection, the proportion of ultrasound energy reflected is measured by the intensity coefficient for specular reflection, defined by:

$$\alpha_R = \frac{I_R}{I_I} = \left(\frac{Z_1 - Z_2}{Z_1 + Z_2} \right)^2 . \tag{1.3}$$

In this expression, I_I and I_R are the intensities – i.e. energy per unit area – of the incident and reflected ultrasound waves, respectively. Z_1 and Z_2 are acoustic impedances of the media on either side of the interface, as indicated in Figure 1.6. If the wavelength of the incident ultrasound is small relative to the roughness of the interface, diffuse reflection occurs, and ultrasound is reflected in many different directions, as illustrated in Figure 1.6(c). This phenomenon is called scattering and is important in the context of medical ultrasound imaging since it is the dominant mechanism through which ultrasound is returned to the transducer; ultrasonic waves are frequently described as having been 'backscattered' to the transducer.

Soft tissue is a complex, composite material made up from many different types of biological material containing scatterers of different sizes. Therefore, scattering occurs continually as ultrasound propagates through tissue, resulting in the attenuation of both transmitted and reflected waves. Consequently, received pulses have a much smaller amplitude than the initial transmitted pulses. Hence, sensitive receiving instrumentation with high amplification is required to detect them. Scattering is responsible for the characteristic 'grainy' texture of ultrasound images, known as speckle, that is produced by constructive-destructive interference of scattered ultrasonic waves.

A third important interaction is refraction. This is depicted in Figure 1.6(d) and occurs when the direction of propagation of incident ultrasound is not perpendicular to an interface. In this case, the direction of the transmitted ultrasound is different to that of the incident ultrasound resulting in a change of direction of the propagation path, and is governed by Snell's law.

Modern ultrasound scanners employ a transducer made up from a series of synthetic piezoelectric transducer elements which convert electrical energy into mechanical vibrations during the transmit phase, and vice versa during the receive phase. Figure 1.7 below shows the structure of a simple, single element ultrasound transducer assembly. A pulse is generated by 'kicking' the transducer using a short electrical pulse, which causes the surface of the transducer to vibrate at its natural

frequency (determined by its thickness). The backing block is necessary to absorb energy emitted from the upper part of the piezoelectric element which may otherwise lead to unwanted echoes from the transducer housing. It also damps the vibrations so that a short pulse is emitted containing only a few cycles. A thin matching layer provides optimal acoustic energy transfer into tissue, and, during a scan, water-based gel is used to provide further acoustic coupling between the transducer surface and the skin.

Nowadays, ultrasound transducers transmit as well as receive. This is possible because of the inverse piezoelectric property, whereby vibrations at the surface of the piezoelectric element generate an electrical voltage across the electrodes. A transducers will alternate between transmit and receive periods, which imposes a limit on the rate at which pulses are transmitted, known as the pulse repetition frequency (PRF). The PRF is important since it determines the maximum depth from which echo pulses can be received without range (i.e. depth) ambiguity.

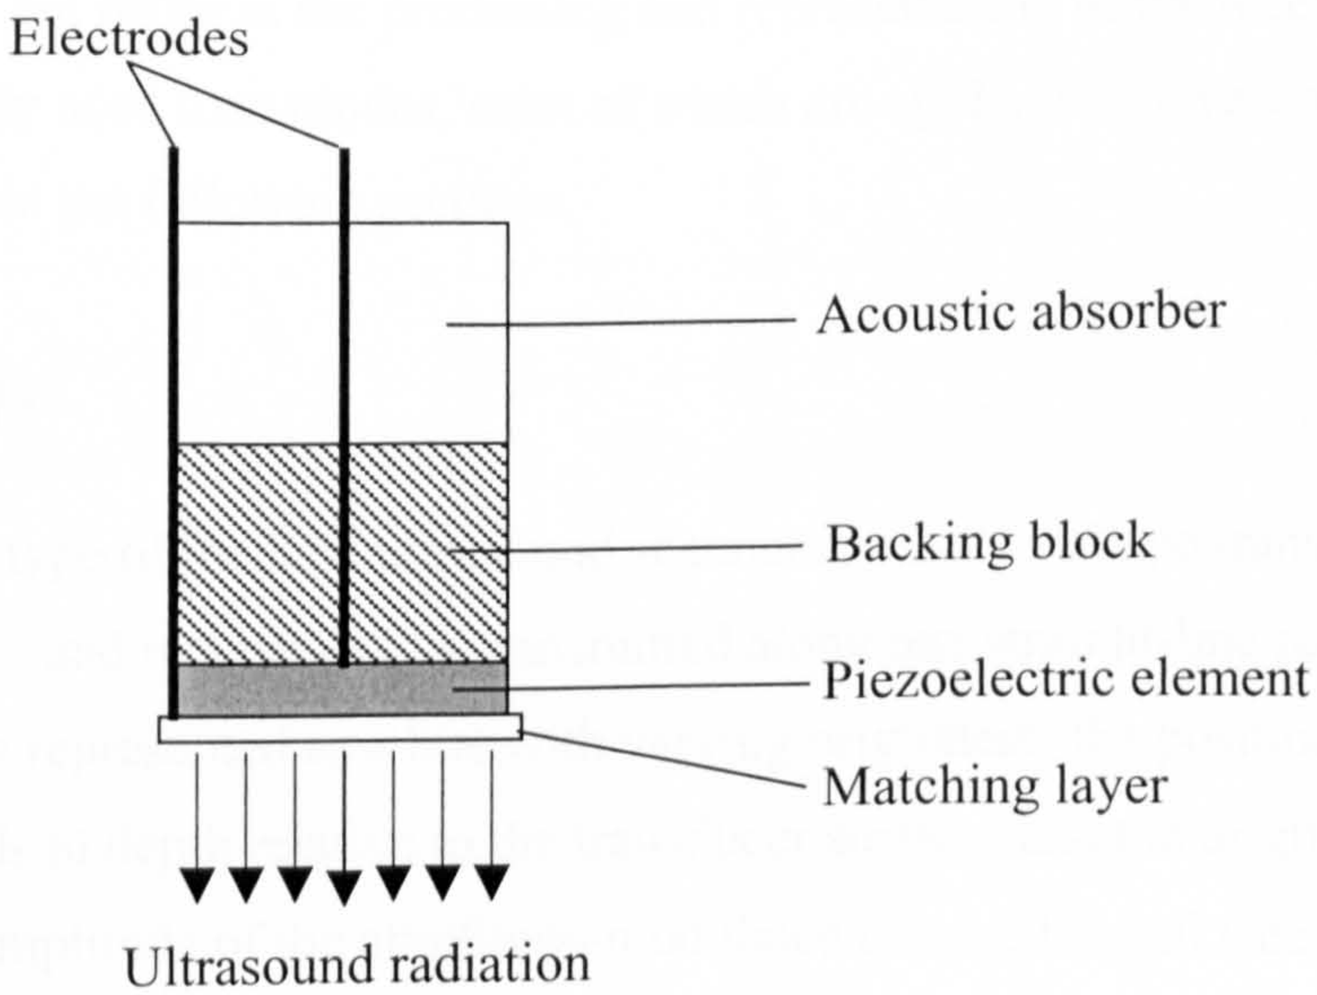


Figure 1.7 A simple ultrasound transducer assembly.

When a series of echo pulses are received a short time after transmitting a pulse, the signal contains the relative timings of the pulses, their amplitudes, and phase information. The depth, d, from which each echo pulse is received is found by converting the time, t, to receive the pulse into distance using the simple formula:

$$d = ct/2, \tag{1.4}$$

and assuming a constant value of *c*. The amplitudes of the pulses are found by rectifying and low-pass filtering the received signal following amplification. This process, called amplitude modulation, retains the envelope of the original high frequency signal, and is the most common method for generating soft tissue images using ultrasound as the phase information is very rarely utilised.

As noted above, ultrasound pulses are attenuated as they propagate through tissue. Hence, received echoes from deep structures are very weak compared with those from more superficial structures. To compensate for this effect, the amplitudes of signals from different depths are equalised by applying a depth-, or, equivalently, a time-dependent gain such that the echoes from deeper structures are amplified more than those from superficial structures. This process is called time-gain compensation (TGC), and is normally adjusted on the scanner by the operator.

Many ultrasound scan-modes have been devised which utilise the basic pulse-echo principle, but differ in the processing and representation of the received signal. The most widely used scan modes, most of which are applicable to vascular imaging, are introduced in the following sections.

1.3.2 A-Mode

The earliest type of medical ultrasound scanners possess a single transducer element to transmit and receive pulses transmitted along one straight-line path. Echo signals are then represented as a line with varying brightness; the position along the line corresponds to depth relative to the transducer surface, and the brightness corresponds to the amplitude of the amplitude-modulated received signal (see Figure 1.8(a)). Such a line is known as an amplitude line, or A-line, and the corresponding imaging mode is called A-mode. A-mode ultrasound has now been superseded by B-mode, described below, where multiple A-lines are combined to form a 2D image with width as well as depth.

1.3.3 M-Mode

M-mode ultrasound, where 'M' stands for 'motion', provides a representation of moving structures by combining A-lines acquired over time. The A-lines for each time interval are displayed side-by-side in the form of a digital image as illustrated in

Figure 1.8(b). The horizontal distance across the image represents time, whereas the vertical distance represents depth. M-mode has very high temporal resolution and is still useful for studying motion of structures, such as the heart valves. The technique is, however, limited to imaging along a single line.

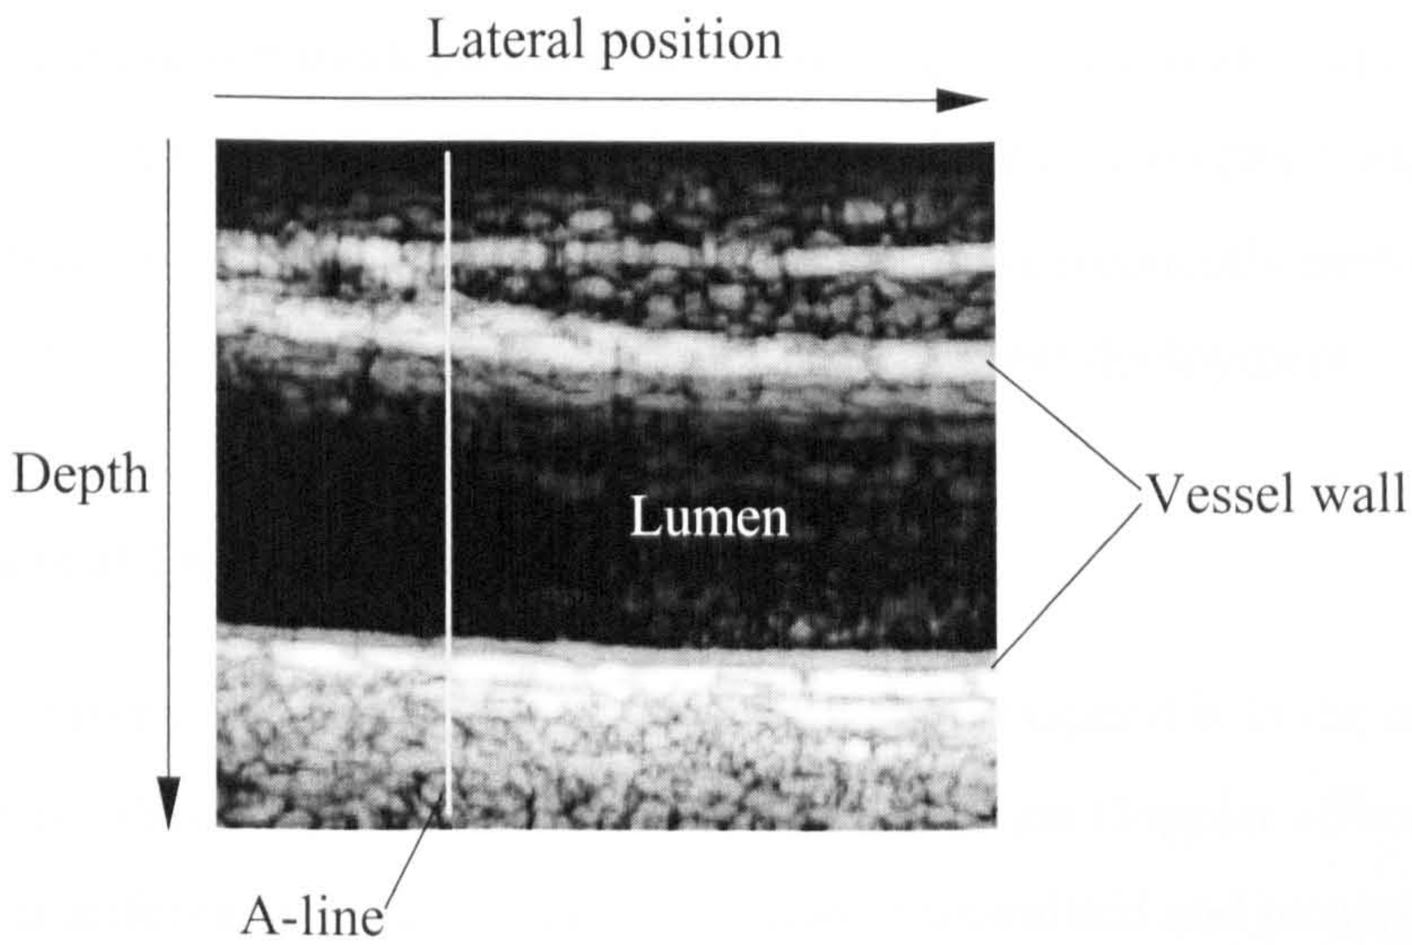
1.3.4 B-Mode

A brightness-mode, or B-mode, image is obtained by generating a series of A-lines, simultaneously or in close succession, at different transducer positions. The lines are then displayed together to form an image, such as the one shown in Figure 1.8(a), which represents a thin tomographic slice through the body. In a B-mode image the vertical direction represents depth and the horizontal direction represents lateral position. The pixel brightness is normally encoded using an 8-bit grey-level intensity scale, with values in the ranging from 0 (black) to 255 (white).

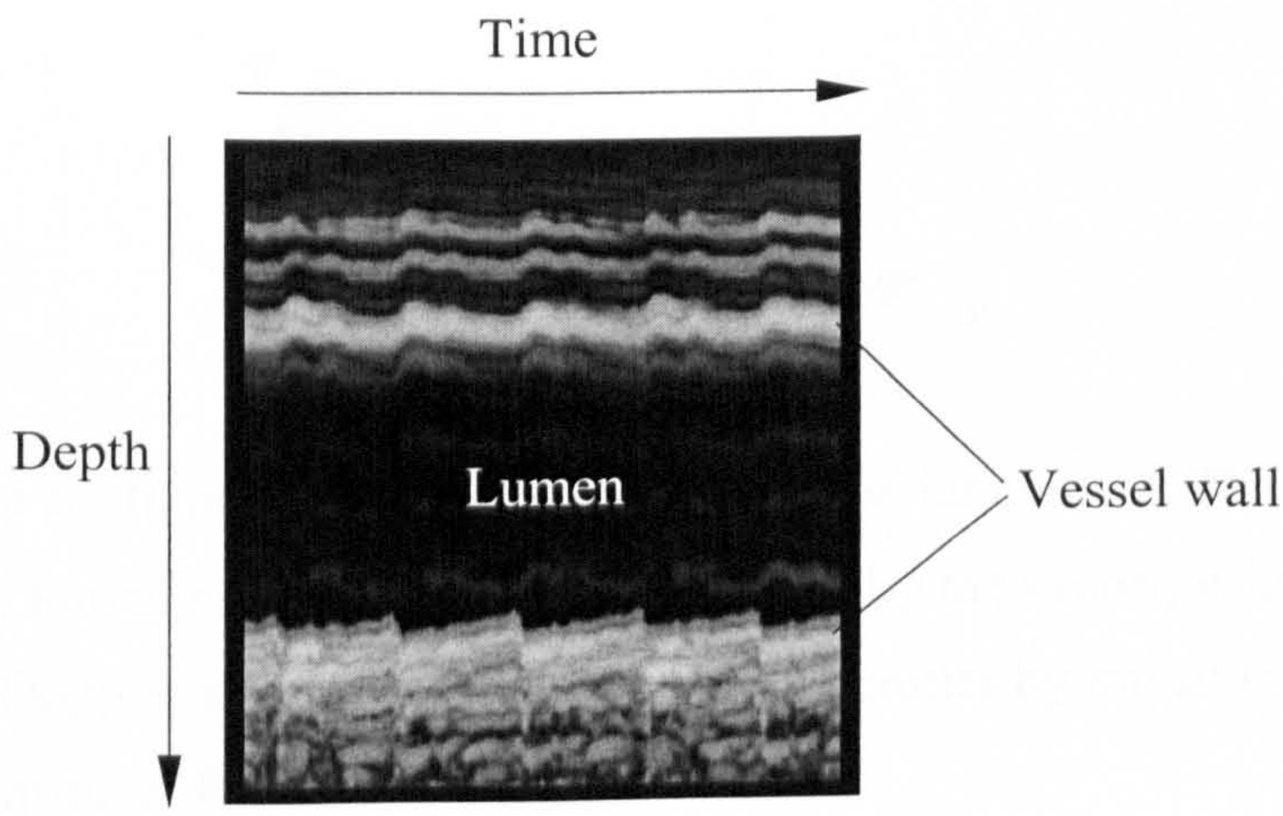
Early B-mode devices obtained A-lines by mechanically moving a single transducer within the scan-probe casing, either laterally, or in an angular motion to produce a sector scan. More recent technology utilises linear or curvilinear arrays of transducers which are activated in sequence electronically. Phased-array transducers activate transducer elements according to a preset pattern of delays, enabling the ultrasound beam to be steered, focused and shaped. The ability of modern scanners to rapidly acquire, process, and display B-mode images means that real-time imaging is feasible, with frame-rates in excess of 20 Hz readily achievable. The high temporal resolution of ultrasound is a significant advantage over alternative imaging methods, and makes it particularly suitable for investigating moving structures, such as the heart.

Intravascular B-mode imaging uses a transducer mounted on an intra-arterial catheter. Images are reconstructed from radial, outwardly projecting A-lines. The A-lines are produced at different angles by either mechanically rotating the transducer, rotating an acoustic mirror, or using phased-array technology to electronically steer the ultrasound beam [Liu 1999]. Due to the close proximity of the transducer to the vessel wall, and the low attenuation coefficient of blood, it is possible to use much higher frequencies than percutaneous imaging (typically around 30 MHz). This, in

(continued on page 37)



(a)



(b)

Figure 1.8 An example of a B-mode and an M-mode ultrasound image. (a) A B-mode image of the CCA formed from a series of A-lines. (b) An M-mode image which depicts the pulsatile motion of the artery wall. Each vertical line corresponds to the A-line highlighted on the B-mode image in (a) at different time steps.

turn, enables images with very high spatial resolution to be obtained.

Intravascular ultrasound has the drawback of being an invasive technique. Nevertheless, the technique provides a valuable alternative to x-ray angiography, where arteries are difficult or impossible to image using non-invasive ultrasound. A good example is the coronary arteries, where catheterisation is frequently performed during interventional procedures, such as angioplasty or stent deployment.

1.3.5 Spectral Doppler

Another extremely useful application of medical ultrasound is in the estimation of blood velocity. This is accomplished by making use of the Doppler effect, which is manifest as a difference in the frequency between transmitted and received ultrasound, known as a Doppler shift. A shift occurs when ultrasound is backscattered by moving blood particles.

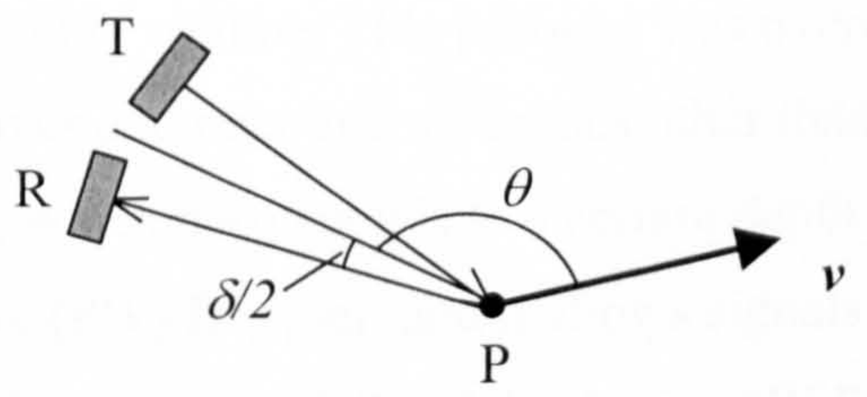


Figure 1.9 The configuration of a Doppler ultrasound instrument. The velocity of the scatterer, P, moving with a true velocity denoted by the vector, $\boldsymbol{v}$, is estimated by measuring the Doppler shift between the signal transmitter by transducer T and that received by transducer R.

For the configuration shown in Figure 1.9, with transmit transducer, T, and receive transducer, R, the Doppler shift, f_S, for the moving scatterer, P, is given by:

$$f_S = f_R - f_T = \frac{-2vf_T \cos\theta \cos(\delta/2)}{c}, \tag{1.5}$$

where f_R is the frequency of the received ultrasound and f_T is the frequency of the transmitted ultrasound. The angle θ is known as the Doppler angle, and, with reference to Figure 1.9, is the angle between the velocity vector $\boldsymbol{v}$ and the bisector of TP and RP. The variable v is the magnitude of vector $\boldsymbol{v}$ and is the speed of the scatterer, P, and c is the speed of sound in blood. Rearranging Equation (1.5) gives:

$$v = \frac{-cf_S}{2f_T \cos\theta \cos(\delta/2)}. \tag{1.6}$$

Since δ is known, and c and θ can be estimated, the speed of blood (v) can be determined by measuring the Doppler shift, f_S, and applying Equation (1.6) above.

Early Doppler ultrasound instruments used two transducers, as shown in Figure 1.9. In this configuration, one transmitted and one received continuously, in contrast to the pulse-echo sequence indicated earlier in Section 1.3.1. Such devices are called continuous wave (CW) Doppler instruments, and these provided the first ultrasonic method for estimating the speed of blood. The sample volume of a CW device, from which back-scattered ultrasound is detected, is the region where the transmit and receive beams overlap. Given the comparatively large size of this sample volume, it is impossible to discriminate between Doppler shifts caused by moving blood at different sites within the sample volume. This problem was overcome by transmitting pulsed ultrasound waves and rejecting all echoes other than those received within a narrow time window, which corresponds to a certain depth range. This technique is known as pulsed wave (PW) Doppler, and it allows signals to be received from a very small sample volume at a predefined depth. For PW Doppler only a single transmit-receive transducer is required, and, in this case, Equation (1.6) reduces to

$$v = \frac{-cf_S}{2f_T \cos\theta}. \tag{1.7}$$

Received ultrasound signals undergo both amplification and demodulation: demodulation typically involves low-pass filtering the signal obtained by multiplying the received signal by a sinusoidal reference signal with the same frequency as the transmitted ultrasound. Two reference signals are used: the in-phase and quadrature signals, with a phase difference of 90°. The signals resulting from the multiplication form the real and imaginary parts of a complex Doppler signal which encodes the direction of the Doppler shift; that is, either towards or away from the transducer [Fish 1999].

A high-pass filter is applied to the demodulated Doppler signal to remove large unwanted components produced by moving vessel walls. The cut-off for this filter, known as the wall filter, is typically about 200 Hz, and can usually be adjusted by the operator. Although the wall filter is necessary to remove high amplitude wall signals, there is the undesirable side-effect that low amplitude Doppler shifts, corresponding to slow near-wall velocities, are also removed.

The spectrum of a Doppler signal is displayed by grey-scale encoding the amplitude in a similar manner to that already discussed for B-mode images (see Figure 1.10). Doppler spectra from different time samples are displayed successively to provide a scrolling display which mirrors the temporal velocity profile of blood moving through the sample volume. This display is commonly referred to as a spectral Doppler trace, and is used when measuring blood velocity at a particular point in the cardiac cycle.

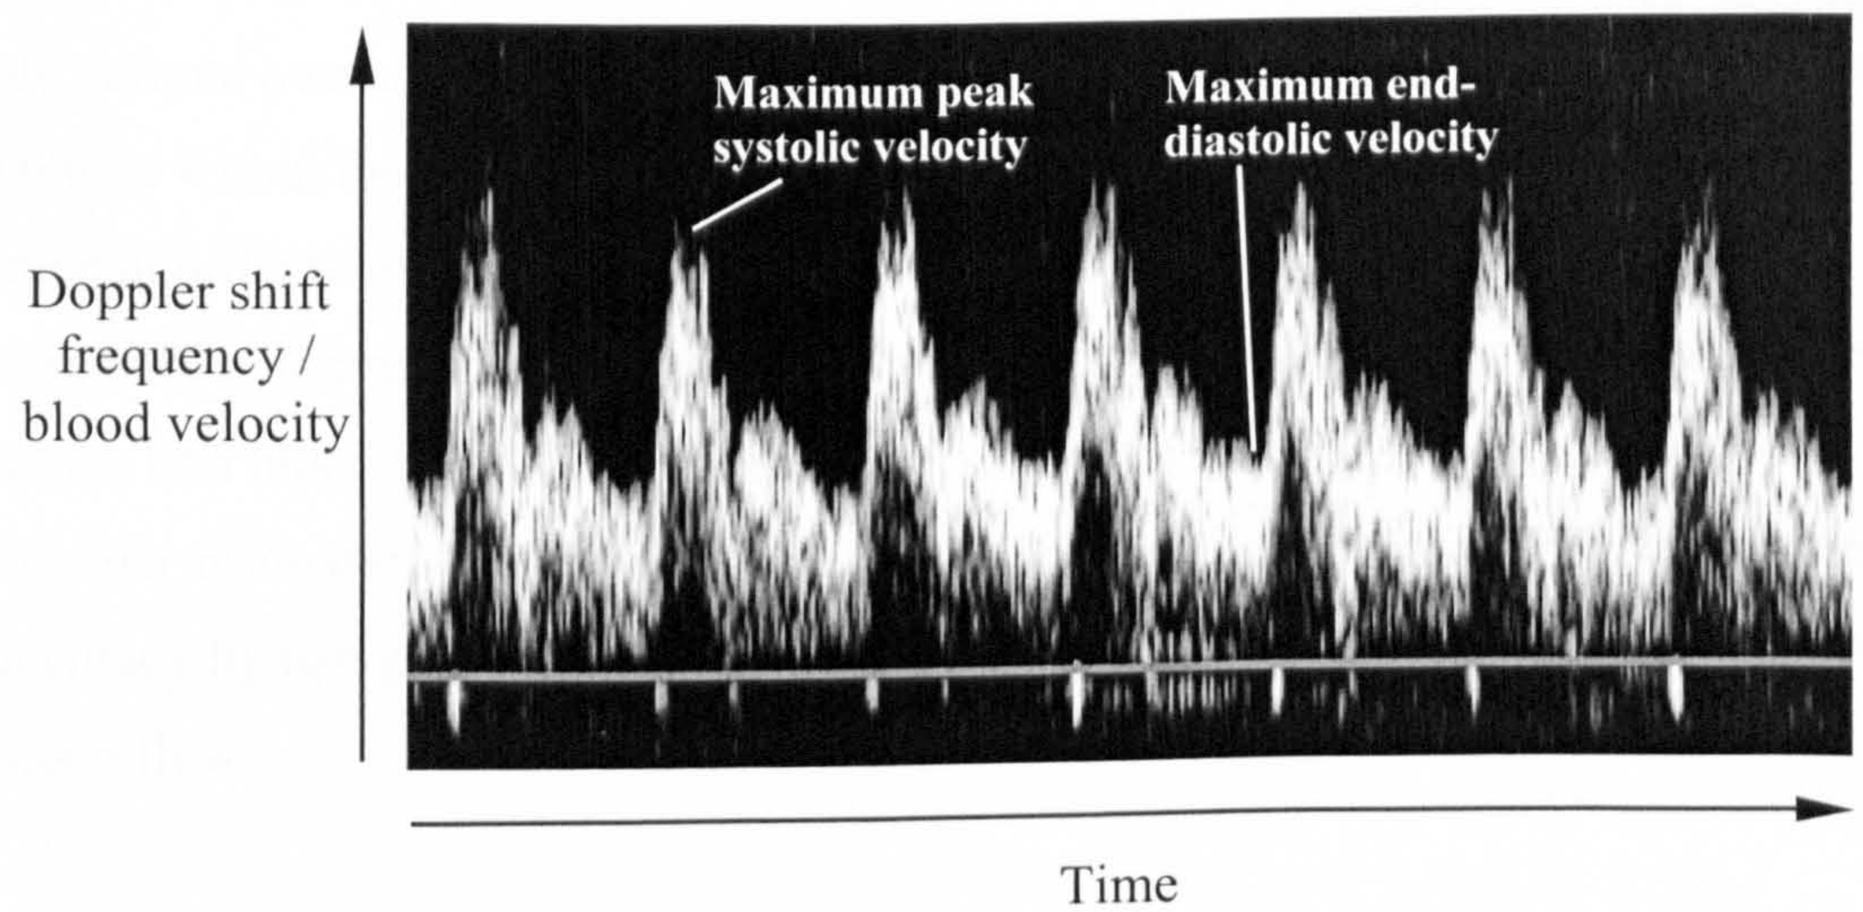


Figure 1.10 A spectral Doppler trace for the common carotid artery. The vertical axis corresponds to Doppler shift frequency (proportional to speed of blood), and the horizontal axis is time.

It can be seen from Figure 1.10 that the maximum blood velocity varies considerably through the cardiac cycle, with the maximum velocity occurring at peak systole. The velocity then falls to the end-diastolic velocity. Inspection of the Doppler trace reveals that the spectral Doppler trace is broad in the vertical direction, indicating that measured Doppler shifts span a range of frequencies. This phenomenon is known as spectral broadening and occurs for the following reasons: firstly, and most intuitively, the blood flowing through the sample volume consists of many scattering

particles, each with a different velocity. Secondly, the finite transit time of the ultrasound to and from the sample volume means that the blood particles have moved a small distance before the next sample. Thirdly, the true Doppler angle between the transducer and each scattering particle is slightly different, partly due to the finite size of the sample volume, and partly due to the finite width of the transducer. The first two mechanisms give rise to what is sometimes called extrinsic spectral broadening, whereas intrinsic spectral broadening is attributable to the third reason [Hoskins 1999a; Hoskins 1999b]. Spectral broadening is a well-recognised and important source of error in velocity measurements [Hoskins 1999a], and precautions to minimise its effect when measuring the elevated blood velocities associated with atherosclerotic disease are considered in Chapter 2.

Most modern ultrasound scanners support Duplex imaging – often simply referred to as 'Duplex' – where a 2D B-mode image, or, more recently, a colour or power Doppler image (see Section 1.3.6), is displayed with the position of the Doppler sample volume overlaid as shown in Figure 1.11. The operator then positions the sample volume within the desired vessel in the image and switches to a spectral Doppler trace in order to estimate the blood velocity. The display of the sample volume is very useful because it allows the sample volume to be placed accurately within the lumen of a blood vessel, and resized if necessary. In addition, the presence of the Doppler insonation line enables the Doppler angle to be estimated with reasonable accuracy by using the direction of the vessel walls as an indication of the direction of blood flow.

Accurate estimation of the Doppler angle is important as inspection of Equation (1.7) reveals that even small errors in the measured Doppler angle may lead to significant errors in the estimated velocity, particularly if the true Doppler angle is large. This point is illustrated by Figure 1.12 in which the percentage error in the estimated velocity, calculated using Equation (1.7), is plotted as a function of the true Doppler angle for errors in the measured Doppler angle from –5 to 5°. It can be seen that, when the true Doppler angle is 60°, a modest over-estimation of 3° in the measured Doppler angle (equal to an error of 5%) results in a velocity error in excess of 10%.

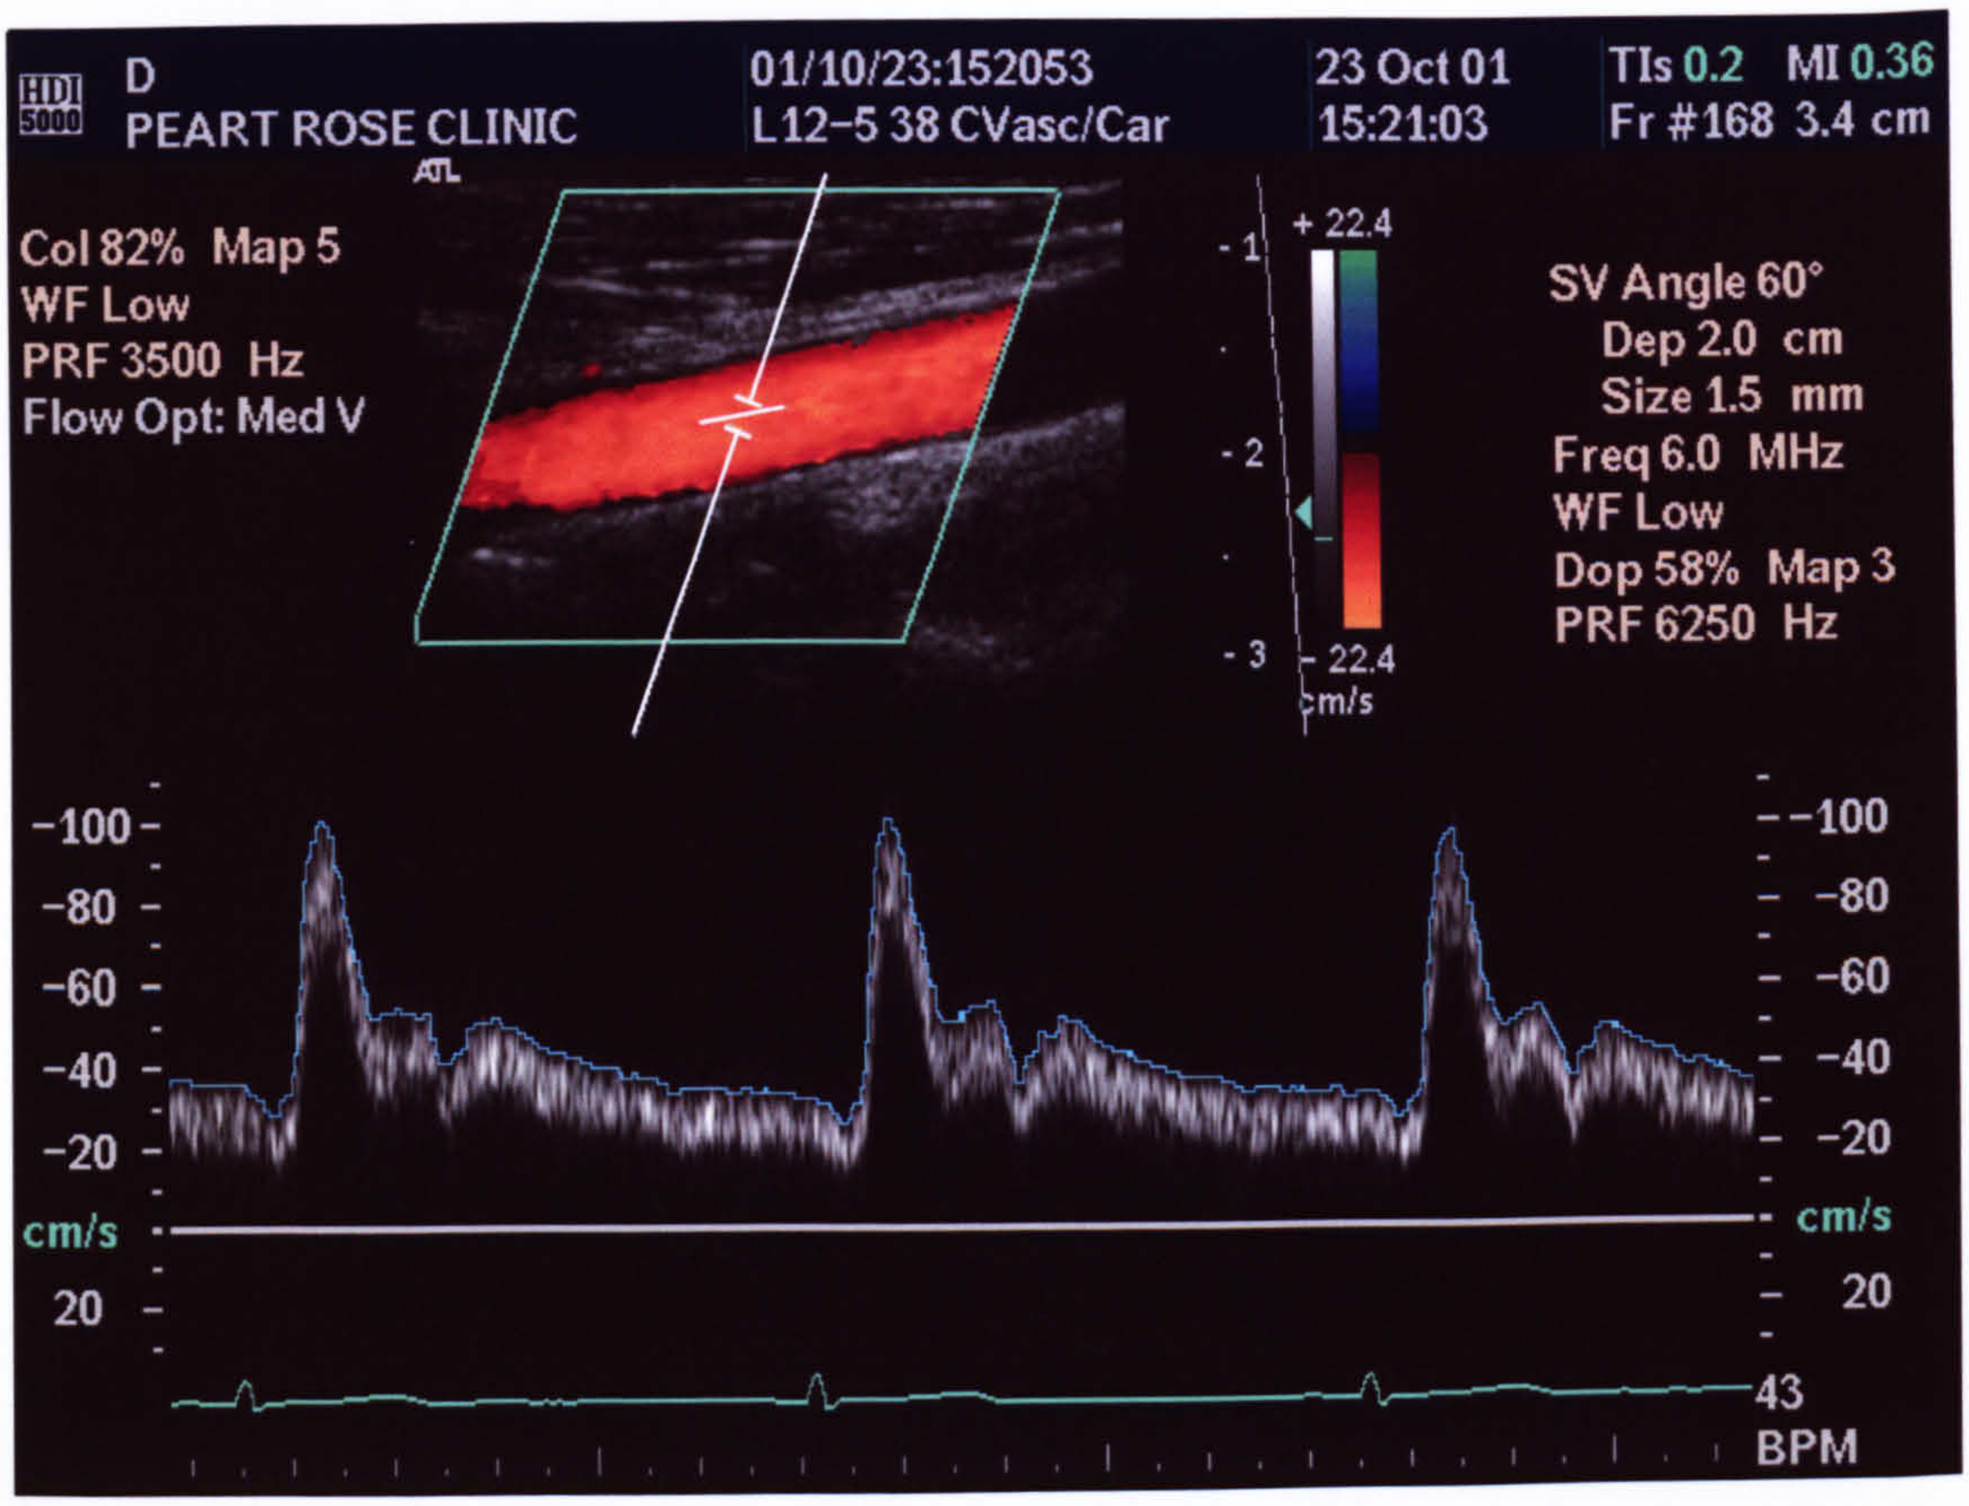


Figure 1.11 A Duplex ultrasound image of the CCA. The time-dependent blood velocity in the centre of the vessel is estimated by placing the sample volume within the lumen in the colour Doppler image.

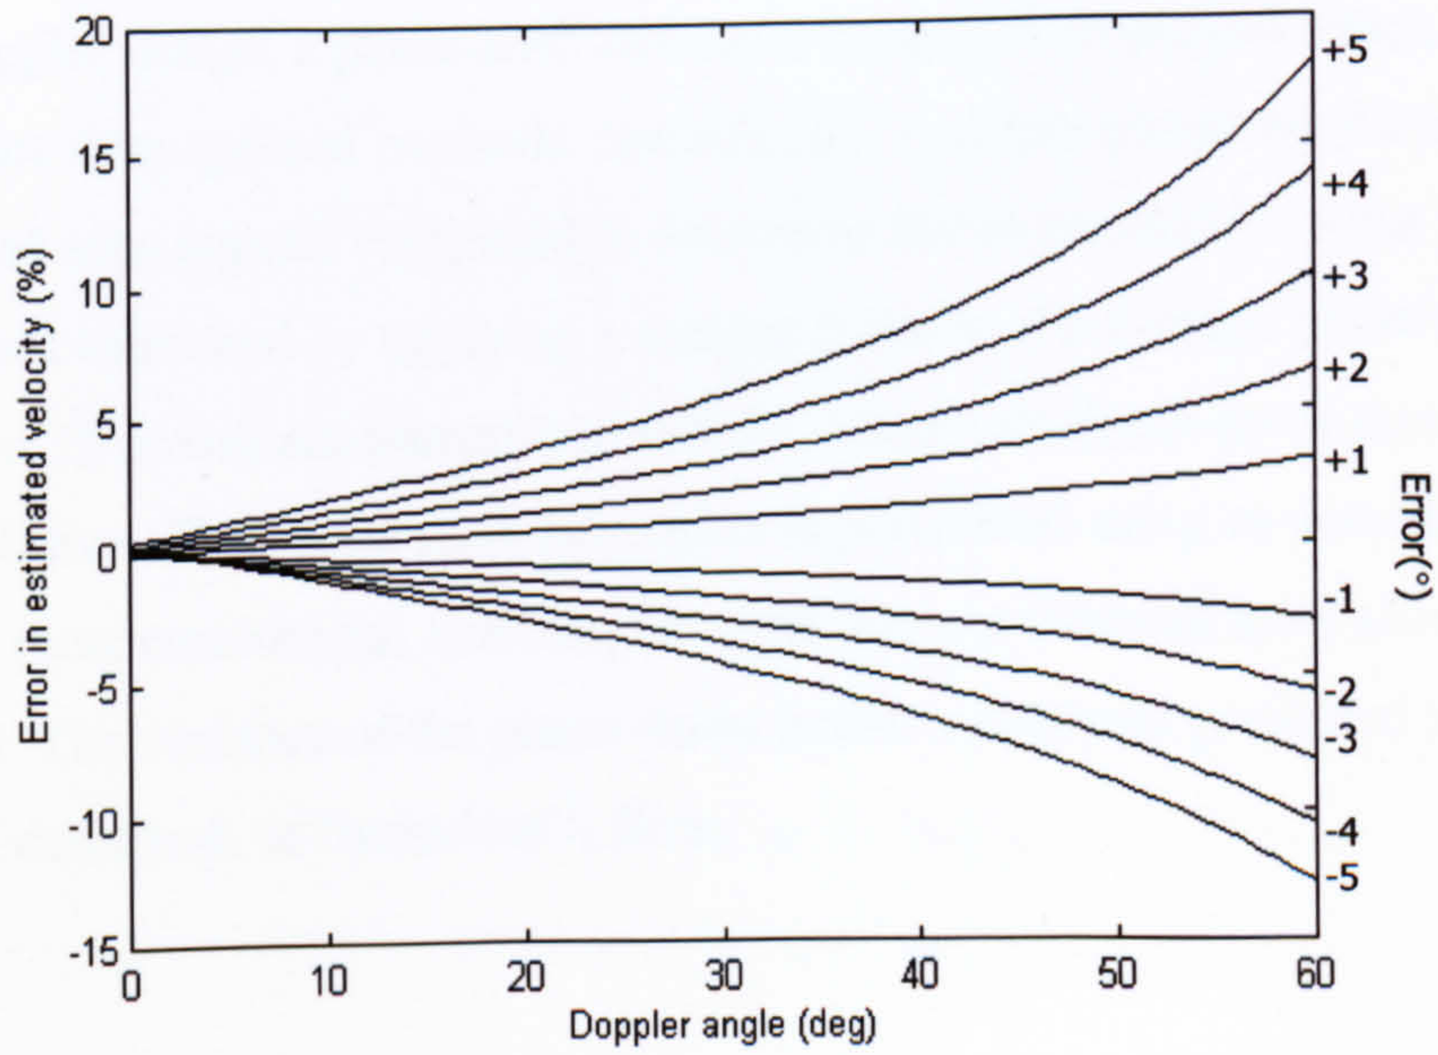


Figure 1.12 Percentage error in velocity estimates due to errors in Doppler angle measurement in the range –5 to 5°.

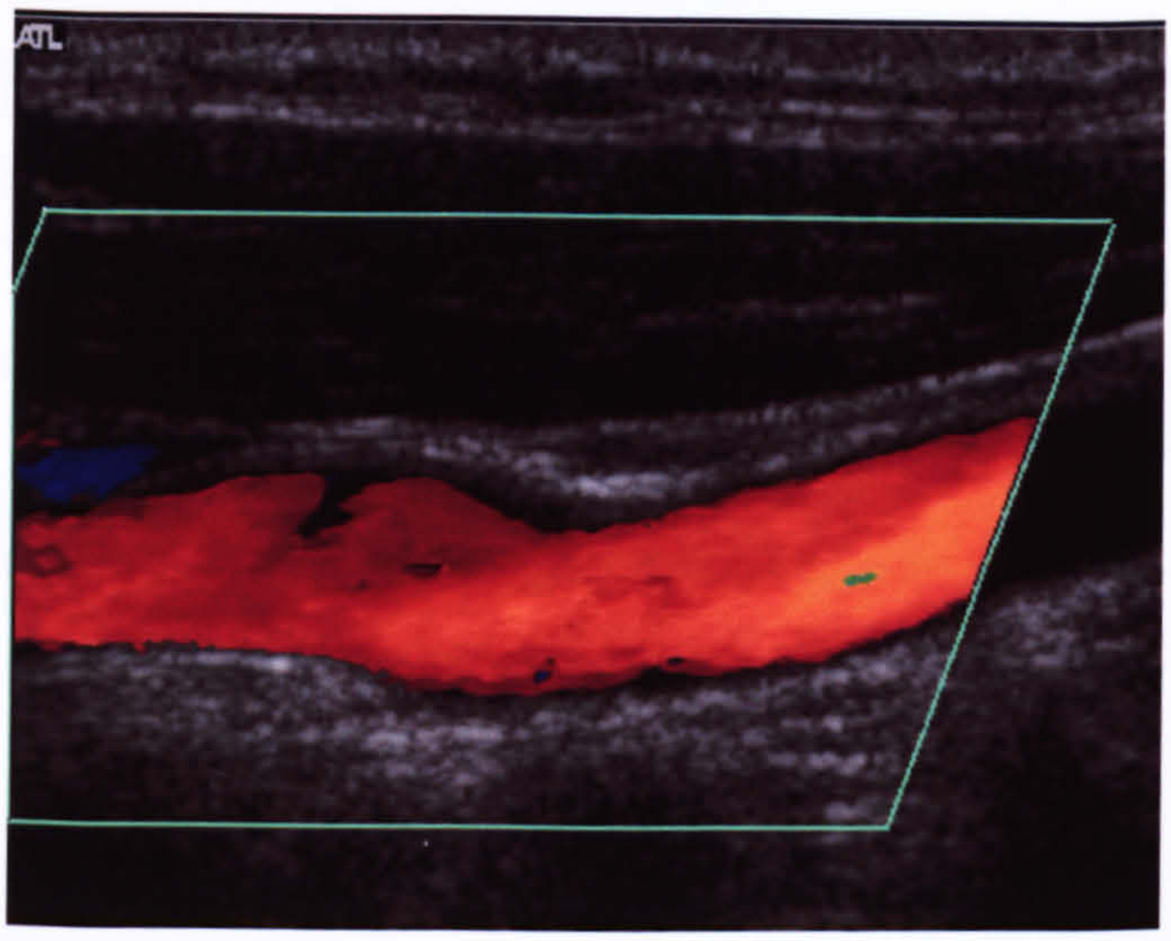
Errors in Doppler angle estimation are unavoidable in practice, since the vessel walls only act as a guide for the direction of blood flow, and blood often exhibits complicated flow patterns, particularly distal to atherosclerotic disease. Further uncertainty in the measured Doppler angle occurs because standard commercial scanners only estimate a one component of the true 3D velocity vector. In isolation, this factor will cause the true velocity to be underestimated. So-called vector Doppler systems which measure two or three components of the velocity vector have recently been developed [Dunmire 2000]. However, no such systems are available commercially as yet.

1.3.6 Colour and power modes

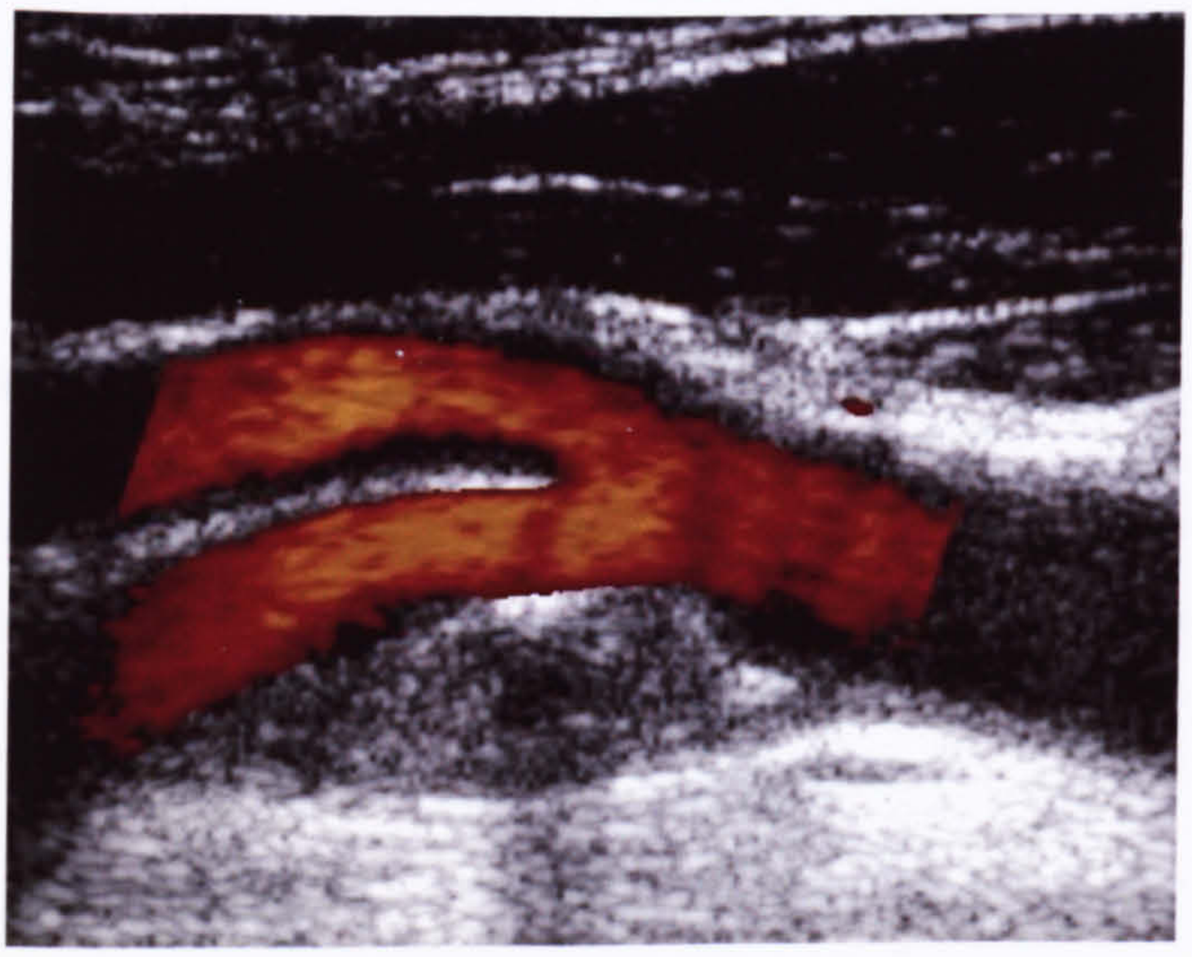
Colour Doppler imaging (sometimes called colour flow imaging) is a method in which blood velocities are estimated from many sample volumes and displayed as a colour-coded image [Ferrara 1997; Hoskins 1997]. Figure 1.13(a) shows an example of a colour Doppler image of the carotid bifurcation, where colour pixels overlay a B-mode image. The colour map is graded according to the velocity magnitude and direction with, say, red indicating flow away from the transducer, and blue indicating flow towards the transducer.

Due to the large number of sample volumes required to create a reasonably sized colour Doppler image, a phase-shift detection method is employed which is faster but less accurate than spectral methods. Specifically, multiple pulses are transmitted and the received echo signals compared to determine the phase shift between them. Velocity is then estimated by applying a scaling factor to the average phase-shift. Typically, about 10 pulses are transmitted and the echoes are depth-gated into, say, 256 sample volumes. The phase shift estimation is performed using an autocorrelation method in most commercial systems, but time-domain methods have also been developed. The variance of the phase shifts is also sometimes presented to indicate regions of disturbed, or 'turbulent'[1], flow.

[1] The phrase 'turbulent flow' is used commonly in the context of colour Doppler imaging to describe blood which is exhibiting very complex flow patterns, with elements flowing in multiple directions. It is, however, usually used inappropriately *vis-à-vis* the strict fluid mechanical sense. For this reason, and to avoid confusion, the term 'disturbed flow' has been adopted here as a substitute.



(a)



(b)

Figure 1.13 Examples of colour and power mode images: (a) a colour Doppler image, and (b) a power Doppler image of the carotid arteries.

Not surprisingly, due to the longer acquisition and computation time of Doppler data compared with B-mode, colour is usually confined to a region of interest (ROI) smaller than the size of the B-mode image. The maximum frame rate achievable using colour Doppler is also somewhat less than B-mode imaging alone. It is possible to use the same transmit pulses to generate both colour and B-mode images, but this is disadvantageous since, in general, it is desirable to have longer pulses for colour Doppler than for B-mode in order to improve the phase-shift estimate. In contrast, shorter pulses are desirable for B-mode to maximise axial resolution.

Another way of representing Doppler data is to colour code the power of the backscattered signal, rather than the frequency shift (see Figure 1.13(b)). This technique is called power Doppler imaging [Hoskins 1997]. Unlike colour Doppler, this mode is essentially independent of the angle of insonation, and has the advantage of being much more sensitive to slow flow than colour Doppler. Consequently, it is excellent for imaging flow in small or tortuous vessels. Recently, some commercial scanners have retained directional information in order to allow some 'directionalisation' of the power Doppler image. Again, power Doppler images are normally combined with a B-mode image to maximise information content.

1.3.7 Other ultrasound modes

Several new ultrasound imaging modes have emerged over the last ten years, with many of them directly applicable to vascular imaging. The most popular of these are introduced briefly below.

1.3.7.1 Contrast imaging

Ultrasound contrast imaging uses a contrast agent comprising microbubbles in a suspension [Hennerici 1999; Correas 2001; Frinking 2000]. The contrast is injected intravenously and passes through the circulation before being removed by the pulmonary vascular bed. Microbubbles are generally gas-filled with a synthetic shell, and resonate when insonated. This property greatly increases the scattering cross-section of individual bubbles; in other words, the strength of the backscattered Doppler signal increases. Therefore, microbubble contrast agents are particularly useful for imaging vessels that normally only provide weak echoes due to their depth, or highly attenuating surrounding tissue.

1.3.7.2 Harmonic imaging

A property of some microbubbles is resonance at the second harmonic of the transmit frequency of the interrogating ultrasound. This makes the microbubbles uniquely identifiable, and, hence, it is possible to produce Doppler images containing considerably less noise than conventional modes. The principle of this method, called harmonic imaging, may also be applied to tissue by transmitting pulses with high amplitude and taking advantage of the non-linear scattering effects of tissue [Burns 1997].

1.3.7.3 Tissue Doppler imaging

Tissue Doppler imaging (TDI) is a way of imaging tissue motion using the Doppler effect [Hoskins 1997]. TDI targets the opposite information to blood Doppler systems which filter out high amplitude, low frequency-shifted echoes from vessel walls. The method has become popular in cardiac applications, where certain abnormalities of the heart are detectable by investigating ventricular wall motion [Price 2000]. It has also been applied to tracking the motion of artery walls and atherosclerotic plaque [Schmidt-Trucksäss 1998; Bonnefous 2001].

1.3.7.4 Compound B-mode

Two limitations of conventional B-mode imaging are the dependence of the image on the direction of insonation, and the presence of speckle. Both these can be reduced by combining multiple B-mode images obtained from multiple directions to form a single, composite image. This technique is known as compound B-mode imaging and can be neatly implemented by electronically steering the ultrasound beam through a range of angles, and averaging pixels in the overlapping region. This approach allows very high co-registration accuracy and, although the frame rate is inevitably compromised compared with standard B-mode, high frame rates are feasible nonetheless. Compound B-mode has recently become available on at least one commercial ultrasound scanner (ATL-Phillips' SonoCT™ on the HDI-5000). In certain cases, compound images result in a marked reduction in direction-dependent artefacts, such as shadowing (see Section 1.3.7.1), as well as enhanced visualisation of complex structures. Figure 1.14 demonstrates the advantages of this mode for visualising carotid plaque.

1.3.7.5 B-flow

B-flow imaging uses advanced B-mode technology to induce and detect back-scattered ultrasound from blood as opposed to soft tissue [Deane 2000; Weskott 2001]. This is achieved by using digitally encoded pulse-echo sequences [O'Donnell 1992]. B-flow imaging allows significantly higher spatial and temporal resolution to be achieved, compared with colour or power Doppler. It is also not subject to artefacts characteristic of these modes (see Section 1.3.8.9). Since the technique is purely qualitative at present, and cannot be used to estimate blood velocity, applications have been confined to plaque surface visualisation, for example. This is particularly useful when imaging echolucent lesions, where the boundary is sometimes impossible to estimate without the aid of colour or power Doppler imaging.

1.3.8 Artefacts

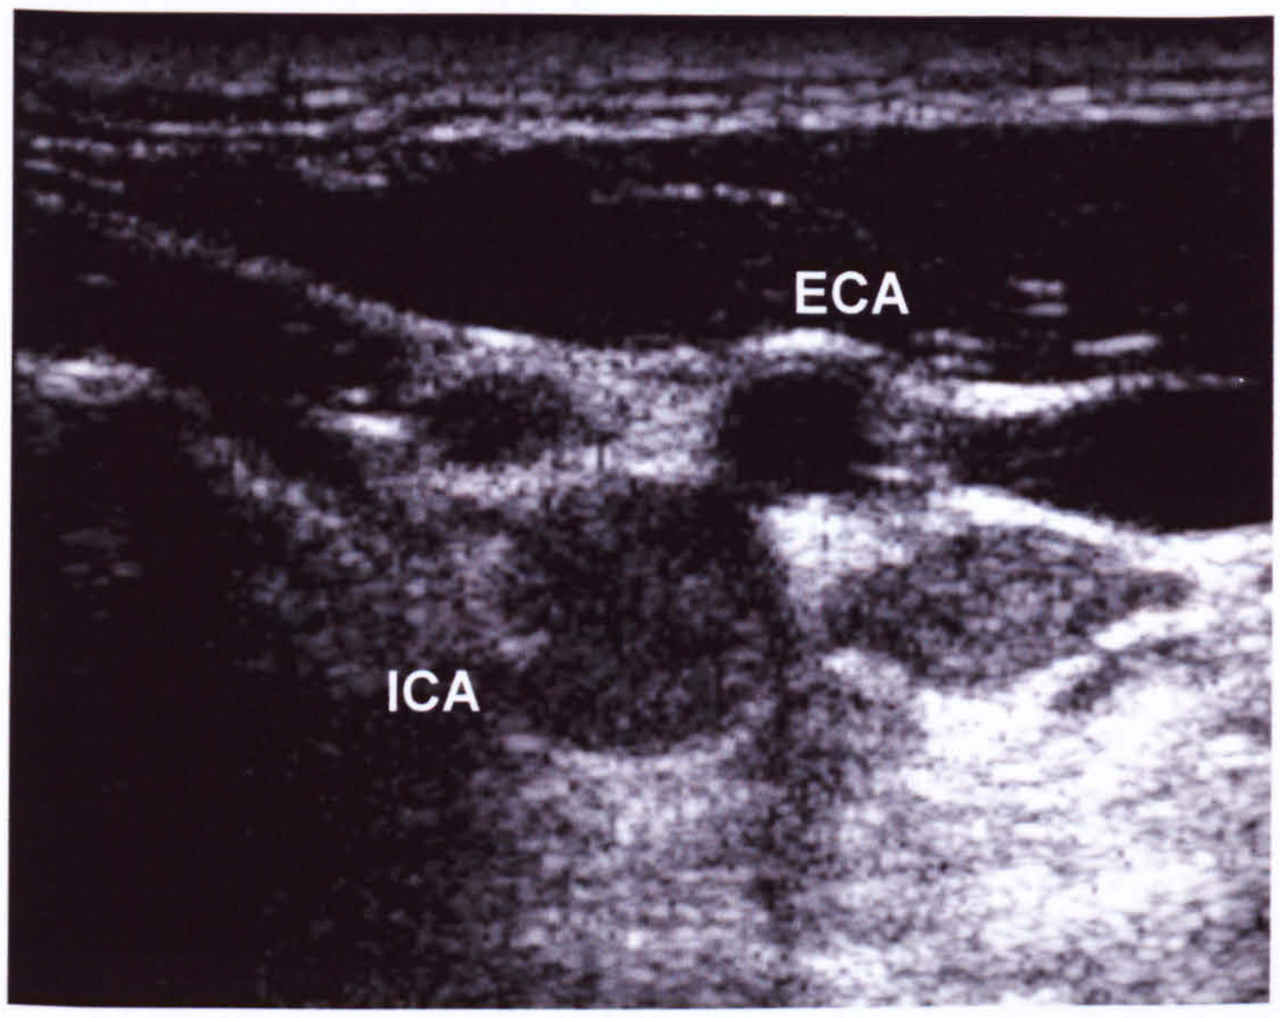
Ultrasound images are potentially subject to a wide variety of artefacts. A detailed explanation of common artefacts encountered can be found in one of the many textbooks on medical ultrasound. The principal artefacts are summarised below.

1.3.8.1 Shadowing

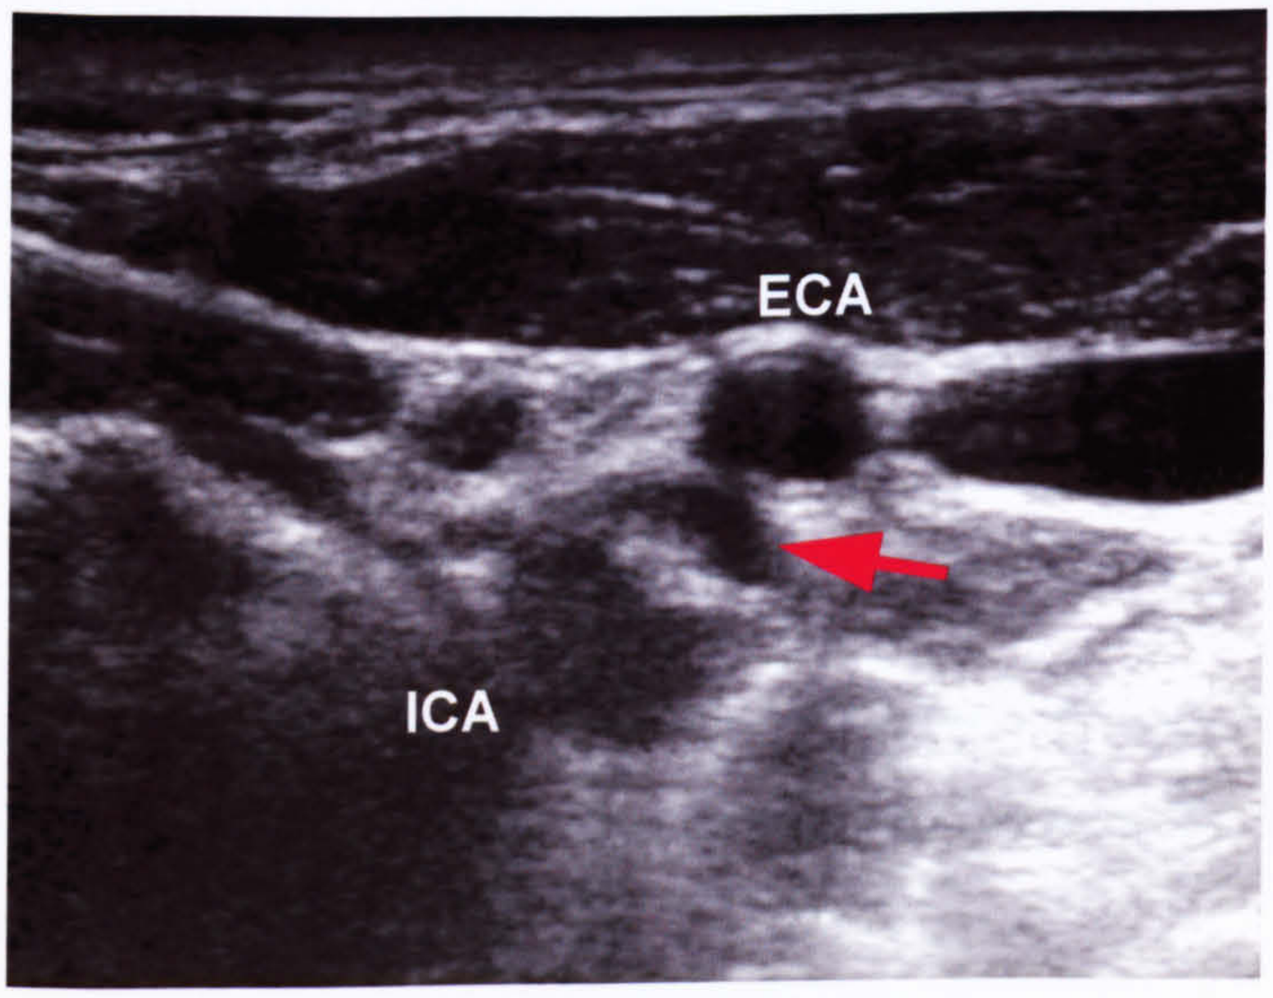
Shadowing is perhaps the most common ultrasound artefact, and is characterised by a dark region lying immediately below a structure. Shadowing can occur on B-mode as well as colour and power Doppler images, and is caused by highly reflective and/or attenuating structures which permit only a very small proportion of the incident ultrasound energy to be transmitted to deeper levels. Hence, echoes from these levels are weaker than those from similar depths elsewhere in the image.

1.3.8.2 Enhancement

Acoustic enhancement is the opposite effect to shadowing whereby ultrasound energy propagates through a region that is either poorly reflective or weakly attenuating, or both. This results in higher intensity waves propagating into deeper levels than those in the surrounding tissue at the same depth. The signal from this region is, therefore, strong. Enhancement is manifest, for example, as a bright area immediately below a structure in a B-mode image.



(a)



(b)

Figure 1.14 Comparison of normal and compound B-mode images. (a) shows a carotid plaque imaged using conventional B-mode. (b) shows the equivalent image obtained using compounding. Note how the ICA lumen, indicated by the arrow, is more clearly visualised in the compounded image. (The images were acquired by the author using the SonoCT™ feature on the ATL HDI-5000 scanner).

1.3.8.3 Reverberation

Reverberation artefacts are caused by multiple reflections from the same structure. There are several ways in which this can occur: for example, the transducer surface may act as a strong reflector of returning echoes, producing a 'ghost image' below a strongly reflective structure. Such artefacts are usually minimised by providing adequate acoustic coupling between the transducer and the tissue. Reverberation can also occur when multiple strong reflectors lie below an object, giving rise to the characteristic 'comet-tail' artefact. Another reverberation artefact, which takes place due to the resonance of gas bubbles in abdominal imaging, is called 'ring-down', and appears as a fuzzy tail.

1.3.8.4 Side- and grating-lobe artefacts

Side-lobe artefacts are produced by echoes arising from peripheral ultrasound energy travelling in a direction different to that of the main beam. This radiation forms so-called side-lobes around the main beam and is common to all transducers. A particular problem arises with array transducers which produce grating lobes. The size of the grating lobes is decreased by either increasing the number of transducer elements in the array, or by weighting the signals from the transmit and receive elements (a method known as apodization [Wells 1993]).

1.3.8.5 Mirror artefact

The mirror artefact occurs when a strong reflector lying below an object acts as an acoustic mirror. This leads to an additional artefactual image of the object appearing below the true image.

1.3.8.6 In-plane misregistration

The principal causes of misregistration within an ultrasound image are direct consequences of the breakdown of assumptions intrinsic to the formation of the image. Two fundamental assumptions made in ultrasound imaging are that a pulse and its echo travel along straight-line paths through tissue, and that the speed of propagation is constant. The first assumption is not always the case because the ultrasound beam can be refracted (see Section 1.3.1). The implication of this is that the path-lengths of pulses are sometimes longer than the assumption presupposes, and that echoes re-

ceived by the transducer may actually arise from structures lying away from the straight-line path. The second assumption of constant speed of sound is also subject to error: for example, the speed of sound in fat is approximately 1440 m/s, whereas the speed in muscle is 1570 m/s. Thus, tissue inhomogeneity may lead to misregistration. The speed of sound is also dependent of the transmit frequency of the transducer, but this effect is negligible.

1.3.8.7 Artefacts due to finite beam thickness

For a standard linear array transducer, the thickness of the ultrasound beam in the elevational direction (perpendicular to the image plane) is much larger than the axial or lateral resolutions. Typical ultrasound beams can have a thickness ranging from less than 1 mm in the focal zone, to several millimetres in the far-field, away from the transducer [Skolnick 1991]. It is, therefore, possible for structures to appear in an image even if they lie outside the ideal imaging plane. The implications of finite beam thickness on 3DUS imaging are considered in Chapter 4.

1.3.8.8 Noise

The quality of ultrasound images is degraded significantly by the presence of noise. The principal source of noise is speckle produced by constructive-destructive interference of back-scattered ultrasound. Another significant source of image noise is electronic signal processing instrumentation.

1.3.8.9 Doppler artefacts

Doppler-based images are subject to many of the artefacts listed above, in particular, shadowing, enhancement, misregistration, and noise. In addition, a few mode-specific artefacts exist. One of these, spectral broadening, has already been discussed in Section 1.3.5. A very common artefact in colour and power Doppler images occurs due to the dependency of the image on the angle between the ultrasound beam and the vessel, resulting in different images for different beam-vessel angles. Another colour/power Doppler artefact is manifest as over- or under colour-filling of the vessel lumen. This can often be corrected by adjusting the signal gain. A third artefact, termed flash artefact, can occur when the scan-probe is moved quickly across the skin. This artefact is more of a problem for power Doppler than colour Doppler because of its greater sensitivity. Several manufacturers now incorporate anti-flash

detection software into scanners to reduce its effect. A further artefact is aliasing whereby high velocities appear as low velocities. This is particularly common when high velocities are present that are above the detectable range of the scanner. Aliasing is far more noticeable in spectral Doppler traces and colour Doppler images than in power Doppler images, and can usually be remedied by increasing the PRF as this increases the maximum velocity detectable by the system.

1.4 Magnetic resonance imaging

MRI encompasses a collection of techniques which take advantage of the paramagnetic properties of atomic nuclei – specifically, protons and neutrons. Hydrogen is the most frequently used target for MRI due to its abundance in biological tissue and its relatively strong magnetic moment. Magnetic resonance angiography (MRA) is a potential alternative to DSA due to the fact that ionising radiation is not used [Yucel 1999]. The long-term biological effects of MRI, if any, are, however, unknown.

MRA produces images with high contrast between blood and background signals, with the blood appearing as either dark (signal absence) or bright (signal present). Traditional spin-echo techniques produce 'black blood' images due to blood flowing out of a slice-plane between successive radio-frequency (RF) pulses [Edelman 1990]. Bright blood images, on the other hand, are generated using time-of-flight (TOF) [Keller 1992] or phase contrast [Dumoulin 1995] techniques. Both can be used non-invasively, without a contrast enhancing agent, although the best images are obtained using an agent such as gadolinium [Prince 1996; Serfaty 2000]. Phase contrast methods also enable blood velocities to be quantified [Walker 1988]. The main artefact occurring in TOF and phase contrast images is signal loss in regions of slow or complex flow, such as the carotid bulb [Fürst 1995].

Soft tissue MRI has been used to image atherosclerotic plaque [Yuan 1998], but this requires the use of a specialised coil. The technique does, however, have potential for non-invasive determination of plaque composition. MRI is currently undergoing rapid development and has shown large potential in many aspects of vascular imaging. However, a major drawback so far is cost, which limits the availability of the technique. The slow image acquisition time is improving, but this is still unacceptably long in many cases. MRI scans have been unpopular with some patients because

of the requirement to remain still in a rather claustrophobic tunnel for prolonged periods. This issue is being addressed by manufacturers, and new scanners are designed to be more patient friendly with reduced scan times.

1.5 3D imaging

The vascular imaging techniques described in the previous sections are predominantly used to acquire 2D images. However, the same techniques are extendable to 3D. Among these, 3D MRA and x-ray CT are already available to many clinicians. 3D angiography is perfectly suited to the visualisation of complex arterial anatomy, for example, the cerebral circulation, and many techniques exist to process and display 3D data. The development of 3DUS has lagged behind the other methods. Reasons for this include the relatively poor quality of ultrasound images, which makes 3D visualization a difficult process, and the interactive nature of traditional (2D) ultrasound imaging. Presently, 3DUS remains largely a research tool. Several dedicated commercial 3D system have become available, however, and the technique is slowly becoming accepted for routine clinical use, particularly in the field of obstetrics. A comprehensive review of 3DUS is given in Chapter 3.

Chapter 2

Quantification of carotid artery disease

This chapter outlines current methods for the quantification of carotid artery disease. The literature on this subject is reviewed, and the merits and disadvantages of each method are considered. Following on from these considerations, the motivations for this thesis are explained. The chapter concludes with a list of aims of the thesis.

2.1 Degree of stenosis

Conventionally, the severity of atherosclerotic disease has been assessed by quantifying the degree of stenosis in an artery. Stenosis is a Greek word meaning a narrowing of a passage, and a quantitative measure of degree of narrowing of the vessel lumen may be defined in a number of different ways. It is normally expressed as a percentage index with respect to some reference indicative of normality; for example, the diameter in a disease-free portion of artery. For clarity, a generalised method for calculating arterial stenosis is presented below. In clinical practice, stenosis indices are usually determined from one-dimensional (1D) – i.e. linear distance – measurements. The method may, however, be extended to 2D (cross-sectional area) or 3D (volumetric) measurements, as explained shortly.

Indices of carotid stenosis have conventionally been based on lumenal diameter measurements. If the lumen diameter at the point of maximal narrowing is d_{lumen} (see Figure 2.1a) then the percentage degree of stenosis, S, is defined as:

$$S = 100 \times \left(1 - \frac{d_{\text{lumen}}}{d_{\text{REF}}}\right), \tag{2.1}$$

where d_{REF} is some reference diameter. Common choices for d_{REF} are: the diameter of a distal, normal section of the ICA (adopted by the North American Symptomatic Carotid Endarterectomy Trial (NASCET)); the estimated diameter of the carotid bulb (adopted by the European Carotid Surgery Trial (ECST)); the diameter of the CCA [Bladin 1995]; or the inner vessel wall diameter of the vessel where d_{lumen} is measured [Beebe 1999].

Inspection of Equation (2.1) reveals that S is 100% when d_{lumen} is zero. Therefore, total occlusion is indicated by a stenosis of 100%. Choosing d_{REF} to be, say, the diameter of the vessel wall where d_{lumen} was measured ensures that S is 0% when d is equal to d_{REF}. Intuitively, 0% stenosis implies a disease-free artery. However, adopting one of the other definitions for d_{REF} mentioned above does not ensure that this is true. Indeed, in some cases, negative values of S may arise for mild disease.

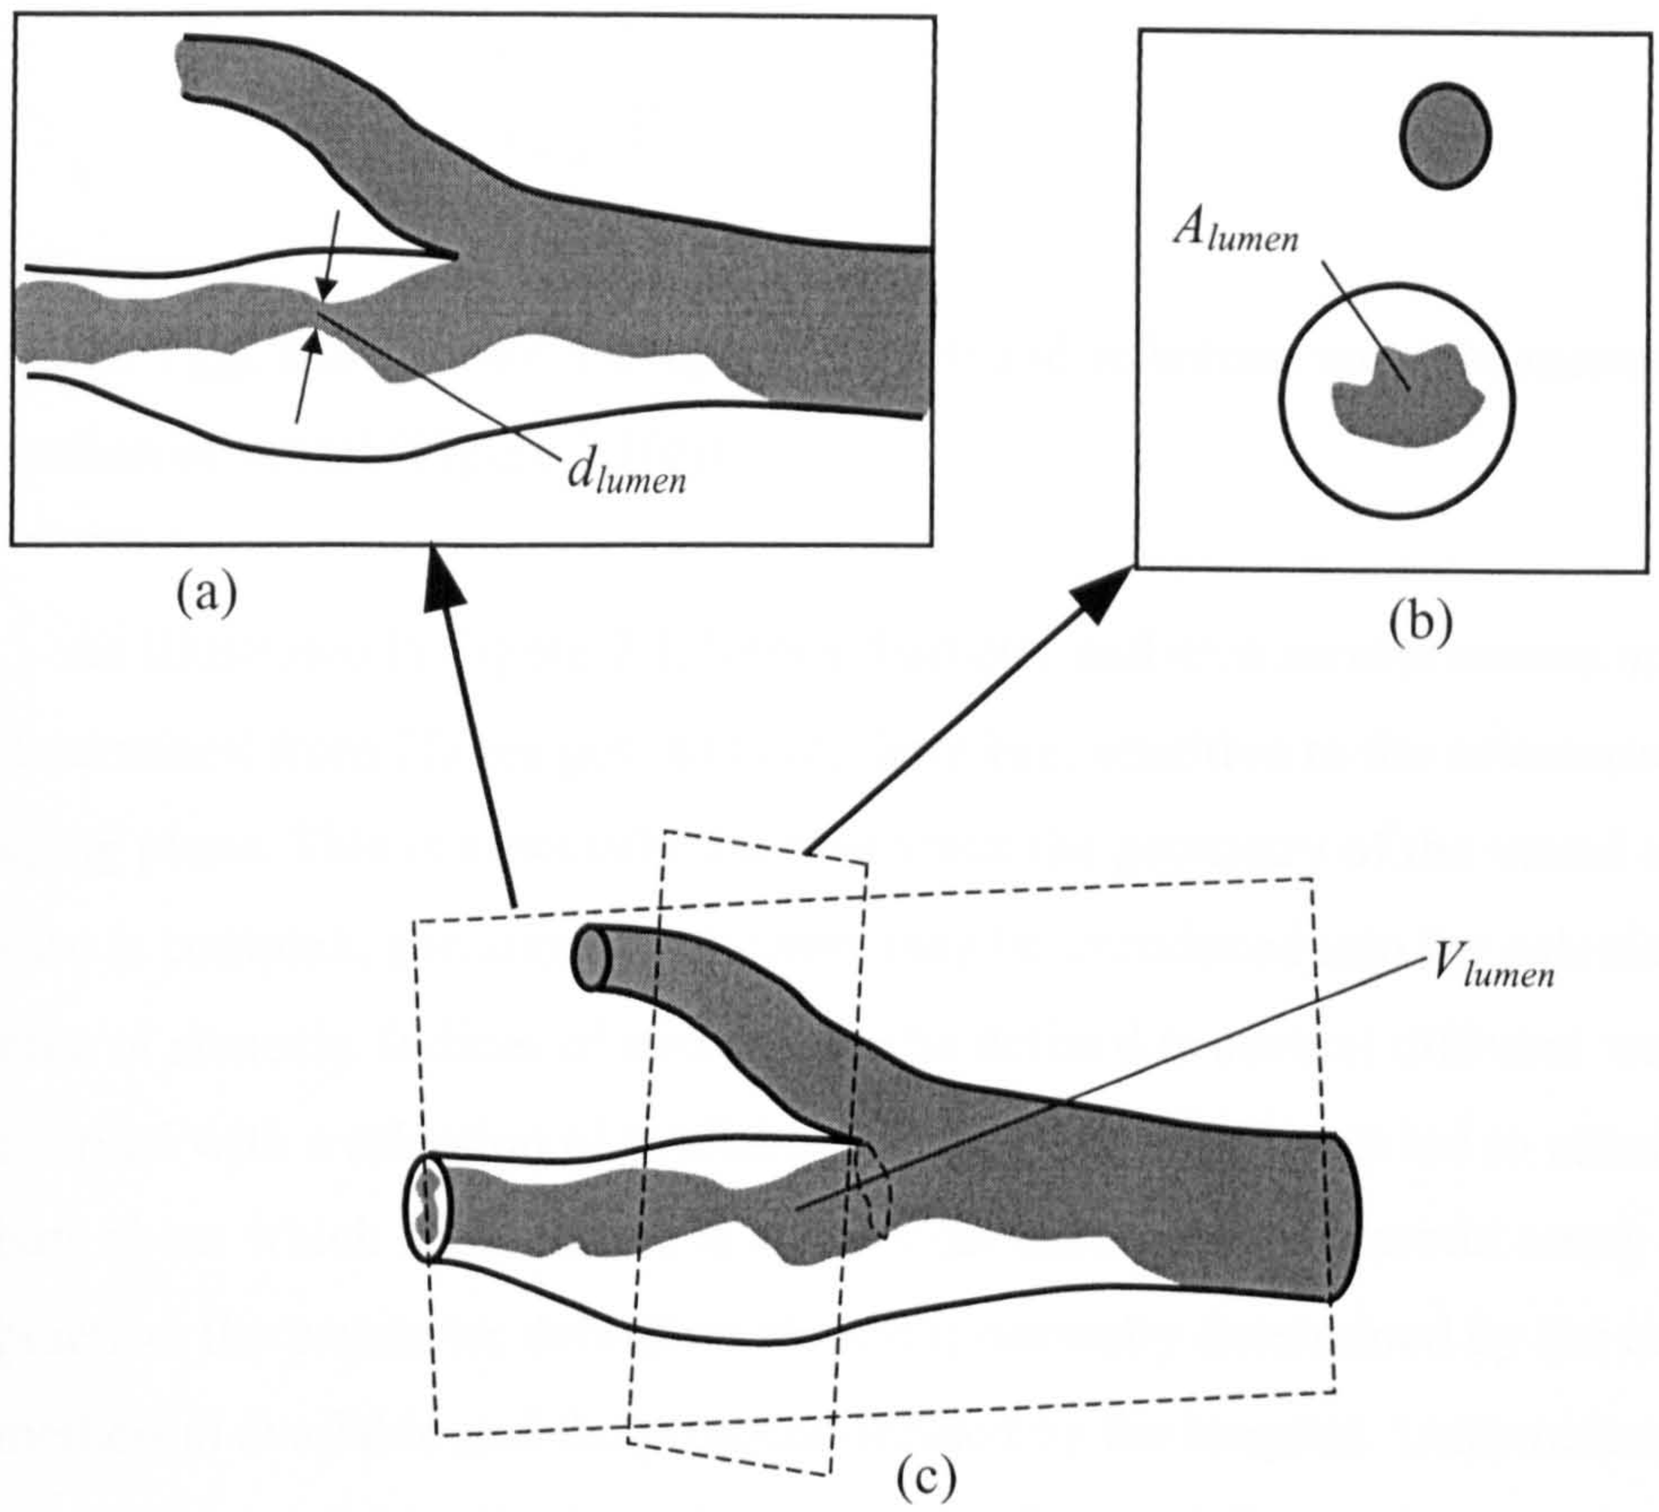


Figure 2.1 Measurement of degree of stenosis in the internal carotid artery. (a) The lumen diameter, d_{lumen}, is measured at the site of maximal narrowing from a longitudinal 2D image. (b) The lumen cross-sectional area, A_{lumen}, is measured from a transverse image. (c) Lumen volume, V_{lumen}, is measured from a 3D image.

An analogous definition of degree of stenosis using cross-sectional area measurements is:

$$S = 100 \times \left(1 - \frac{A_{\text{lumen}}}{A_{\text{REF}}} \right), \tag{2.2}$$

where A_{lumen} is the cross-sectional area at a point along the vessel, and A_{REF} is a reference area, such as the cross-sectional area to the inner boundary of the vessel wall. Area-based stenosis may be readily calculated from measurements made from a transverse image, as shown in Figure 2.1(b). Extending the definition of stenosis to 3D, we arrive at the volumetric stenosis index, defined by:

$$S = 100 \times \left(1 - \frac{V_{\text{lumen}}}{V_{\text{REF}}} \right), \tag{2.3}$$

where V_{lumen} and V_{REF} are appropriate lumen and reference volumes measured within a section of vessel (Figure 2.1(c)).

As illustrated in Figure 2.1, lumen diameter and area measurements are typically determined from 2D images, and are, therefore, sensitive to the orientation of the imaging plane. This is especially the case when the geometry of the vessel and/or disease is complex, and significant errors may be introduced into the calculation of degree of stenosis. Indices of stenosis can be defined in several different ways. This, coupled with a selection of available imaging techniques, has led to considerable debate about which combination is best for the assessment of carotid artery disease. In practice, the particular definition chosen is normally determined by the imaging method(s) available and the protocol decided by the hospital. Unfortunately, there is at present no standardisation of assessment of carotid disease between centres [Byrd 1998], and it is left to individual centres to establish their own interpretation of the general guidelines for the quantification and subsequent management of carotid disease. This situation makes comparisons between research findings very difficult to establish, and raises important questions about whether the results of the quantitative assessment of carotid disease carried out in different centres, and their interpretation, with respect to established criteria for clinical management, are consistent [Zwiebel 1999; Toole 1994].

The standard stenosis indices are examined in detail in the following sections. To aid in the assessment of these indices, it is helpful to construct a hypothetical model of the ideal stenosis index (ISI) as a benchmark for comparisons. In one sense, this should be considered to be the true Gold Standard as it provides a level against which practical indices may be compared. The characteristics of the ISI proposed here are as follows:

- The ISI should be based on one or more measurable, well-defined, physical quantities, and take a value between 0 and 100%.
- The ISI should have high accuracy: i.e. its error compared with the imaginary 'true' value is small. (The accuracy may be calculated from the estimated error

bounds of the dependent physical measurements.)

- The ISI should have high precision, i.e. repeated measures should be close in value. Often, an operator is involved in the measurement process, in which case the precision is reflected by the intra- and inter-observer reproducibility.

- The ISI should be useful as a screening test: in this context it should be strongly predictive of risk of stroke (or TIA) related to the carotid territory. To this end, it should be both specific (low number of false-positives) and sensitive (low number of false-negatives).

- It is desirable for the ISI to have some sensible zero level; in particular, its value should be 0% for a disease-free artery.

2.2 Influential clinical trials

Surgical removal of the plaque, a procedure known as carotid endarterectomy (CEA) has become a relatively common vascular operation. There are, however, inevitable risks associated, including peri-operative stroke and complications due to (general) anaesthesia during surgery [Ferguson 1999]. Medical complications can also occur post-surgery, but these are less common [Paciaroni 1999].

To date, both the methods for quantifying degree of stenosis and consequential patient management have been heavily influenced by the NASCET and ECST trials [ECST Collaborators 1991; ECST Collaborative Group 1998; Ferguson 1999; NASCET Steering Committee 1991; NASCET Collaborators 1991]. The principal aim of both of these large clinical trials was to determine the optimal treatment (surgical versus prophylactic) in order to prevent stroke in patients with symptomatic carotid artery disease. In both studies, different stenosis indices were calculated using diameter measurements made from x-ray angiograms (see Section 2.3), but both were shown to be a strong predictor of risk of ipsilateral ischaemic stroke. Moreover, the results of the NASCET trial suggested that CEA was beneficial for a stenosis of greater than 70% [North American Symptomatic Carotid Endarterectomy Trial Collaborators 1991], whereas the ECST trial concluded that surgery is beneficial for patients with a stenosis of 80% or above [European Carotid Surgery Trialists' Collaborative Group 1998]. It is important to note that the stenosis indices adopted by the NASCET and ECST trials were different. It will be shown shortly that, at least super-

ficially, the thresholds for CEA determined by each trial are not particularly inconsistent. It is probably wise, however, to consider the findings of each trial separately, and to adopt the measurement methods and criteria for one or the other.

A third clinical trial, the Asymptomatic Carotid Atherosclerosis Study (ACAS), investigated the benefit of CEA for asymptomatic patients with carotid disease, concluding that CEA was indicated for a degree of stenosis above 60 % (using the NASCET method specified in Section 2.3 below) [Executive Committee for the Asymptomatic Carotid Atherosclerosis Study 1995]. The European equivalent of ACAS – the Asymptomatic Carotid Surgery Trial (ACST) [Halliday 1994] – is still ongoing and no interim results have yet been published.

In addition to the stenosis criteria quoted here, many other indications and contra-indications must be taken into account. Consensual guidelines for CEA, including details of these, are given in Moore et al. [1995].

2.3 Angiography-based measurement of degree of stenosis

The NASCET, ECST, and ACAS trials all used images of the carotid bifurcation acquired using x-ray angiography as a basis for stenosis measurements. For the NASCET and ACAS trials, the stenosis index was calculated using the formula given in Equation (2.1) with d_{REF} equal to the measured diameter of a 'normal' section of the ICA distal to the stenosis. In the protocol for the NASCET study two orthogonal angiograms were acquired for each patient, with measurements made on the one with the narrowest observed lumen diameter. For the ECST trial, d_{REF} is the estimated diameter of the carotid bulb.

In an attempt to directly compare the NASCET and ECST stenosis indices, some studies have sought to establish a mathematical relationship to allow conversion between them [Eliasziw 1994; Rothwell 1994]. Using 1001 angiograms from the ECST trial, Rothwell et al. [1994] established the linear relationship:

$$S_{ECST} \approx 0.6 S_{NASCET} + 40. \tag{2.4}$$

A smaller study (112 angiograms) by Eliasziw et al. [1994], on the other hand, found a parabolic relationship between S_{ECST} and S_{NASCET}:

$$S_{ECST} \approx 0.002 S_{NASCET}^2 + 0.290 S_{NASCET} + 55.160 . \quad (2.5)$$

The relationships defined by Equations (2.4) and (2.5) are shown graphically in Figure 2.2.

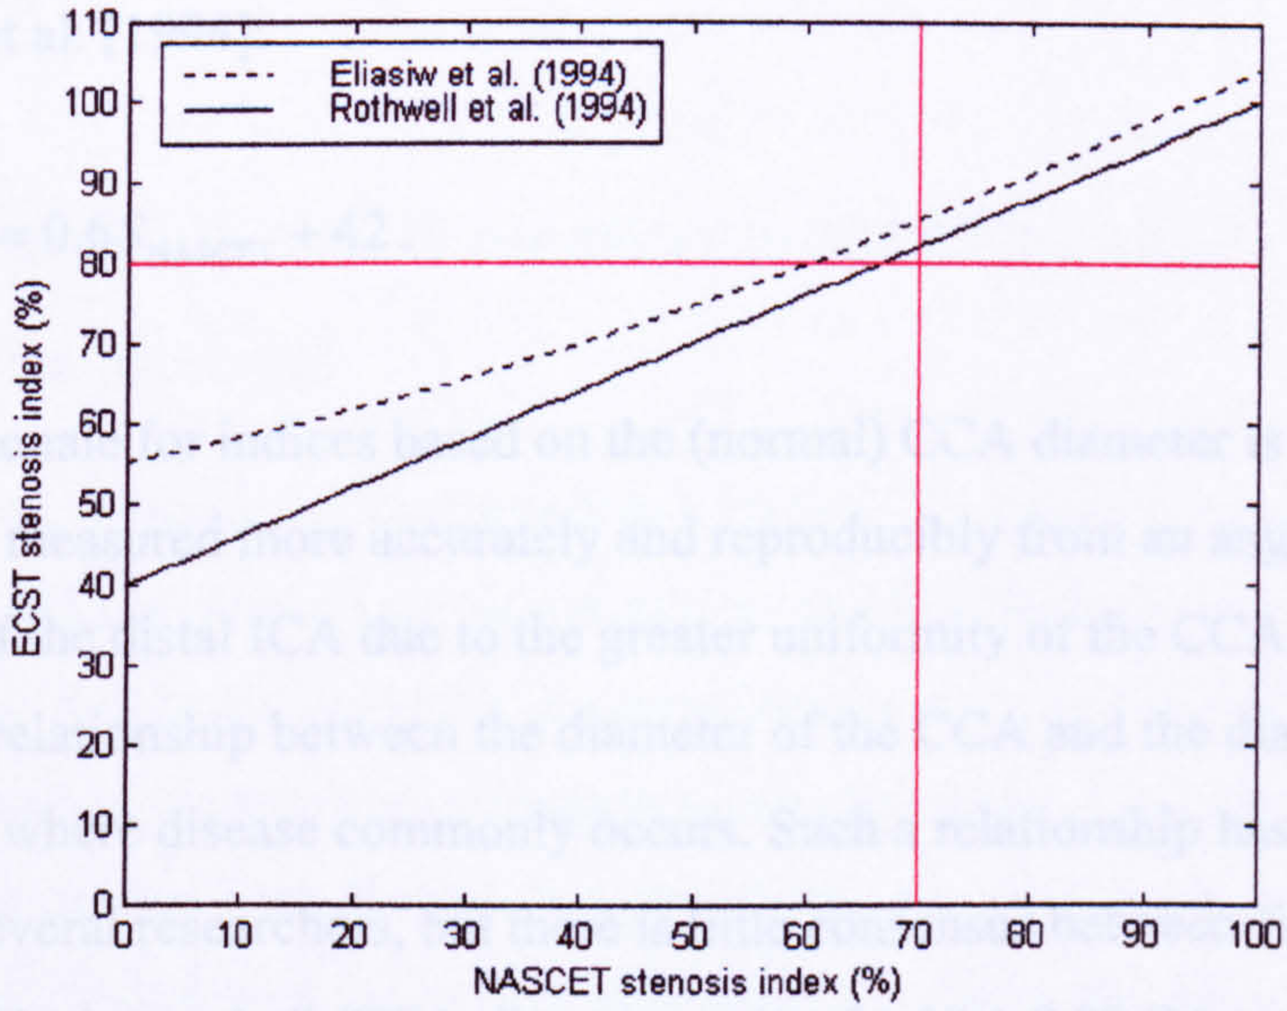


Figure 2.2 Conversion from NASCET to ECST stenosis indices from two published studies. The red lines indicate the clinically significant levels according to the NASCET and ECST trials.

Some general observations can be made from inspection of Figure 2.2: firstly, and most importantly, the conversion formulae appear to be fairly consistent (within 5%) above the threshold values (70% for the NASCET and 80% for the ECST). However, there is wider discrepancy between the converted values well below these thresholds. This may be due to the difference in numbers of angiograms analysed in the two studies. A rather disturbing observation is the reclassification after conversion of 'moderately diseased' artery, measured using the ECST method, as a 'mildly diseased' or 'disease-free' artery, according to the NASCET method. The clinical classifications 'mild', 'moderate', and 'severe' are important given lack of precision of many measurement methods, and reclassification under different stenosis indices is highly confusing.

An alternative stenosis index has been proposed in which the diameter of the normal CCA is used as the reference [Rothwell 1994; Eliasziw 1994]. A similar index, called the Carotid Stenosis Index, has also been put forward by Bladin et al. [1995], where the reference diameter is calculated by applying a scaling factor of 1.2 to the measured CCA diameter. This index is supposedly equivalent to the ECST measure, assuming a simple relationship between the bulb and CCA diameters. Bladin et al. [1995] found a very similar linear relationship to that discovered by Rothwell et al. [1994]:

$$S_{\text{CSI}} \approx 0.6 S_{\text{NASCET}} + 42 . \tag{2.6}$$

The rationale for indices based on the (normal) CCA diameter is that this parameter may be measured more accurately and reproducibly from an angiogram than the diameter of the distal ICA due to the greater uniformity of the CCA. Also, there is a purported relationship between the diameter of the CCA and the diameter of the carotid bulb, where disease commonly occurs. Such a relationship has been investigated by several researchers, but there is little consensus between findings: Williams et al. [1987] report a bulb/CCA diameter ratio of 1.19 ± 0.09 (Mean ± SD); Rothwell et al. [1994] found a mean ratio of 1.0 with a range of 0.7-1.4; whilst Eliasziw et al. [1994] report a mean ratio of 1.31. The inconsistency between these results might be accounted for by methodological or population differences between the studies, or by errors inherent in the estimation of bulb diameter from angiographic projections.

Conversion between vessel diameters, and in turn stenosis indices, presumes a relationship between the diameters at points along the vessel. Given the anatomical variability of the human carotid, this may not necessarily be true. Even if this is the case, measurements may have significant uncertainty for diseased vessels where asymmetrical disease or smooth plaque surface is present, or remodelling of the vessel wall has occurred. All of these factors can lead to misidentification as 'normal' an artery on an angiogram, and, hence, underestimation of true vessel diameter [Alexandrov 1993].

X-ray angiography is frequently termed the 'Gold Standard' technique for imaging and estimating the severity of carotid disease. This is due to the fact that the tech-

technique is well-established, widely available, and was used in the major clinical trials cited above. However, the term 'Gold Standard' is a somewhat misleading since it implies that the method is superior to other methods in terms of measurement error (i.e. accuracy and precision). Clearly, this is not the case, for rasons discussed in Section 1.2. Further, the limitations in terms of reproducibility of angiography-based measurement of carotid stenosis are well-documented [Chikos 1983; Murie 1986; O'Leary 1987; Rothwell 1998; Vanninen 1994]. Factors influencing reproducibility include poor vessel visualization, and subjectivity of diameter measurements.

2.4 Ultrasound-based measurement of stenosis

2.4.1 Direct methods

Ever since the potential of ultrasound as a method for imaging vascular disease was realised, it has provided a safe, non-invasive, and inexpensive alternative to x-ray angiography for the assessment and quantification of carotid stenosis [Sidhu 1997]. A direct approach to stenosis estimation using ultrasound is to make diameter (or area) measurements from B-mode images [Beebe 1999]. This can be done in a similar way to measurements made from angiograms whereby the lumen diameter is measured at predefined positions along the artery. Applying the NASCET stenosis measure directly to ultrasound can be problematic, however, since the field-of-view distal to the bifurcation is limited by the presence of the mandible. This means that for a proportion of patients the distal ICA where the vessel is normal cannot be imaged. The wall of the carotid bulb and CCA can usually be visualised using B-mode ultrasound, and, therefore, using the ECST or CCA-based stenosis indices is more appropriate. Indeed, the measurement of the diameter of the vessel wall using ultrasound is intrinsically more reliable than estimation from angiograms because it is visualised directly. In common with x-ray angiography, measurements from 2D ultrasound images are dependent on the imaging plane, particularly in the longitudinal view. For this reason, it is preferable to make linear measurements from transverse rather than longitudinal images, in order to minimise errors introduced by plane dependency.

Since the vessel wall can usually be visualised using B-mode ultrasound, a localised stenosis index can be calculated by using the local diameter of the vessel at the

point of maximal stenosis as the reference diameter. Similarly, a local area stenosis index can be defined (Equation (2.2)) as both the lumen and the vessel wall cross-sectional areas may be measured readily from a transverse B-mode image. In practice, where the geometry of carotid disease is asymmetrical and complex, all planar stenosis indices will be dependent on the imaging plane, which is selected subjectively by the ultrasonographer. It should also be noted that, in some circumstances, acoustic shadowing due to dense material within plaque can obscure part of the vessel, making accurate measurements impossible.

Several researchers have attempted to use colour or power Doppler information to quantify stenoses directly with varying success [Griewing 1996; Ranke 1992; Arbeille 1999; Steinke 1997; Steinke W 1996]. This practice cannot be recommended for quantitative purposes, however, due to the relatively poor spatial resolution of these modes, and the ubiquity of artifacts, especially over- and under-filling of the vessel (see Section 1.3.8.9).

2.4.2 Indirect methods

The most common ultrasound method for assessing carotid disease is to employ Doppler ultrasound to estimate the degree of stenosis indirectly from intra-stenotic blood velocity measurements. This is possible because as the lumen area decreases, conservation of mass for the flowing blood dictates that the mean velocity through the stenotic region increases. The mean velocity is often difficult to measure reliably [Evans 2000], so it is common practice to use the peak velocity or frequency shift. Most velocity-based stenosis criteria – frequently referred to as *Doppler parameters* in the literature – are based on peak systolic velocity (PSV), or end-diastolic velocity (EDV), or a combination of both. Original work by [Spencer 1979] used CW Doppler, but modern PW Duplex systems allow the velocities to be measured at precise locations within the vessel lumen (see Section 1.3.5).

During a Doppler ultrasound assessment, the maximum blood intra-stenotic velocity is measured [Zwiebel 1999]. This occurs either at the point where the lumen diameter is narrowest, or, where a high-grade stenosis is present which produces a 'flow jet', a short distance immediately distal to the stenosis, and within the jet. Elevated velocities are only recorded for reasonably high degrees of stenosis: an *in vitro*

model of an axisymmetric stenosis in the ICA predicts that the maximum intra-stenotic PSV rises steadily above about a 50% lumen area reduction (30% diameter reduction for this model), as the cross-sectional area of the lumen decreases [Langlois 1984]. At about 85% (equivalent to 60% diameter reduction) the PSV begins to fall sharply, as the increased flow resistance impedes the flow rate, ending in zero velocity when complete occlusion is reached. From this idealised model, the steep slope of the PSV-stenosis curve lies between about 65 and 85%. In addition, the double valued relationship between stenosis and PSV may lead to significant inaccuracy in stenosis estimation. Furthermore, in the more realistic case of an asymmetrical stenosis, this relationship will be distorted.

There are many studies in the literature reporting a correlation between PSV (and/or EDV) and a reference stenosis index, usually the NASCET or ECST indices [Grant 1999; Grant 2000; Carpenter 1996; Hunink 1993; Alexandrov 1997; Moneta 1995]. An important discriminating threshold is the velocity above which CEA has been shown to be beneficial. Hence, particular attention has been paid to successfully classifying stenoses as either above or below 60% or 70% stenosis using the NASCET method, or 80% using the ECST criterion. Thirteen published PSV (or peak systolic frequency shift) thresholds for determining greater than 70% stenosis using the NASCET index are summarised in Table 2.1.

It can be seen from Table 2.1 that there is a wide variation in the published velocity criteria for determining whether a stenosis is >70% according to the NASCET method, with the reported PSV threshold ranging from 125 to 325 cm/s. The differences are attributable to a number of factors: Doppler velocity measurements are well known to be subject to significant inter-equipment and inter-observer variability [Schwartz 1997; Howard 1991; Meyer 1997; Mikkonen 1996; Ranke 1999; Fillinger MF 1996]. This is due largely to differences in the characteristics of hardware and post-processing software between ultrasound machines, differences in operator-adjustable settings used, and errors inherent in the technique. Various sources of measurement error in Doppler velocities are well recognised [Ranke 1992; Hoskins 1996; Scissons 1997; Hoskins 1999a; Arbeille 1999]. These include angle dependent overestimation of velocity due to spectral broadening [Scissons 1997; Hoskins 1999a], and errors due to signal loss and single component measurements of the true

3D velocity vector. The former effect suggests that the Doppler angle should remain at a fixed value during velocity measurements.

Study	PSV threshold	EDV threshold	Velocity ratio threshold
Derdeyn 1995	125	-	-
Hood 1996	130	100	-
Browman 1995	175	-	-
Grant 1999	175 (200)[1]	-	2.5 (3.0)[1]
Carpenter 1996	210	70	3.0
Faught 1994	210	100	-
Withers 1990	225	-	-
Hunink 1993	230	-	-
Neales 1994	270	110	-
Elgersma 1998	270 (210)[2]	-	-
Moneta 1993	325	130	4.0 – 4.2
Wilterdink 1996	7 kHz	-	-
Ranke 1999	-	-	5.0[3]
Dhanjil 1997	-	-	15.0[4]

Table 2.1 Velocity (frequency) thresholds for discriminating > 70% stenosis using the NASCET stenosis index. [1] The value in brackets is the threshold relevant to a subgroup of asymptomatic patients. [2] The value in brackets refers to the threshold determined for a different time period. [3] The ratio (Mean $PSV_{INTRA\text{-}STENOTIC\text{-}ICA}$)/(Mean $PSV_{DISTAL\text{-}ICA}$) was used in this study. [4] The ratio $PSV_{INTRA\text{-}STENOTIC\text{-}ICA}/EDV_{CCA}$ was used in this study.

A further reason for discrepancies between reported results is the choice of criteria used to assess the performance of Doppler parameters and, hence, select a threshold. Methods commonly used include analysis of accuracy (with respect to the reference index), sensitivity, specificity, receiver-operator curves, cost-benefit, or clinical outcome. Often, a combination of such measures are considered. However, there is evidence that it may be difficult to satisfy, for instance, high sensitivity and high specificity requirements simultaneously [Hunink 1993].

Lastly, differences in the population sample used and bias introduced by the preselection of patients based on angiographic findings. A study by Grant et al. [1999] concludes that asymptomatic and symptomatic patients should be treated as separate groups after finding different optimal PSV and ratio thresholds for each group (see Table 2.1). Another study by Elgersma et al. [1998] found that the optimal PSV threshold changed over time (Table 2.1), and the authors therefore concluded that:

"Individual laboratories should assess duplex accuracy regularly and adjust adopted criteria if necessary to keep diagnostic performance optimal".

A serious criticism of assessment of carotid disease using a single velocity measurement is that it does not account for inter-patient variations arising from differences in blood pressure, cardiac output, heart rate, and local arterial geometry and compliance. Intra-stenotic velocities are also influenced by arrhythmia, the extent of disease within the extracranial ICA, severe ipsilateral atherosclerotic disease upstream (intracranial ICA or cerebral vessels) or downstream (aorta, proximal CCA, innominate artery) from the carotid bifurcation, or the presence of a contralateral stenosis [Beckett 1990; AbuRahma 1995; Henderson 2000]. All of these factors influence the rate of blood flow through the ICA. Inter-operator and inter-equipment variations also contribute. It is not surprising, therefore, that plots of, say, intra-stenotic PSV, versus a reference stenosis index typically exhibit a wide dispersion for any particular degree of stenosis.

In an attempt to normalise for the inherent variability in *in vivo* Doppler velocity measurements, velocity ratios have been proposed as an alternative [Moneta 1993; Carpenter 1996; Ranke 1999; Dhanjil 1997; Grant 1999]. In general, a velocity ratio is defined as:

$$\frac{v}{v_{\mathrm{REF}}}, \tag{2.7}$$

where v is an intra-stenotic velocity (usually the PSV), and v_{REF} is a reference velocity. Popular choices for the reference velocity are the PSV or the EDV in the CCA. Some published ratio thresholds are included in Table 2.1. Again, a wide range of values have been reported with values from 2.5 to 15.0, depending on which ratio is adopted. A recognised problem with some velocity ratios is the choice of a CCA velocity as the reference [Hunink 1993; Ranke 1999]. This is because some high-grade ICA stenoses are known to cause significant reduction and variation in the velocity in the CCA. As a consequence, the CCA velocity is not independent of the degree of stenosis. The CCA velocity is also affected by the proportion of blood flowing into the ECA, which is also subject to considerable variation due to anatomical and physiological differences between individuals. Recognising these deficiencies, Ranke

et al. [1999] propose the distal ICA velocity as more suitable reference. Questions remain, however, concerning the feasibility of measuring post-stenotic ICA velocities given the limited field-of-view for many patients, as well as the tortuosity of the distal ICA. Problems will also arise when disease extends far distally in the ICA.

One other method for quantifying degree of stenosis relies on the assessment of the Doppler spectrum in the post-stenotic region [Kalman 1985; Kassam 1985; Arbeille 1995]. The method is based on the presence of disturbed flow immediately distal to a stenosis, which is manifest as spectral broadening in the Doppler trace. The appearance of this spectral broadening changes with disease severity, and attempts have been made to categorise these changes and relate them to standard stenosis indices. This is usually done qualitatively by inspecting the Doppler spectrum. Hence, the method is highly observer dependent and prone to errors arising from spectral artifacts. The relationship between degree of stenosis and spectral broadening is also very general. Consequently, the method is only suitable for the purpose of broad classification, and is not appropriate for precise quantification of carotid stenosis on its own. Improvements in B-mode images, and the advent of PW Doppler, and colour and power modes, mean that post-stenotic spectral analysis is now rarely used. The method may, however, be helpful as a supplement to other Doppler parameters, for instance, when flow disturbances provide the only evidence of a stenosis obscured by shadowing.

It was noted above that the various Doppler parameters are only applicable to high-grade stenoses. This does not significantly undermine their usefulness as a substitute to x-ray angiographic methods for determining clinically significant stenoses. It does, however, place limits on when Doppler parameters should be applied, and how they should be interpreted. The implication here is that direct measurement methods are the method of choice for quantifying low to medium grade stenoses, as well as plaque progression or regression. There are also circumstances where a stenosis exhibits considerable lumenal narrowing, but intra-stenotic velocity measurements are barely elevated [D'Alotto 2000]. This type of 'paradoxical stenosis' usually occurs within the carotid bulb, which can become very dilated.

2.5 3D quantification of carotid disease

There are several reports in the literature of 3D quantification of carotid artery disease [Palombo 1998; Delcker 1994a; Delcker 1996; Allott 1999; Allot 1998; Meairs 2000; Griewing B 1997; Schminke U 2000; Yao 1998; 1998; Arbeille 1995]. To date, the majority of studies have concentrated on the measurement of plaque volume, particularly using 3D ultrasound (see Chapter 3), although Arbeille et al. [1995] propose a method of plaque estimation based on planar area and length measurements. This method only provides a rather crude estimation of the true plaque volume, however. Yao et al. [1998] also determined plaque length and area and diameter stenosis indices from parallel transverse slices. 3D analysis is usually performed on 3DUS images reconstructed from 2D B-mode images, but Schminke et al. [2000] advocate the value of power Doppler to improve plaque visualisation.

Plaque volume is ideal as a gross measure of disease severity. It is also well suited to monitoring plaque progression or regression, for example, as a result of drug therapy [Hennerici 1991]. This is the case because plaque progression/regression is always associated with a change in volume, but not necessarily a change in degree of diameter (or area) stenosis. Plaque volume alone, however, is likely to have limited utility as a risk factor, since it confers no information on the distribution or geometry of disease. Unlike, stenosis indices, the measure is also not normalised for inter patient variations in vessel size. Therefore, a volumetric stenosis index (Section 2.1) may be more useful as this better reflects both disease severity and 'plaque burden'.

Recently, serial cross-sectional area measurements from 3DUS data have been used to quantify the distribution of atherosclerotic disease in lower extremity bypass grafts [Leotta DF 2001], and the 3D geometry of abdominal aortic aneurysms [Leotta 2001]. To the best of the author's knowledge, however, this approach has so far not been applied to carotid artery disease.

2.6 Motivations for the thesis

From the considerations above it is clear that many limitations exist in current clinical methods for assessing the severity of carotid artery disease. This is due partly to limitations in imaging techniques, and partly to limitations in the definitions of

commonly adopted stenosis indices. Such indices conceptually provide an intuitive measure of disease severity, but many indices have been proposed and several imaging techniques may be employed to estimate them. This has led to a great deal of inconsistency in quantification of carotid disease, and very little standardisation.

Of particular concern is the lack of consensus in the literature on velocity-based criteria measured using Doppler ultrasound. This has provoked intense debate about whether ultrasound velocity measurements are sufficiently accurate and reproducible to be used in place of x-ray angiography. Doubtless, this debate will continue. Although ultrasound is undeniably effective as a screening tool [Derdeyn 1995], it would seem that the sources of error are numerous enough to recommend that conventional velocity-based parameters should probably be used only to indicate the clinical significance of disease in the ICA (e.g. > 70% using the NASCET stenosis index), rather than a precise measure of degree of stenosis [Grant 2000]. Such a view implies that direct measurements from ultrasound images are more reliable from the point of view of precise quantification [Beebe 1999]. Fortunately, direct plaque measurements are becoming easier as a result of advances in ultrasound technology.

In keeping with the general trend of technological advancement, several new methods for estimating degree of stenosis have been introduced within the past 20 years, often claiming distinct advantages over existing methods. This has resulted in a large number of publications on imaging and the quantification of carotid stenosis in the literature. In an important recent paper, Rothwell et al. [2000] reviewed a total of 40 published studies on imaging and carotid stenosis measurement – 20 published before1991 and 20 published between 1993 and 1997. The results found "many basic methodological deficiencies" in the majority of the studies examined, and in conclusion the authors write:

> "The majority of published studies are not of a sufficient standard to enable the results to be used to inform clinical practice. The utility of future studies could be improved considerably by better adherence to 9 simple methodological guidelines."

The nine guidelines referred to here are as follows:

"1. Prospective rather than retrospective study design.
2. Patient selection based on a consecutive series or a random sample.
3. Adequate detail of study population (age, sex, clinical presentation, and indications for investigation).
4. Adequate detail of imaging techniques (sufficient for the study to be repeated).
5. Inclusion of all investigations, i.e. patients with poor-quality imaging were not excluded.
6. Blinded assessment of images.
7. Adequate detail of derivation of measurement of stenosis from images or data (sufficient for the study to be repeated).
8. Adequate data on the reproducibility of measurements of stenosis (data on either intra-observer or interobserver agreement were accepted).
9. Study powered according to a sample-size calculation."

A fundamental limitation of planar imaging methods is that a 2D representation of 3D anatomy is produced, resulting in an inevitable loss of information. For x-ray angiography, this reduces to the projection problem highlighted in Section 1.2.1. Ultrasound measurements are highly dependent on the imaging plane, but the guided movement of a hand-held probe gives some scope for optimal selection of the plane(s) from which measurements are made. Transverse vessel views are also readily accessible using ultrasound. With experience, ultrasonographers become adept at mentally reconstructing 3D anatomy from a number of planar views. However, such a process is still subjective, and measurements become increasingly difficult as the 3D anatomy becomes more complicated.

Although these problems can be reduced by multiplanar acquisition, this may not be effective in eliminating measurement errors. 3D imaging techniques are, therefore, preferable. 3DUS is comparatively new technology which combines the advantages of ultrasound with the ability to perform volumetric imaging. This technique was adopted for this thesis as it overcomes many of the fundamental limitations of conventional, planar imaging. Since the soft tissue of the vessel wall and plaque are imaged in addition to the lumen, in theory, more accurate and descriptive quantitative measures of atherosclerotic disease may also be determined from 3DUS reconstructions.

2.7 Aims of the thesis

The overall aim of this thesis is to investigate methods for quantifying carotid artery disease using 3DUS imaging. The specific aims are as follows:

- to develop and validate a low-cost, flexible 3DUS acquisition system which takes advantage of the range of imaging modes (present and future), and digital image transfer facility, available on the ATL HDI-5000 ultrasound scanner.
- to develop a method for reconstructing the wall and lumen of the diseased carotid artery bifurcation directly from 2D images acquired using the 3DUS system.
- to develop algorithms to automatically quantify carotid artery disease from the 3D reconstruction. The quantitative measures are to include clinically accepted (planar) measures, as well as new 3D-specific measures.
- to assess the accuracy of the vessel reconstruction and quantification algorithms using anthropomorphic ultrasound phantoms.
- to assess the clinical feasibility of the techniques developed using data from a group of patients with carotid artery disease.

Chapter 3

3D ultrasound imaging

This chapter presents an overview of the current state-of-the-art in 3D ultrasound imaging. The literature on the subject is reviewed including clinical applications of the technique, together with the technical considerations for each of the stages involved in the acquisition and processing of three-dimensional ultrasound data.

3.1 Applications of 3D ultrasound

Three-dimensional ultrasound (3DUS) is a relatively recent advance in ultrasound imaging. The idea of constructing a 3D ultrasound image was conceived as early as the 1960's, but did not become realisable for many years due to the limitations of available technology. Potential applications are as extensive as those of conventional 2D ultrasound. 3DUS has shown great potential in the detection of anatomical abnormalities, especially in the foetus.

To date, by far the most popular applications have been in the area of obstetrics [Devonald 1995; Kelly 1992; Klein 1993; Lee 1995; Steiner 1994]. The most common application is prenatal diagnosis of congenital fetal malformations, and other abnormalities of the foetus [Nelson TR 1992; Merz 1995a; Merz 1995b]; 3DUS imaging has been demonstrated for the fetal heart [Deng 1996; Chang 1997; Sklansky 1999; Nelson TR 1996], liver [Laudy 1998], cranium [Pretorius 1994], and skeleton [Steiner 1995; Nelson 1995; Lee 1999].

Other areas amenable to 3DUS investigation are the adult heart [Arbeille 2000; Belohlavek 1993; Binder 1996; Coppini 1995; Ghosh 1982; Handschumacher 1993; Levine 1992; Moritz 1983; Sarti A 1993; Seward 1995], the peripheral vasculature [Downey 1995a; Fenster 1998; Griewing B 1997; Picot 1993; Ritchie 1996], the prostate [Sehgal 1993; Downey 1995b; Richard 1993; Tong 1996], the uterus [Athanasiou 1997; Feichtinger 1993; Weinraub 1996], the bladder [Riccabona 1996a], the gall bladder [Fine 1991; Pauletzki 1996], the kidneys [Matre 1999; Gilja 1995], the breast [Rotten 1999; Weismann 2000a], the eye [Downey 1995c; Downey 1996], and the anal canal [Gold 1999; Hünerbein 1997].

Examples of more specialised applications include: the guidance of biopsy needles and endoscopic and surgical instruments [State 1996; Kavic 1996; Weismann 2000b]; the analysis of tongue motion [Stone 1990] and lip motion [Deng 2000]; planning of maxillofacial surgery [Hell 1993]; and muscle imaging [Delcker 1997a].

Dynamic and functional imaging of the heart has been the focus of recent attention, since the heavy technological demands of real-time 3D - or four-dimensional

(4D) – imaging are now being met by technical advances [Arbeille 2000; Delabays 1995; Seward 1995].

Numerous studies have investigated the use of 3D reconstruction of colour or power Doppler data for blood vessel visualisation [Rickey 1991; Guo 1995; Pretorius 1992; Pretorius 1998; Griewing B 1997; Downey 1995a; Ritchie 1996; Picot 1993; Li 1997]. In addition to the peripheral vasculature, this technique has been applied to the cerebral vasculature [Delcker 1997b; Lyden 1997], the placental vasculature [Pretorius 1998], and to the assessment of tumour vascularity [Downey 1995b].

3DUS images of blood vessels have also been generated using intravascular transducers [Lees 1991; Kovalski 2000; von Birgelen 1996; Gil 1996; Chandran 1996; von Birgelen 1997]. In the past, 3D intravascular images have been generated by 'stacking' together 2D image slices, acquired regularly as the catheter is pulled back through the vessel. This approach does not use any additional means to accurately register the acquired images in 3D space, and, therefore, there is an implicit assumption that consecutive slices are parallel is implicit. Consequently, the geometric accuracy of images generated in this way cannot be guaranteed. So-called 'forward looking' 3D intravascular transducers do offer anatomical accuracy, but a relatively limited field-of-view [Ng 1994].

In general, 3DUS has been successful for applications where a 3D image enhances the appreciation of anatomy, particularly where normal anatomy is complex. The other immediate advantage of 3DUS is the accurate quantification of volume [Gilja 1999]. Traditionally, the volume of large organs has been estimated by calculating the product of linear dimensions obtained from three orthogonal (or near orthogonal) 2D views, and multiplying by a scaling factor. This method is, however, susceptible to very large errors [Pauletzki 1996; Griffiths 1986; Hughes 1996]. 3DUS allows volumetric measurements to be made without relying on simplifying assumptions about the shape of the volume.

3.2 Generation of a 3D ultrasound image

The purpose of a 3DUS acquisition system is to record and register in 3D space ultrasonic echo data received from a tissue volume. Although the feasibility of volu-

metric imaging using matrix-array transducers has been demonstrated (see Section 3.3.3), most 3DUS acquisition systems utilise conventional transducer technology to obtain 2D images in conjunction with external image capture and position tracking devices.

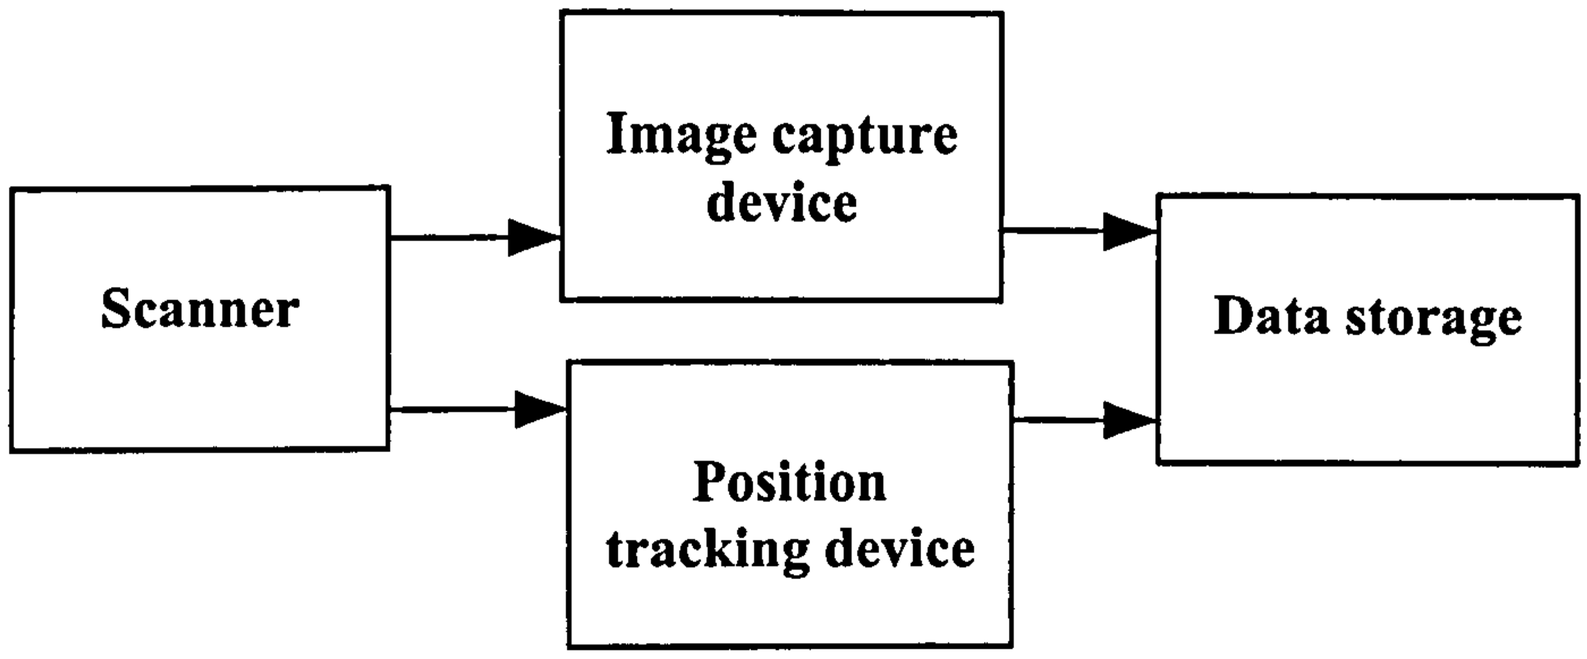


Figure 3.1 Block diagram of a generalised 3DUS acquisition system. 2D ultrasound images are acquired with current transducer technology, and captured using an image capture device. The 3D position of captured ultrasound images is measured by a position tracking device. Both images and positional data are stored on appropriate data storage devices.

Figure 3.1 shows a block diagram of a generalized 3DUS acquisition system which illustrates the architecture of most 3DUS systems. An image capture device stores ultrasound image data produced by the scanner, and a position tracking device registers the 3D location of each captured image. Common image capture devices include video recorders and analogue frame grabbers. Images captured using a video recorder are stored on video tape, whereas images captured using a frame grabber are digitised and stored on any high capacity digital medium. A wide variety of position tracking devices have been used, and these are reviewed in Section 3.3.

For some applications it is necessary to synchronise image capture with an event in a physiological signal; for instance, the peak of the R-wave (i.e. QRS complex) of an electrocardiogram (ECG). Such as procedure is known as gating and is used to minimise artefacts which might otherwise arise due to tissue motion within the scan volume.

Gating may be performed either during acquisition or retrospectively. In the for-

mer, image data is only captured at a particular point (or points) in time corresponding to a detectable physiological event. This approach reduces the amount of initial data that is acquired, but relies on reliable and automatic event-detection method. Where this is not possible, it is necessary to adopt a retrospective gating approach. Then, images are acquired at a regular rate, and a subset of the acquired images is selected off-line according to their acquisition time in relation to the physiological signal of interest.

The two most important types of gating are ECG gating and respiratory gating. ECG gating is implemented on most commercial ultrasound scanners as the peak of the R-wave can be identified automatically using modern signal processing techniques. Respiratory gating may be implemented straight forwardly, but is rarely employed in practice.

3DUS acquisition systems may be divided into those which use a dedicated 3D transducer, and those which use a conventional 2D imaging transducers with the scanner adapted so that data may acquired from a tissue volume. They may also be classified as either constrained or freehand systems: in a constrained acquisition system, the motion of the scan-probe is constrained. A special case is a system in which the transducer is moved within the probe casing, and the scan-probe remains stationary during acquisition (see Section 3.3.1.1). Freehand acquisition, on the other hand, allows data to be acquired using unconstrained motion of the scan-probe as for standard (2D) ultrasound imaging.

In general, constrained systems have a more limited field-of-view; that is, the size of the volume which can be imaged during a single acquisition is smaller than a freehand system. However, constrained systems have the advantage that regularly-sampled ultrasound data are acquired, with a simple spatial relationship between consecutive images. Therefore, acquisition, post-processing and display are both simpler and more efficient than for freely-acquired 3DUS data.

The acquisition of 3DUS data is the first in a series of stages required to generate a 3D image – or, more generally, a 3D representation – of the scan-volume. Figure

(*continued on page 76*)

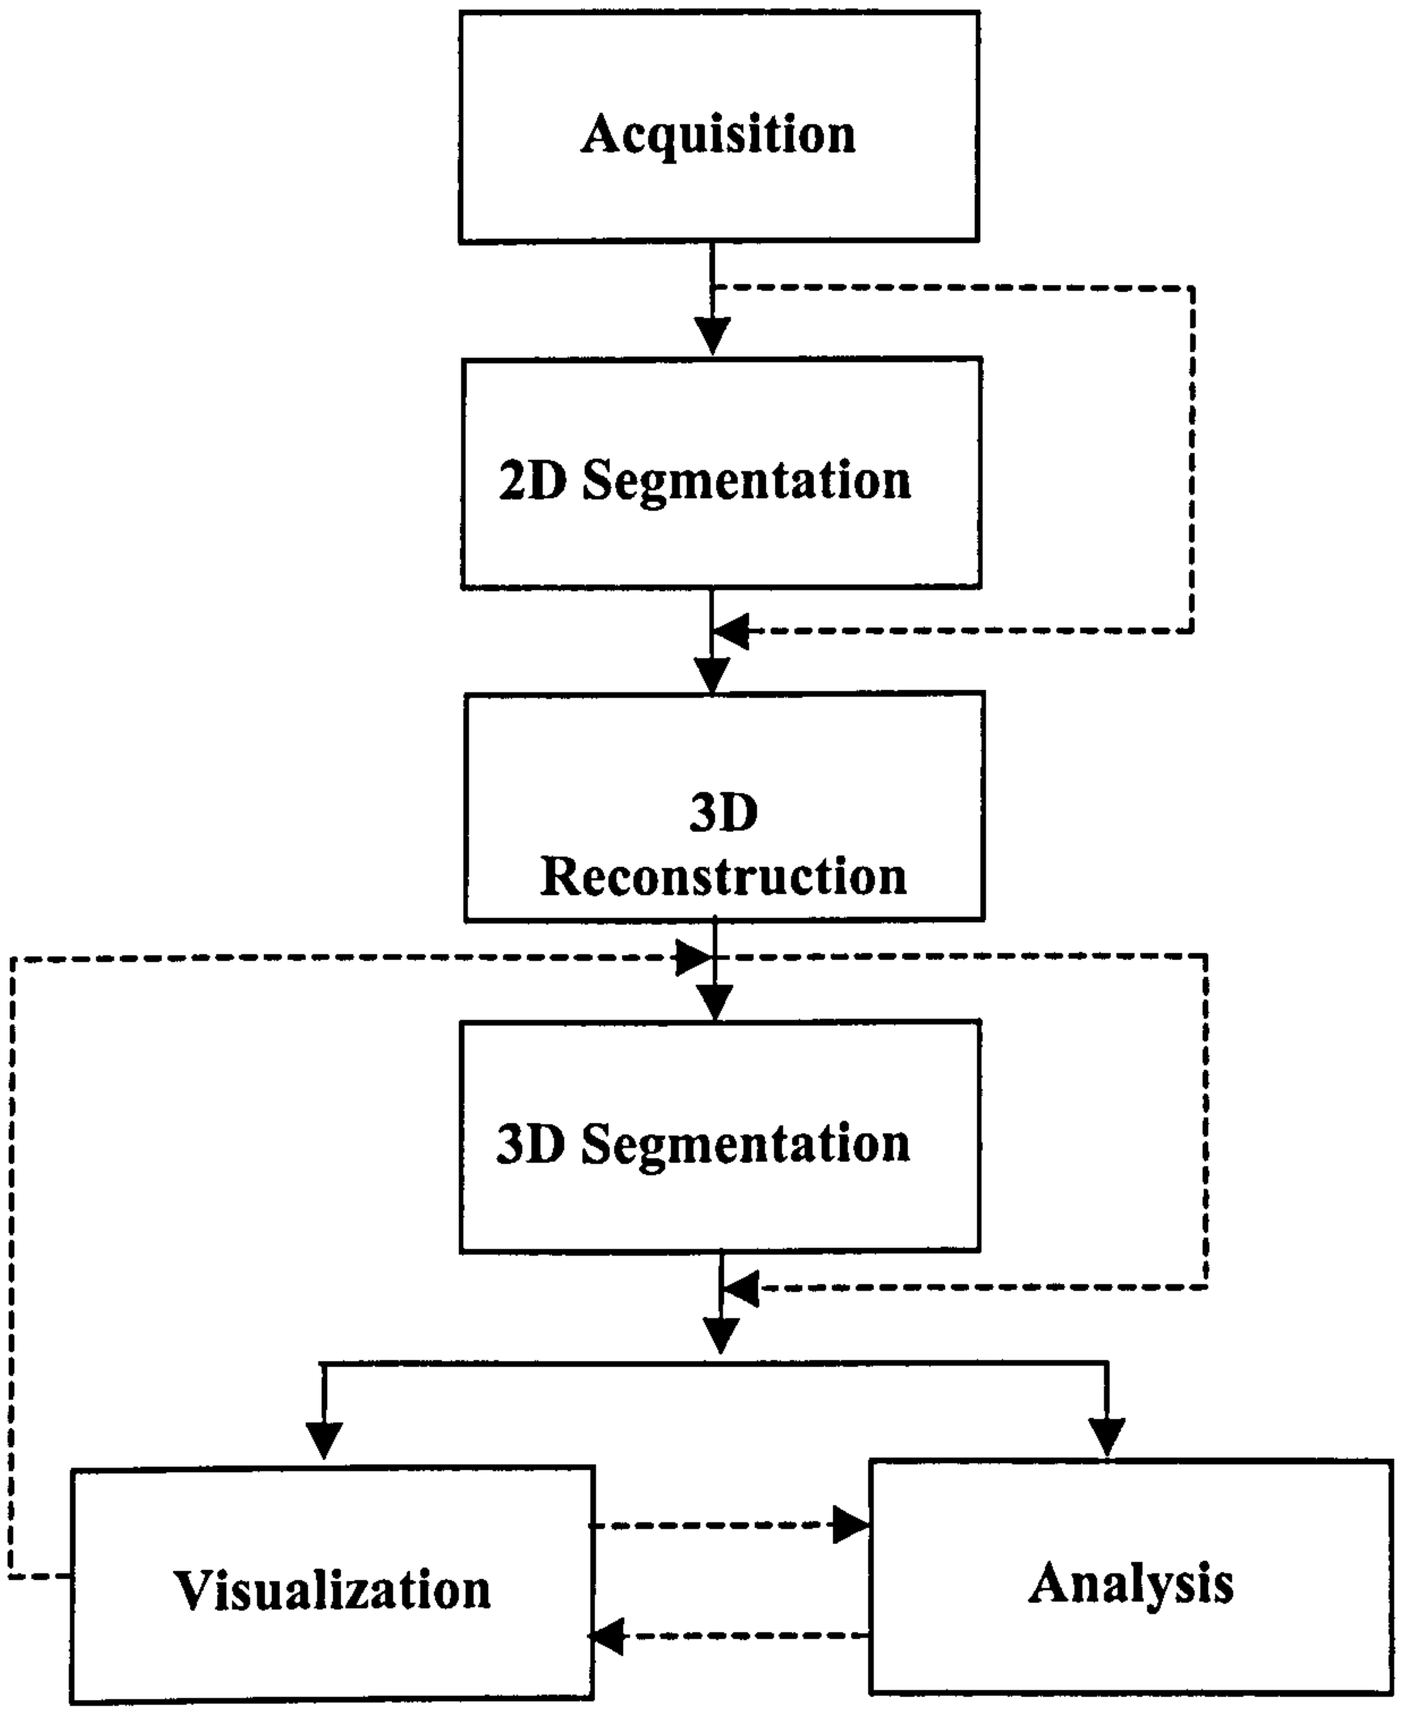


Figure 3.2 Stages in the generation and analysis of a 3DUS image.

3.2 shows the principal stages in 3DUS imaging. For clarity, each of these is introduced briefly below. More detailed descriptions of specific techniques associated with each stage form are presented later in the chapter.

(i) Segmentation

Following the acquisition of ultrasound data from a tissue volume, image data are usually processed to highlight features of interest. This process is known as *segmentation*, and might be performed manually, semi-automatically (i.e. involving some user-interaction), and automatically (i.e. involving no or minimal user-interaction). A simple example of automatic segmentation is retaining only colour pixels in a 3D colour or power Doppler image. A direct by-product of segmentation is data reduction. The process may also improve the quality of a 3D image by removing unwanted structures that may otherwise obscure features of interest. As indicated in Figure 3.2, segmentation may be performed either before or after 3D reconstruction, or both. For instance, each of a set of acquired 2D ultrasound images may be segmented prior to 3D reconstruction. In contrast, segmentation following reconstruction is applied directly to a 3D representation, and for the remainder of this thesis this will be distinguished by the term '3D segmentation'.

(ii) 3D reconstruction

The next stage is *3D reconstruction*: this describes the process of generating a convenient 3D representation of the scanned volume from acquired ultrasound and positional data.

(iii) 3D visualisation

3DUS data is frequently presented by displaying multiple 2D images corresponding to cut-planes through the 3D data set. Such a technique is called multi-plane slicing and the 2D images are sometimes called C-planes. Multi-plane slicing can be implemented simply and efficiently, but modern graphics technology enables 3D data sets to be displayed on a 2D screen with the illusion of three dimensions. Unlike a hologram, the image presented to the viewer in this manner is not truly 3D. Rather, inspection from different physical viewpoints (as opposed to virtual viewpoints altered by 'rotating' the 3D data set) does not allow the viewer to see behind objects. It is in fact a depth-dependent projection of a 3D scene onto a 2D display. 3D visualisation

typically involves extensive user interaction to enhance the perception of 3D structures.

(iv) 3D analysis

In order to make diagnoses or clinically meaningful measurements, an analysis of the reconstructed 3D data set is required. In its simplest form, analysis is qualitative involving, for example, visual interrogation of a 3D image by an expert. Quantitative analysis must be used to precisely measure dimensions, volume, and other parameters of interest, however.

3.3 3D acquisition systems

3.3.1 Constrained acquisition

3.3.1.1 Mechanical carriage systems

Many early 3DUS acquisition systems employ a mechanical carriage device to translate or rotate the scan-probe by predefined displacements. The scan-probe is secured in the carriage device, and images are captured at regular intervals as it is moved across the skin surface. The driving mechanism may be controlled manually, or by a motor. Such devices permits limited motion of the scan-probe, as shown in Figure 3.3(a).

The principal advantages of systems based on a mechanical carriage device are that they are straightforward to implement, ensure accurate registration of the captured images, and, in the case of motor-driven devices, enable rapid acquisition. The captured slices are regularly-spaced, and this simple spatial relationship between slices significantly reduces the complexity of a 3D image reconstruction (see Section 3.4).

In practice, the constraint on the scanning motion limits the applicability of mechanical carriage systems. Conventional (2D) ultrasound imaging involves collecting a range of different views of the anatomy of interest, each obtained from different scan-probe position and orientation. Constrained systems limit the range of views possible, and this may lead to a portion of a volume being imaged satisfactorily. A further problem arises from the fact that maintaining sufficient contact between the scan-probe surface and the skin can be difficult or impossible. Some researchers have

overcome this problem by applying gel liberally or, more practically, using a deformable, fluid-filled sack which can be moulded to the local contour of the skin. Anecdotal evidence suggests that linear carriage systems have also been unpopular also because they are often bulky and cumbersome to use. In spite of their limited applicability, linear systems have been used for imaging the heart [Raichlen 1986; Ghosh 1982], the prostate [Hastak 1982], the foetus [Klein 1993], and the carotid arteries [Delcker 1994a; Delcker 1994b; Griewing B 1997; Hennerici 1991; Picot 1993; Rosenfield 1992].

2.2.1.2 Electromechanically-driven transducers

A few commercial 3DUS scanners (for example, the Kretztechnik Voluson 730 system) incorporate a small driving motor within an enlarged scan-probe casing. This is used to translate or rotate the transducer within the casing as ultrasound slices are captured at regular positions/angles, in the same manner as bulkier carriage systems described above. Loss of acoustic coupling is not a problem because the fluid-filled scan-probe is held steady, in contact with the skin. A popular configuration is to rotate the transducer about an axis parallel to the direction of the linear array so that a fan shaped set of images is acquired, as shown in Figure 3.3(b).

Electromechanically-driven transducers require a housing larger than conventional scan-probes. They are, nevertheless, small enough to be handled in much the same way and provide a neat solution to 3DUS acquisition with the same advantages as mechanical carriage systems. One disadvantage is that the motion of the transducer is usually quite limited enabling only a small volume to be imaged. There is also no control over the orientation of captured images slices, and this may degrade the 3D reconstructed image if sub-optimal views are used.

The scan-probe must be held stationary during acquisition so that all images are acquired such that their relative positions and orientations are consistent with those expected by the 3D reconstruction algorithm. This condition may be difficult to satisfy in some applications, for example, when pronounced skin surface motion arises from respiration.

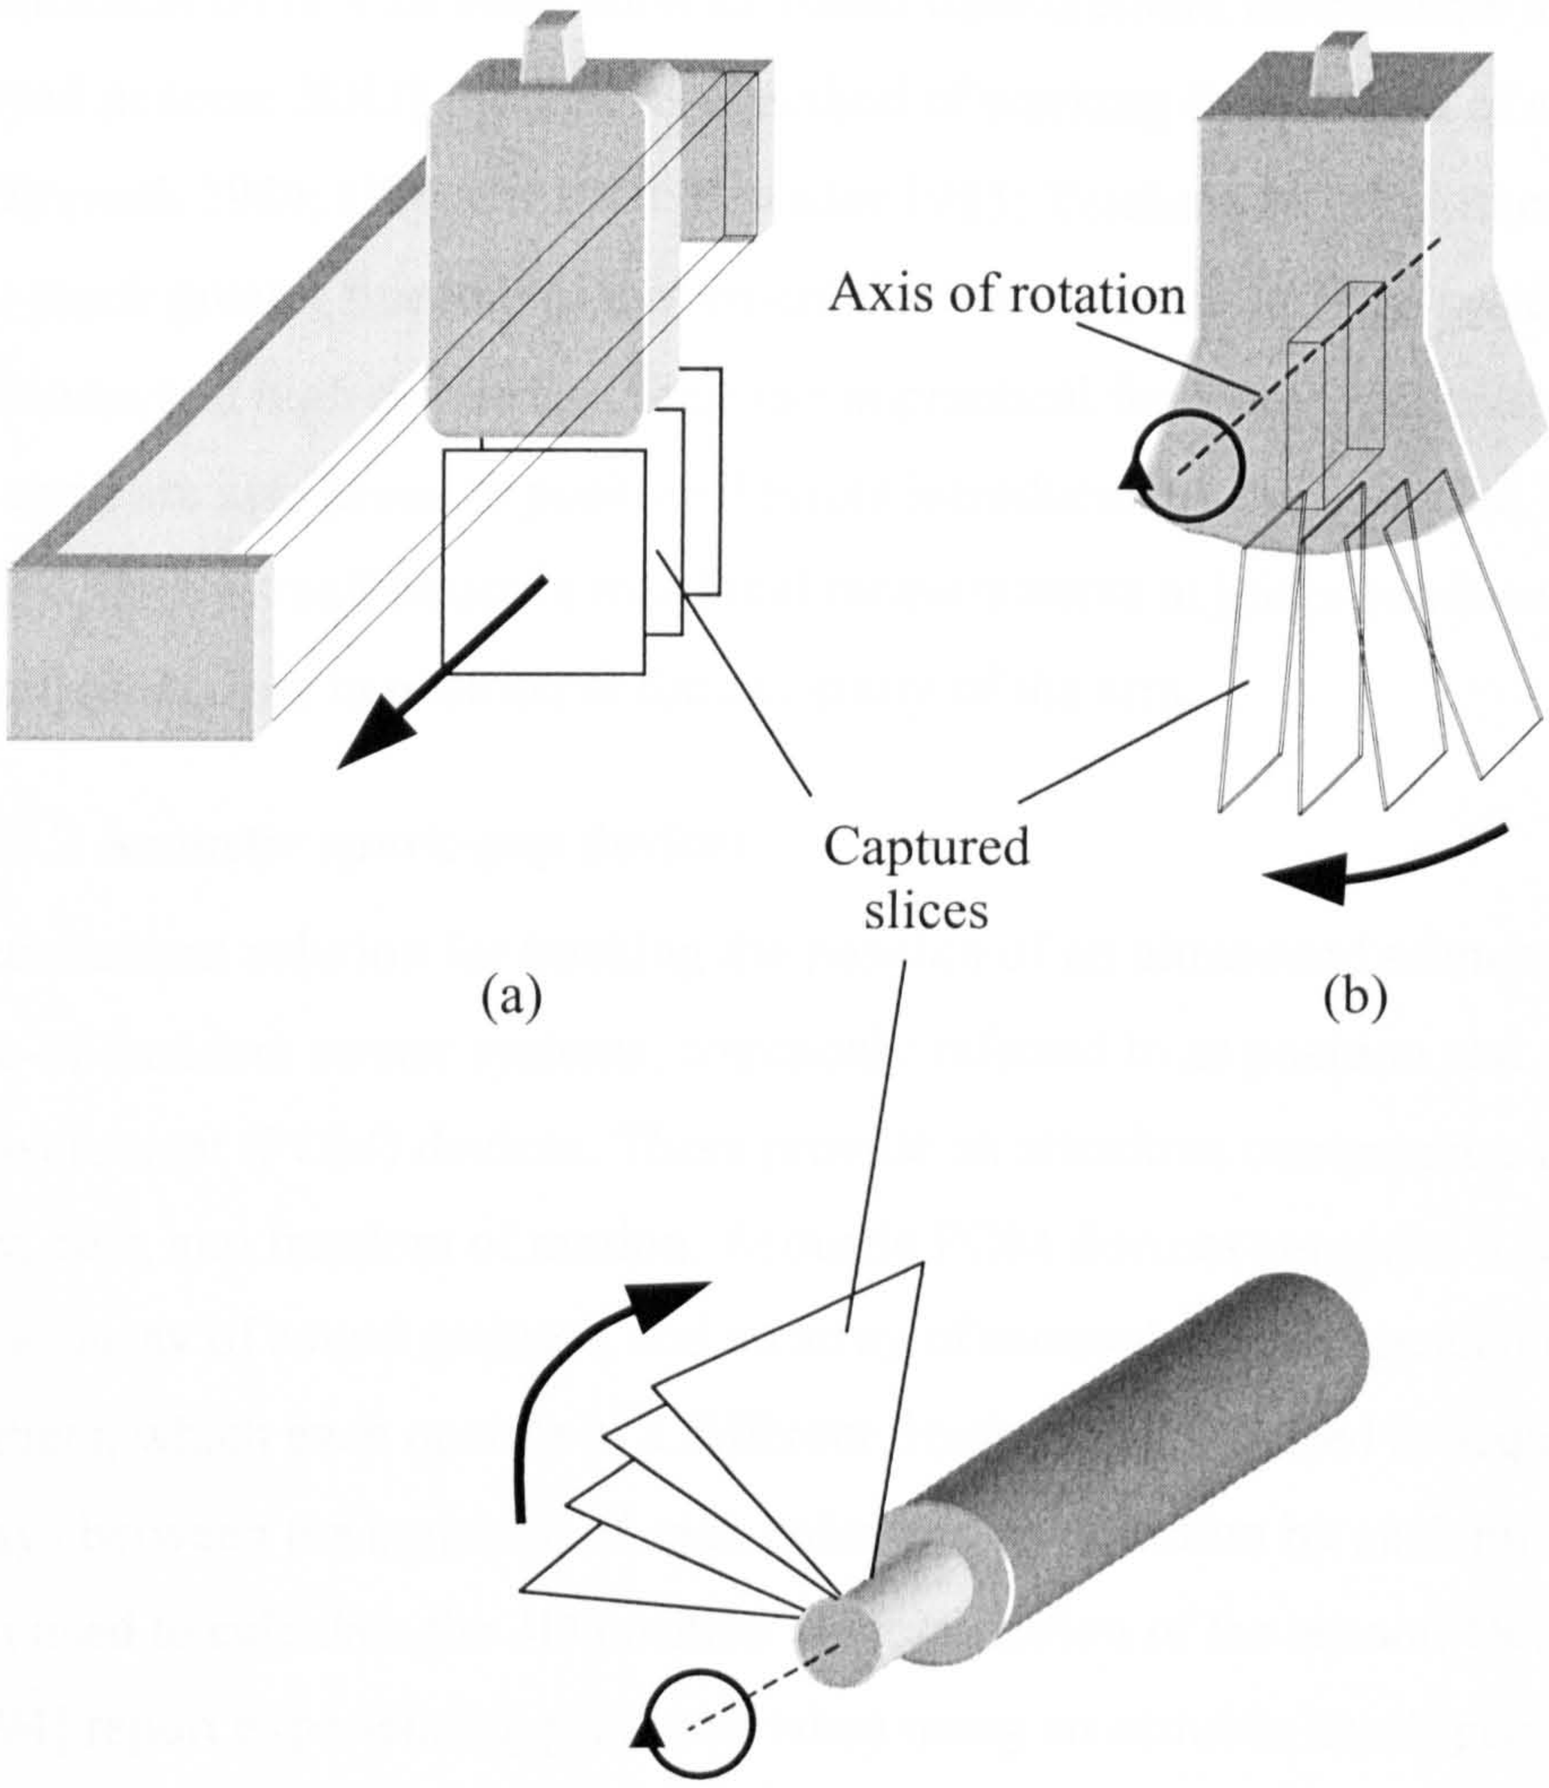


Figure 3.3 Three different types of constrained 3DUS acquisition system: (a) a linear 3DUS scan using a mechanical carriage system to translate a conventional (2D) scan-probe; (b) fan shaped image acquisition achieved by tilting a linear-array transducer within the scan-probe casing; (c) fan shaped acquisition using a rotating, catheter-mounted transducer.

3.3.2 Freehand acquisition

3.3.2.1 Mechanical arms

Mechanical arms with potentiometer-based measurement instruments have been employed in some 3DUS systems as a method of tracking the position of a scan-probe [Nikravesh 1984; Ohbuchi 1990; Sawader 1983; Teicholz 1974]. Although arms allow much greater freedom of movement than mechanical carriage systems, their bulkiness and high cost makes their use impractical for many applications. Mechanical arms are also prone to positional errors introduced by the so-called lever-arm effect, whereby small errors in rotational measurements at joints combine to produce significant errors in position at the extremity of the arm.

3.3.2.2 Acoustic spark-gap devices

A convenient solution for tracking the position of an ultrasound scan-probe are 6 degree-of-freedom sensor systems, commonly referred to as position and orientation measurement (POM) devices. These provide an attractive compromise between accuracy, cost, and freedom of motion. Acoustic POM devices comprise a sensor containing an array of sound emitters, and an array of microphones suspended above. The emitters, which each operate at a different frequency, are pulsed in sequence. The delays between the emission of each pulse and its detection by each microphone are then used to calculate the 3D position and orientation of the sensor. Detmer et al. [1994] report experiencing problems when using an acoustic spark-gap device within a 3DUS acquisition system. For the system in question, developed by Moritz et al. [1983], they report:

> "...the examiner had to ensure that the spark-gap array always faced toward the microphone array above the examining table which limited the scanhead's orientation to range of about ±70 degrees around the vertical."

There are also indications in the literature that some patients are unable to tolerate the loud sounds and sparks produced by some acoustic POM devices. Further, the accuracy of these devices is, to some extent, sensitive to the temperature and humidity of the operating environment.

3.3.2.3 Optical methods

The use of multiple video cameras to record a routine (2D) ultrasound scan has been investigated [Trobaugh 1994]. The technique involves processing movie sequences off-line to determine the position and orientation of the scan-probe. This stage is computationally intensive, and there is the restriction that the scan-probe must be within the field-of-view of all the cameras during acquisition. Alternative optical methods employ optical POM devices, typically based on light-emitting diodes or lasers. Examples of commercial devices suitable for 3DUS applications include the Northern Digital *Polaris* and *Optotrak* systems (Northern Digital Inc., Ontario, Canada), which have RMS positional accuracies of 0.35 mm and < 0.15 mm, respectively. The Ascension *LaserBird* has a positional accuracy of 1 mm RMS, and a rotational accuracy of 0.5° RMS (Ascension Technologies Corp., Vermont, USA). To date, optical POM devices, have been little used for 3DUS imaging in spite of their high accuracy specifications, mainly due to their high cost. However, they are becoming less expensive. Another factor which may have contributed is the necessity of maintaining an unobstructed line-of-sight between the sensor and fixed reference points.

3.3.2.4 Electromagnetic devices

Currently, electromagnetic POM (EPOM) devices are the most popular solution for freehand 3DUS systems described in the literature [Delcker 1996; Delcker 1998; Gardener 1991; Hodges 1994; Hughes 1996; Kelly 1994; Kelly 1992; Lees 1991; Lees 1993; Leotta 1997; Nelson 1993; Pretorius 1994; Raab 1979; Barry 1997]. They are also used in at least one commercial 3DUS system (Tomtec Imaging Systems GmbH, Munich, Germany). In general, EPOM devices consist of one or more small, lightweight sensors (receivers), a single transmitter, and an electronic controller. Both the receiver and transmitter contain three orthogonal coils. The coils in the transmitter are used to generate three orthogonal magnetic fields, which induce small currents in the receiver coils. The relative magnitudes of the receiver coil currents are then used to calculate the 3D position and orientation of the receiver relative to the transmitter [Raab 1979]. The transmitter and receiver are activated electronically by the controller, which also performs signal processing operations necessary for the determination of low noise position/orientation measurements.

Two types of electromagnetic POM device are available commercially: the first generate smoothly changing fields, generated by an alternating current (AC), and are, therefore, known as AC devices. An example is the Polhemus *Fastrak* (Polhemus Inc., Vermont, USA). The second type of EPOM device is the *Bird* range of EPOM devices, manufactured exclusively by the Ascension Technologies Corp. These devices generate a pulsed direct current (DC), and are known as *pulsed-DC devices*.

3D position and orientation measurements rely on sampling the magnetic field generated by the transmitter. Therefore, they are susceptible to errors caused by distortion of the transmitter field. Certain types of metal and extraneous magnetic fields can produce field distortion. Common sources of electromagnetic interference are cathode-ray-tube (CRT) monitors, voltage sources, and high power electrical equipment.

A survey of EPOM devices was carried out by Meyer et al. [1992]. Independent studies suggest that pulsed-DC devices are less susceptible to interference from metals with low permeability than AC-devices [Bull 1998; Birkfellner 1998a]. The reason behind this is that pulsed-DC fields induce eddy currents in the metal only for a short period of time, when the field is changing. Thereafter, the field is in a 'quasi-static' state, and no eddy currents are induced. Hence, measurements made during this period are not subject to field distortions arising from the interaction of the transmitter field and those produced by eddy currents. In contrast, AC devices produce eddy currents throughout the measurement cycle because the transmitter field is constantly changing.

EPOM devices are small, accurate and inexpensive. In principle, there are also no restrictions on movement of the scan-probe, other than ensuring that the receiver remains within a certain distance from the transmitter. This poses few practical limitations, since even the close range systems enable measurements to be performed accurately within approximately 1 m of the transmitter. The specifications and operational details for the pulsed-DC device used in this thesis are given in Chapter 4.

3.3.2.5 Image correlation methods

A novel method of registering acquired 2D images is to correlate adjacent slices based upon statistical properties of the speckle pattern. The method is essentially the

same as that used for extended field-of-view imaging [Kroger 1998], and is currently employed in the Siemens 3D-Scape™ system. The accuracy of this acquisition technique has yet to be fully established in independent tests, but it appears that it is probably unsuitable for applications requiring accurate 3D quantification. A specific criticism is the inability of the technique to account for non-uniform interslice separations and orientations.

3.3.3 Matrix-array transducers

In the same way that advances in transducer technology have led to the widespread used of linear- and phased-array transducers, it is likely that future advances in matrix-array transducers will make volumetric scan-probes a standard feature of commercial ultrasound scanners. Currently, most modern commercial scan-probes contain a one-dimensional (1D) transducer array made up of $1 \times N$ transducer elements, where $N > 100$. Matrix-array transducers are a natural progression of this technology where the transducer is made up from a 2D array of $M \times N$ elements, with $M,N > 100$ [Light 1998; Smith 1992].

A first step towards this goal is offered by so-called 1.5D arrays, which contain up to 8 elements in elevational direction and more than 100 in the lateral direction [Wildes 1997; Zhang 1996]. 1.5D arrays are superior to 1D arrays because they enable focusing in the elevational direction, resulting in a smaller slice thickness than a 1D array. One effect of this is a reduction in artefacts due to the otherwise poor resolution in the elevational direction. Synthetic aperture imaging using a 1.5 or higher dimension array also improves the within-plane resolution, and enables correction for phase aberration across the transducer surface due to small, random changes in the speed of sound in tissue [1991]. The main limitation of 1.5D arrays is that they do not contain enough elements in the elevational direction to steer or focus the beam electronically in that direction; an $M \times N$ array, where $M, N > 100$ and $M \sim N$ is required in order to achieve this. In practice, beam steering is essential to interrogate a volume large enough to be clinically useful.

Presently, the major challenge to successful fabrication of matrix-array transducers is the very large number of electrical connections required to activate a fully sampled array. This problem has been partly overcome by applying techniques de-

veloped by the microelectronics industry. However, a fully-sampled matrix-array with elements separated by one half wavelength (as for 1D transducers) is not yet realisable. Instead, sparse transducer arrays must be used where only a subset of the elements are used in the transmit and receive cycles. The relatively large inter-element separation also imposes a limit on the operating frequency, with frequencies limited to less than 2.5 MHz. The speed of sound in tissue imposes a more fundamental limitation on matrix-arrays: this restricts the rate at which ultrasound data can be collected, and, therefore, limits the size of the volume which can be interrogated within a time period. Thus, there is an inevitable trade-off between volume size and acquisition rate.

Despite the significant technical challenges, matrix-array technology has progressed considerably in the last 10 years. Indeed, it is very likely that they will supersede other methods of 3DUS acquisition discussed here when the performance of matrix-arrays can meet the demands of most clinical applications. To date, matrix-array transducers consisting of up to 256 × 256 elements have been constructed [Light 1998; Arbeille 2000]. One commercial system (the 3-D Volumetrics Model 1) uses a 16 × 16 mm, 256 × 256 element array, and can acquire up to 25 pyramidal volumes per second (each pyramidal volume is formed from 4096 lines and has an angle 64° and a depth of 15 cm).

In principle, matrix-array transducers have considerable advantages over other 3DUS acquisition methods in that they are compact, easy to use, capable of high acquisition rates, and acquire accurately registered and regularly-sampled ultrasound data.

3.4 Segmentation

The purpose of image segmentation is to extract features of interest from an image. In general, such features may be defined by a region (i.e. a set of pixels or voxels (see Section 3.5.2)) or a boundary (a curve or surface that represents an interface). Segmentation is effectively a binary classification of an image into two constituent parts. The motivations for image segmentation are: (i) the amount of data required to represent the image is reduced through elimination of unwanted information; and (ii) the quantification of size and other morphometric parameters is simplified. In the

simplest case, this is because the operation is often confined to a sub-region of the original data set.

There are many standard methods for image segmentation, ranging from delineating regions/boundaries by hand to fully automatic algorithms. Ultrasound images present a significant challenge for automatic segmentation; the combination of high noise content, poor contrast between dissimilar tissues, and multifarious artefacts means that most conventional segmentation algorithms perform badly, even after pre-processing. In particular, B-mode ultrasound images are generally too noisy for traditional edge detection operators, such as the Roberts and Sobel filters [Gonzalez 1992], to be applied successfully. A more robust approach to the edge/surface detection problem is to use deformable models: in this method, an initially chosen contour or surface deforms subject to global and local image-determined forces until equilibrium is reached. The two main types of deformable model are snakes, which are 2D active contours, and balloons, which are 3D active surfaces [Eviatar 1996; Kass 1988; Mao 2000].

Alternative approaches include: region-based methods, for example threshold filtering, texture analysis, region growing, or isosurface extraction [Gonzalez 1992]; region competition, which is a unification of edge- and region-based techniques proposed by [Zhu 1996]; and the semi-automatic live-wire and live-lane methods, which have been proposed as a good compromise between user-interactivity and fully automatic approaches for noisy images [Falcao 1998].

The main drawbacks of many automatic and semi-automatic segmentation methods are that they can be highly dependent on initial parameters, and they are often unable to handle complicated shapes. It is also usually necessary to pre-process images before segmentation algorithms can be successfully applied. An interesting algorithm for enhancing edges in B-mode images uses *stick operators* and has been described by Czerwinski et al. [1993].

The 3D segmentation problem may be treated in two ways: either by applying operations which work directly on 3D data sets, or reducing the problem to a series of 2D segmentations. The latter strategy has the advantage that there are many more

segmentation algorithms for 2D images, although many of these can, in principle, be extended to 3D.

3.5 3D reconstruction

3D image reconstruction is the process of converting input data – i.e. back-scattered ultrasound intensities and associated 3D position and orientation data – into a form which is both a convenient and a complete representation of the imaged volume. Once the relevant parameters have been set, reconstruction is typically performed without any further user-interaction.

There are two general approaches to reconstruction: the first is to reconstruct surfaces from input data that correspond to anatomical surfaces; for example, the outer surface of an organ, or the wall of a blood vessel. The second method generates a discretized volumetric representation whereby the volume is divided into small cuboid volume elements called voxels. Each voxel is assigned a value corresponding to its ultrasonic grey-level intensity or colour in an analogous way to pixels in a 2D digital image.

The complexity of a 3D reconstruction algorithm depends on the form of the input data, the method used for reconstruction, and the particular information required from the final analysis. The latter is usually determined by clinical questions, such as: "Is an anatomical abnormality present?" or "What is the exact size?" The efficiency of a reconstruction algorithm is also an important consideration since 3D data sets are invariably very large.

As input data is sampled, interpolation is normally required to maintain a sufficiently high voxel resolution in the reconstruction. This is particularly the case when an external POM device is used, since, typically, the inter-slice separation will be significantly greater than the within-slice pixel dimensions. Interpolation methods are considered in Section 3.5.2.

3.5.1 Surface reconstruction

In general, a 3D surface may be reconstructed from a set of sampled points which lie on that surface. The points may be sampled randomly or have some order; for exam-

ple, they may form a series of planar, cross-sectional contours through the surface. This is frequently the case when ultrasound data are acquired using conventional 2D scan-probes.

The accuracy of the surface reconstruction is dependent partly on the surface point sampling rate: the higher the sampling rate, the finer the detail of the surface point representation. Thus, the sampling rate must be sufficiently high to preserve features down to a certain scale. An anatomically faithful surface representation also depends on the accurate determination of interfaces. As discussed in Section 2.3, these may be found by either processing an acquired set of 2D images, or a reconstructed 3D image. In practice, noise is present in the sampling process, and this may give rise to artefactual surface detail. Such surface irregularities may be removed by filtering. Interpolation will also be required in order to 'fill in the gaps' in the acquired data set.

The standard method for reconstructing a continuous surface from a set of surface points is to generate a surface mesh. Surface meshes may be thought of as being made up from number of small, flat surface primitives, called facets. A facet is defined by a set of vertices, or, alternatively, by a set of edges and has a regular shape (e.g. a triangle). Each facet interconnects with adjacent facets. One popular technique for surface mesh generation is called triangulation. As the name suggests, a surface is created from triangular facets, as depicted in Figure 3.4.

Meshes provide an efficient geometric representation of a physical surface, with the implication that they may be displayed rapidly using standard surface rendering techniques introduced in Section 3.6.1. Some meshing algorithms accept a randomly sampled set of input surface points. A simpler problem, which is relevant to 3DUS imaging, is the reconstruction of surfaces from a series of planar, cross-sectional contours extracted from a set of image slices. This is discussed in depth in Chapter 6.

3.5.2 Voxel-based reconstruction

Voxel-based reconstruction methods map ultrasound echo data into a regularly sampled 3D array of voxels, sometimes referred to as a *cuberille*. This is illustrated in Figure 3.5 below.



Figure 3.4 An example of a triangulated surface mesh which forms the shape of a heart.

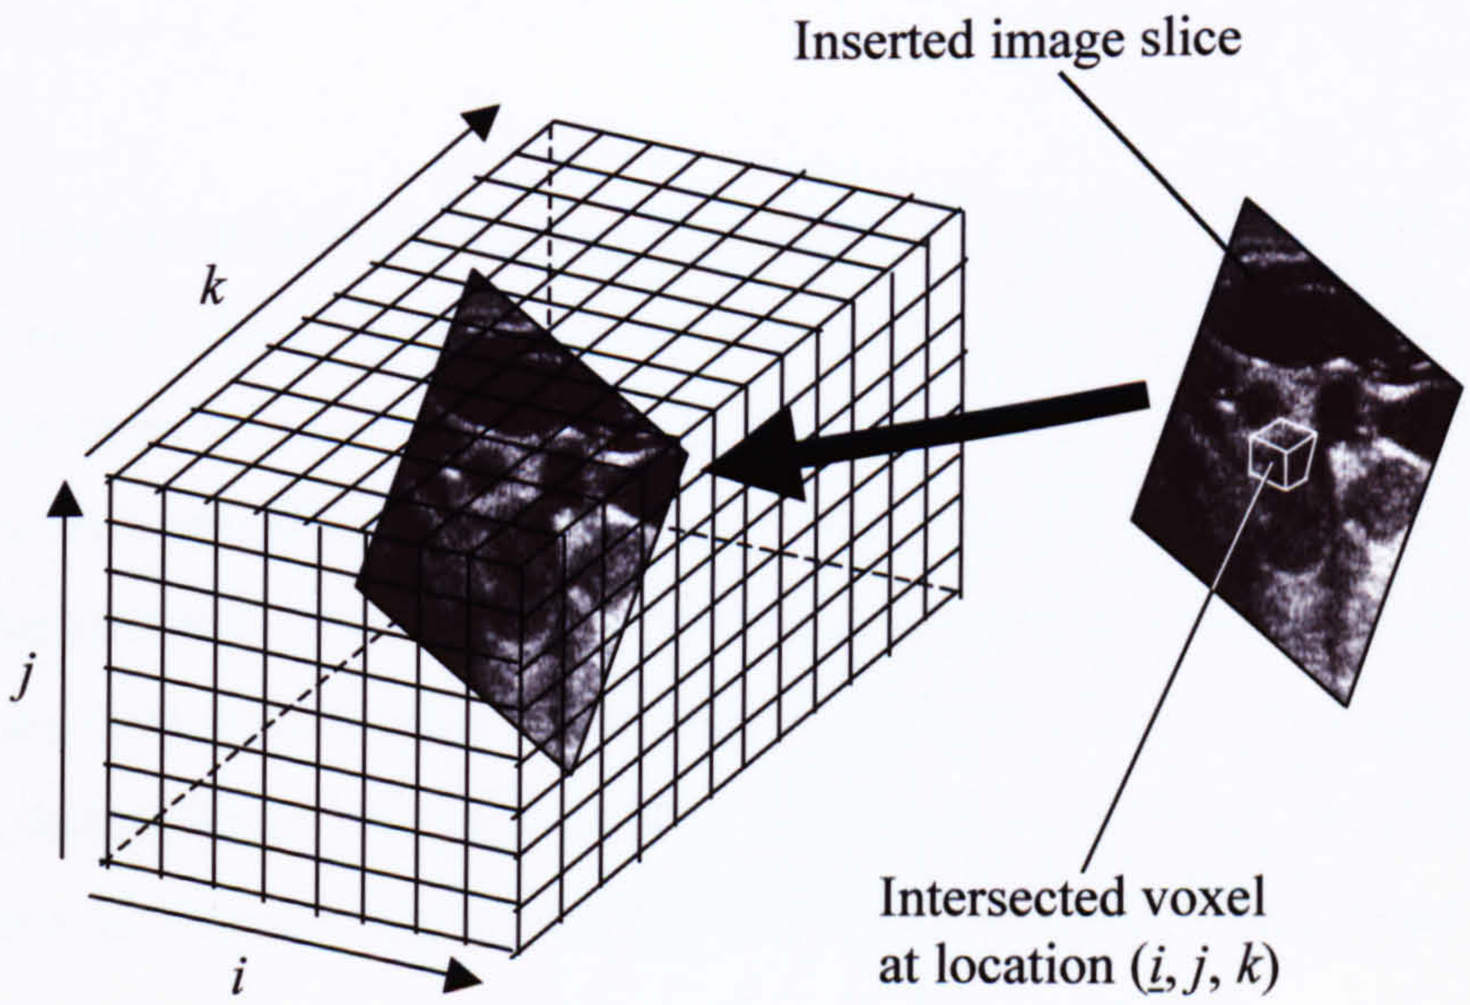


Figure 3.5 The mapping of a 2D ultrasound image into a 3D voxel array. Each voxel is identified by the indices (i, j, k) and assigned a value that represents its grey-level intensity or colour. The value of each sample voxel is found by evaluating a function of the grey-level/colour values of pixels in the inserted image slice which lie within the voxel intersected by the image (indicated on the right).

Currently, voxel-based reconstruction is the most popular representation for 3D ultrasound data. The main reason for this is that it can be performed automatically following acquisition, with little or no intermediate user-interaction. As shown in Figure 3.5, each pixel value is used to determine the value of the voxel with indices (i, j, k) in which it lies when registered within the voxel array. Using this scheme, if the dimensions of individual voxels are greater than those of pixels in the input image, more than one input pixel may lie within a particular voxel. This is a common scenario for three reasons: firstly, memory limitations dictate that voxels cannot be made arbitrarily small because the number of voxels in an array would become unmanageably large, leading to unacceptably slow reconstruction and post-processing times. Secondly, the resolution in elevational direction of 2D ultrasound images is much lower than the within-image directions. This implies that a voxel dimension larger than the within-image dimensions of pixels is appropriate for the 3D reconstruction. Thirdly, volumes are often sparsely sampled: for example, the inter-slice separation is commonly much larger than the image pixel dimensions. Hence, increasing the voxel size reduces the amount of interpolation required to fill empty voxels.

Voxel interpolation is essential, particularly for irregularly-sampled volumes acquired using a freehand system. Constrained acquisition ensures that the spatial relationship between acquired images is regular, and interpolation then becomes a straightforward operation. For instance, a parallel set of image slices may be interpolated in the direction of acquisition – i.e. perpendicular to the image planes. The operation then reduces to a series of 1D interpolations. For freehand acquired images, however, interpolation is not trivial since images are not parallel, and, in the extreme case when a volume is imaged from multiple directions, they will intersect one another.

A generalised method for interpolating freely-acquired, non-parallel input slices is to adopt weighted nearest-neighbour interpolation. In this scheme, the value assigned to an interpolated voxel is a normalised, weighted summation of the intensities of the sample voxels within a so-called volume of influence surrounding the voxel. This is expressed mathematically as:

$$I_{\text{int}} = \frac{\sum_{V} w.I_{\text{samp}}}{\sum_{V} w} \quad , \tag{3.1}$$

where $I_{\text{int}}(i, j, k)$ is the value assigned to an interpolated voxel; $I_{\text{samp}}(i, j, k)$ is the value of a sample voxel; $w(i, j, k)$ is a weighting function (dependent on the location of I_{samp}); and V is the volume of influence – i.e. the set of sample voxels surrounding I_{int} whose intensities are considered in the interpolation.

A special case of weighted nearest-neighbour interpolation has been proposed by Carr [1996] and is depicted in Figure 3.6. The algorithm is intended for use with 'swept volumes', i.e. acquired 2D image sets in which consecutive slices are semi-parallel. For an empty voxel lying between the n^{th} and $(n+1)^{\text{th}}$ slices, the centre of the voxel is projected onto each slice. The pixels, p_n and p_{n+1}, coinciding with the within-plane projection points are then used to determine the value of the voxel. The weighting function is inversely proportional to the distances between each of these pixels and the interpolated voxel (d_n and d_{n+1} in Figure 3.6) .

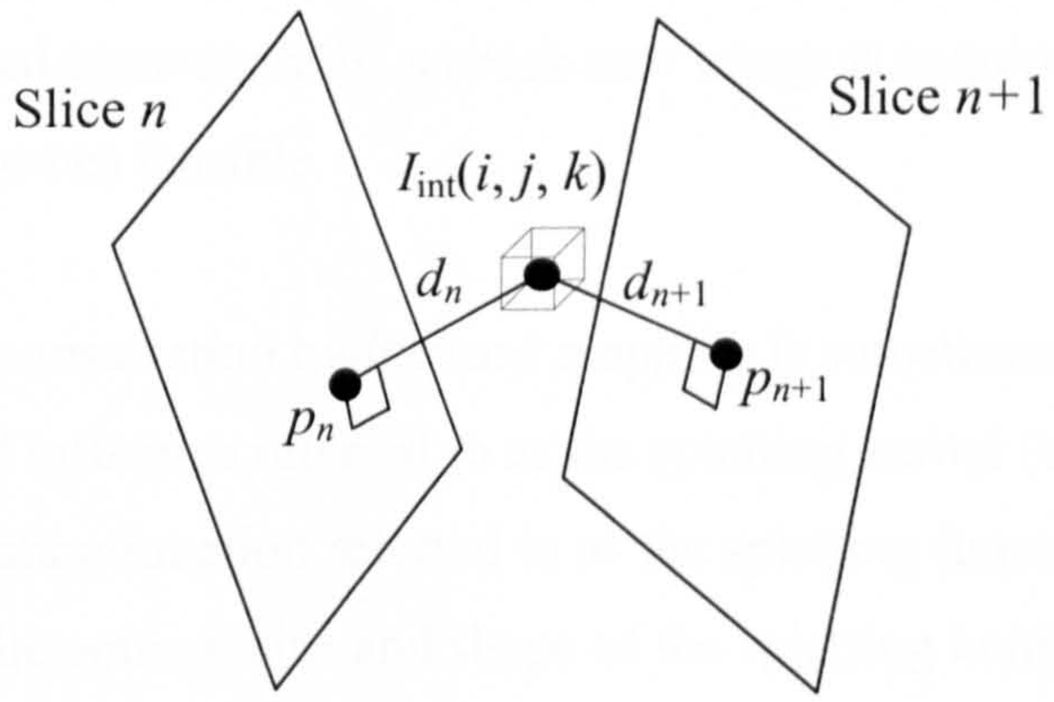


Figure 3.6 Weighted nearest-neighbour interpolation for a freely-acquired swept volume.

If λ_n and λ_{n+1} are the intensities of pixels p_n and p_{n+1}, respectively, then Equation (3.1) becomes:

$$I_{\text{int}} = \begin{cases} f & \text{if } d_n, d_{n+1} \geq \delta \\ p_n & \text{if } d_n < \delta \\ p_{n+1} & \text{if } d_{n+1} < \delta \end{cases} \tag{3.2}$$

where

$$f = \frac{\frac{\lambda_n}{d_n} + \frac{\lambda_{n+1}}{d_{n+1}}}{\frac{1}{d_n} + \frac{1}{d_{n+1}}} = \frac{d_{n+1}\lambda_n + d_n\lambda_{n+1}}{d_{n+1} + d_n}. \tag{3.3}$$

The value of δ is small and sets a distance threshold below which a pixel value dominates. The algorithm automatically handles intersecting slices and only the previous two acquired slices are required to incrementally reconstruct the volume. Hence, the algorithm is suitable for real-time implementation.

Voxel reconstruction algorithms can be implemented using a direct approach; that is, the interpolated value of each non-sample voxel is computed from the values of all sample voxels. This method is called backward mapping. A much more efficient implementation is forward mapping where the values of each of the (non-sample) voxels within the volume of influence surrounding a sample voxel are adjusted, subject to some weighting function, as each sample voxel is filled. The equivalence of forward and backward mapping implementations can be shown [Pitt 1998]. In addition to reduced computation time, an advantage of forward mapping is that a volume can be reconstructed incrementally, as each new image is acquired, and so real-time reconstruction becomes feasible.

Voxel-based reconstruction by forward mapping is sometimes called splatting, with the volume of influence referred to as the splatting kernel (or convolution kernel), and the weighting function referred to as the splatting function. Largely unresolved issues are the optimal size and shape of the splatting kernel, and the optimal form of the splatting function. Several different combinations have been reported in the literature. For example, Barry et al. [1997] use a spherical kernel and a inverse square splatting function; Meairs et al. use an ellipsoidal kernel and a Gaussian splatting function; and Ohbuchi and Fuchs [1990] use an octahedral kernel based on the approximate point-spread function of the ultrasound beam in the elevational direction. Recently, Prager et al. proposed a 'thick slice' method, equivalent to using a uniform splatting function and rectangular splatting kernel with a length equal to the assumed thickness of the slice. This approach obviates the need for voxel reconstruction altogether prior to rendering.

Most of the voxel-based reconstruction methods described in the literature were developed primarily for reconstructing B-mode ultrasound date, but the same techniques can be applied to the reconstruction of colour or power Doppler data [Guo 1996; Ritchie 1996; Picot 1993].

3.5.3 3D compounding

A straightforward approach to reducing the amount of interpolation required to reconstruct a volume fully is to increase the amount of sample data. In the context of 3DUS imaging, this is achieved by performing multiple scans of the same volume, preferably using a large number of arbitrarily oriented views. This technique is known as 3D compounding [Allott 1999; Rohling 1998]. The principal attractions of 3D compounding are that the signal-to-noise ratio (SNR) is increased by reducing speckle [Rohling 1998], and that, in some cases, the visualisation of structures is enhanced.

3D compounding is a logical extension of single-sweep acquisition, but its success relies on accurate co-registration of acquired images (or volumes). Otherwise, interfaces in the image may become blurred. Methods for compensating for inadequate co-registration have been described by Rohling et al. [1998] and Meyer et al. [1999]. However, these methods are dependent on unambiguous identification of anatomical landmarks.

A fundamental problem encountered in 3D compounding is the appearance of image artefacts in the reconstructed volume which arise from the presence, or absence, of particular structures in different views. In other words, when the visual content of images acquired from a broad range of views is inconsistent. It is also important to note that the averaging effect of compounding is just as likely to degrade strongly visible features as enhance poorly visible ones. In spite of this, compounding can be effective in enhancing features when the quality of the original acquired images is poor. A compromise aimed at overcoming some of the problems of 3D compounding arising from inadequate co-registration, that retains the benefits, is to perform a single swept-volume acquisition using compound B-mode imaging (see Section 1.3.7.4).

3.6 3D visualisation

3.6.1 Surface rendering

Surface rendering is the display of a 3D scene, specified by one or more 3D surface meshes, as a 2D image. The rendering operation can be repeated using a new viewpoint, and this enables a volume to be rotated on a graphics workstation, thus conferring a sense of three-dimensionality in the displayed image. As explained in Section 3.5.1, a surface mesh is made up from a set of small, interconnecting facets. Surface properties are defined for each facet, and the whole surface may be rendered using standard algorithms that support a number illumination and shading schemes. A number of surface properties govern the overall appearance of a surface, including colour, reflection coefficients, opacity, and texture. Reflection coefficients determine how reflective a surface appears when lit; for example, to make a surface appear shiny. Opacity is a measure of how opaque a surface is: a high opacity (equivalent to a low transparency) allows little or none of the surfaces behind to be seen through the object. Conversely, a low opacity (high transparency) allows the viewer to see through the surface.

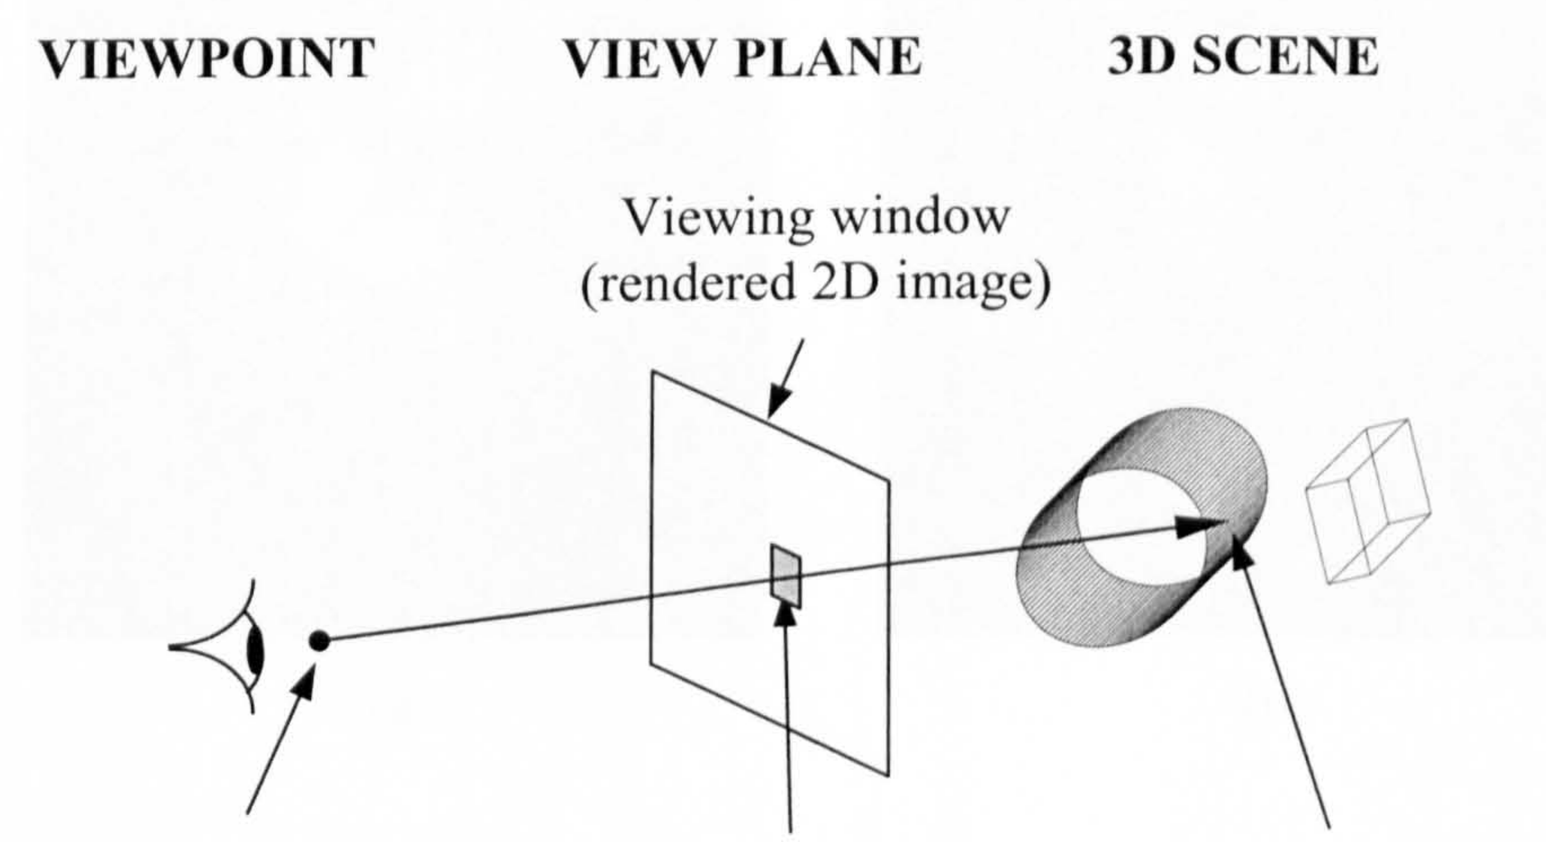
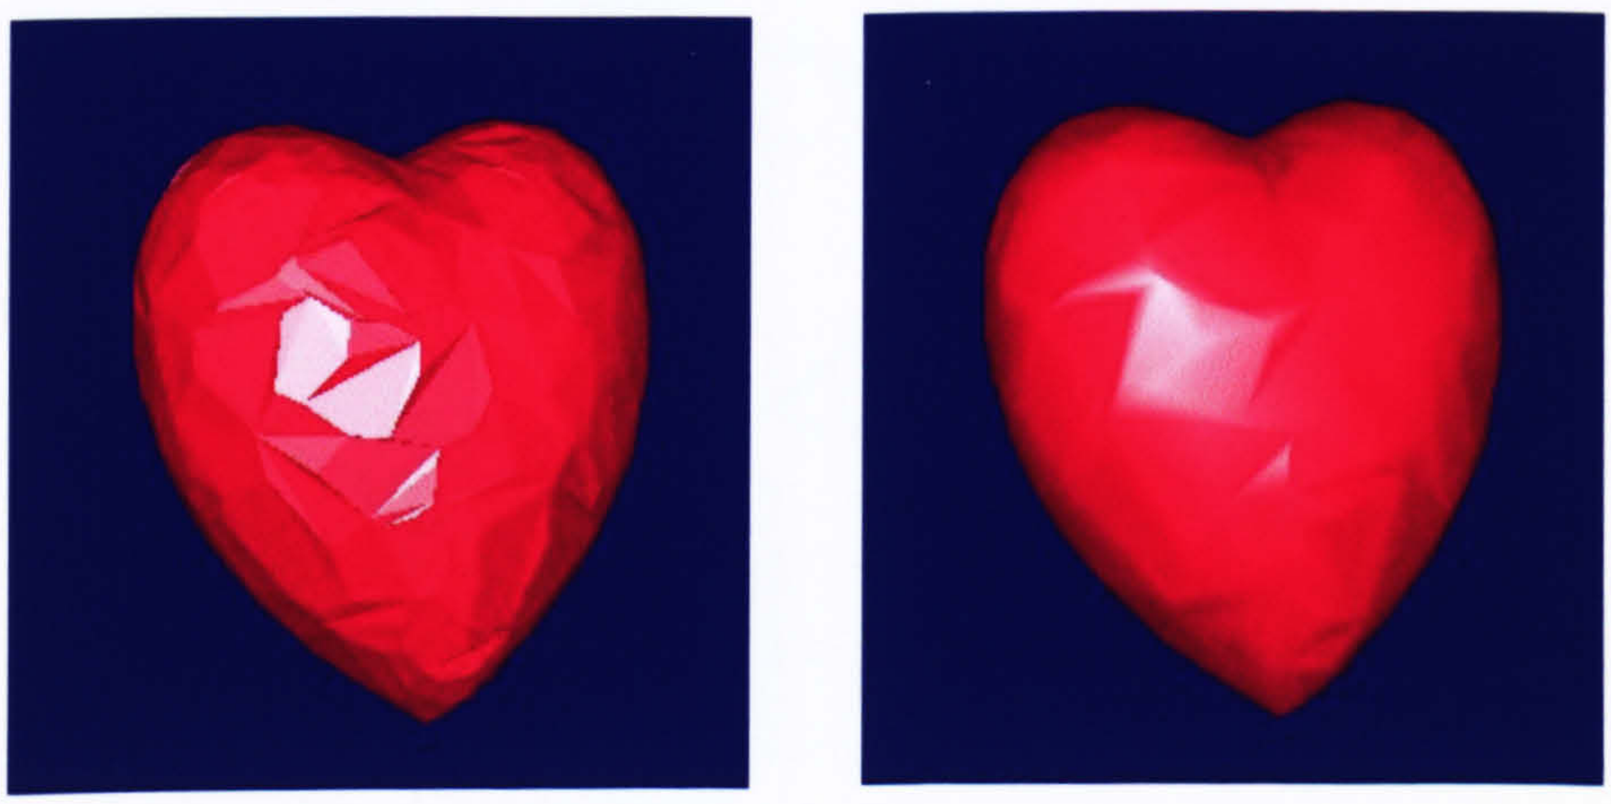


Figure 3.7 Ray-tracing for 3D surface rendering.

A widely used surface rendering technique is ray-tracing [Foley 1995]. Figure 3.7 illustrates the principle underlying the ray-tracing algorithm. A linear ray is cast from the centre of projection, coincident with the position of a viewer's eye, through the viewing window and into the 3D scene. The window directly corresponds to the dis-

play area of the rendering. The pixel in the window which is intersected by the ray is coloured according to the surface properties at the point where the ray first intersects a surface in the scene. In practice, much of the algorithm is dedicated to finding the ray intersections with the surfaces in the scene. Consequently, many methods of making this process more efficient have been proposed [Foley 1995].

In order to render surfaces realistically, the surface normal vector and surface properties must be found at the point where the surface is intersected by a ray. If each facet in a polygon mesh is shaded independently using a flat shading scheme then individual facets become visible in the rendering, resulting in a 'blocky', unrealistic-looking surface (see Figure 3.8(a)). Increasing the number of facets used to approximate surfaces does not solve this problem because of the Mach Band Effect which accentuates surface discontinuities. Instead, Gouraud or Phong shading is used to render a smoothly shaded surface [Foley 1995]. The improvement gained from using Gouraud shading is illustrated in Figure 3.8(b).



(a) (b)

Figure 3.8 The heart mesh shown in Figure 3.4 surface rendered using (a) flat shading, and (b) Gouraud shading models.

Surface rendering is significantly faster than volume rendering, introduced in the next section. This is because, in general, fewer data points are required in a surface representation. Rapid rendering enables 3D images to be displayed in real-time. Therefore, surface rendered images can be updated quickly in response to user-interactive tasks, such as, rotation and magnification. This feature enhances 3D visualization significantly. Surface rendering does, however, depend on the availability

of surfaces extracted from the original volumetric data. As discussed in Section 3.3, invariably this is a difficult process because ultrasound data are noisy. Nevertheless, if a surface representation can be created for structures of interest, surface rendering provides a very efficient and easily interpretable method of visualising anatomical data in 3D.

3.6.2 Volume rendering

Volume rendering involves projecting a 3D voxel array directly onto a 2D display, without first generating a surface representation [Levoy 1990; Sabella 1988; Sarti 1993; Watt 1992; Drebin 1988; Sakas 1995; Elvins 1992; Kajiya 1984]. In common with surface rendering, the concept of ray-tracing is used in the vast majority of volume rendering algorithms [Elvins 1992]. The method of volume rendering by ray-tracing is illustrated in Figure 3.9.

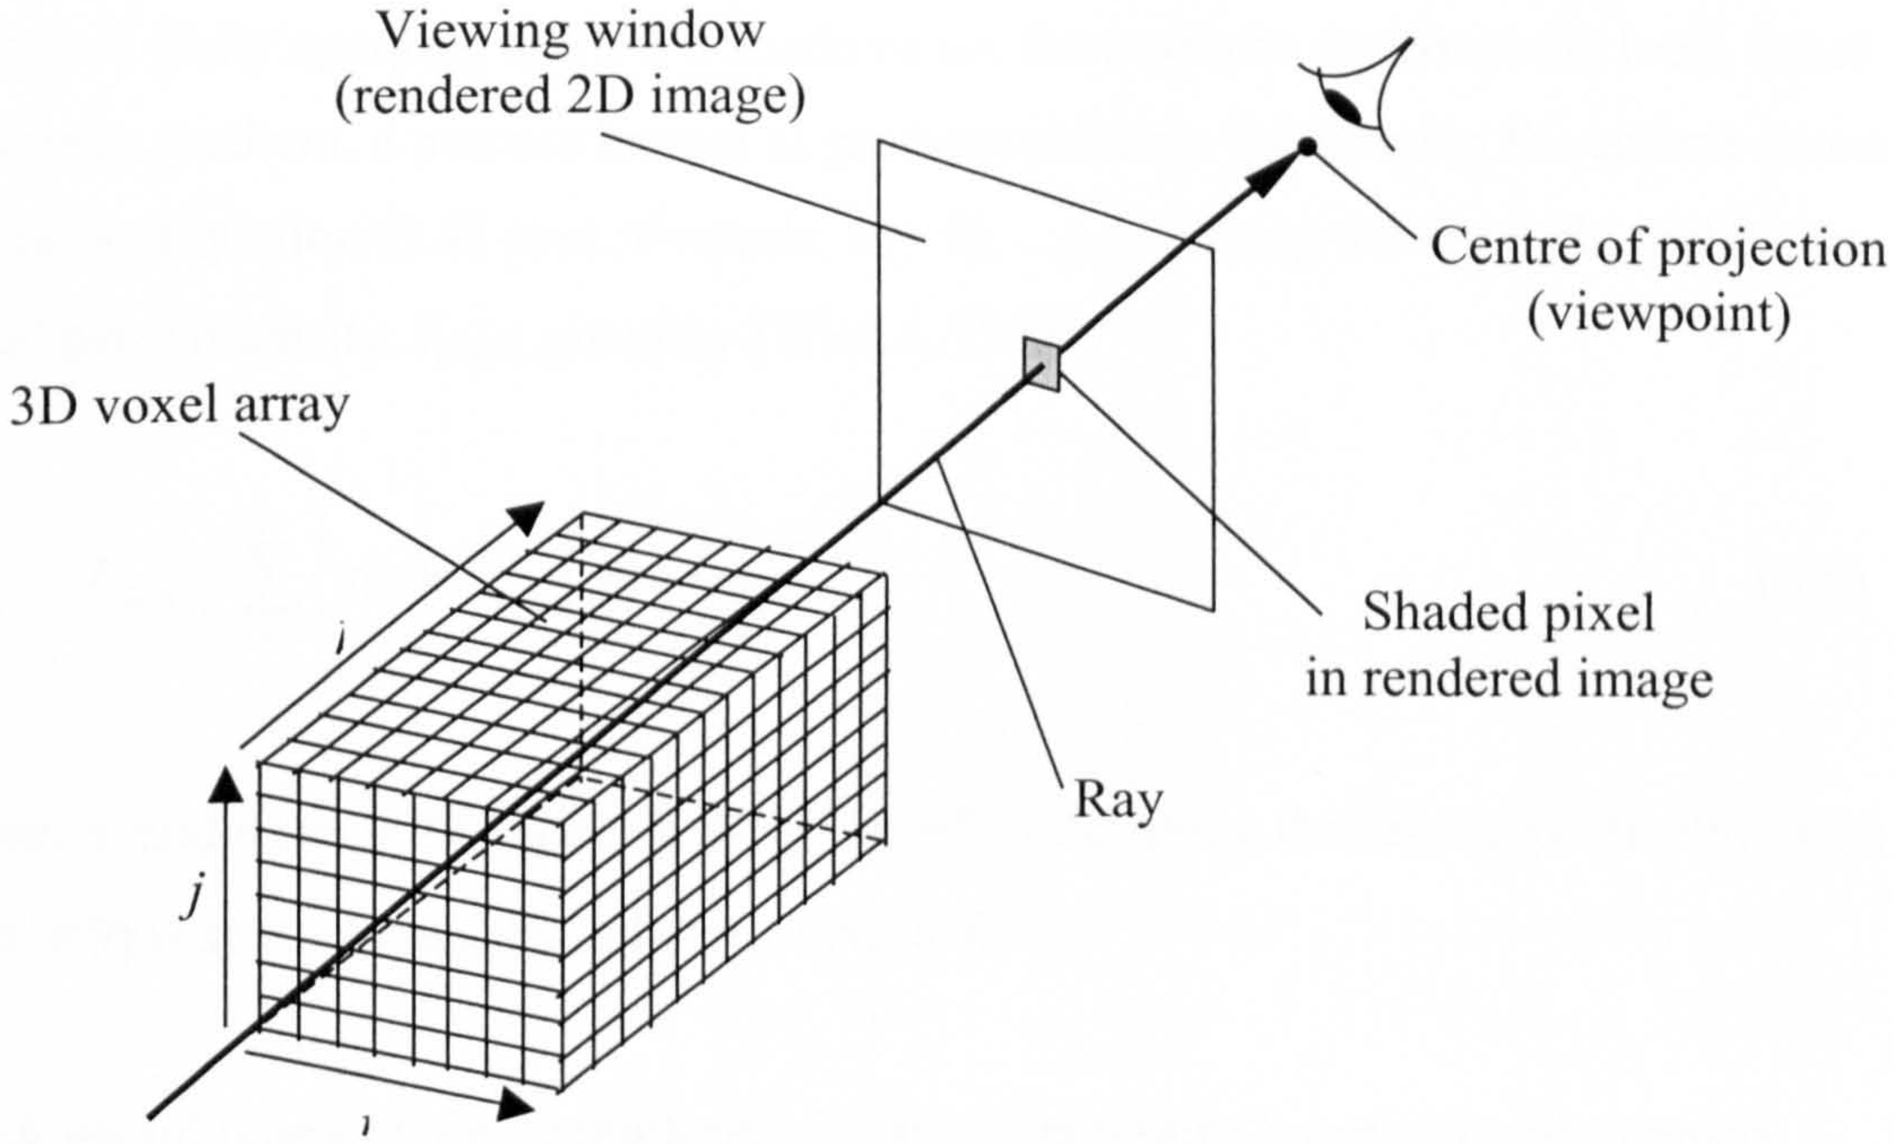


Figure 3.9 Ray-tracing for 3D volume rendering. A ray is projected through the 3D voxels array. The intensity of the pixel in the rendered image at the intersection with a ray is determined by a function of the grey-level/colour intensities and opacities of all the voxels intersected by the ray.

With reference to Figure 3.9, each ray passes through the volume and a pixel in the viewing window to the centre of projection (coincident with the viewpoint). As

the ray passes through the volume, it may be thought of as carrying a value – the ray intensity – which changes as it progresses. The ray intensity at any point is dependent on the grey-level (or colour) intensities and opacities of all the voxels intersected along its path up to that point. The value of the ray intensity when it emerges from the volume determines the grey-level (colour) intensity of the pixel intersected by the ray in the viewing window. One ray is cast for each pixel within the viewing window so that the complete voxel array is rendered.

The function which is used to compute the ray intensity is sometimes called a blending function and has the general form [Watt A 1992]:

$$I_{\text{out}} = I_{\text{in}}(1-\alpha)+s\alpha, \tag{3.4}$$

where I_{out} is the ray intensity as it emerges from a voxel, I_{in} is the ray intensity entering the voxel, α is the opacity of the voxel with a value in the range of 0 (transparent) to 1 (fully opaque), and s is a shade value, for example, based on the local voxel intensity gradient, a process known as gradient shading. Summating the contributions given by Equation (3.4) over N voxels, V_1, V_2, …, V_N, along the ray path, yields the final pixel intensity, I_{pixel}, given by [Watt A 1992]:

$$I_{\text{pixel}} = \sum_{n=0}^{N}\left\{s(n)\alpha(n)\prod_{m=n+1}^{N}[1-\alpha(m)]\right\}. \tag{3.5}$$

Here, m and n ($m, n > 0$) specify the m^{th} or n^{th} voxel along the ray, respectively. For $n = 0$, $s(0)$ is set to a background value and $\alpha(0) = 1$.

A useful type of volume rendering is called maximum (or minimum) intensity projection (MIP). In this scheme, only the maximum (minimum) intensity of the voxel lying on the ray path is used to calculate pixel intensity. Hence,

$$I_{\text{pixel}} = \max\{V_1, V_2, ..., V_N\} \tag{3.6}$$

for maximum intensity projection, or

$$I_{\text{pixel}} = \min\{V_1, V_2, ..., V_N\} \tag{3.7}$$

for minimum intensity projection. MIP methods have the advantage of being straight-forward to implement and produce good results for many applications [Pretorius 1994; Nelson 1992].

A number of additional techniques are used in practice to enhance the visualization of volume rendered images. These are necessary because the basic methods of ray-tracing do not account for more subtle features, such as depth perspective, shadows, and reflections, which enhance the perception of 3D [Nelson 1993]. One useful technique for improving the perception of depth in an image is to employ depth shading, implemented by weighting voxel intensities according to their distance from the centre of projection.

One of the drawbacks of volume rendering algorithms is that they are computationally intensive. This can cause difficulties if the volume needs to be visually analysed in real-time. Often there is a trade-off between image detail (governed by the voxel size) and the ease with which volume rendered images can be manipulated. This situation is improving, however, as the speed of graphics workstations increases, and the increasing use of parallel computer architectures [Edwards 1998][†].

A specific problem encountered with volume rendered 3DUS images is difficulty in differentiating structures due to artefacts, speckle, and poor tissue contrast. Ironically, visual interpretation can be hindered by the very large information content of these images. Modifying the opacity values or the blending function may help, but this can become a rather tedious and subjective process. Another approach is to segment volumetric data either before rendering, or as part of the rendering process [Steen 1994]. In this way, only features of interest are retained in the rendered image.

Volume rendering is widely used to display 3DUS data, and several software packages are commercially available to perform basic rendering and manipulation tasks. An example of a volume rendered 3DUS image is shown in Figure 3.10. In this example, image reslicing has been used to expose structures of interest.

[†] Ray tracing lends itself to parallelization since each ray can be processed independently.

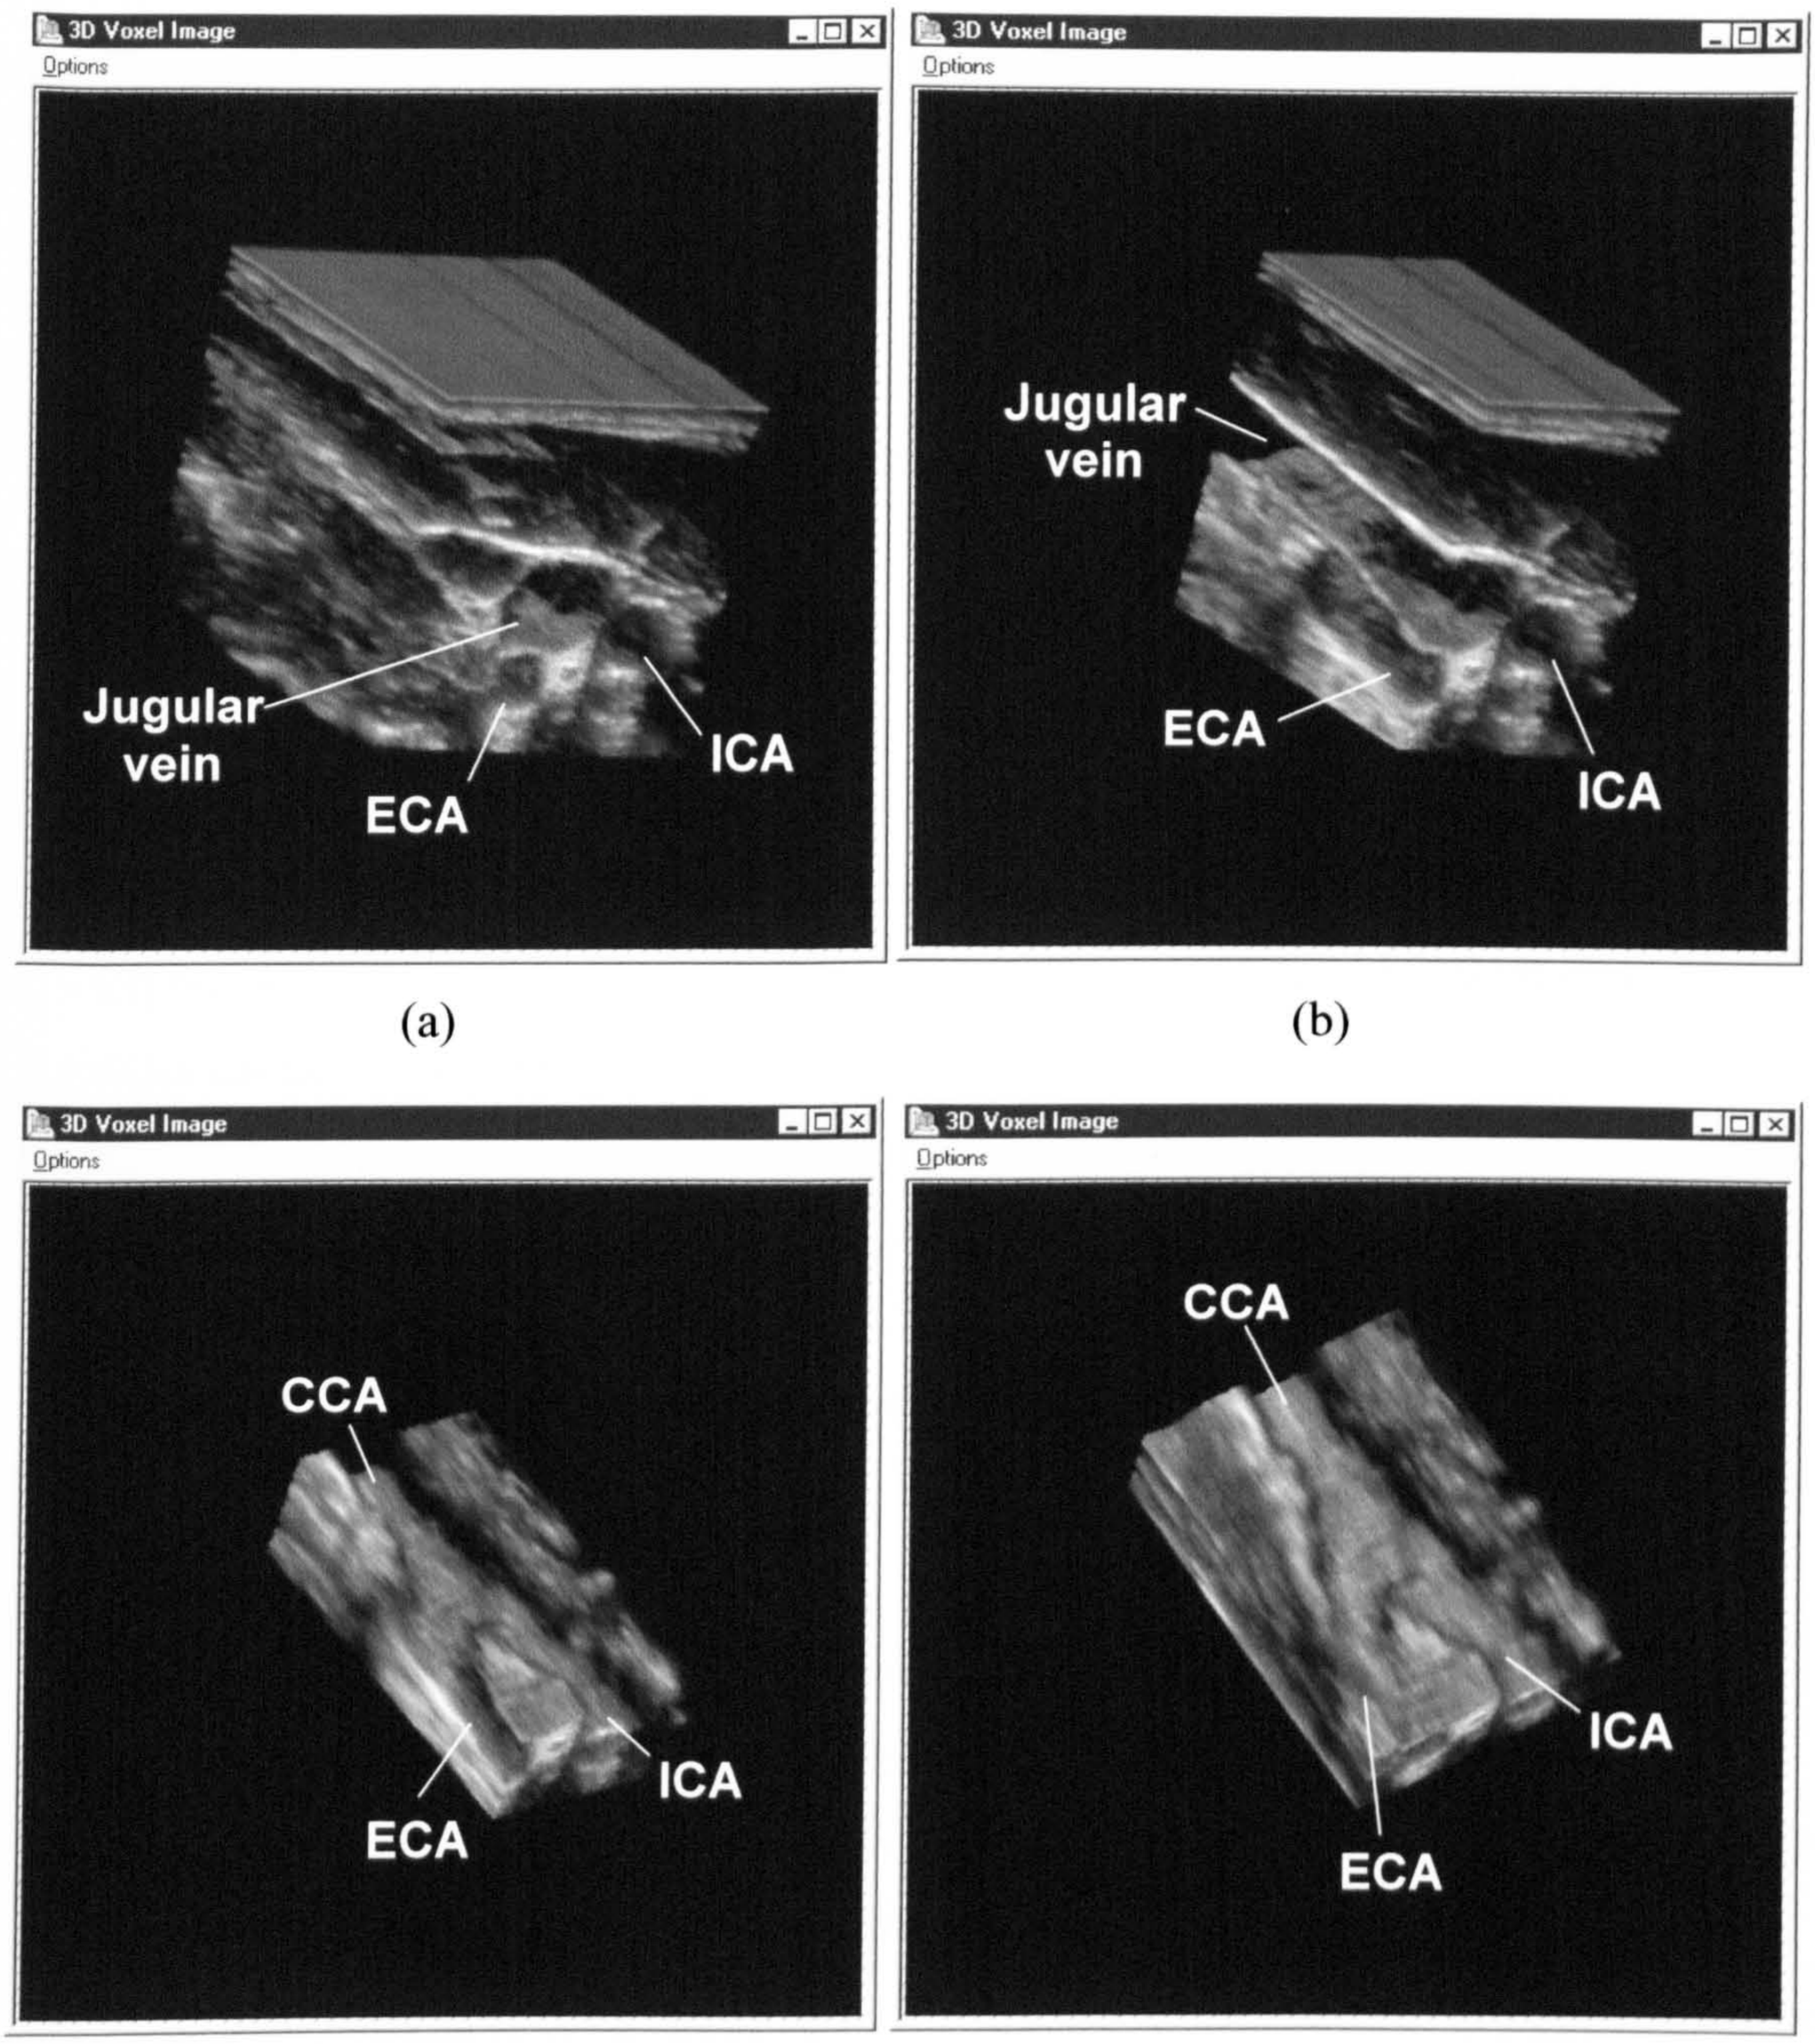


Figure 3.10 3D ultrasound images of a carotid bifurcation rendered using VoxBlast (Fairfield Imaging Ltd, Tunbridge Wells, UK): (a) the complete volume; (b) the volume after slicing in a sagittal plane, exposing the jugular vein and the ECA; (c) and (d) the volume following further slicing along a coronal plane, exposing the carotid bifurcation. All images were produced by the author.

3.6.3 Interactive visualisation

Interactive visualisation of images is an extremely important aspect of 3D medical imaging. The process requires a user-interface between the analyser, who issues commands, and a computer, which executes the commands and generates the appropriate image. A simple user interface may take the form of task-specific controls ('buttons') on a display. Typically, these are selected using a mouse and pointer, or by touching the screen. Examples of common controls are: rotation (or, more correctly, changing the rendering viewpoint), slicing, opacity changes, and thresholding to remove portions of a particular grey-level or colour. This kind of interaction greatly enhances the perception of 3D data and is a critical factor in the clinical acceptance of 3DUS imaging. One of the reasons why 3DUS is still slow to be accepted in routine clinical practice is the current lack of effective tools for interactive visualisation of scanned volumes.

There are several essential attributes for an effective user-interface for interactive visualisation of 3DUS data. These are listed below (adapted from Nelson and Elvins [1993], and Nelson and Pretorius [1998]):

(i) An intuitive command process

The way in which commands are issued by the user should be intuitive and simple to use. More advanced user-interfaces have been explored with which adopt virtual reality techniques to create a more natural interaction [Fuchs 1996]. Such methods rely on a sophisticated sensory feedback mechanisms so that information is relayed back to the user, thus increasing the level of interaction. Virtual surgery is a good example [Taylor 1996]. Sensors on a glove convert hand movements into electrical signals. Information is fed back to the user via a 3D display of virtual instruments and organs, and actuators in the glove provide resistive forces. One can readily conceive a similar system which allows an 3D anatomical image to be 'explored' in a virtual reality environment.

(ii) Ease of use

The user-interface should be easy to use with minimal training. In common with most commercial software, such as word processors and spreadsheets, there is an inevitable compromise between functionality – i.e. how many functions (or features)

are available – and 'usability': it is easy to overwhelm the user with a myriad of functions, many of which they will seldom use. One way of resolving this problem is to hide functions which are less commonly used, allowing immediate access only to a pre-configured set of commands. Ease of use can also be a significant factor in the reproducibility and repeatability of image analysis because a complex sequence of commands is less likely to be repeated exactly by a different user, or by the same user on a different occasion. Therefore, there is sometimes good reason for limiting the number of commands available.

(iii) Provision of clinically useful tools

A key role of a visualisation system is to permit the user to extract clinically useful information. This is, to a large extent, dependent on all of the earlier stages in the 3D imaging process, but it is essential that a clinician is able to *retrieve* useful information from an image via the user-interface. Tools which enable this require varying degrees of user-interaction, and many may require expert knowledge – for example, a familiarity with a specific anatomical region or disease – to be used effectively.

Recently, a general purpose tool for interactive visualisation of 3DUS data was described by Robert et al. [2000]. The method creates a surface rendered display from three orthogonal boundary contours defined by the user. The advantage of this segmentation approach is that it can be quickly applied, but only simple shapes are can be reconstructed accurately.

3.7 3D quantification

Quantitative measurements, derived from 3D images, fall into one of two categories: physical or functional measurements. Physical measurements include measures of size, such as volume, and morphological measures characterised by shape, greyscale intensity, colour, or texture. Functional measurements, on the other hand, are based on physical measurements but relate directly to the functionality of an organ, for example, perfusion or cardiac performance.

3.8 3D image artifacts

Because 3DUS images are generated from data acquired using the same basic princi-

ples as conventional 2D ultrasound, the same artefacts as those discussed in Section 1.3.8 can arise. As indicated in Section 3.5.3, some of these artefacts, such as shadowing, enhancement, and speckle, may be less pronounced if 3D compounding is employed. There are, however, several artefacts which are specific to 3DUS images [Pretorius 2000]. These vary depending on the type of acquisition system and post-processing algorithms used. Misregistration of acquired 2D images is a common cause of artefact, particularly when an external device is used to locate acquired images in 3D space. Misregistration is manifest as the misplacement of structures or the blurring of tissue interfaces. Motion artefacts are also common due to the need for much longer acquisition times than for 2D imaging.

It is important to remember that, following acquisition, artefacts may arise from any of the subsequent stages. Probably the most frequent source of artefacts is the interpolation of poorly registered or 'inconsistent' input data during the reconstruction stage. Inconsistent data – i.e. images of the same structure with significantly different visual characteristics – is generally caused by the dependency of ultrasound on the orientation of the scan-view, in addition to artefacts, such as shadowing.

3.9 4D ultrasound imaging

With many of the technical challenges associated with 3DUS acquisition and display overcome, there has been growing interest in dynamic 3D, or four-dimensional ultrasound (4DUS) for a variety of applications [Deng 1996; Meairs 2000; Seward 1995; Raichlen 1986; Wollschlager 1989; Roelandt 1994; Delabays 1995]. In general, because near real-time volumetric acquisition is difficult to achieve, one solution has been to gate acquisition, and generate a series of 3D reconstructions for different time-steps. This method is both cumbersome and slow, however. 4DUS imaging using an electromechanically-driven transducer has been demonstrated on some commercial systems. For example, the Kretztechnik Voluson 730 (Kretztechnik AG, Zipf, Austria) allows up to 16 volumes per second to be acquired. As stated in Section 3.3.3, real-time matrix-array transducers have also been used for cardiac applications.

Chapter 4

An electromagnetic 3D ultrasound acquisition system

This chapter describes the development of a freehand 3DUS acquisition system based on an electromagnetic tracking device. The system was designed to provide a low-cost solution for high quality, digital 3DUS data acquisition using a commercial ultrasound scanner used for routine clinical vascular investigations. Details are given on the architecture, components, and operation of the system, and associated software. This is followed by a description of the system calibration, including a comparison with the results of calibration methods used by other researchers. The chapter concludes with a brief discussion of the advantages and limitations of the system.

4.1 System architecture

A schematic of the electromagnetic 3DUS acquisition system developed for this thesis is shown in Figure 4.1. The system is a freehand acquisition system, but differs from other systems in that an external image capture device is not used. Instead, the frame buffer – or 'cineloop' – on the ultrasound scanner stores digital, pre-scan-converted ultrasound image data. Each component of the system is considered in the sections below.

4.1.1 Ultrasound scanner

The ultrasound scanner used in the 3DUS system was the ATL HDI-5000 (ATL-Phillips Ltd, Bothell, MA, USA). The manufacturer provided software called *HDI-Lab* to enable pre-scan-converted image data stored in the scanner's cineloop to be downloaded to a personal computer (PC) via a network connection. The software also allows downloaded images to be viewed and manipulated off-line. Two high frequency scan-probes – the L10-5 and L12-5 (5-10 MHz and 5-12 MHz broadband linear-array transducers, respectively) – were used to acquire high resolution B-mode and power Doppler images of the carotid arteries.

4.1.2 Electromagnetic tracking device

The EPOM device used in this study was the Ascension *pcBIRD* (Ascension Technology Inc, Vermont, USA). The receiver has dimensions 25.4 × 25.4 × 20.3 mm and the cubic transmitter has sides approximately 96 mm in length. The receiver was attached rigidly to the scan-probe using a thermoform plastic mount and nylon bolts, as illustrated in Figure 4.2. The distance between the centre of the receiver and the surface of the transducer was approximately 30 mm. The *pcBIRD* transmitter was supported by a Perspex shelf secured to the head of the examination couch using brass screws (see Figure 4.3). The *pcBIRD* is a pulsed-DC device, and was chosen because of its low sensitivity to metal in the environment compared to equivalent AC devices. It also has the advantage that the controller board has an ISA interface which enables it to be installed neatly inside a PC.

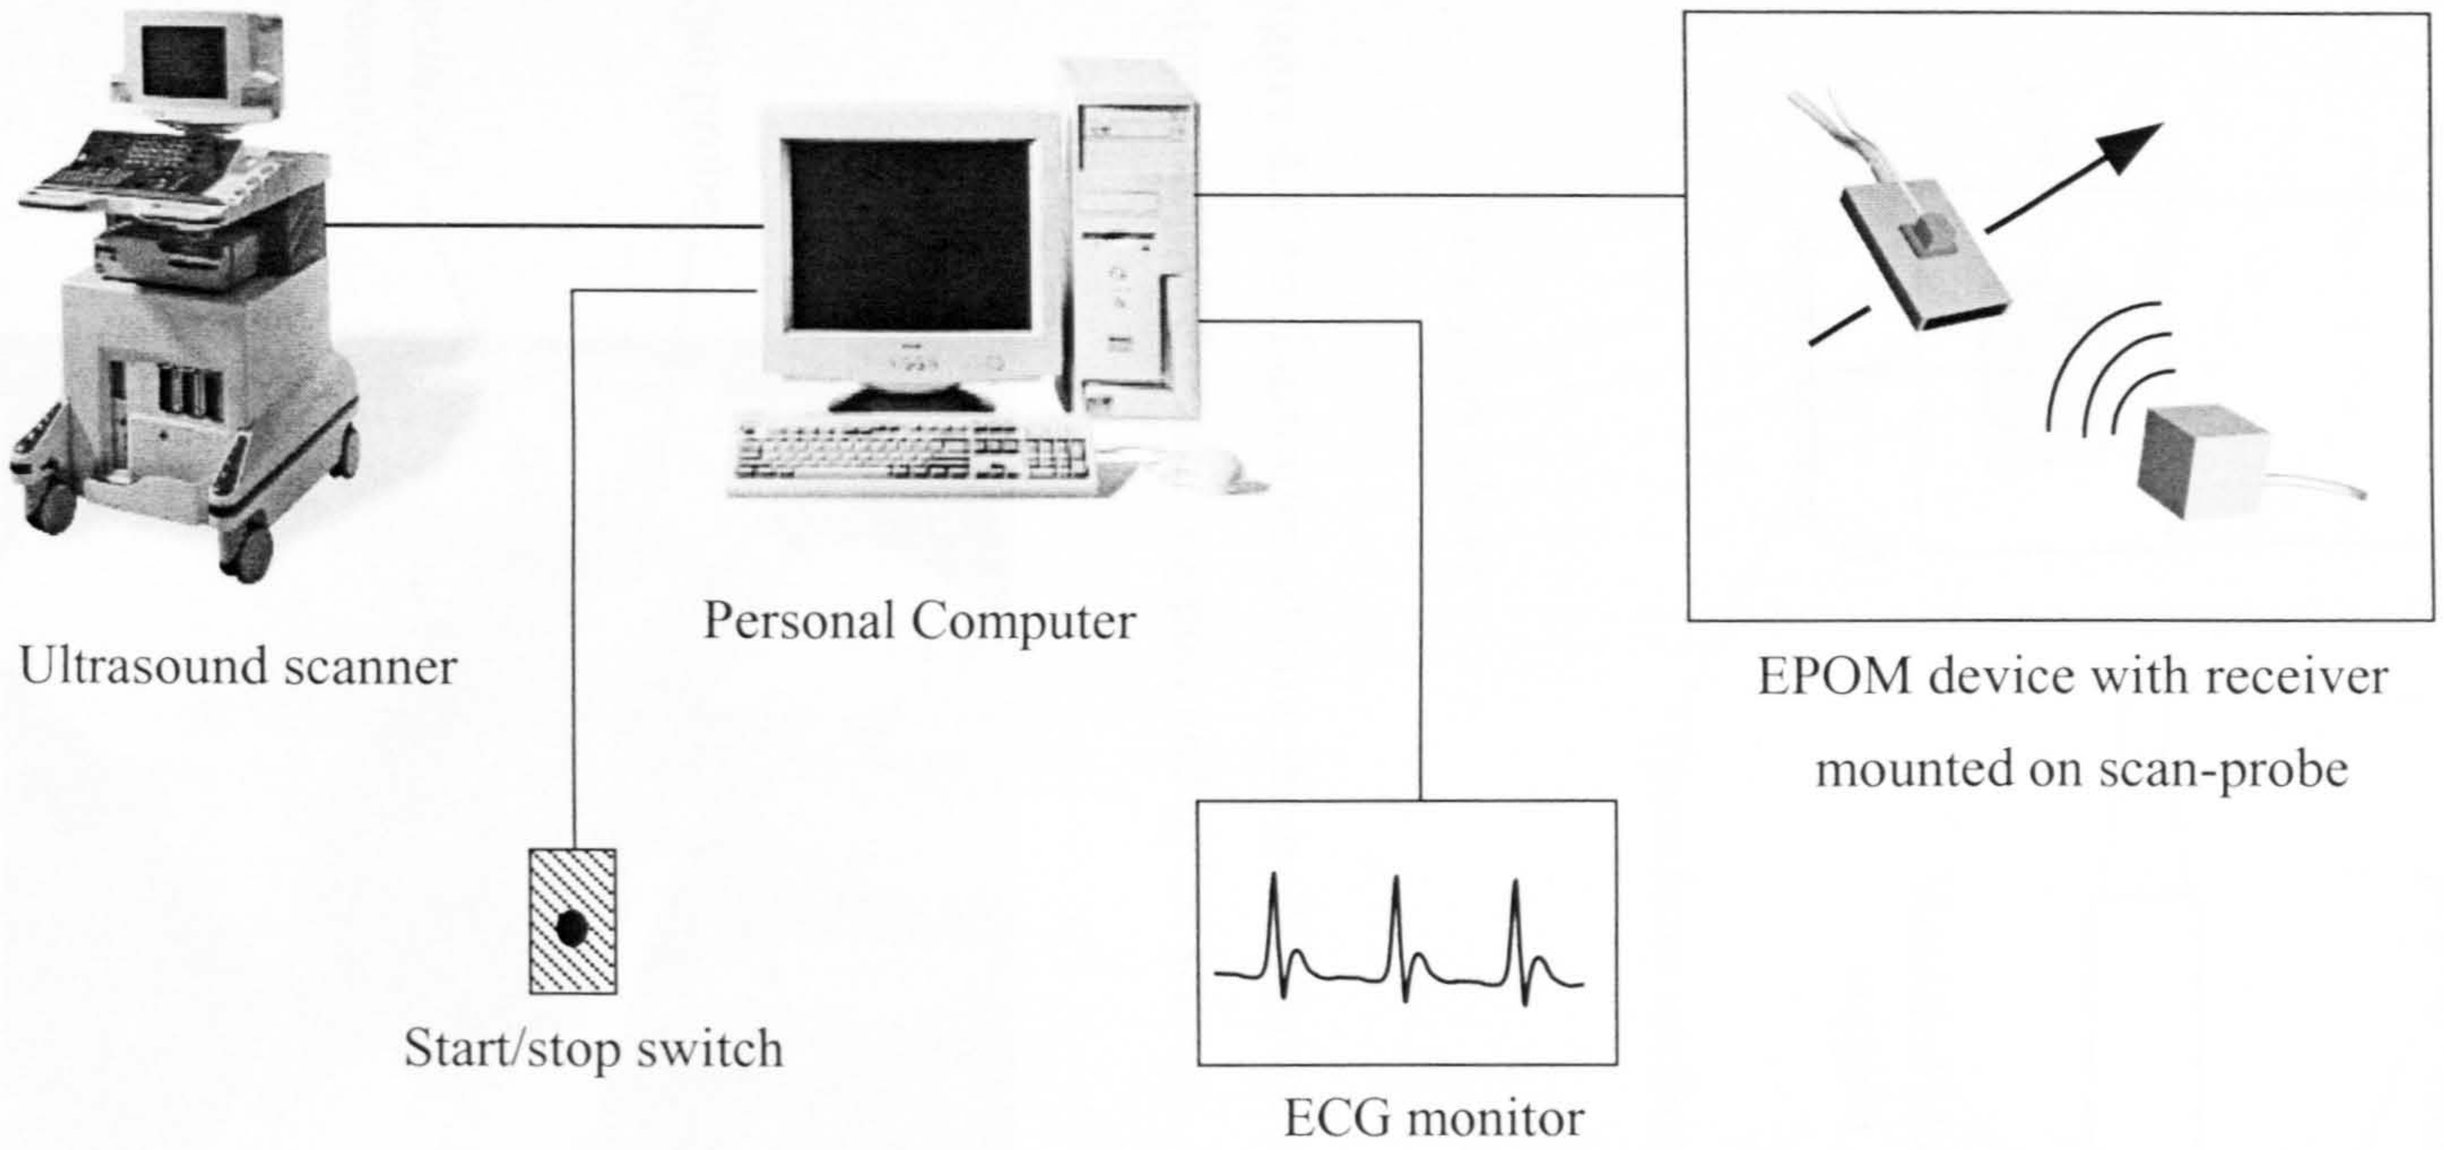


Figure 4.1 The electromagnetic 3DUS acquisition system developed for this project. During acquisition, ultrasound image data are stored in the scanner's cineloop (frame buffer) upon receiving a trigger signal from the personal computer. At the same time, the position and orientation of the electromagnetic receiver mounted on the scan-probe, relative to the fixed transmitter, are measured and recorded.

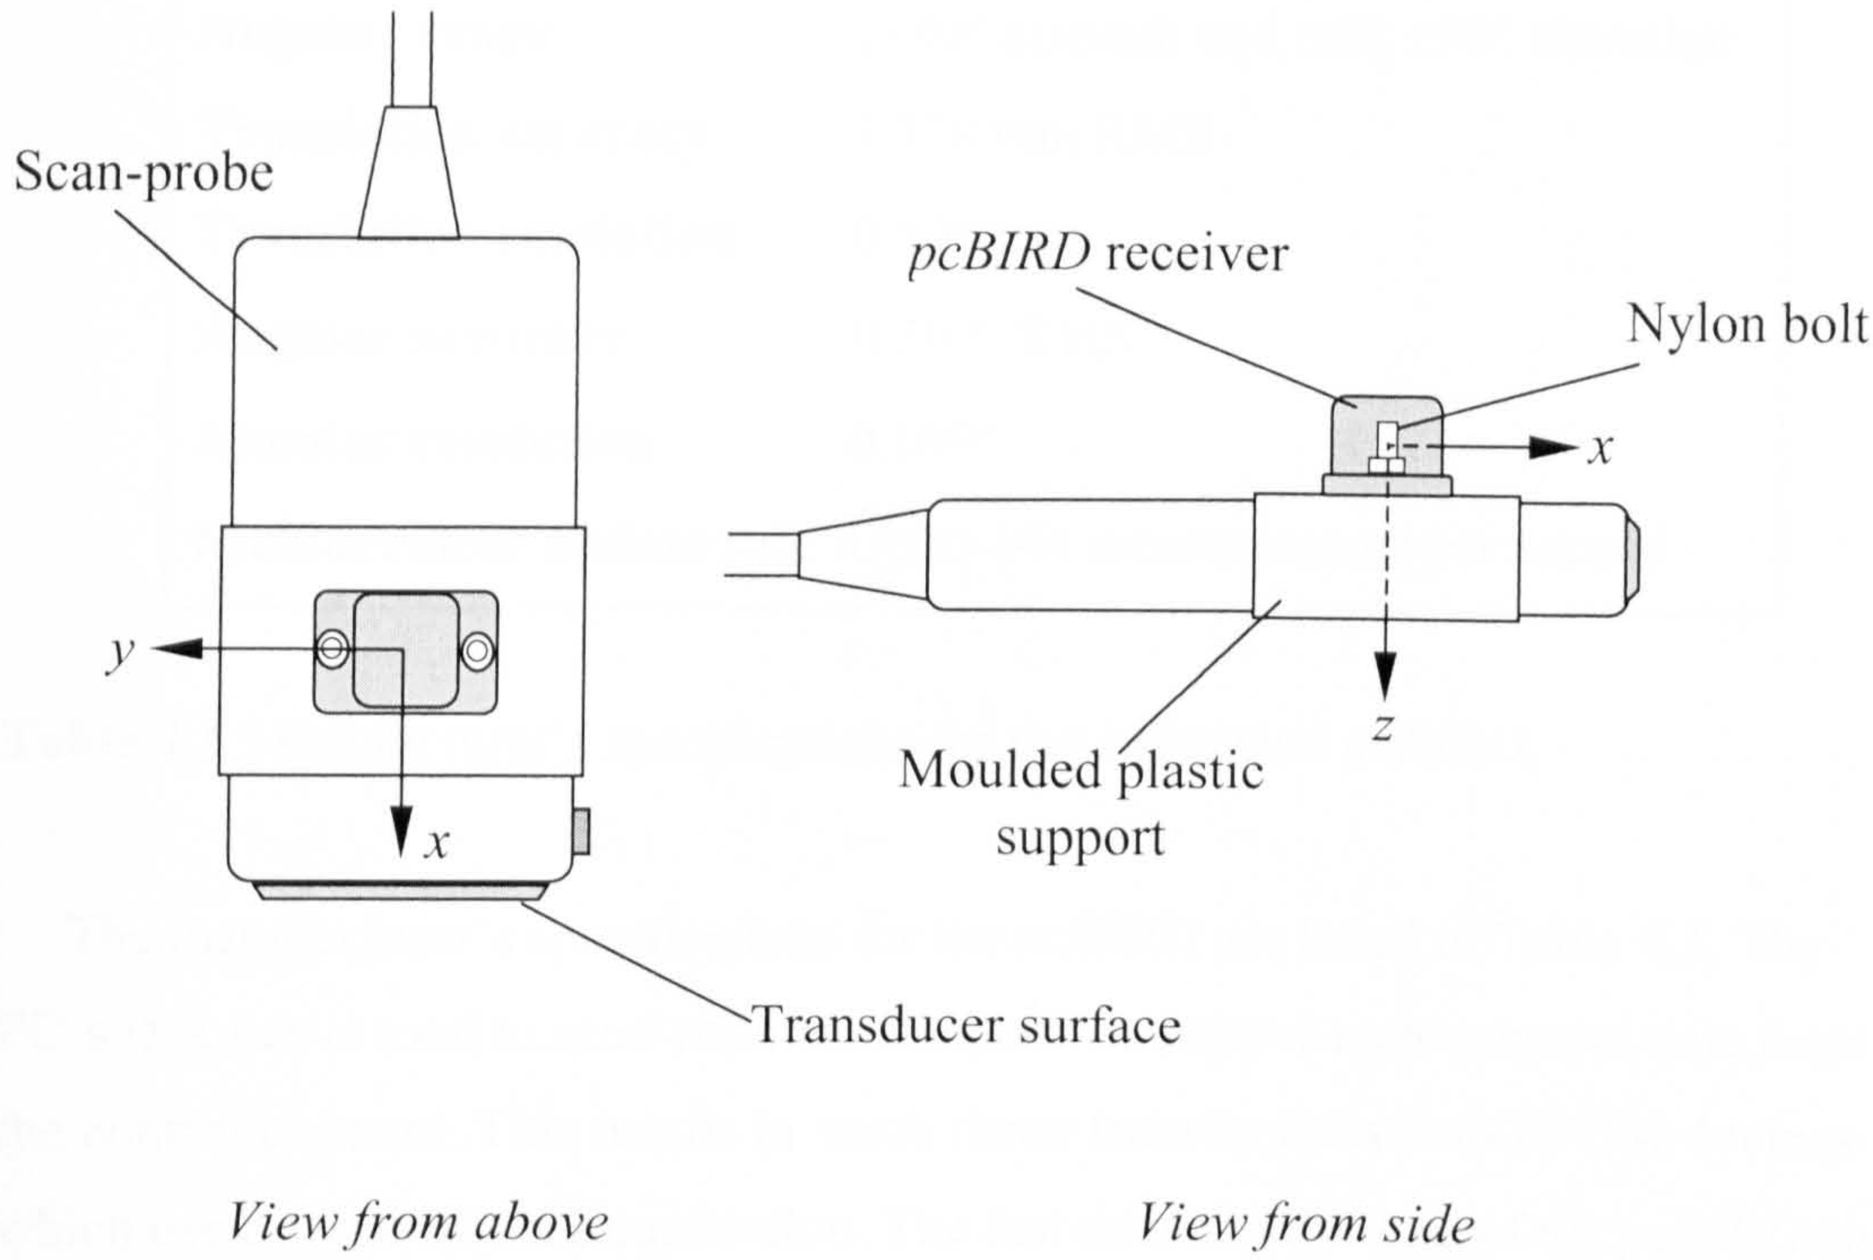


Figure 4.2 An illustration of how the *pcBIRD* receiver is mounted on a scan-probe using a moulded plastic support and nylon screws.

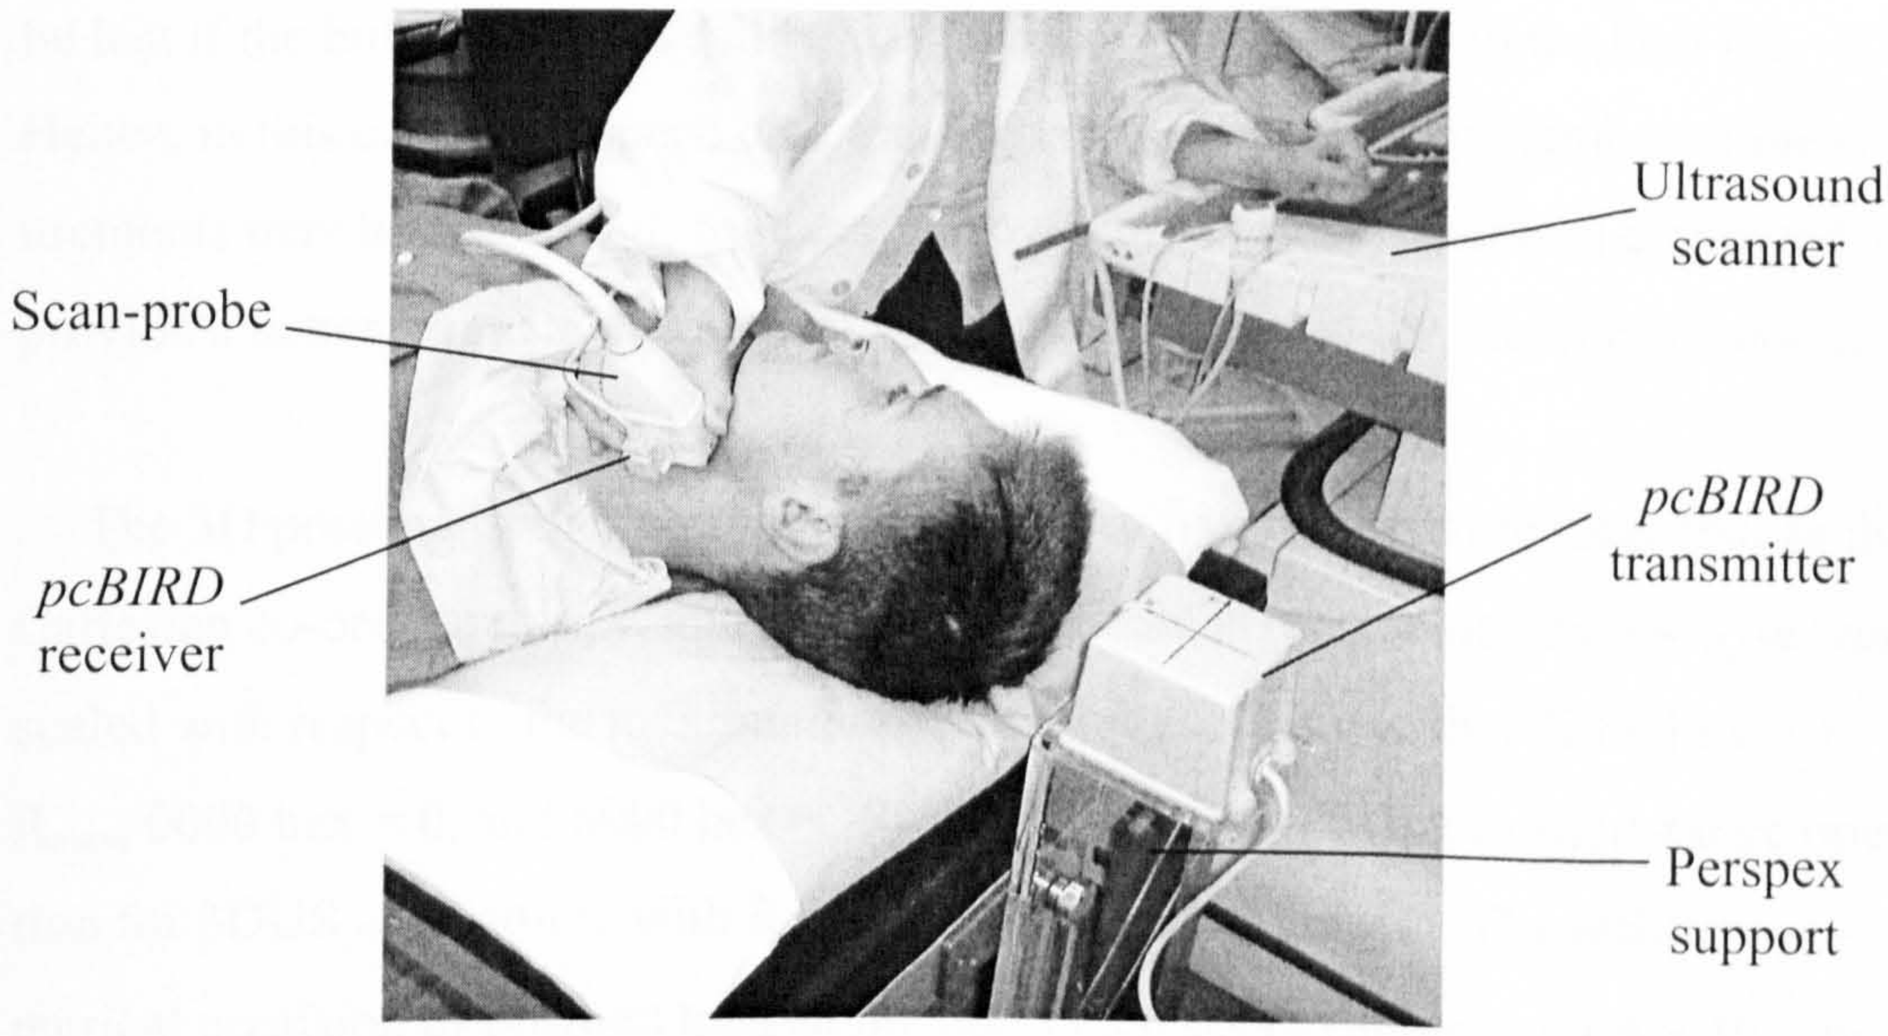


Figure 4.3 The configuration of the *pcBIRD* transmitter and receiver for bilateral carotid scanning.

Translation range	±1.219 m in any direction
Angular range	±180° azimuth and roll; ±90° elevation
Translation accuracy	1.778 mm RMS
Translation resolution	0.508 mm
Angular accuracy	0.500° RMS
Angular resolution	0.100°
Measurement update rate	Up to 144 measurements per second

Table 4.1 Manufacturer's specifications for the Ascension *pcBIRD*.

The manufacturer's specifications for the *pcBIRD* are listed in Table 4.1. The PC's ISA bus is used to send commands to, and to receive measurement data from the controller board. This results in much faster transfer rates than EPOM devices which use a serial RS-232 connection. The fast data transfer rate of the *pcBIRD* enables acquisition software to record measurements more quickly than an equivalent RS-232 device operating at the same measurement rate. In most cases, this is not a significant advantage since measurements are buffered. However, if a large number of measurements is required, made at a high measurement rate, measurements may be lost if the buffer becomes full before data can be transferred to the host PC. Hence, in this case, high speed data transfer is desirable. A large number of measurements may be required if, for example, averaging measurements is employed to provide a better estimate of true position and orientation in the presence of noise.

The 3D position of the receiver relative to the (fixed) transmitter is output as three Cartesian co-ordinates, x, y, and z. Each co-ordinate is represented by a 4-byte word, scaled with respect to the maximum operating range, R_{max}, so that 7FFF hex = + R_{max}, 0000 hex = 0, and 8000 hex = -R_{max}. The *pcBIRD* was set to short-range operation for 3DUS acquisition, with R_{max} = 914.40 mm (36 inches). The resultant numerical precision of position measurements is then approximately $\pm 1.4 \times 10^{-2}$ mm.

The orientation of the receiver is defined by three fixed-axis rotation angles – α, β and γ – and orientations were output as rotation matrices. In this representation, each orientation is described by a 3 × 3 matrix of the form:

$$\begin{bmatrix} \cos\beta\cos\gamma & \cos\beta\sin\gamma & -\sin\beta \\ -\cos\alpha\sin\gamma+\sin\alpha\sin\beta\cos\gamma & \cos\alpha\cos\gamma+\sin\alpha\sin\beta\sin\gamma & \sin\alpha\cos\beta \\ \sin\alpha\sin\gamma+\cos\alpha\sin\beta\cos\gamma & -\sin\alpha\cos\gamma+\cos\alpha\sin\beta\sin\gamma & \cos\alpha\cos\beta \end{bmatrix}. \quad (4.1)$$

This matrix, which is derived in Appendix 1, represents a counter clockwise rotation, α, about the transmitter z-axis; followed by a counter clockwise rotation, β, about the transmitter y-axis; followed by a counter clockwise rotation, γ, about the transmitter x-axis. The form has the advantage of providing a more stable representation of orientation than Euler angles or quaternions when β approaches 90°.

Like the position measurements, each element of the rotation matrix is represented by a 4-byte word with the following encoding: 7FFF hex = + 1, 0000 hex = 0, and 8000 hex = -1. Consequently, the numerical precision for each rotation matrix element is approximately $\pm 1.5 \times 10^{-4}$ over the interval [-1, 1].

A number of parameters of the *pcBIRD* can be modified. This is a useful feature as it allows the performance of the device to be optimised in noisy environments. The two primary user-adjustable parameters are the *measurement rate* and the *digital filter* setting. Increasing the measurement rate of the device improves dynamic performance because the electromagnetic field is sampled more quickly, resulting in more measurements being made per second. However, the measurement noise may also increase if metal is present in the environment. This is because eddy currents induced in nearby metals have less time to subside before a measurement is taken. Measurement rates of up to 144 Hz are available, with the default being 103 Hz.

Applying additional filtering can reduce measurement noise, but may also reduce the dynamic performance of the device. The filters can be set to any the following: (i) no filtering; (ii) an adaptive, low-pass DC filter (DC); (iii) a wide-notch AC filter (ACW); (iv) a narrow-notch AC filter (ACN); (v) a DC filter in combination with a wide-notch AC filter (DC + ACW); and (vi) a DC filter in combination with a narrow-notch AC filter (DC + ACN) . The frequency response characteristics of all the filters can also be changed but this was not attempted. The default setting for the *pcBIRD* is DC+ACW.

pcBIRD measurements may be output in one of two modes: POINT mode, where a single set of measurements (e.g. a position and a rotation matrix) are made immediately after a command has been issued, or STREAM mode, where one set of measurements is made regularly each measurement cycle.

4.1.3 Personal Computer

A 200MHz Pentium PC was used in the system. The function of the PC is to coordinate the operations of the other components within the system via custom-written software (detailed in Section 4.1.6). The PC houses the pcBIRD controller card, and a 12-bit, 50 kHz ISA data acquisition card (DAS802, Computer Boards Ltd, Massachusetts, USA). This enabled ECG-gated acquisition, and a mechanism for triggering digital image capture on the scanner, as explained below.

4.1.4 ECG monitor

A clinical ECG monitor with an external R-wave trigger option enabled (LifePulse 10, HME Ltd, Hertfordshire, UK) was used to provide ECG-gated acquisition. A short TTL trigger pulse was output by the monitor on every peak R-wave. This was detected by software via the data acquisition board. The trigger pulse was used to synchronise image capture and pcBIRD measurements with a regular point in the cardiac cycle, thus reducing artefacts due to vessel motion. The way in which ECG-gating was implemented is described in Section 4.1.6.

4.1.5 Start-stop switch

Acquisition was started and stopped comfortably using a hand-held switch connected to the PC printer port. This enabled the PC to be placed in a convenient position away from the scanning region.

4.1.6 Software

3DUS acquisition software was written in Microsoft Visual C++ v. 4.0, and runs as a Windows 95/98 application on the PC. Its purpose is to synchronise pcBIRD measurements with the capture of ultrasound image data. A delay is introduced by the acquisition software so that both events are initiated at a preset time after the detection

of an ECG R-wave. This allows data acquisition to be synchronised to a regular point within the cardiac cycle.

Figure 4.4 shows a screen-shot of the graphical user interface (GUI) for the acquisition software. Once an acquisition cycle has been started by pressing the start-stop switch, the 3D position and orientation of the current ultrasound imaging plane is indicated by a filled red rectangle. Green rectangles are used to represent the images captured so far within the current acquisition. Therefore, during acquisition, the region enclosed by the green rectangles indicates the current total volume scanned so far. The OpenGL graphics library (Silicon Graphics Inc., USA) was used to generate this real-time 3D display.

The *pcBIRD* is operated in POINT mode throughout acquisition, and positional co-ordinates and rotation matrices are stored in PC memory. At the end of acquisition, the position and orientation matrices are converted into a single 4 × 4 homogeneous matrix (see Appendix 1) and saved in an ASCII file.

The *HDILab* software is used to download images for subsequent processing on a PC. Image data are stored in a proprietary file format indicated by the '.xif' file extension. Each file contains pre-scan-converted ultrasound data together with information on the scan protocol used, and the ECG signal. Figure 4.5 shows the GUI for *HDILab*. A powerful feature of the software is that ultrasound images can be manipulated offline using a variety of functions. A subset of these, such as TGC controls and display map selection, are equivalent to those available on the ultrasound scanner. Others are related to post-processing operations which change the appearance of the displayed ultrasound image. The software therefore provides a versatile platform for viewing and processing ultrasound images. The capability of the software is also extendable by adding custom operations via software plug-ins.

4.2 System operation

Figure 4.6 summarises the operation of the 3DUS system. During acquisition, the scan-probe is slowly swept across the surface of the skin. During the sweep, R-wave

(*continued on page 113*)

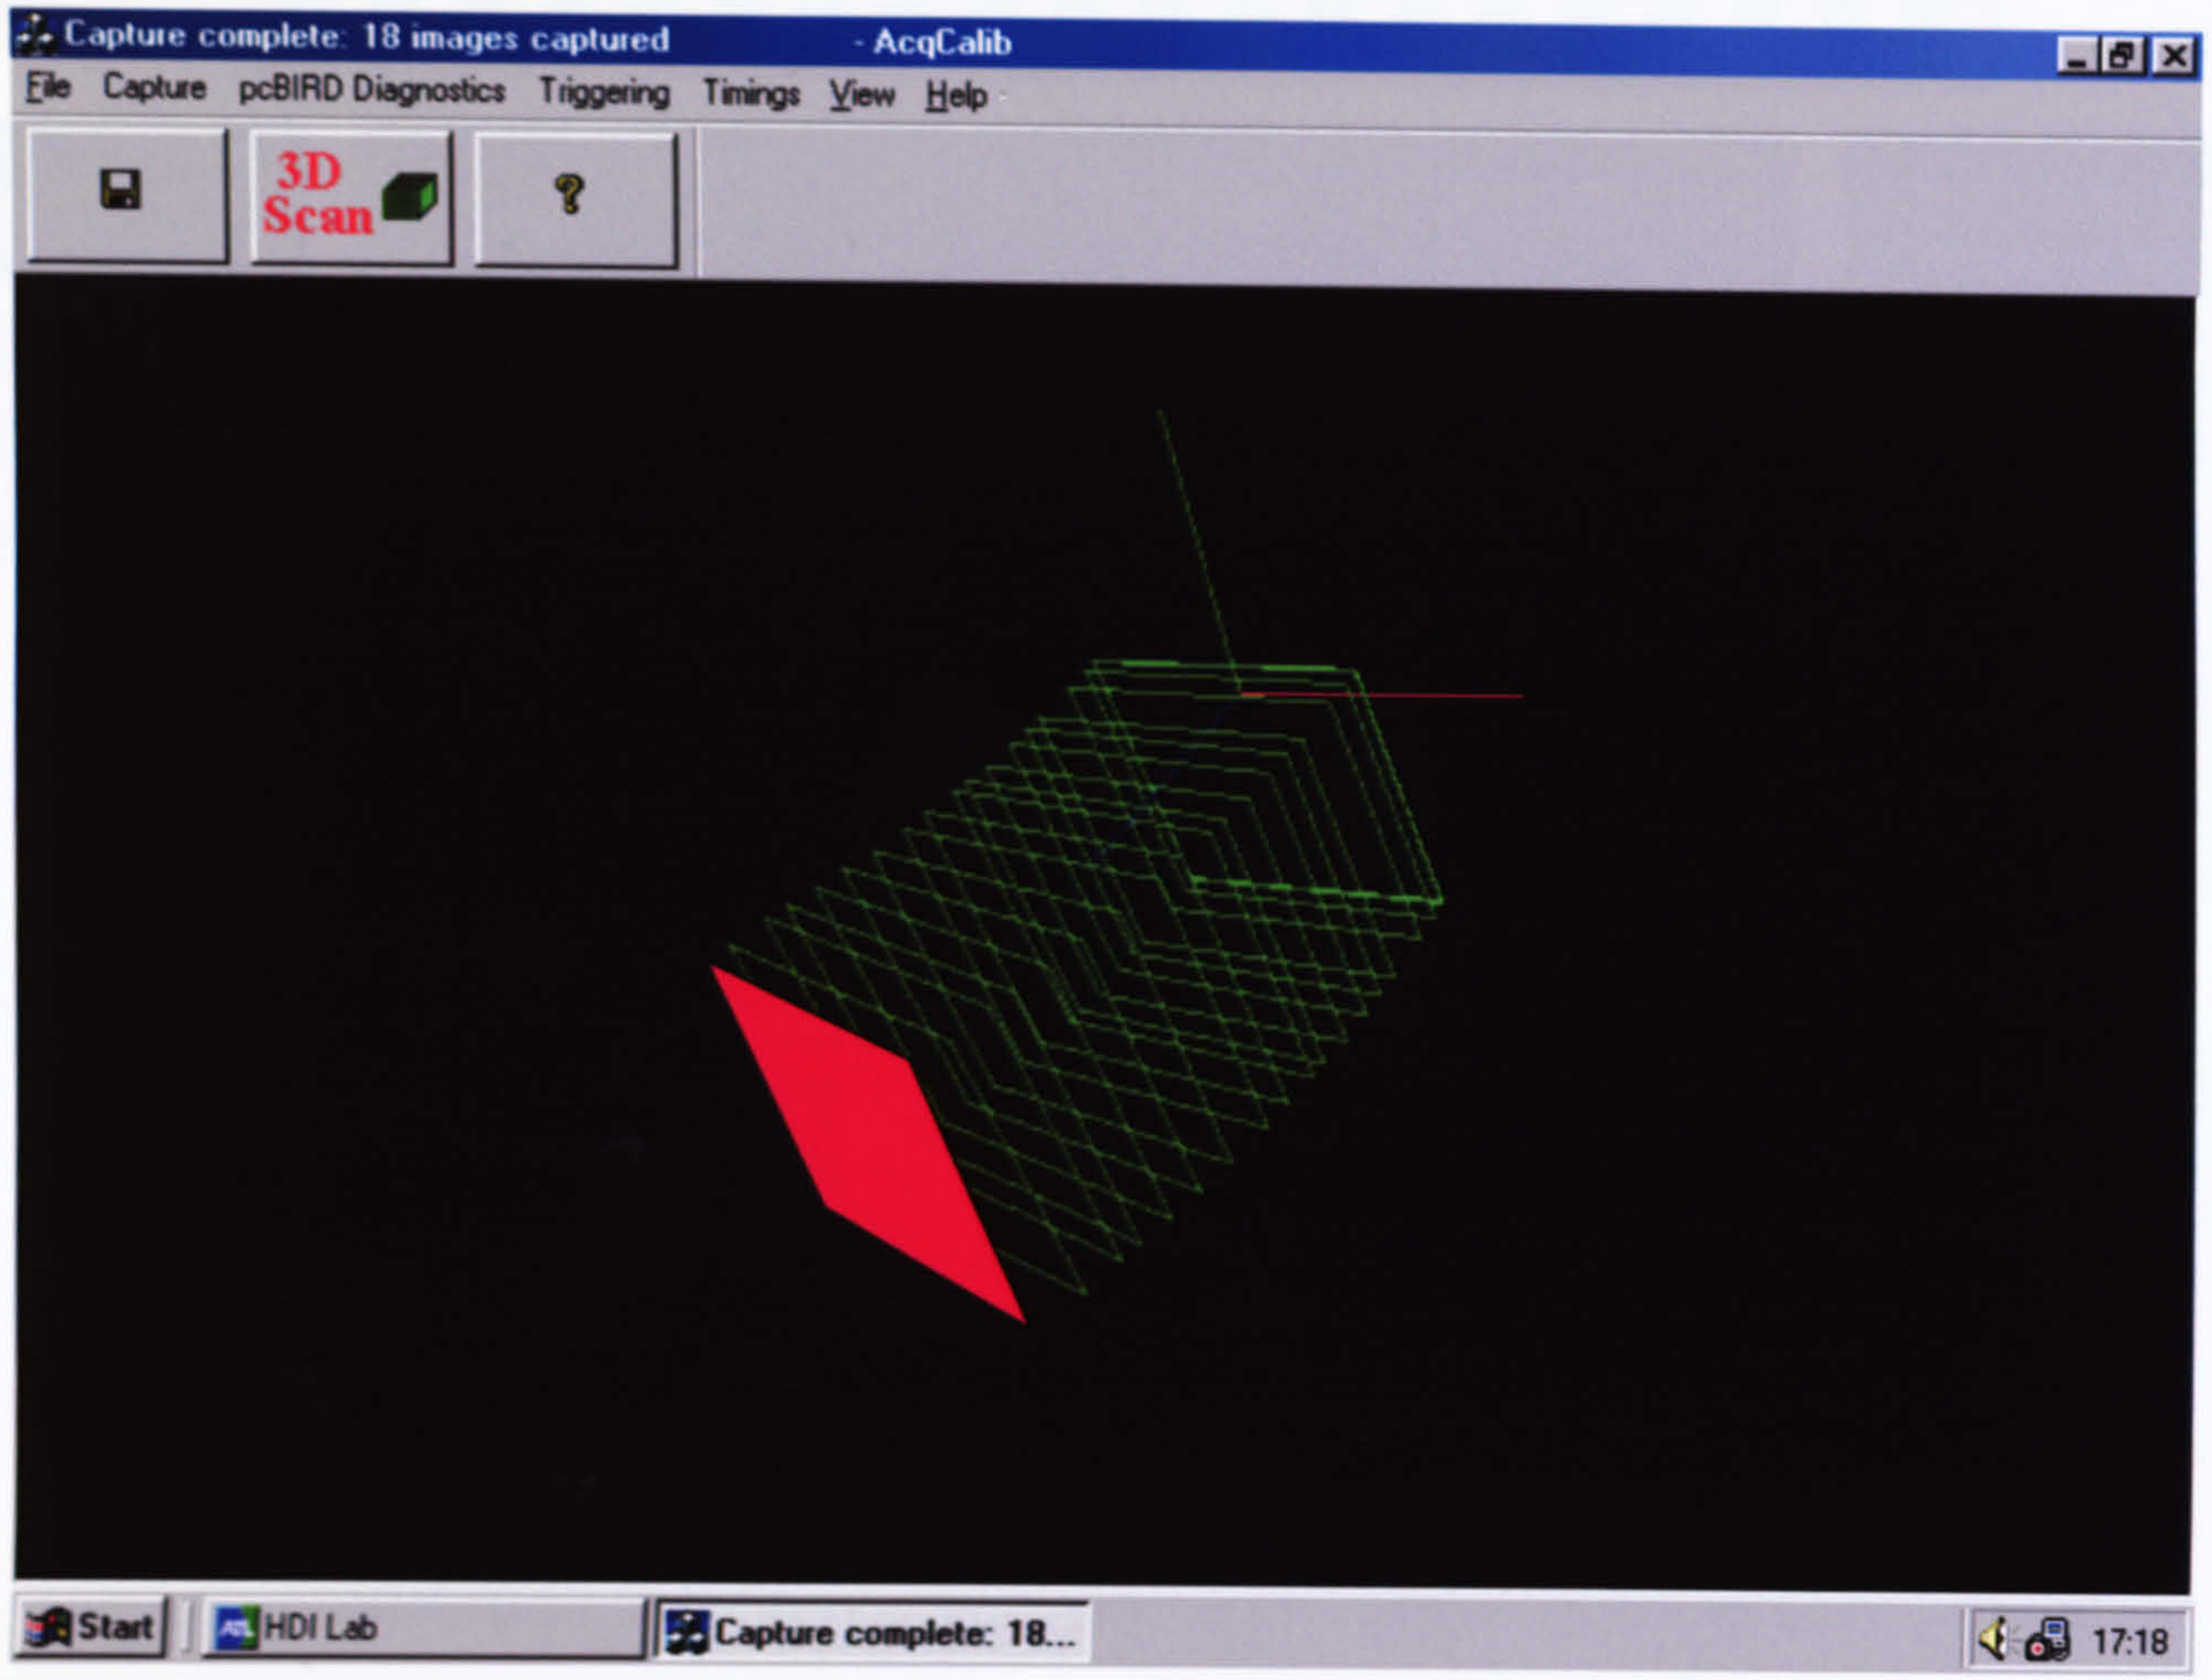


Figure 4.4 Graphical user interface for the 3DUS acquisition software. Captured images are represented by green rectangles, whereas the current image is represented by a filled red rectangle.

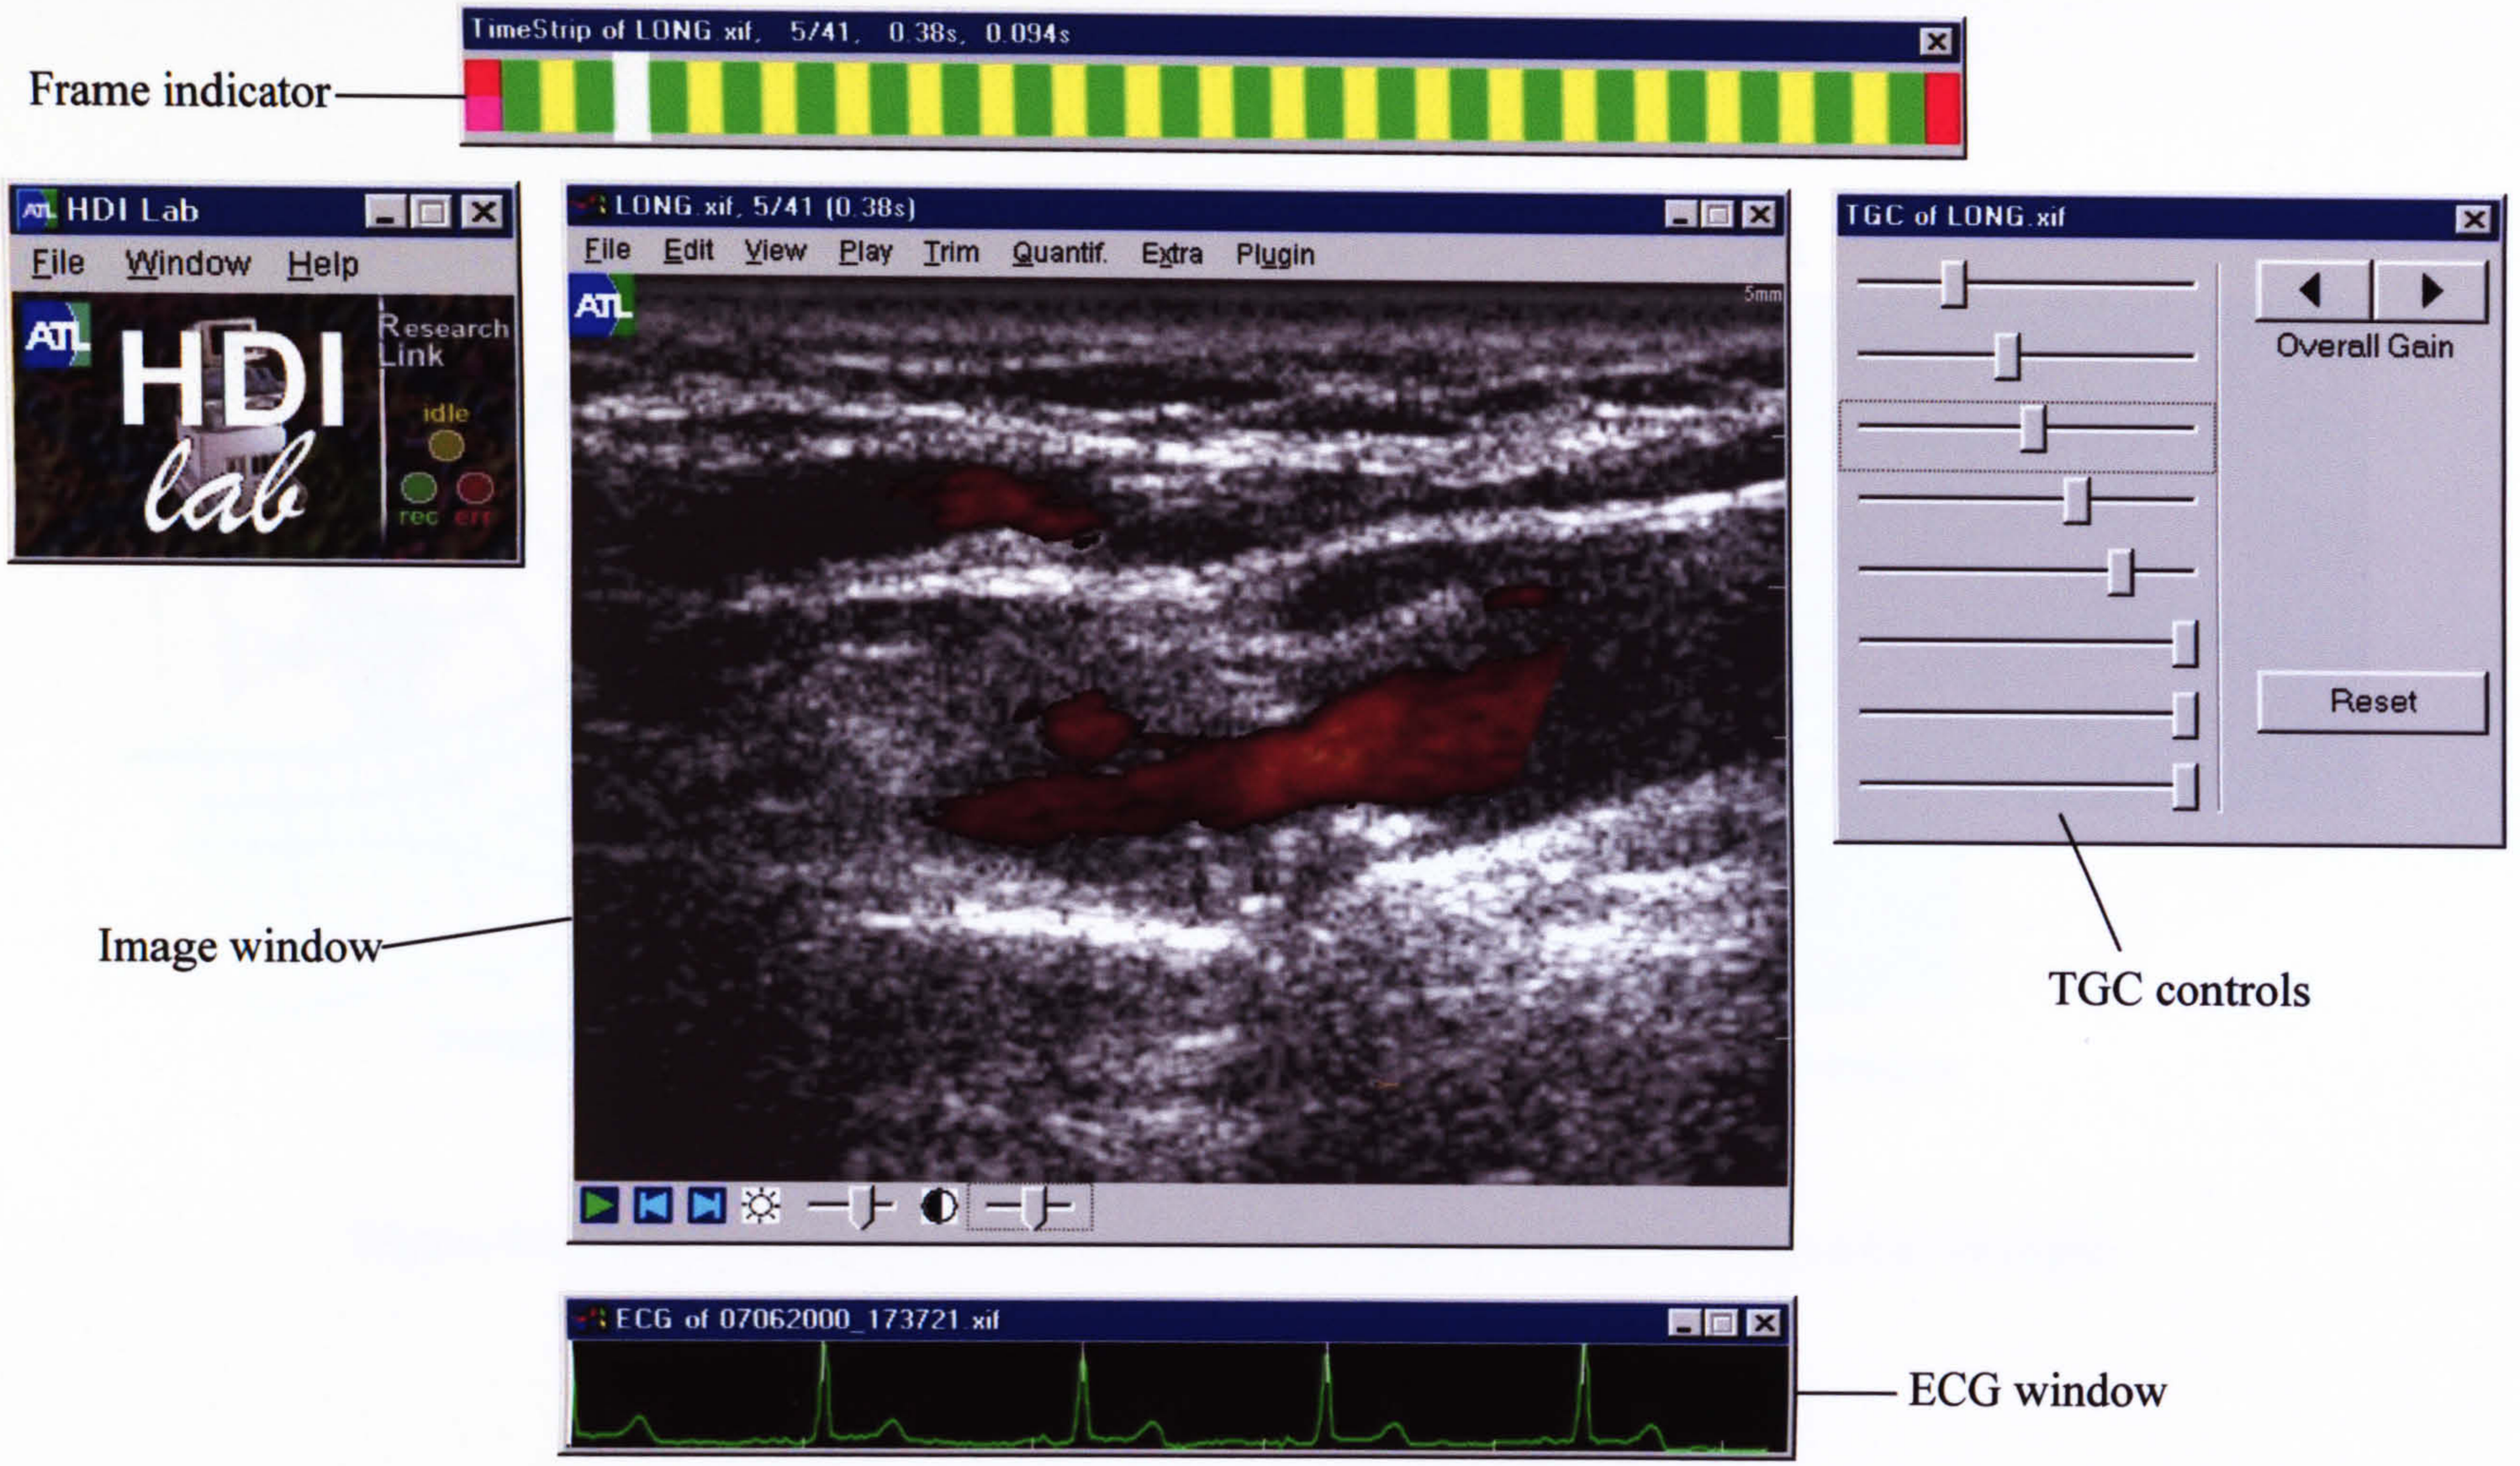


Figure 4.5 Graphical user interface for the *HDILab* software. The image displayed is frame 5 in a set of 41 power Doppler images (as shown by the frame indicator at the top). Off-line TGC controls are shown on the right, and the ECG signal is displayed in a separate window underneath the image window.

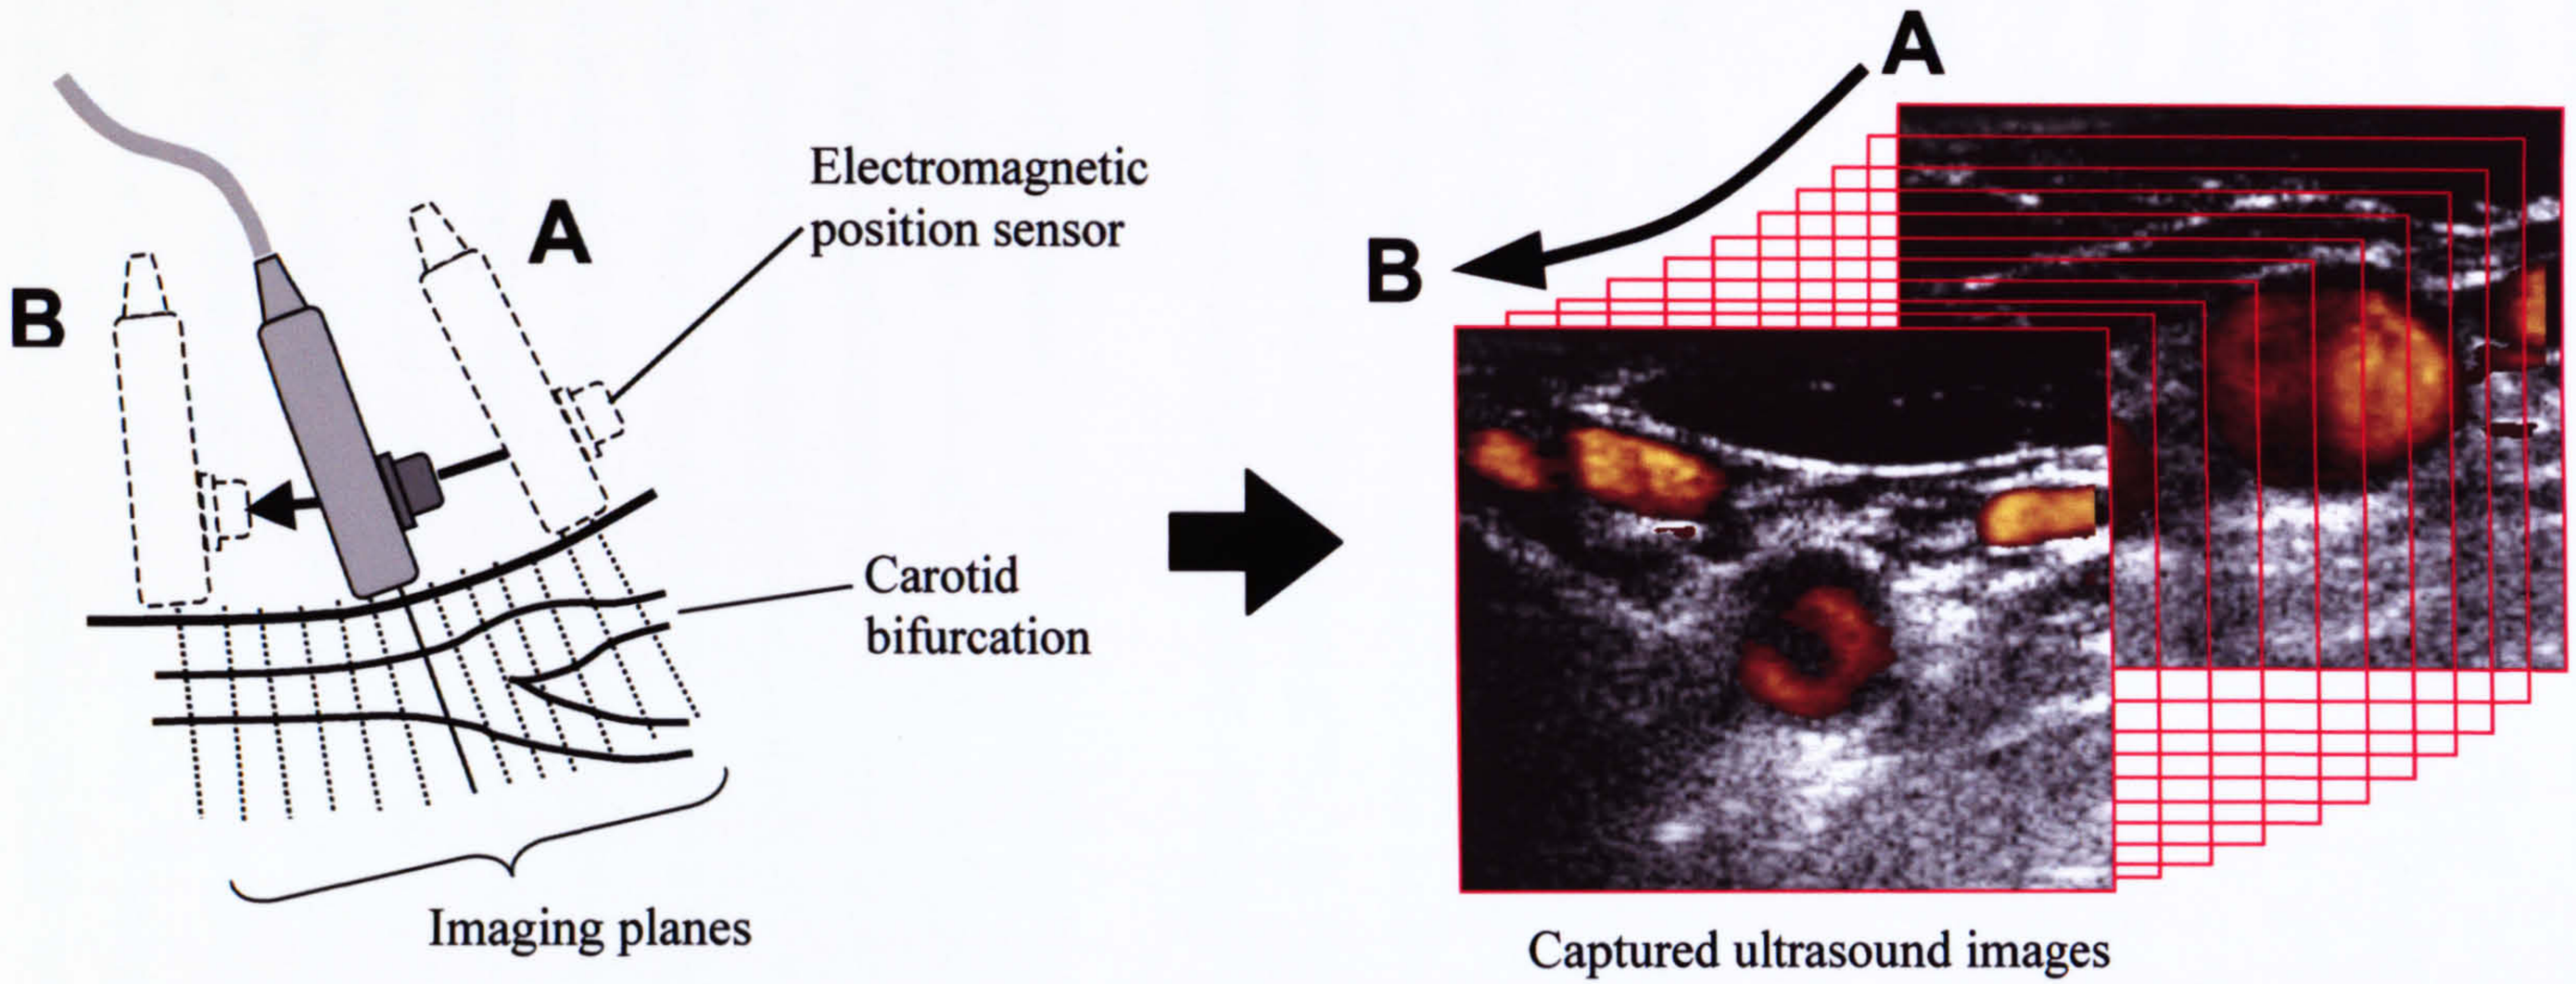


Figure 4.6 3DUS data acquisition using the electromagnetic system developed for this thesis.

trigger signals are generated periodically by the ECG monitor. These are detected by the acquisition software via an analogue to digital conversion channel of the data acquisition board. The detection of an ECG trigger is followed by a preset delay, after which a trigger pulse is sent to the scanner via its external ECG port. When the trigger is detected, the image currently being formed is transferred into the cineloop as soon as it is ready to be displayed. A request for a position and orientation measurement is sent to the pcBIRD at the same time as the image capture trigger is sent. Once received, pcBIRD data are stored in PC memory until the end of the scan. The beginning of an acquisition is signified by a short burst of 5 closely spaced trigger pulses. This enables the first captured image in a sequence to be identified during off-line analysis. The sequence of events in an acquisition cycle is depicted in Figure 4.7.

The configuration of this system has the advantage that only those images used in the 3D reconstruction are captured. Therefore, there is no need to filter out images, manually or otherwise, at a later stage. The 3DUS system described in this thesis is fully digital in the sense that the pre-scan-converted ultrasound data used to generate each image are stored in the cineloop of the scanner during acquisition, and later downloaded to the computer for further processing offline.

The size of the scanned volume is limited by the maximum number of images which can be stored in the cineloop, the sweep speed, and the size and type of captured images – colour images require more memory than B-mode images. Larger volumes can be imaged by increasing the sweep speed, but this may give rise to image misregistration or flash artefacts if colour or power Doppler modes are used. A second way to increase the scan-volume size is to acquire images less frequently than once every cardiac cycle, while maintaining the normal sweep speed. This can be implemented very simply by keeping a count of the number the R-wave triggers detected since the previous image capture. The inter-slice separation will, however, increase. These considerations suggest that there is a trade-off between the size of the scan-volume and the quality of the acquired data. Presently, data for up to about 140 power Doppler images can be stored on the HDI-5000 scanner, using a penetration depth of 4 cm. This was found to be sufficient for the majority of carotid scans, provided the heart rate is less than 100 bpm.

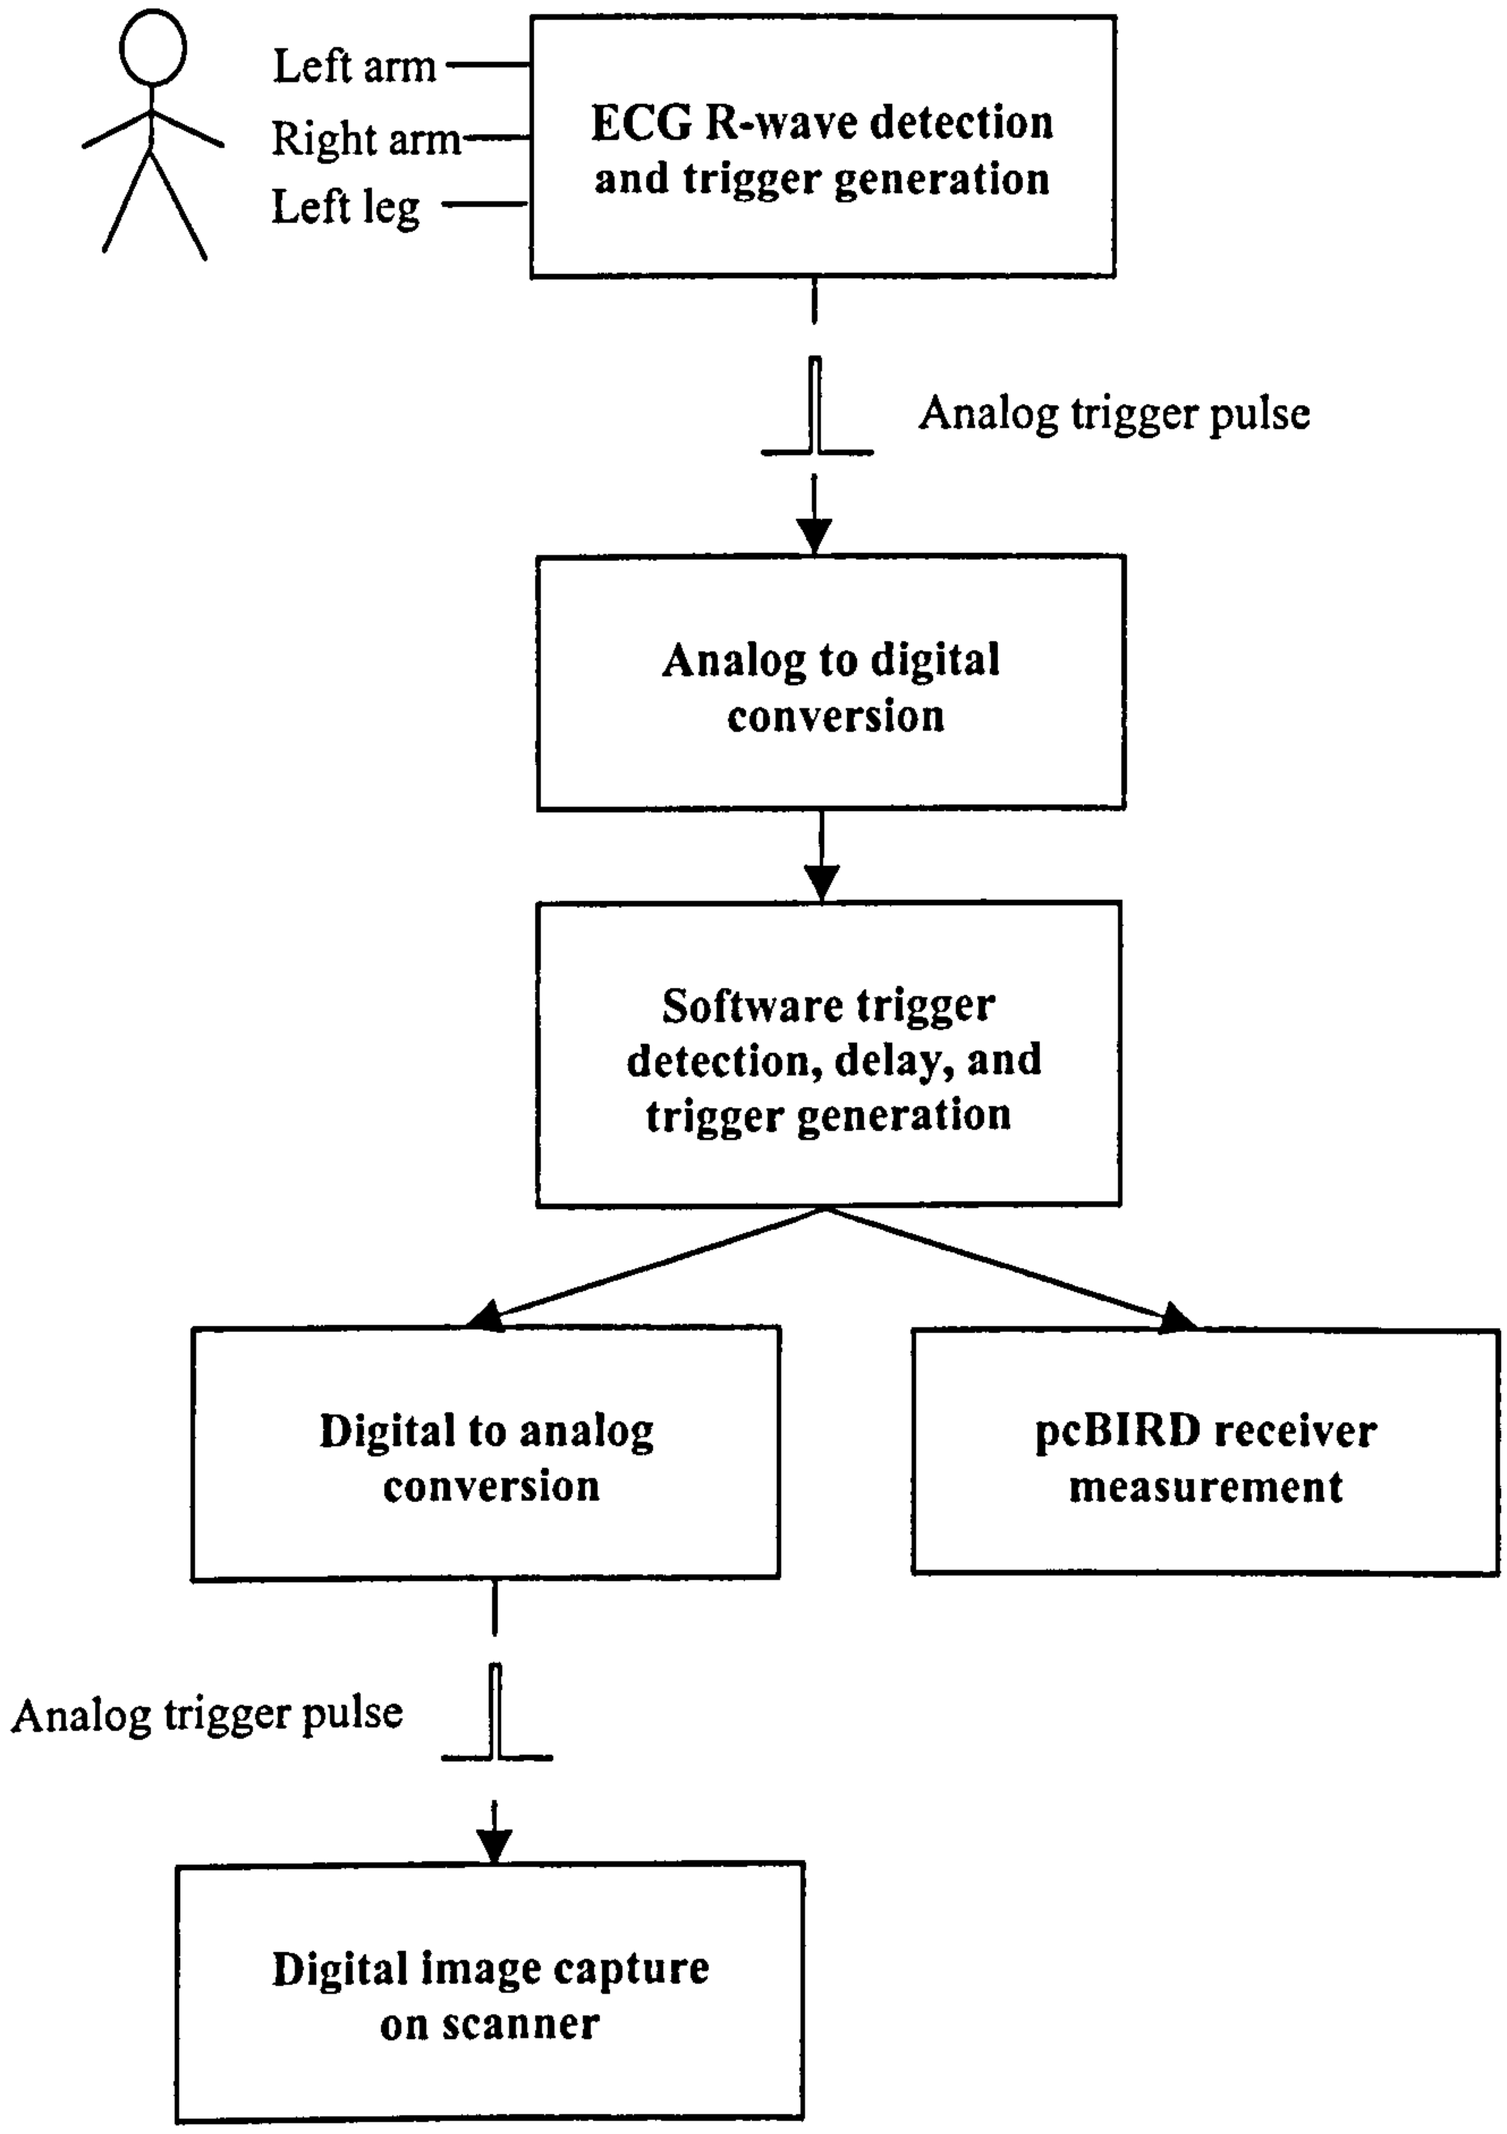


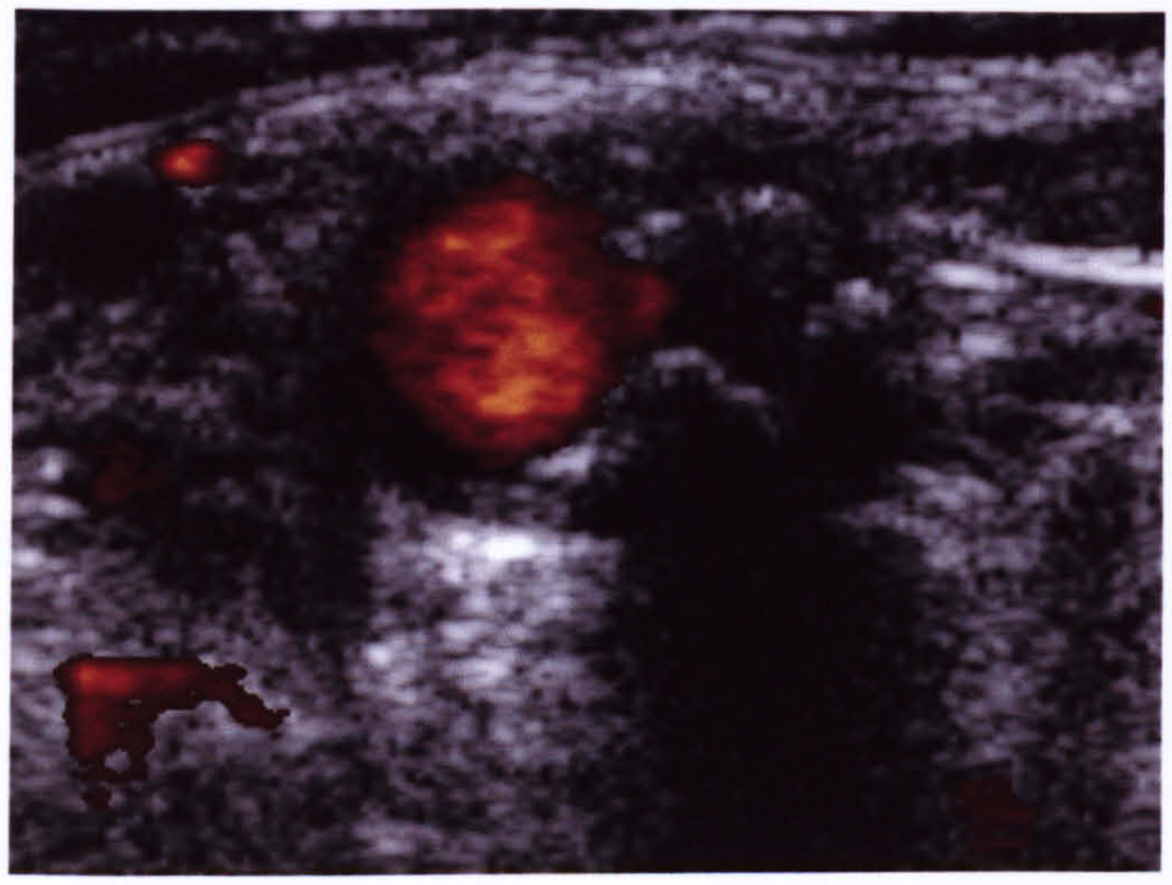
Figure 4.7 Sequence of events following the detection of an ECG R-wave during a 3DUS acquisition.

The use of the scanner's cineloop to store captured ultrasound image data is a novel feature of the 3DUS system described here, and removes the need for a VCR or analogue frame grabber to digitise images from the video output port of the scanner. Video-based image capture has been used in many other systems, but this method is inherently subject to some image degradation. Video timing irregularities in commercial ultrasound scanners can also lead to the misregistration of image capture and positional measurements [Meairs 2000]. However, such systems can, in principle, be used in conjunction with any scanner that has a video output – a feature which is standard on commercial scanners.

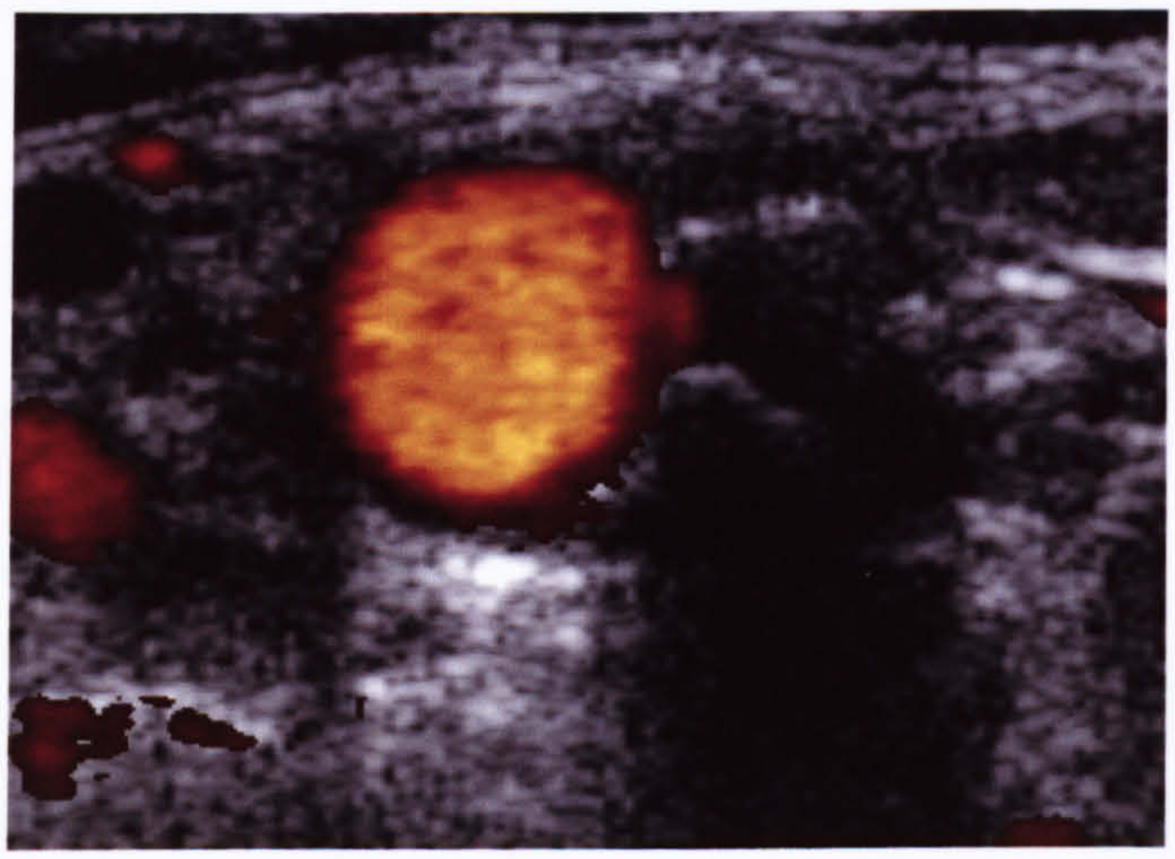
An important issue relating to the implementation of ECG-gating is the choice of the delay, Δ, between receiving an R-wave trigger and initialising an image and pcBIRD measurement. Setting $\Delta = 0$, so that acquisition takes place immediately following the peak of the R-wave, is the simplest and most reliable way to implement ECG-gating. If $\Delta > 0$, the image capture trigger occurs at point later in the cardiac cycle. Provided the R-R interval is relatively constant, setting $\Delta > 0$ enables images to be captured at approximately the same point within each cardiac cycle. This cannot be guaranteed for patients with irregular cardiac rhythms, however. If image slices are not captured at the same point in the cardiac cycle then some motion artefact may be introduced, since the gross position and dilatation of the vessel will be at a slightly different phase when each image is captured. Even if ECG-gating is employed, small differences in the point of image capture within the cardiac cycle may lead to appreciable motion artefact. The magnitude of this effect is dependent on which region of the cardiac cycle gating is applied: when the artery is undergoing rapid motion (a short time after the peak of the ECG R-wave), consistency between images is less likely. Therefore, in order to minimise motion artefact, Δ should be set so that image acquisition coincides with a phase of the cardiac cycle where the vessel wall is moving least, namely, diastole.

An additional benefit of using a non-zero value for Δ is that the timing of image acquisition may be adjusted so that maximum colour filling of the vessel occurs when colour or power Doppler modes are used. This is useful as it enhances the

(continued on page 117)



(a)



(b)

Figure 4.8 (a) A transverse power Doppler image of the CCA captured immediately after the peak R-wave. (b) An equivalent image captured after a delay of 400 ms after the peak R-wave. In both cases, the image persistence was set to zero on the ultrasound scanner.

visualization of the vessel lumen. In clinical scanning, it is normal practice to set a non-zero persistence for both colour and B-mode images as this allows consecutive frames to be averaged, thus reducing noise in the displayed image and sustaining colour filling. This is undesirable for 3D acquisition, however, since it increases uncertainty in the synchronization of captured images and positional measurements. For this reason, the persistence of both the B-mode and the colour/power Doppler components of an image were set to zero in the system described in this thesis. This has the side-effect that the colour filling of the vessel is very sporadic during the cardiac cycle. This is particularly noticeable for colour Doppler which is not as sensitive to low flow velocities as power Doppler. Hence, selecting a suitable value for Δ will also ensure optimal colour filling. Figure 4.8 illustrates the effect of changing Δ on colour filling in a power Doppler image of the CCA. In this case, maximum colour filling was found to occur when $\Delta \approx 400$ ms, corresponding to the end of the ECG T-wave. Fortunately, this is consistent with the requirement that gating should occur during diastole. The 3DUS acquisition software was adapted to include this delay by default. An option to further modify the value of Δ for individual patients, if necessary, was also added.

4.3 Calibration

4.3.1 Problem Definition

In order to correctly register acquired 2D ultrasound images in a reconstructed volume, it is necessary to determine, with some accuracy, the position and orientation of each image with respect to a fixed 3D co-ordinate system placed within the volume. Only then can a 3D representation be reconstructed which faithfully portrays anatomical structures, and from which accurate quantitative measurements may be made. The co-ordinate system of the reconstructed volume specifies a convenient 'real world' frame of reference for anatomical structures imaged within the scan-volume. During acquisition, the position and orientation of the *pcBIRD* receiver, mounted on the scan-probe, are measured. Because the receiver is offset relative to the surface of the transducer, *pcBIRD* measurements must be corrected to obtain the 3D location of each acquired image relative to the reference frame. This is achieved by means of a mathematical transformation which converts – or transforms – the (2D) co-ordinates of image pixels to 3D co-ordinates in the reconstructed volume.

This transformation, which may be expressed mathematically as a 4×4 matrix (see Appendix 1), is found by calibrating the 3DUS imaging system.

For the 3DUS acquisition system described here, calibration essentially involves determining the spatial relationship between the pcBIRD receiver co-ordinates and the co-ordinates of pixels within an ultrasound image. Figure 4.9 shows the conventions adopted for 3D co-ordinate systems. The co-ordinate system for the pcBIRD receiver, the (fixed) *pcBIRD* transmitter, and the ultrasound image are denoted by ***R***, ***T***, and ***I***, respectively. A fourth co-ordinate system, ***V***, is also defined as the fixed co-ordinate system for the reconstructed volume. The position and orientation of ***V*** are arbitrary, but they are generally chosen to provide a convenient frame of reference for the anatomical area under investigation. For example, ***V*** may be chosen so that the *z*-axis is approximately parallel to the longitudinal axis of an organ, and its origin is close to the scanned volume.

For ***I*** the x- and y-axes define the plane of the ultrasound image. A *z*-axis is defined in ***I*** for consistency, but it is redundant as, for the purposes of the present discussion, 2D ultrasound image slices are considered to have zero thickness, and image pixels lie in the *x*-*y* plane. In reality, images may be considered to have a finite thickness corresponding to the thickness of the ultrasound beam in the elevational direction. The justifications for assuming zero thickness here are that the actual thickness of the slice is much smaller than the axial and lateral dimensions, and that it is impossible to determine with any accuracy the z-position of a small object which lies within the thickness of the slice; it will produce a reflection, but could be at any position within the thickness of the beam. This point, and the effect of finite slice thickness on the accuracy of the 3DUS system, are considered later in Section 5.1.3.

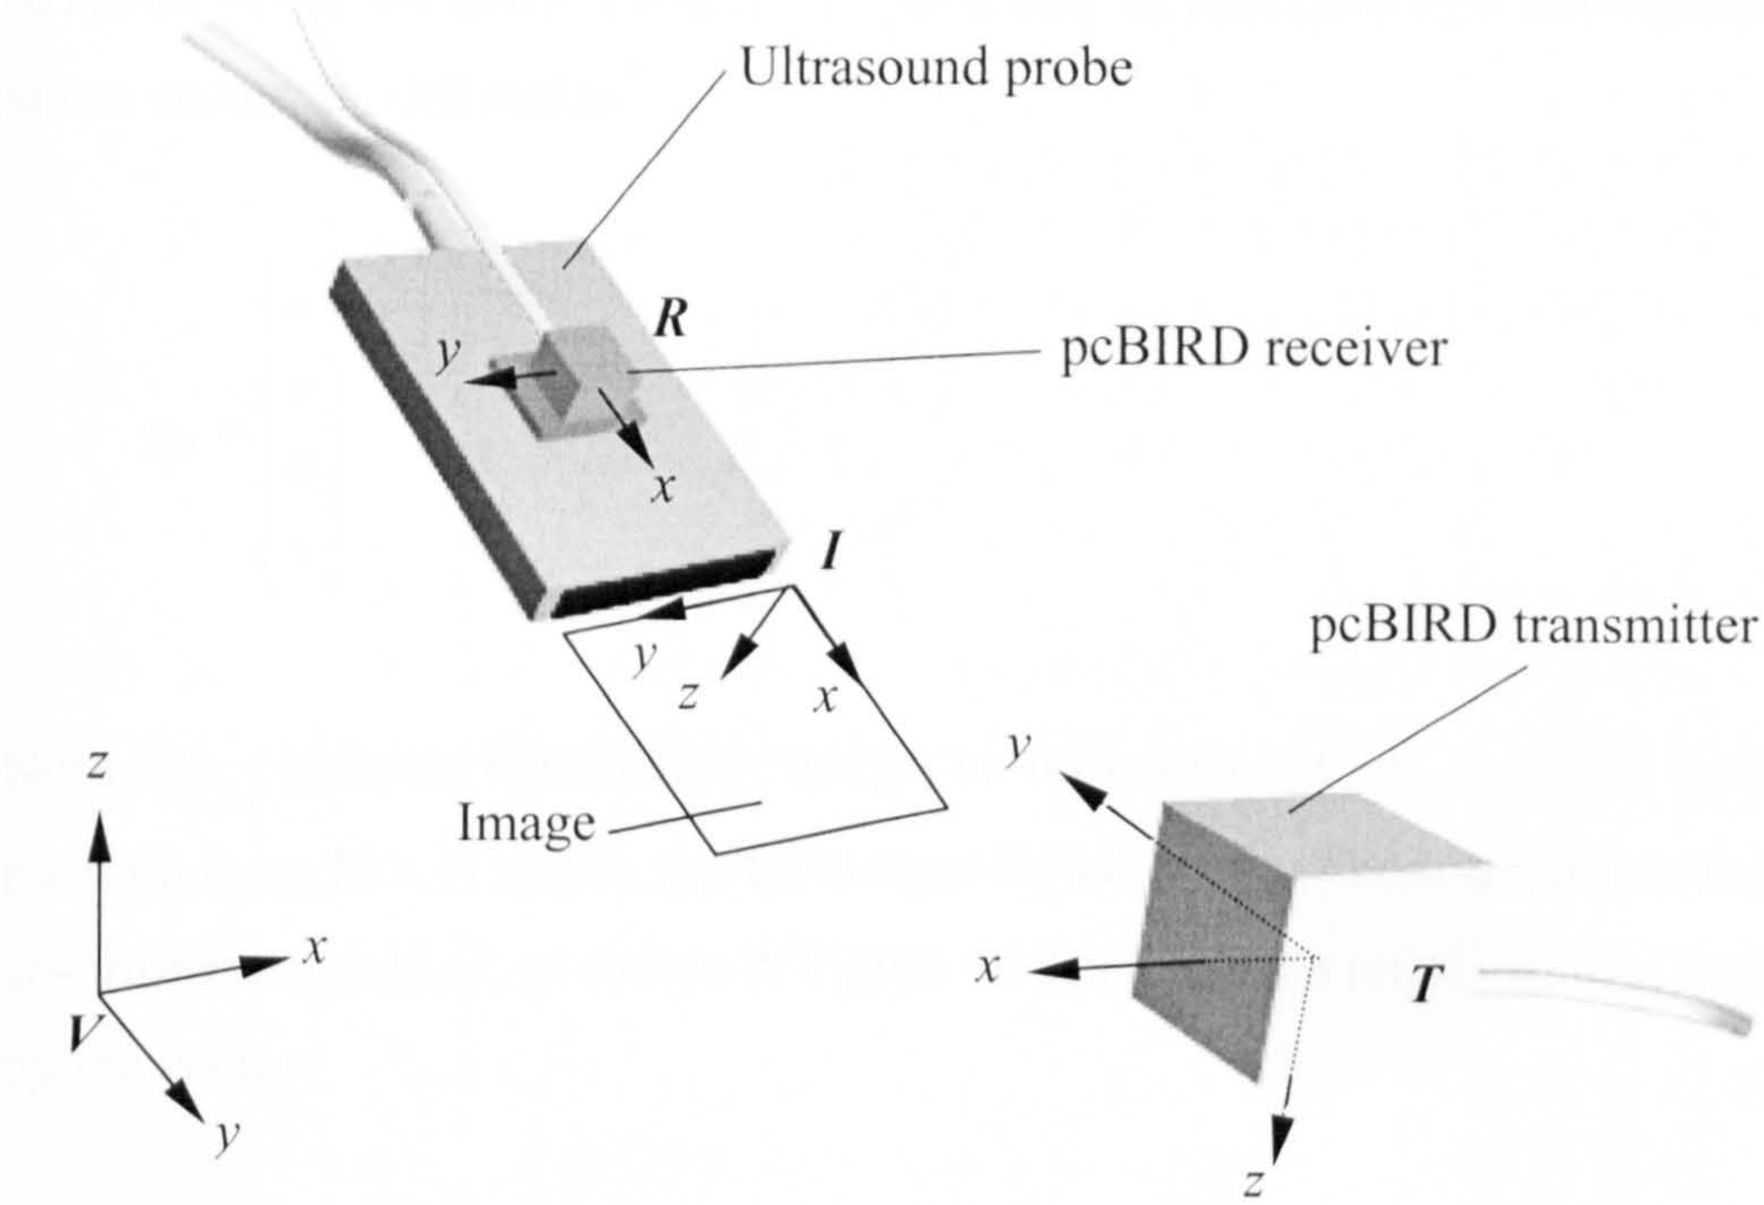


Figure 4.9 The 3D co-ordinate systems used for freehand 3DUS imaging. The co-ordinate systems for the pcBIRD receiver, the pcBIRD transmitter, the ultrasound image, and the reconstruction volume are denoted by ***R***, ***T***, ***I***, and ***V***, respectively (the *z*-axis for ***R*** is not shown but points downwards through the scan-probe).

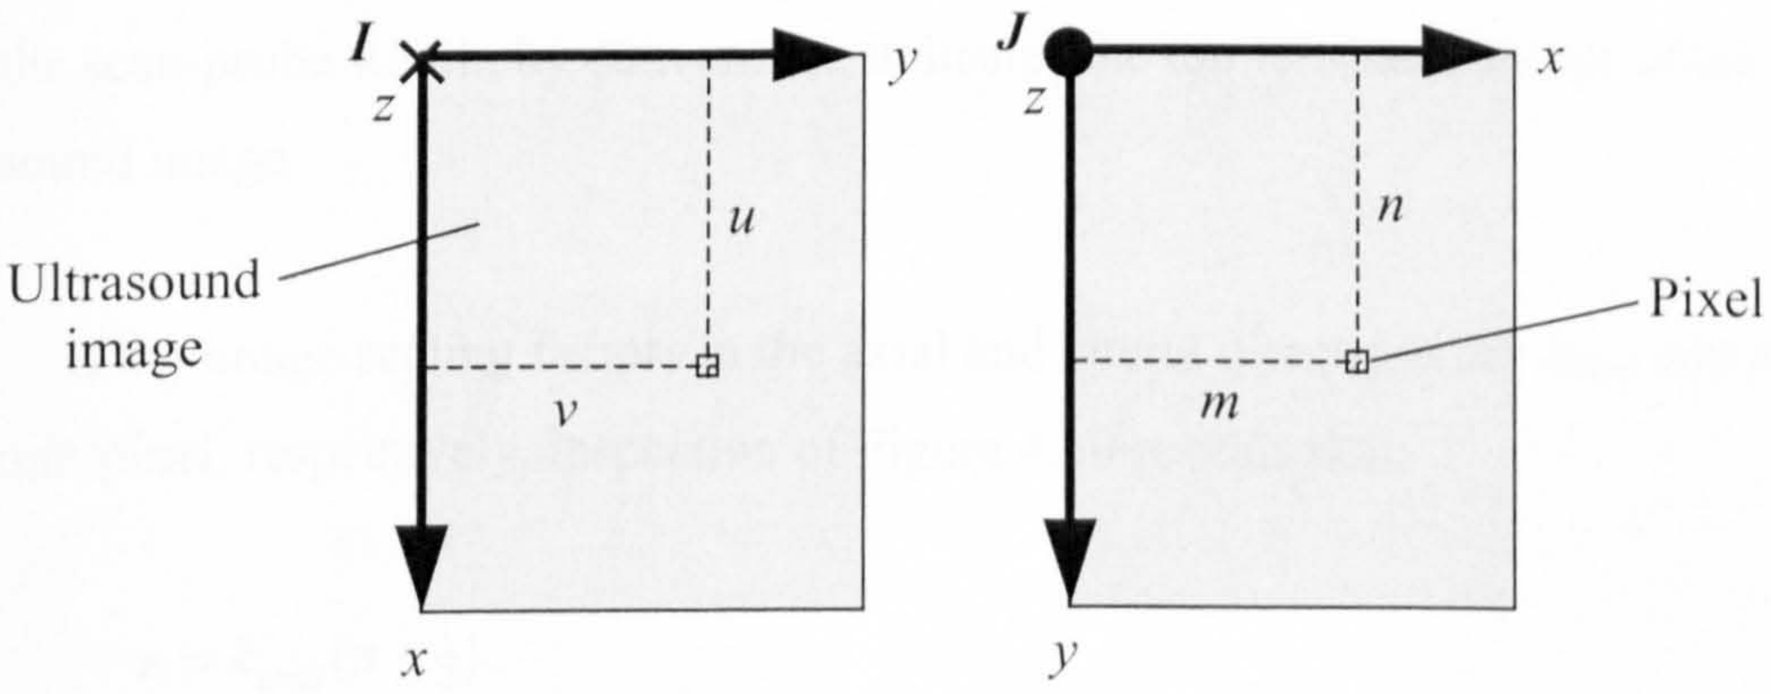


Figure 4.10 Two different co-ordinate systems for acquired images: the system on the left (***I***) is oriented to coincide roughly with the co-ordinate system for the pcBIRD receiver (***R***), whereas the system on the right (***J***) is more conventional for digital images. The z-axes for ***I*** and ***J*** point out of and into the page, respectively, indicating that both are left-handed co-ordinate systems (see Appendix 1).

The position of an image pixel, represented by the co-ordinates (u, v) with respect to $\boldsymbol{I}$, and where the units of u and v are mm, may be specified by a homogeneous position vector, $\mathbf{p}_I$, defined as[†]:

$$\mathbf{p}_I = \begin{bmatrix} u \\ v \\ 0 \\ 1 \end{bmatrix}. \tag{4.2}$$

Normally, pixels are identified by integer co-ordinates, say (m, n), which specifies a position in an M × N matrix representing a digital image. Thus, we define another co-ordinate system, $\boldsymbol{J}$, as shown in Figure 4.10, such that a pixel location is specified by the vector:

$$\mathbf{p}_J = \begin{bmatrix} m \\ n \\ 0 \\ 1 \end{bmatrix}, \tag{4.3}$$

where m and n are integers. It should be noted that, compared with the image shown in Figure 4.9, the image on the left in Figure 4.10 appears inverted. There is no inconsistency here since the origin of $\boldsymbol{I}$ is aligned with the image orientation marker on the scan-probe which, by convention, indicates the top left-hand corner of the ultrasound image.

If the image scaling factors in the axial and lateral directions are k_{axial} and k_{lateral} mm/pixel, respectively, inspection of Figure 4.10 reveals that:

$$u = k_{\text{axial}}\left(n + \tfrac{1}{2}\right), \tag{4.4}$$

and

$$\mathrm{v} = k_{\text{lateral}}\left(m + \tfrac{1}{2}\right). \tag{4.5}$$

[†] The last element of the vector is unity as this (*homogeneous*) form lends itself to a more concise representation of 3D co-ordinate transformations wherein a combined 3D rotation and translation are described by a single 4 × 4 matrix (see Appendix 1).

Combining Equations (4.2) – (4.5) gives:

$$\mathbf{p}_I = \begin{bmatrix} k_{\text{axial}}(n+\frac{1}{2}) \\ k_{\text{lateral}}(m+\frac{1}{2}) \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 & k_{\text{axial}} & 0 & \frac{1}{2}k_{\text{axial}} \\ k_{\text{lateral}} & 0 & 0 & \frac{1}{2}k_{\text{axial}} \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} m \\ n \\ 0 \\ 1 \end{bmatrix} = {}^{I}\mathbf{M}_{J}\,\mathbf{p}_{J}. \tag{4.6}$$

The matrix ${}^{I}\mathbf{M}_{J}$ in Equation (4.6) defines the simple transformation which converts the position vector, $\mathbf{p}_J$, defining a point specified with respect to J, to the position vector, $\mathbf{p}_I$, which specifies the same point with respect to I. Here, the notation ${}^{B}\mathbf{M}_{A}$ is used to denote a matrix transformation *from* co-ordinate system A *to* co-ordinate system B, and $\mathbf{p}_S$ denotes a vector which specifies the location of a point with respect to some co-ordinate system S.

In general, a rigid-body 3D transformation is composed of a 3D rotation and a 3D translation, and, as shown in Appendix 1, may be expressed by a 4 × 4 matrix of the form:

$${}^{B}\mathbf{M}_{A} = \begin{bmatrix} \cos\beta\cos\gamma & \sin\alpha\sin\beta\cos\gamma+\cos\alpha\sin\gamma & \sin\alpha\sin\gamma-\cos\alpha\sin\beta\cos\gamma & X_{AB} \\ -\cos\beta\sin\gamma & \cos\alpha\cos\gamma-\sin\alpha\sin\beta\sin\gamma & \cos\alpha\sin\beta\sin\gamma+\sin\alpha\cos\gamma & Y_{AB} \\ \sin\beta & -\sin\alpha\cos\beta & \cos\alpha\cos\beta & Z_{AB} \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{4.7}$$

where α, β, and γ are counter-clockwise rotations about the x-, y-, and z-axis of A, respectively; and X_{AB}, Y_{AB}, and Z_{AB} are translations in the directions of the x-, y- and z-axis of B, respectively.

A point, $\mathbf{p}_J$, in an ultrasound image corresponds to a point in the reconstructed volume. The 3D location of this point, with respect to the scanned volume co-ordinate system, V, is specified by the vector, $\mathbf{p}_V$, where the relationship between $\mathbf{p}_V$ and $\mathbf{p}_I$ is:

$$\mathbf{p}_V = {}^V\mathbf{M}_J . \mathbf{p}_J . \tag{4.8}$$

${}^V\mathbf{M}_J$ may be rewritten as the matrix product:

$${}^V\mathbf{M}_J = {}^V\mathbf{M}_T . {}^T\mathbf{M}_R . {}^R\mathbf{M}_I . {}^I\mathbf{M}_J , \tag{4.9}$$

where ${}^T\mathbf{M}_R$ in is computed from pcBIRD measurements, but ${}^V\mathbf{M}_T$, ${}^R\mathbf{M}_I$ and ${}^I\mathbf{M}_J$ are unknown and must be found in order that acquired images can be registered within V. As mentioned above, the orientation and position of the origin of V are somewhat arbitrary, and, therefore, V is initially chosen to be identical to T so that ${}^V\mathbf{M}_T$ becomes the identity matrix. Then the calibration problem is reduced to finding ${}^R\mathbf{M}_I$ and ${}^I\mathbf{M}_J$,

4.3.2 Review of calibration methods

The simplest method of calibrating a sensor-based, freehand 3DUS system is to measure the distance between the transducer surface and the physical centre of the POM sensor using a ruler [Hughes 1997; Hughes 1997]. This method is not very accurate, not least because the origin of the sensor co-ordinate system does not necessarily coincide with the physical centre of the sensor. Even if it does, the measurement of distances in the x, y, and z directions can be very difficult to determine accurately. The method also assumes that the axes of the receiver co-ordinate system are perfectly aligned with those of the image co-ordinate system, which may not be the case for the same reasons as above.

Several more advanced calibration methods are described in the literature [Leotta 1997; Barry 1997; Prager 1998; Detmer 1994; Birkfellner 1998b; Meairs 2000; Prager 1998]. All of these methods share a common procedure of scanning a phantom with known geometry from multiple views, and then solving an associated set of equations to estimate ${}^R\mathbf{M}_J$. Three examples of calibration phantoms are depicted in Figure 4.11. Detmer et al. [1994] used the cross-wire phantom (Figure 3.10(b)), while Leotta et al. [1997] used a 1.5 mm diameter brass bead suspended by a thread (Figure 3.10(a)). Meairs et al. [2000] describe a 'calibration box' which contains 4 wire intersections constructed in a tetrahedral arrangement. Prager et al. [1998] have proposed a novel calibration phantom, dubbed the 'Cambridge phantom', which

comprises a thin bar mounted on an axis connecting two wheels. The scan-probe is held in a support with slits on the lateral sides so that its motion is guided by the bar. This design ensures that the bar always appears in the ultrasound image as the phantom is moved around a water bath. The guided motion of the scan-probe also allows the bar to be imaged at varying depths. Another fast calibration method has been recently described by Pagoulatos et al. [2001], and uses a set of fish line strings to form fiducial markers in the shape of the letter 'N'. This method has the advantage that calibration can be performed from a single calibration image, rather than 30 or more typically required using other techniques.

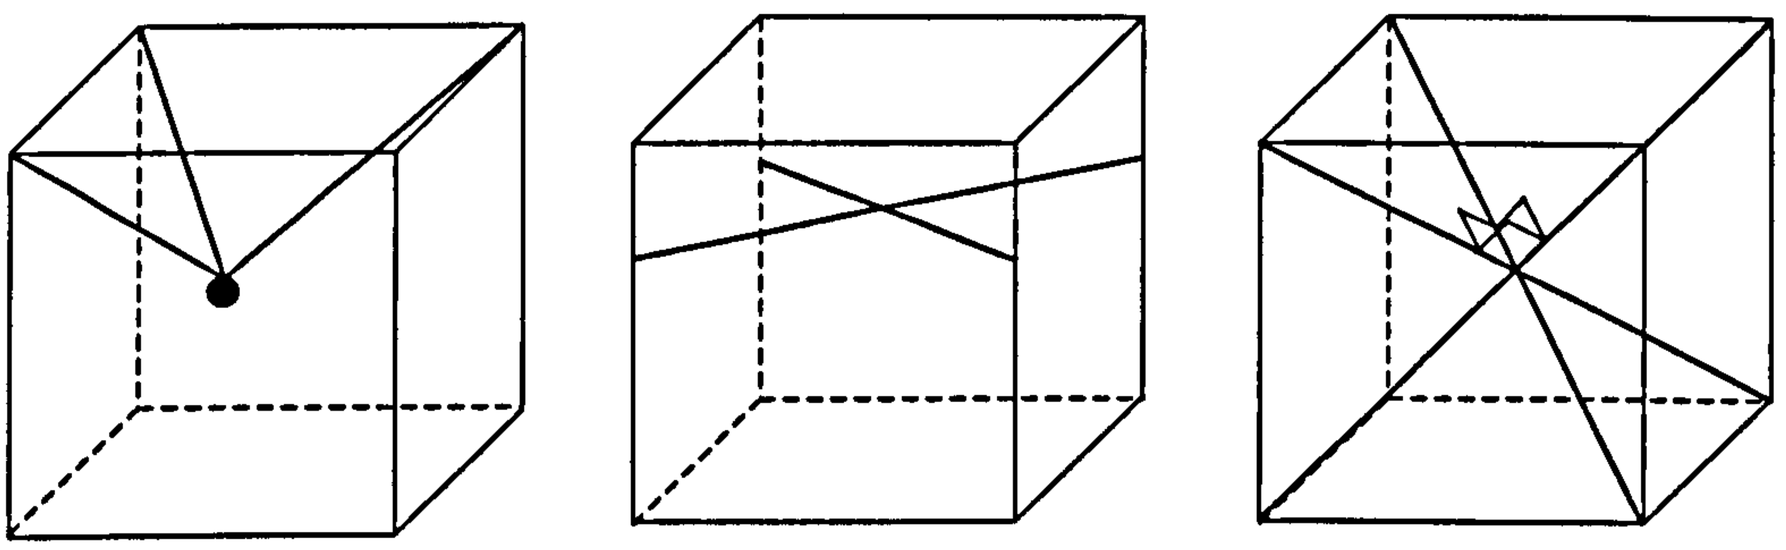


Figure 4.11 Examples of calibration phantoms for 3D freehand ultrasound acquisition systems: (a) a suspended bead; (b) two crossing wires, commonly referred to as the 'cross-wire' phantom; and (c) three orthogonal wires. All phantoms are immersed in a water-based fluid.

It is often difficult to compare the results of calibrations performed by different researchers due to the wide variety of methods and criteria used to assess performance. To address this problem, a study by Prager et al. [1998] compared 4 calibration phantoms using the same criteria. Twenty calibration experiments were carried out for each phantom, apart from the 3-wire phantom, for which 10 were performed. The following measures of performance were adopted:

(i) Quality of calibration data

The quality of calibration data was assessed in two ways: the first was to compare the number of calibration images used to estimate ${}^{R}\mathbf{M}_{J}$ with how many were collected during the calibration experiment. The second was to calculate the RMS error (in

mm) in target location over the B-mode images obtained for an individual calibration experiment. The way in which this is done is explained in Section 4.4.4.

(ii) Calibration precision

The calibration precision of each method was evaluated by taking the estimates of ${}^{R}\mathbf{M}_{J}$ for any pair of calibration experiments – referred to below by the subscripts 'trial 1' and 'trial 2' – and calculating the quantity:

$$|\Delta_R| = \left| \left({}^{R}\mathbf{M}_J\right)_{\text{trial 1}} (\mathbf{p}_J)_{\max} - \left({}^{R}\mathbf{M}_J\right)_{\text{trial 2}} (\mathbf{p}_J)_{\max} \right| , \tag{4.10}$$

with

$$(\mathbf{p}_J)_{\max} = \begin{bmatrix} m_{\max} \\ n_{\max} \\ 0 \\ 1 \end{bmatrix}, \tag{4.11}$$

which is the vector specifying the bottom right-hand pixel of an image, (m_{max}, n_{max}). The quantity $|\Delta_R|$ was calculated for all possible pairs drawn from all the calibration experiments. The results provide an indication of the variability in the position of the bottom right-hand pixel as a result of variability in ${}^{R}\mathbf{M}_{J}$ estimated from different calibration experiments.

(iii) Reconstruction precision

The reconstruction precision was evaluated by using all images of the cross-wire phantom, obtained during one calibration experiment. For each possible pair of images – denoted below by the subscripts 'image 1' and 'image 2' – the following quantity was calculated:

$$|\Delta_T| = \left| (\mathbf{p}_T)_{\text{image 1}} - (\mathbf{p}_T)_{\text{image 2}} \right|, \tag{4.12}$$

where

$$(\mathbf{p}_T)_{\text{image } i} = \left({}^{T}\mathbf{M}_R\right)_{\text{image } i} {}^{R}\mathbf{M}_J (\mathbf{p}_J)_{\text{image } i} , \; i \in \{1,2\}. \tag{4.13}$$

Here, $(\mathbf{p}_J)_{\text{image } i}$ specifies the location of the pixel closest to the centre of the target in image i, and ${}^{R}\mathbf{M}_J$ is the estimate obtained from a particular calibration experiment. In this case, the quantity $|\Delta_T|$ was calculated for each pair of images drawn from 49 cross-wire images; i.e. 1176 values of $|\Delta_T|$ were computed for each calibration trial.

Since the publication of the study by Prager et al. [1998], the three measures above have become regarded as the standard criteria for assessing the performance of a calibration method for freehand 3DUS systems. The results of the original study are summarised at the end of the chapter in Table 4.2 for the purpose of comparison with the measures computed for the system developed in this thesis. Prager an colleagues concluded that, overall, the cross-wire technique produced the best performance in terms of repeatability and accuracy, but the Cambridge phantom had the advantage of allowing a calibration to be performed rapidly with good performance.

The apparatus required to construct a calibration phantom is typically very simple. However, with the exception of the Cambridge Phantom and that described by Pagoulatos et al. [2001], the majority of calibration methods described in the literature are time consuming. This limits the frequency with which calibration can practically be repeated within the clinical environment, and raises the question: how often should a 3DUS system be calibrated? There is no obvious answer to this, but it seems sensible that calibration should be performed as regularly as possible.

4.3.3 Calibration using a pin phantom

The acquisition system was calibrated using the 'pin phantom' shown in Figure 4.12. The phantom was constructed from a stainless steel pin inserted into a rubber bung glued to the bottom of a water bath about 50 mm deep. The pin was held steady in an upright position during the calibration procedure, and the small pinhead formed a fixed point target, producing a bright reflection on a B-mode ultrasound image, as shown in Figure 4.13.

Calibration was performed by acquiring a series of B-mode images of the pinhead from a wide range of views. The position and orientation of the pcBIRD receiver was measured simultaneously with each image capture and recorded.

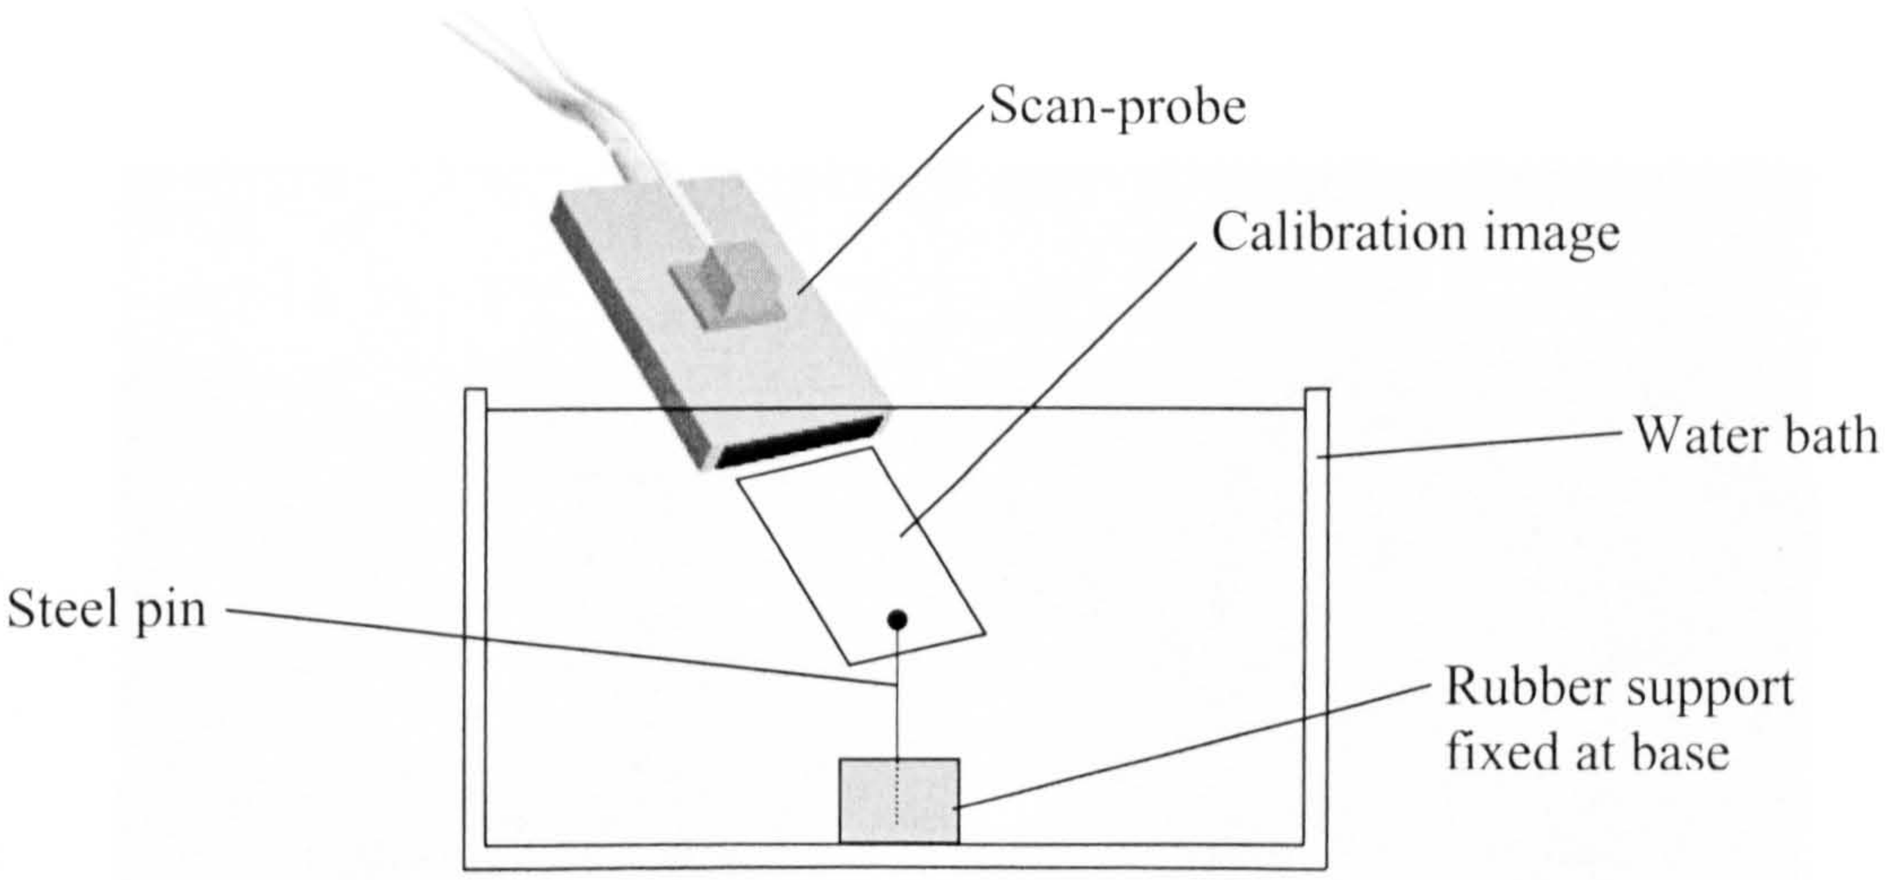


Figure 4.12 Calibration of the 3D ultrasound system using a pin phantom: the stationary pinhead is imaged in a water bath using multiple scan views.

Images of the pinhead, together with the corresponding pcBIRD data, were obtained using both the L10-5 and L12-5 scan-probes, and the following calibration imaging protocol was adopted:

- The scan depth was set to 40 mm with a single focal zone at a depth of 20 mm (these settings were considered to be the most appropriate for carotid scanning).

- 50 B-mode calibration images were obtained during each calibration experiment, using as wide a range of scan-probe positions and orientations as possible.

- the scan-probe was carefully tilted until the strongest reflection from the pinhead was found. This was to ensure that the pinhead was as near as possible to the centre of the ultrasound beam in the elevational direction.

- the probe was held stationary when each image capture was acquired.

For the L10-5 scan-probe, 12 calibration experiments were performed with the pin phantom placed at random positions within the region appropriate to carotid artery scanning, whereas, for the L12-5 scan-probe, 4 calibrations were performed.

A custom-written program was used to manually locate the pixel closest to the pinhead for each calibration image. This was done by placing a cross-hair cursor over the pinhead and recording the pixel co-ordinates, as illustrated in Figure 4.13.

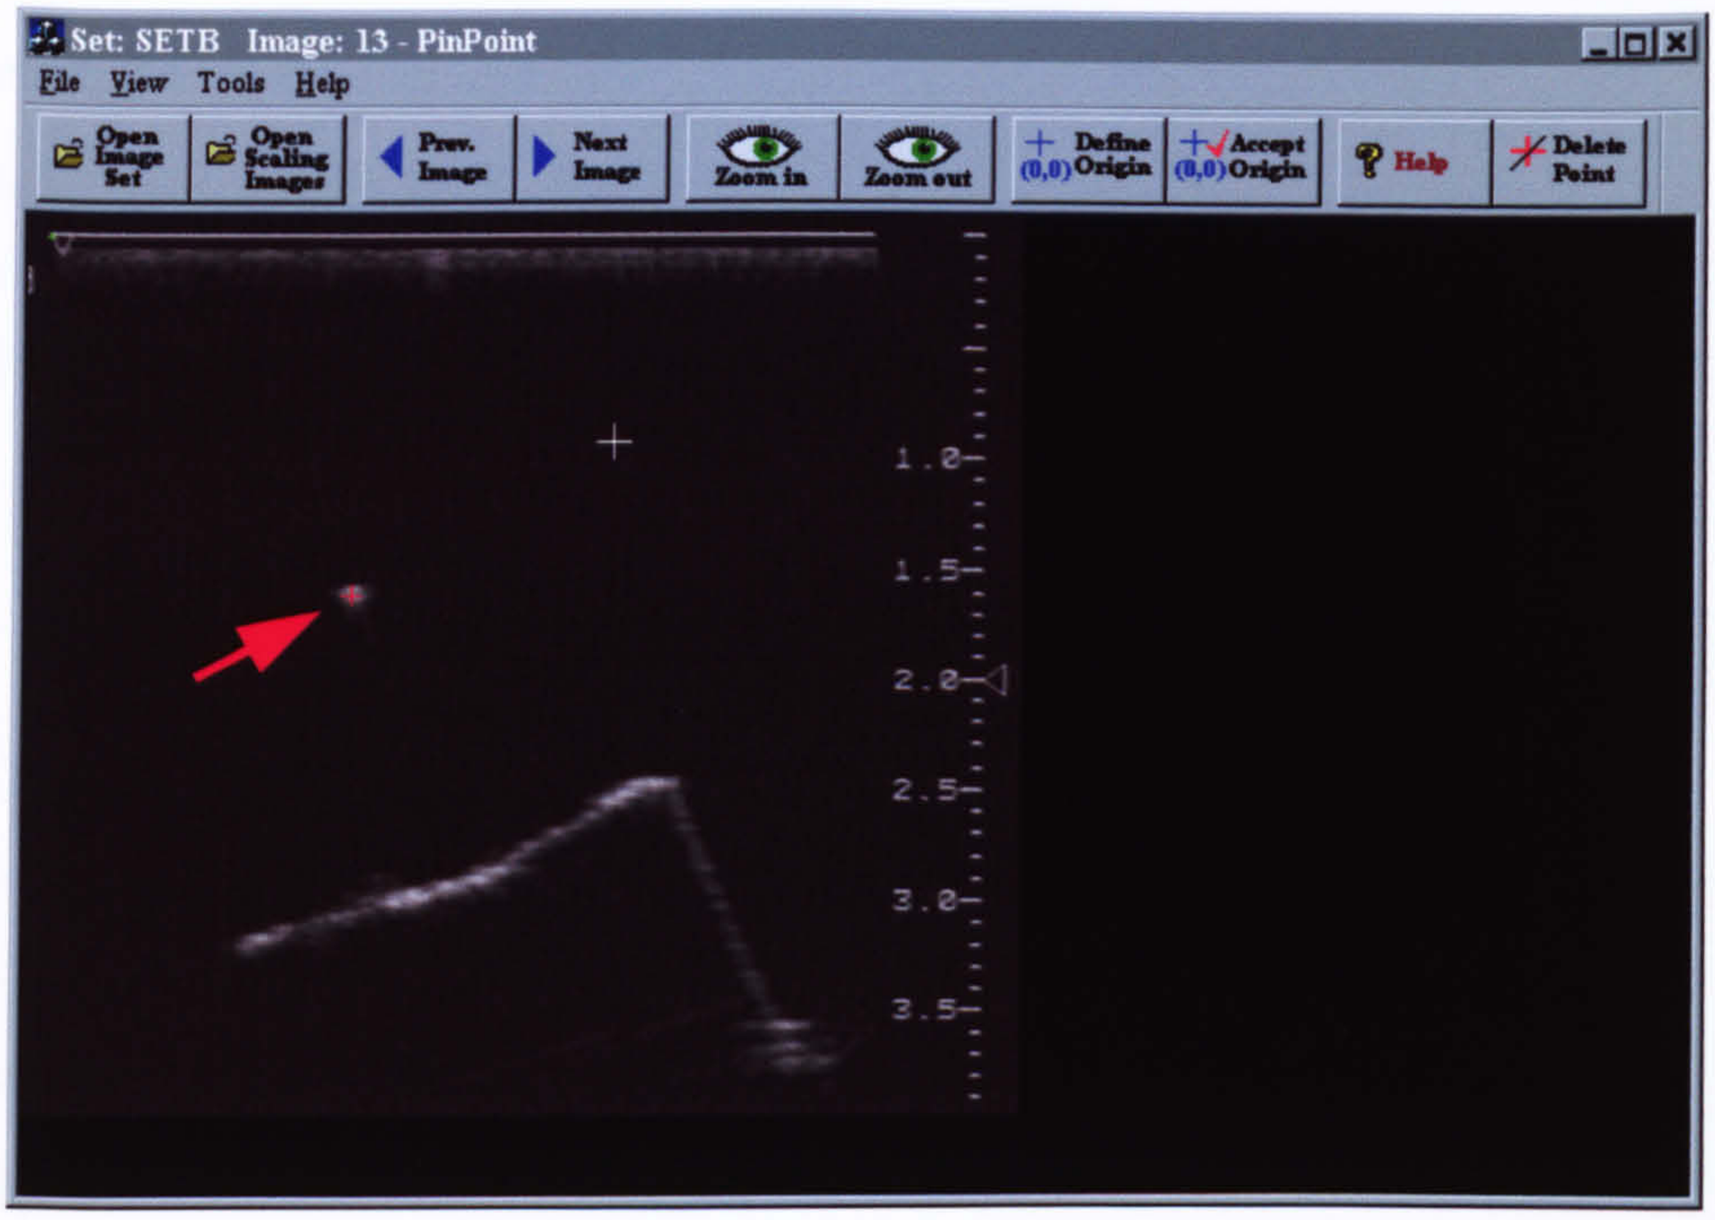


Figure 4.13 A B-mode image of the pin phantom. The position of the pinhead is manually located on each calibration image using the cross-hair cursor. The location pinhead is marked by a small red cross and indicated by the arrow.

The axial and lateral pixel scaling factors (k_{axial} and $k_{lateral}$, respectively – see Equations (4.4-6)) were determined using a custom-written HDILab plug-in to access image scaling information in the downloaded image files. A file containing the within-image pinhead co-ordinates (in mm) was generated for each image. Since the pin phantom was imaged in water, measurements in the axial direction were corrected to account for the difference in the speed of sound in water and soft tissue. The speed of sound in water, c_{water}, was calculated using the following formula [Lubbers 1998]:

$$c_{water} = 1404.3 + 4.7T - 0.04T^2, \tag{4.14}$$

where the units are m/s for c_{water} and °C for T, the temperature of the water bath.

The method used to estimate the parameters of ${}^{R}\mathbf{M}_{I}$ and ${}^{V}\mathbf{M}_{T}$ was similar to that described by Prager et al. [1998]. Combining Equations (4.7) and (4.8) gives:

$$\mathbf{p}_V = {}^{V}\mathbf{M}_T\,{}^{T}\mathbf{M}_R\,{}^{R}\mathbf{M}_I\,{}^{I}\mathbf{M}_J\mathbf{p}_J. \tag{4.15}$$

Unlike the method described by Prager et al., however, the parameters k_{axial} and k_{lateral}, and hence the matrix ${}^{I}\mathbf{M}_J$, were determined separately as explained above. Therefore, it is more convenient to rewrite Equation (4.15) in terms of $\mathbf{p}_I$ rather than $\mathbf{p}_J$ as follows:

$$\mathbf{p}_V = {}^{V}\mathbf{M}_T\,{}^{T}\mathbf{M}_R\,{}^{R}\mathbf{M}_I\mathbf{p}_I. \tag{4.16}$$

For the scans of the pin phantom, Equation (4.16) relates $\mathbf{p}_I$, the vector specifying the location of the pin, and $\mathbf{p}_V$, the physical 3D position of the pin with respect to the co-ordinate system V for the reconstructed volume. Because the pinhead remained stationary within V for each of the views performed during one calibration experiment, $\mathbf{p}_V$ is invariant over a single set of calibration images. If we force the origin of V to coincide with the centre of the pinhead then $\mathbf{p}_V$ becomes the (homogeneous) zero vector, and Equation (4.16) is rewritten as:

$${}^{V}\mathbf{M}_T\,{}^{T}\mathbf{M}_R\,{}^{R}\mathbf{M}_I\mathbf{p}_I = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \tag{4.17}$$

The equation holds for all the images obtained during a calibration experiment, and, therefore, each calibration image yields three non-linear equations. If the number of calibration images is N ($N = 50$ here) then applying Equation (4.17) to each image produces an over-determined system of $3N$ non-linear equations. This system may be solved using standard numerical methods, and the solution provides an estimate of the unknown parameters of ${}^{R}\mathbf{M}_I$ and ${}^{V}\mathbf{M}_T$. Setting the origin of V to coincide with the pinhead simplifies the solution significantly. ${}^{V}\mathbf{M}_T$ is later discarded, since it is dependent on the location of the pinhead relative to the origin of the pcBIRD transmitter for a particular calibration experiment.

For each calibration image, the matrix ${}^{T}\mathbf{M}_R$ is computed from the 3×3 rotation matrix $\mathbf{R}$ and receiver position (x, y, z) output by the pcBIRD as follows:

$$ {}^{T}\mathbf{M}_{R} = \begin{bmatrix} \mathbf{R}_{11} & \mathbf{R}_{21} & \mathbf{R}_{31} & x \\ \mathbf{R}_{12} & \mathbf{R}_{22} & \mathbf{R}_{32} & y \\ \mathbf{R}_{13} & \mathbf{R}_{23} & \mathbf{R}_{33} & z \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{4.18} $$

where $\mathbf{R}_{ij}$ is the element in the ith *row* and jth *column* of **R**. The matrices ${}^{R}\mathbf{M}_{I}$ and ${}^{V}\mathbf{M}_{T}$ have the general form given in Equation (4.6); i.e. each is a function of 3 unknown rotation angles – α, β, and γ – and 3 unknown translations – X, Y, and Z. If the 6 parameters of ${}^{R}\mathbf{M}_{I}$ and ${}^{V}\mathbf{M}_{T}$ are denoted by the subscripts 'RI' and 'VT', respectively, then the parameterisation of ${}^{R}\mathbf{M}_{I}$ and ${}^{V}\mathbf{M}_{T}$ is made explicit by writing:

$$ {}^{R}\mathbf{M}_{I} = {}^{R}\mathbf{M}_{I}(\alpha_{RI}, \beta_{RI}, \gamma_{RI}, X_{RI}, Y_{RI}, Z_{RI}), \tag{4.19} $$

and

$$ {}^{V}\mathbf{M}_{T} = {}^{V}\mathbf{M}_{T}(\alpha_{VT}, \beta_{VT}, \gamma_{VT}, X_{VT}, Y_{VT}, Z_{VT}). \tag{4.20} $$

With its origin co-incident with the pinhead, V can be placed at any orientation and still satisfy Equation (4.17). This means that the rotational parameters of ${}^{V}\mathbf{M}_{T}$ cannot be identified and therefore are set to zero for convenience. This leaves 9 identifiable parameters: α_{RI}, β_{RI}, γ_{RI}, X_{RI}, Y_{RI}, Z_{RI}, X_{VT}, Y_{VT}, and Z_{VT}. These parameters are the calibration parameters for the 3DUS system using a particular scan-probe, pcBIRD receiver mount, and pin-phantom position. For brevity, the angles α_{RI}, β_{RI}, and γ_{RI} will be denoted by α, β, and γ, respectively, for the remainder of this analysis.

Collecting all of the non-linear equations represented by Equation (4.17) and grouping for all N calibration images, we have:

$$ \mathbf{F}(\alpha, \beta, \gamma, X_{RI}, Y_{RI}, Z_{RI}, X_{VT}, Y_{VT}, Z_{VT}) = \mathbf{F}(\Phi) = \mathbf{0}, \tag{4.21} $$

where **F** is a $3N$ element, non-linear vector function, **0** is the $3N$ element zero vector, and

$$ \Phi = [\alpha \quad \beta \quad \gamma \quad X_{RI} \quad Y_{RI} \quad Z_{RI} \quad X_{VT} \quad Y_{VT} \quad Z_{VT}]^{T} \tag{4.22} $$

is a vector containing the calibration parameters. The calibration parameters were estimated by solving Equation (4.21) using the iterative modified-Marquardt algorithm described in Appendix 2. The algorithm rapidly converged to a solution within 5 iterations, indicating that the data from the calibration technique formed a well-conditioned system of equations. A measure of the 'quality' of the calibration (see Section 4.3.2) is the RMS error in pin location over all the calibration images. Expressing Equation (4.21) in terms of the components of **F** we have:

$$\mathbf{F}(\Phi) = \begin{bmatrix} f_1(\Phi) \\ f_2(\Phi) \\ \vdots \\ f_{3N}(\Phi) \end{bmatrix}, \tag{4.23}$$

where f_i, $1 \le i \le 3N$, is a (scalar) function of the calibration parameters. Now, if $\Phi_{\text{Marquardt}}$ denotes the vector containing the calibration parameters estimated using the modified-Marquardt algorithm, the RMS error is given by:

$$\sqrt{\frac{\sum_{i=1}^{3N} f_i(\Phi_{\text{Marquardt}})^2}{3N}}. \tag{4.24}$$

The RMS error will be zero for the hypothetical 'perfect' calibration.

The procedure detailed above was used to estimate the calibration parameters and standard performance measures defined by Prager et al. [1998] for each calibration experiment.

4.3.4 Calibration results

Tables 4.2 and 4.3 show the estimated calibration parameters and associated RMS error (calculated using Equation (4.24)) for the L10-5 and L12-5 scan-probes. Table 4.4 summarises the results for the calibration quality, calibration precision, and reconstruction precision, together with the results of comparable studies by Prager et al. [1998] and Meairs et al. [2000]. It can be seen that the results of calibration using

the pin phantom compare very favourably with those reported by the other researchers using the same performance criteria.

The study by Prager et al. [1998] used a Toshiba SSA-270A/HG scanner, a 7 MHz linear array scan-probe, a scan-depth of 40 mm, and an AC EPOM device (the Polhemus *Fastrak*). The receiver was mounted on a rigid bar approximately 250 mm from the transducer surface. Meairs et al. [2000] used the ATL HDI-5000 and L12-5 scan-probe with a scan-depth of 35 mm (the same as this thesis). A pulsed-DC EPOM device with a similar specification to the *pcBIRD* (the Ascension *Flock of Birds*) was used in this study.

The estimated calibration parameters were averaged over all the trials to produce a composite estimate of ${}^{R}\mathbf{M}_{I}$. The averaged calibration matrices were:

$$ {}^{R}\mathbf{M}_{I} = \begin{bmatrix} 1.00 & 0.03 & -0.02 & 34.25 \\ -0.03 & 1.00 & -0.03 & -21.65 \\ 0.02 & 0.03 & 1.00 & 28.57 \\ 0 & 0 & 0 & 1 \end{bmatrix} $$

for the L10-5 scan-probe, and

$$ {}^{R}\mathbf{M}_{I} = \begin{bmatrix} 1.00 & -0.01 & -0.01 & 27.62 \\ 0.01 & 1.00 & -0.01 & -20.27 \\ 0.01 & 0.01 & 1.00 & 26.97 \\ 0 & 0 & 0 & 1 \end{bmatrix} $$

for the L12-5 scan-probe.

4.4 Discussion

The 3DUS acquisition system in this chapter combines the advantages of freehand acquisition with digital image capture. There is no image degradation in the image capture process, which is a significant advantage over video-based systems, and the fact that pre-scan-converted ultrasound data are stored allows some image param-

(*continued on page 134*)

Parameter ↓	Calibration trial (L10-5 scan-probe)												
	1	2	3	4	5	6	7	8	9	10	11	12	Mean ± SD
X_{VT} (mm)	-241.53	-241.69	-241.32	-241.24	-264.12	-264.30	-264.28	-264.13	-207.03	-207.05	-207.05	-205.62	-
Y_{VT} (mm)	-19.06	-18.73	-18.66	-18.35	-43.34	-43.30	-43.24	-43.31	32.57	32.29	32.80	33.46	-
Z_{VT} (mm)	-88.59	-88.49	-88.60	-89.02	-93.57	-93.72	-93.72	-94.54	-86.35	-86.34	-86.78	-86.69	-
X_{RI} (mm)	34.37	34.85	35.25	35.37	33.97	33.91	34.04	34.90	33.47	33.48	33.60	33.79	**34.25 ± 0.68**
Y_{RI} (mm)	-21.08	-21.48	-22.02	-22.01	-20.97	-21.37	-21.32	-21.28	-22.22	-22.27	-21.72	-22.00	**-21.65 ± 0.45**
Z_{RI} (mm)	28.60	28.68	29.02	27.80	27.22	28.31	28.94	29.23	28.39	28.71	29.66	28.30	**28.57 ± 0.64**
α (°)	-1.12	-1.01	-0.60	-2.30	-2.73	-2.26	-1.19	-1.01	-1.89	-1.56	-0.84	-2.06	**-1.55 ± 0.68**
β (°)	1.10	1.17	-0.10	1.85	3.78	1.28	0.25	1.07	1.17	1.25	-1.48	1.08	**1.03 ± 1.23**
γ (°)	3.15	1.43	0.72	1.03	2.90	2.24	1.31	1.66	2.04	1.55	2.65	2.11	**1.90 ± 0.75**
RMS error (mm)	0.73	0.72	0.72	0.52	0.77	0.70	0.77	0.76	0.69	0.63	0.68	0.61	**0.69 ± 0.08**

Table 4.2 Calibration parameter estimates for the L10-5 scan-probe.

Parameter ↓	Calibration trial (L12-5 scan-probe)				
	1	2	3	4	Mean ± SD
X_{VT} (mm)	-332.43	-332.48	-432.33	-432.46	-
Y_{VT} (mm)	-11.67	-11.56	7.59	7.51	-
Z_{VT} (mm)	-77.53	-78.33	-81.61	-81.44	-
X_{RI} (mm)	26.60	27.56	28.01	28.31	**27.62 ± 0.75**
Y_{RI} (mm)	-20.02	-19.98	-20.19	-20.87	**-20.27 ± 0.41**
Z_{RI} (mm)	26.23	25.86	27.46	28.33	**26.97 ± 1.14**
α (°)	-1.85	-1.81	0.56	1.96	**-0.29 ± 1.87**
β (°)	0.41	2.52	0.39	-0.44	**0.72 ± 1.26**
γ (°)	-0.06	0.28	-0.23	-1.67	**-0.42 ± 0.86**
RMS error (mm)	0.69	0.64	0.74	0.84	**0.73 ± 0.09**

Table 4.3 Calibration parameter estimates for the L12-5 scan-probe.

		This thesis		Prager et al [1998]				Meairs et al [2000]*
		L10-5	L12-5					
	Phantom	**Pin**	**Pin**	**Cross-wire**	**Three-wire**	**Single-wall**	**Cambridge**	**Mannheim**
Calibration quality	**No. trials**	12	4	20	10	20	20	-
	RMS error (mm)	0.69	0.73	0.56	1.04	0.48	0.34	-
Calibration precision	**No. observations**	66	6	190	45	190	190	120
	Mean (mm)	1.36	1.53	1.47	5.37	3.27	0.92	1.81
	Max (mm)	2.63	2.33	3.37	12.09	10.07	2.32	3.37
Reconstruction precision	**No. observations**	13475	3105	23520	11760	23520	23520	15840
	Mean (mm)	1.89	1.82	1.65	2.67	3.43	2.17	2.11
	Max (mm)	5.58	4.41	4.78	10.31	11.67	7.13	6.01

Table 4.4 Standard measures of calibration performance calculated for the L10-5 and L12-5 scan-probes. The results obtained by Prager et al. [1998] and Meairs et al. [2000] are included in the table for comparison. *In the published paper by Meairs et al. [2000] the measures of 'calibration precision' and 'reconstruction precision', as defined originally by Prager et al. [1998], are referred to incorrectly as 'quality of calibration' and 'calibration precision'. This error has been corrected in the data reproduced above.

eters, such as the TGC settings and colour maps, to be modified off-line using the HDILab software. This is a useful feature as it enables acquired images to be optimised during the pre-reconstruction phase by, for example, the removing colour artefacts, or enhancing the contrast between plaque and lumen.

As indicated in Chapter 3, one of the principal advantages of freehand systems is that there is no constraint on the movement of the scan-probe. This allows the view angle to be altered during acquisition in order to optimise visualisation. Many freehand systems also operate using existing transducer technology. The use of an inexpensive EPOM device results in a relatively low-cost adaptation to a standard clinical ultrasound scanner. Currently, 3DUS is not used for routine clinical vascular investigations, and such a low-cost 3DUS capability is highly attractive, particularly as use of the scanner is preserved. Furthermore, since the EPOM receiver attached to the scan-probe is unobtrusive, and may be removed if necessary, switching between routine 2D and 3D imaging is very quick and simple.

Because a clinical scanner was used in this project, emerging 2D imaging modes, such as compound B-mode and tissue harmonic imaging, are supported by the 3DUS system. Such modes are usually not available on current commercial dedicated 3DUS systems, and image quality, particularly in colour or power Doppler modes, is often inferior to that achievable with high-end 2D scanners. Image quality is an important issue as, ultimately, the quality of 3D reconstruction depends heavily on the quality of the original acquired slices.

Another advantage of freehand systems is flexibility – they can be adapted to work with virtually any commercial scanner. This advantage has been sacrificed in the system described here, since digital image capture and transfer is presently only available on a few commercial scanners. However, there is a general trend towards introducing digital image archiving on ultrasound scanners, aided by the standardised DICOM protocol. Therefore, it is anticipated that soon it will possible to take advantage of digital capture facilities on a number of commercial systems. This will extend the compatibility of the 3DUS system described here to other scanners.

The current version of the 3DUS system incorporates an external ECG monitor to enable ECG-gated acquisition. It is hoped that future incarnations of the system will be able to access the ECG signal produced by the scanner. This would obviate the

need for the ECG monitor which would, in turn, reduce the total component cost, and make the system more integrated. The decision to use an EPOM device with a controller housed within the host PC, rather than in an external box, was made in effort to make the system as self-contained as possible. This has distinct practical advantages. It may also prove advantageous to implement respiratory gating in addition to ECG-gating to further reduce motion artefacts in 3DUS acquisitions.

One of the major concerns about electromagnetic 3DUS systems has been accuracy, based on the well-documented sensitivity of EPOM device to interference, particularly due to metal present in the operating environment. This issue is considered in detail in the next chapter.

The pin phantom provides a simple and effective means of calibrating the 3DUS system. The performance of this calibration method compares very favourably with other methods reported in the literature. Although the calibration procedure is quite time-consuming – taking up to 2 hours – this is not prohibitive if it is performed just a few times per year. So far it has not been considered necessary to adopt a faster calibration method, such as that described by Prager et al. [1998] or [Pagoulatos 2001]. Such methods may be implemented in the future, however, as they enable calibration to be practically performed on a much more regular basis.

Chapter 5

Accuracy of 3D ultrasound acquisition

There are several sources of error that can arise in freehand 3DUS acquisition systems. Errors result from the misregistration of acquired 2D ultrasound images, such that pixels are positioned incorrectly in the reconstructed volume. The purpose of this chapter is to examine each of the sources of error relevant to the acquisition system described in the previous chapter. The relative contributions of each error source to the uncertainty in pixel registration are predicted using a simple mathematical model of the reconstruction process. This is followed by a description of the experimental optimisation and evaluation of the accuracy of the electromagnetic tracking device used in the system. Finally, the results of in vitro experiments aimed at estimating the reconstruction accuracy of the system are presented.

5.1 Sources of pixel misregistration

5.1.1 Position and orientation measurement device errors

A primary source of 3D pixel misregistration in freehand 3DUS systems is error in POM device measurements. Table 4.1 gives an indication of the errors expected for the *pcBIRD* under ideal conditions. However, these figures alone do not provide information on what the expected errors in pixel misregistration are. To obtain this information, further analysis is required.

Position measurement errors propagate through the imaging system in a different way to orientation measurement errors: position measurement errors produce an identical error in all pixels in an acquired image. The error may be expressed as the 3D vector:

$$\mathbf{e} = \begin{bmatrix} e_x \\ e_y \\ e_z \end{bmatrix}_T \tag{5.1}$$

where e_x, e_y and e_z are the component errors in the x-, y- and z- directions, respectively, and the subscript 'T' indicates that the vector is measured with respect to the EPOM transmitter co-ordinate system. The resultant error vector in the reconstructed 3D pixel position is then also **e**. Therefore,

$$\mathbf{p}_V = (\mathbf{p}_V)_{\text{true}} + \mathbf{e}, \tag{5.2}$$

where $(\mathbf{p}_V)_{\text{true}}$ specifies the correct 3D position of a pixel.

Rotational errors in orientation measurements, on the other hand, do not affect all pixels in an image in the same way, and result in a different misregistration error depending on where a pixel is within the image. The reason for this is the action of the lever-arm effect, illustrated in Figure 5.1. In the figure the points O and P are connected by a rigid rod with length r. The orientation of the rod OP is defined by the angle θ, relative to the horizontal. If r is known, but measurement of the angle θ is subject to an error, say $\delta\theta$, then the measured angle, $\theta' = \theta + \delta\theta$, specifies the incorrect orientation OP'. So, if we use the angle θ' to then calculate the position of P, we

arrive at the incorrect location P'. The misregistration error is the distance PP', denoted by δP. Using simple geometry, δP is given by:

$$\delta P = 2r\sin\left(\frac{\delta\theta}{2}\right) \approx r\delta\theta\,, \tag{5.3}$$

when $\delta\theta$ is small. Inspection of this expression indicates that the misregistration error is proportional to r – the distance between the true point P and the centre of rotation O.

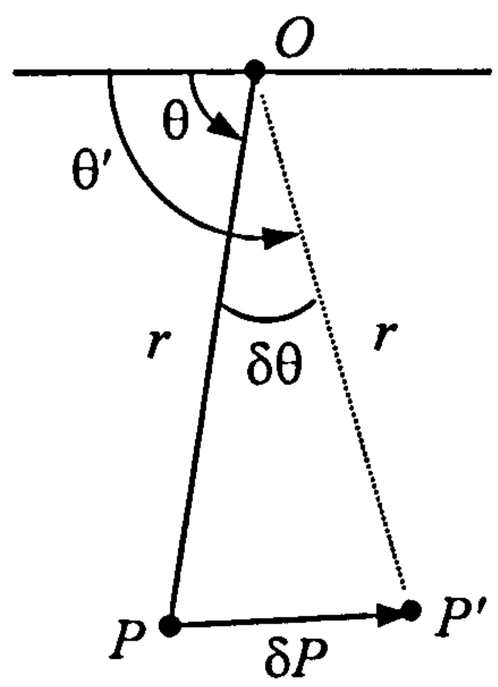


Figure 5.1 The lever-arm effect: a small error, $\delta\theta$, in the measurement of the orientation of OP results in the true angle θ being measured as $\theta' = \theta + \delta\theta$. If θ' is then used to calculate the position of the point P (connected to O by a rigid bar as shown), the point is located incorrectly at P'. The registration error for P is then equal to the length of PP' (δP).

As a direct consequence of the lever-arm effect, a small error in the orientation measurement of the EPOM receiver will produce a misregistration error proportional to the perpendicular distance between the axis of rotation and an image pixel. It is easiest to visualise the effect for a rotation about the receiver's y-axis, which gives rise to an error in the elevational direction and increases with distance in the axial direction. However, it is important to note that the lever-arm effect acts about *all three axes of rotation*, resulting in an error vector with components in the lateral, elevational, and axial directions of the image. The error is also dependent on the position of the receiver relative to the transducer surface. It is because of the lever-arm

effect that the EPOM receiver was mounted on the scan-probe as close as possible to transducer surface without disrupting normal scanning technique.

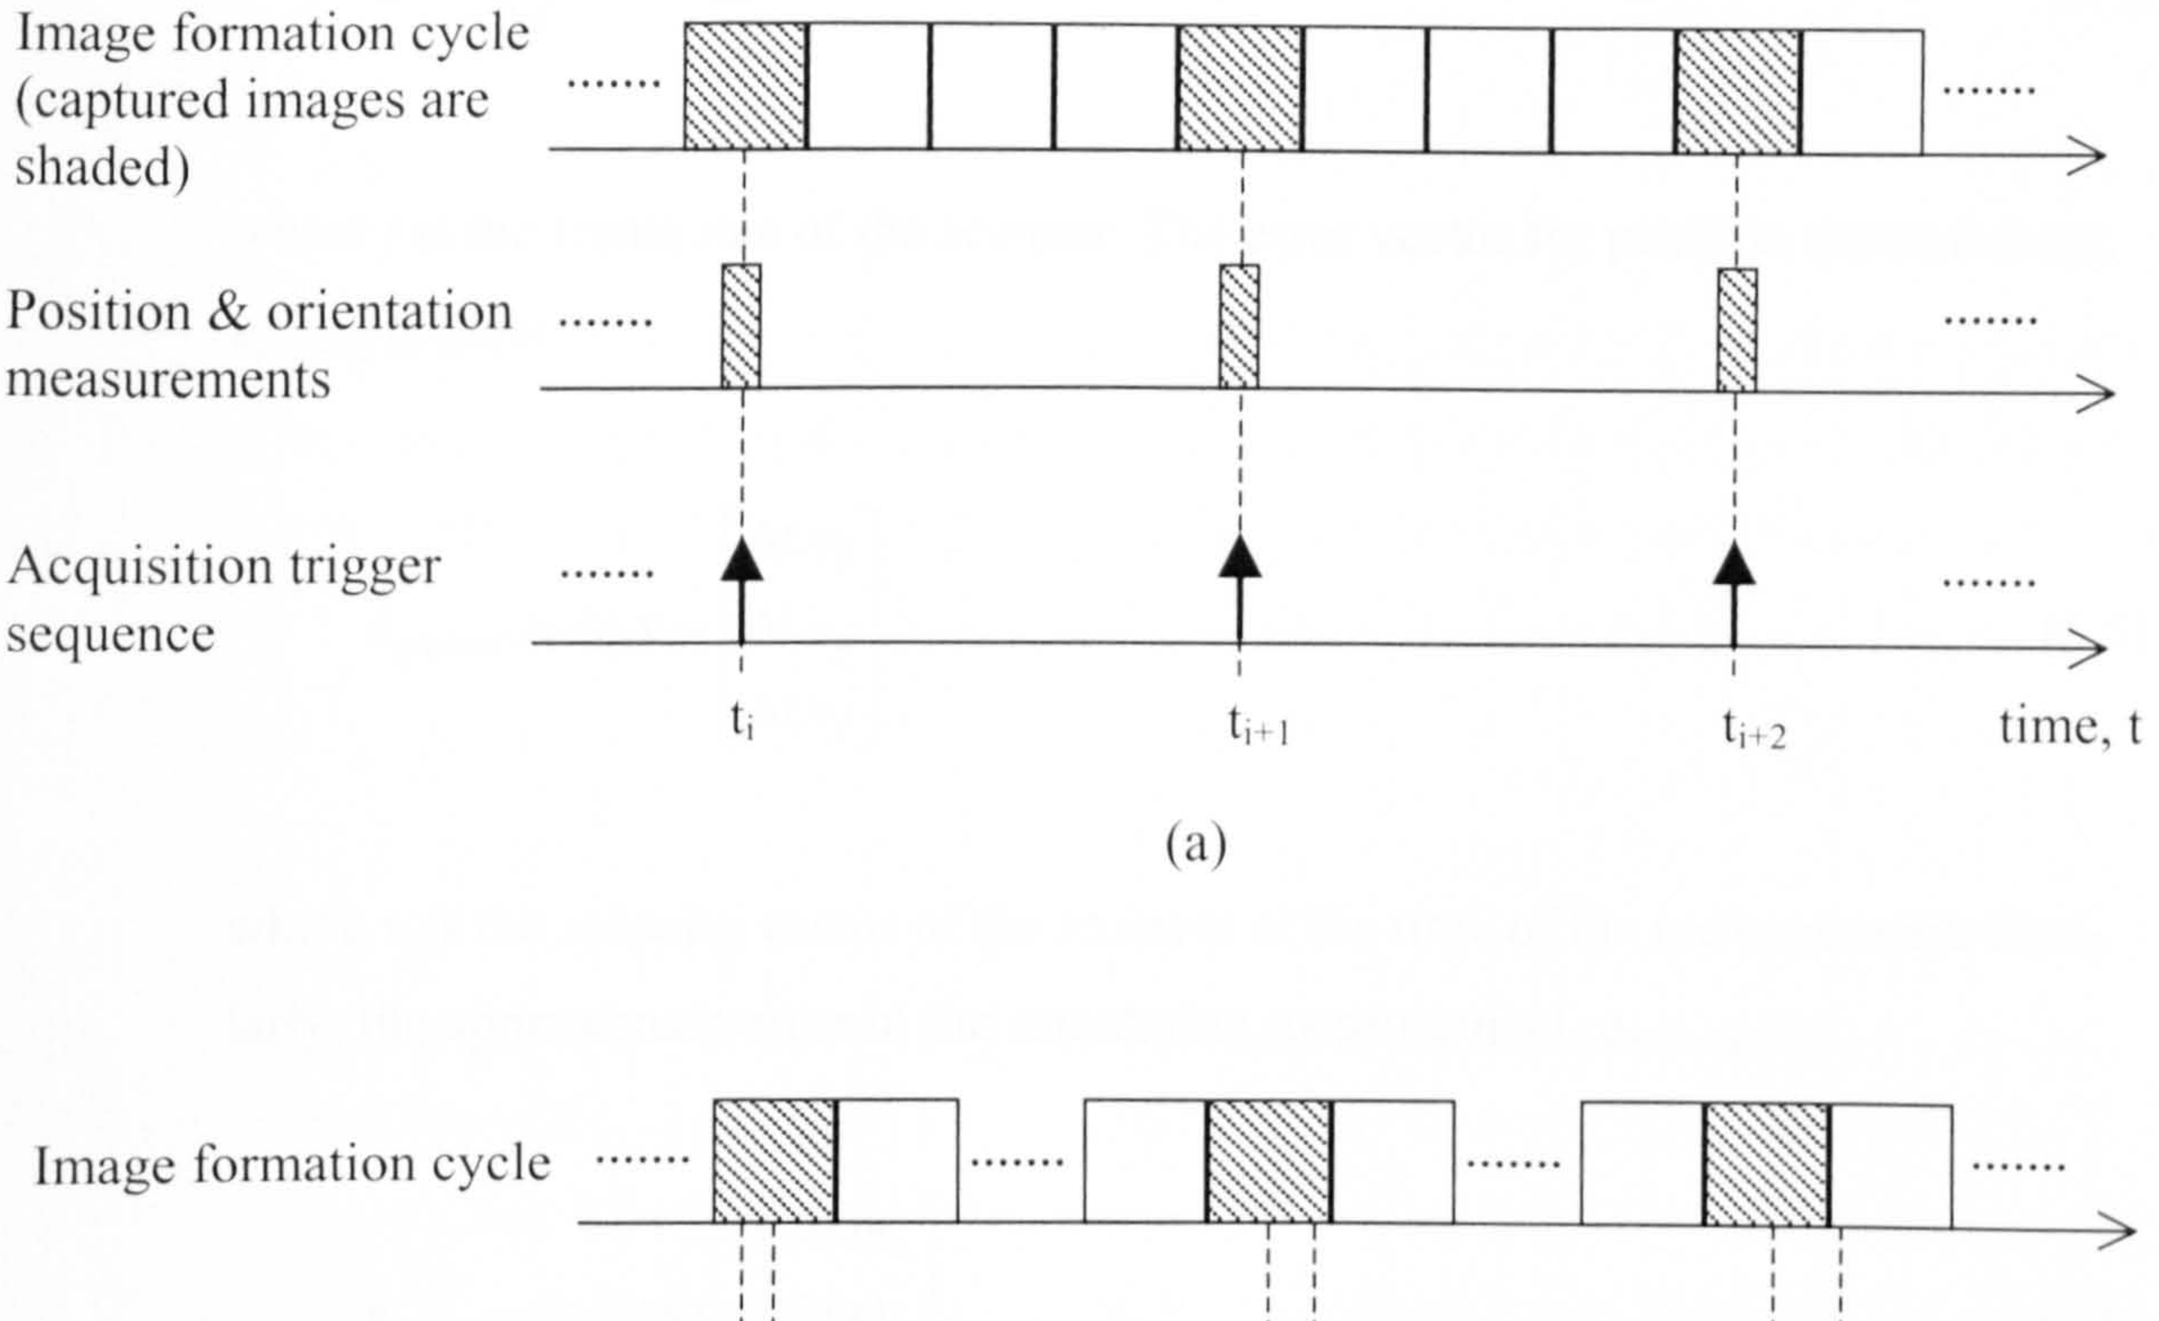
5.1.2 Temporal misregistration

In the 3DUS system described in Chapter 4, the capture of an ultrasound image was synchronised with the measurement of the current position and orientation of the EPOM receiver. Synchronisation is necessary to ensure that the EPOM device measurement is a good estimate of the true position and orientation of the scan-probe when an image is captured. For the purposes of this discussion, it will be supposed that these two events are synchronised perfectly if the 'effective' time of the position and orientation measurement – i.e. the midpoint of the measurement period – coincides with the midpoint of the period between the display of consecutive images[†]. Such perfect synchronisation is illustrated in Figure 5.2(a). When ECG-gating is employed in practice, this idealised synchronisation cannot occur because there is the ECG signal, and hence the acquisition trigger cycle, and the continuous image formation cycle are uncorrelated. It can be seen from Figure 5.2(b) that the EPOM device measurement and image acquisition are no longer synchronised in the sense defined above; in other words, there is a degree of asynchrony between the two events, which varies with each acquired image. This phenomenon is known as *temporal misregistration* and will be quantified here by the random variable, Δt.

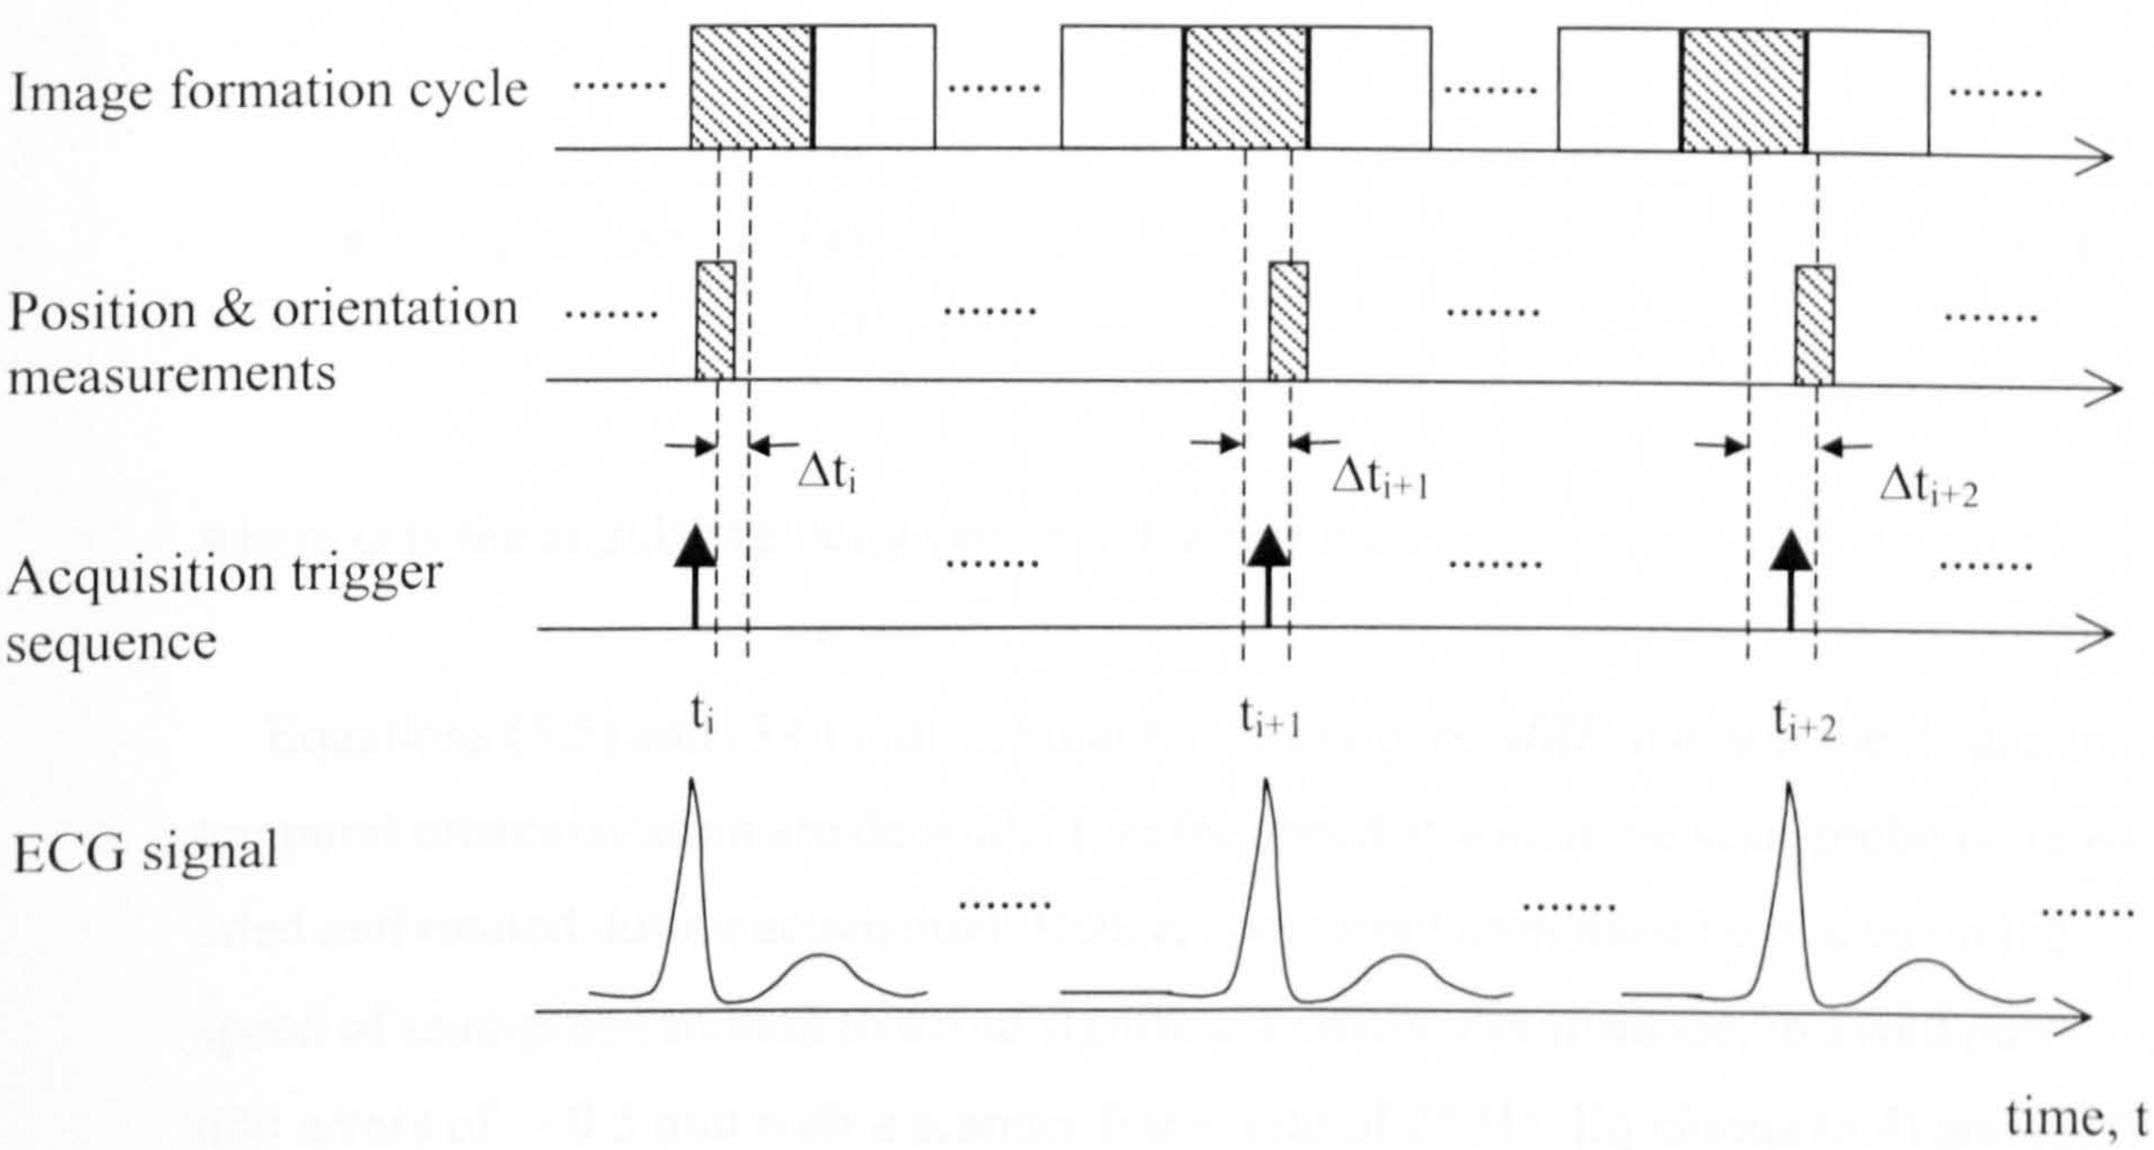
Because the scan-probe is moving as images are acquired, temporal misregistration also gives rise to spatial misregistration in acquired images. The value of Δt is dependent on the timing of the acquisition trigger within the image formation cycle and the finite delay between the initiation of an EPOM device measurement and the effective time of the measurement. Data supplied by the device's manufacturer shows that this delay is < 10 ms for the *pcBIRD*. Therefore, a worst case measure of temporal misregistration is given by:

(*continued on page 141*)

[†] Modern ultrasound scanners use multi-channel processing so that most of the time between the display of consecutive images is spent collecting echo data.



(a)



(b)

Figure 5.2 Relative timings of the 3DUS acquisition cycle: (a) the idealised case where image capture and position and orientation measurement are perfectly synchronised; (b) the real case where acquisition is ECG-gated. The random variable Δt provides a measure of the asynchrony between image capture and position and orientation measurement. For simplicity, the peak of the ECG R-wave is used as the gate.

$$\Delta t < \frac{1}{2f} + 10\,\text{ms}, \tag{5.4}$$

where f is the frame rate of the scanner. The error vector for position measurement, $\mathbf{e}_{\text{position}}$, is then:

$$\mathbf{e}_{\text{position}} \approx \Delta t.\mathbf{v} = \begin{bmatrix} \Delta t.v_x \\ \Delta t.v_y \\ \Delta t.v_z \end{bmatrix}, \tag{5.5}$$

where $\mathbf{v}$ is the velocity vector of the receiver at the time of the measurement. Similarly, the approximate error in the orientation measurement, $\mathbf{e}_{\text{orientation}}$, is:

$$\mathbf{e}_{\text{orientation}} \approx \Delta t.\boldsymbol{\omega} = \begin{bmatrix} \Delta t.\omega_x \\ \Delta t.\omega_y \\ \Delta t.\omega_z \end{bmatrix}, \tag{5.6}$$

where $\boldsymbol{\omega}$ is the angular velocity vector of the receiver.

Equations (5.5) and (5.6) indicate that the errors in *pcBIRD* measurements due to temporal misregistration are dependent on the speed at which the scan-probe is translated and rotated during acquisition. Hence, some limitation must be placed on the speed of scan-probe motion to avoid significant errors; For instance, to avoid position errors of > 0.5 mm with a scanner frame rate of 25 Hz. Equations (5.4) and (5.5) suggest that:

$$\Delta t.|\mathbf{v}| = 0.03|\mathbf{v}| \leq 0.0005\,\text{m}. \tag{5.7}$$

Rearranging gives:

$$|\mathbf{v}| \leq 0.0167\,\text{m/s} = 16.7\,\text{mm/s}. \tag{5.8}$$

Such limitation on the speed of scan-probe translation does not present a difficulty for normal scanning, and, with practice, it is quite feasible to perform slow sweeps at

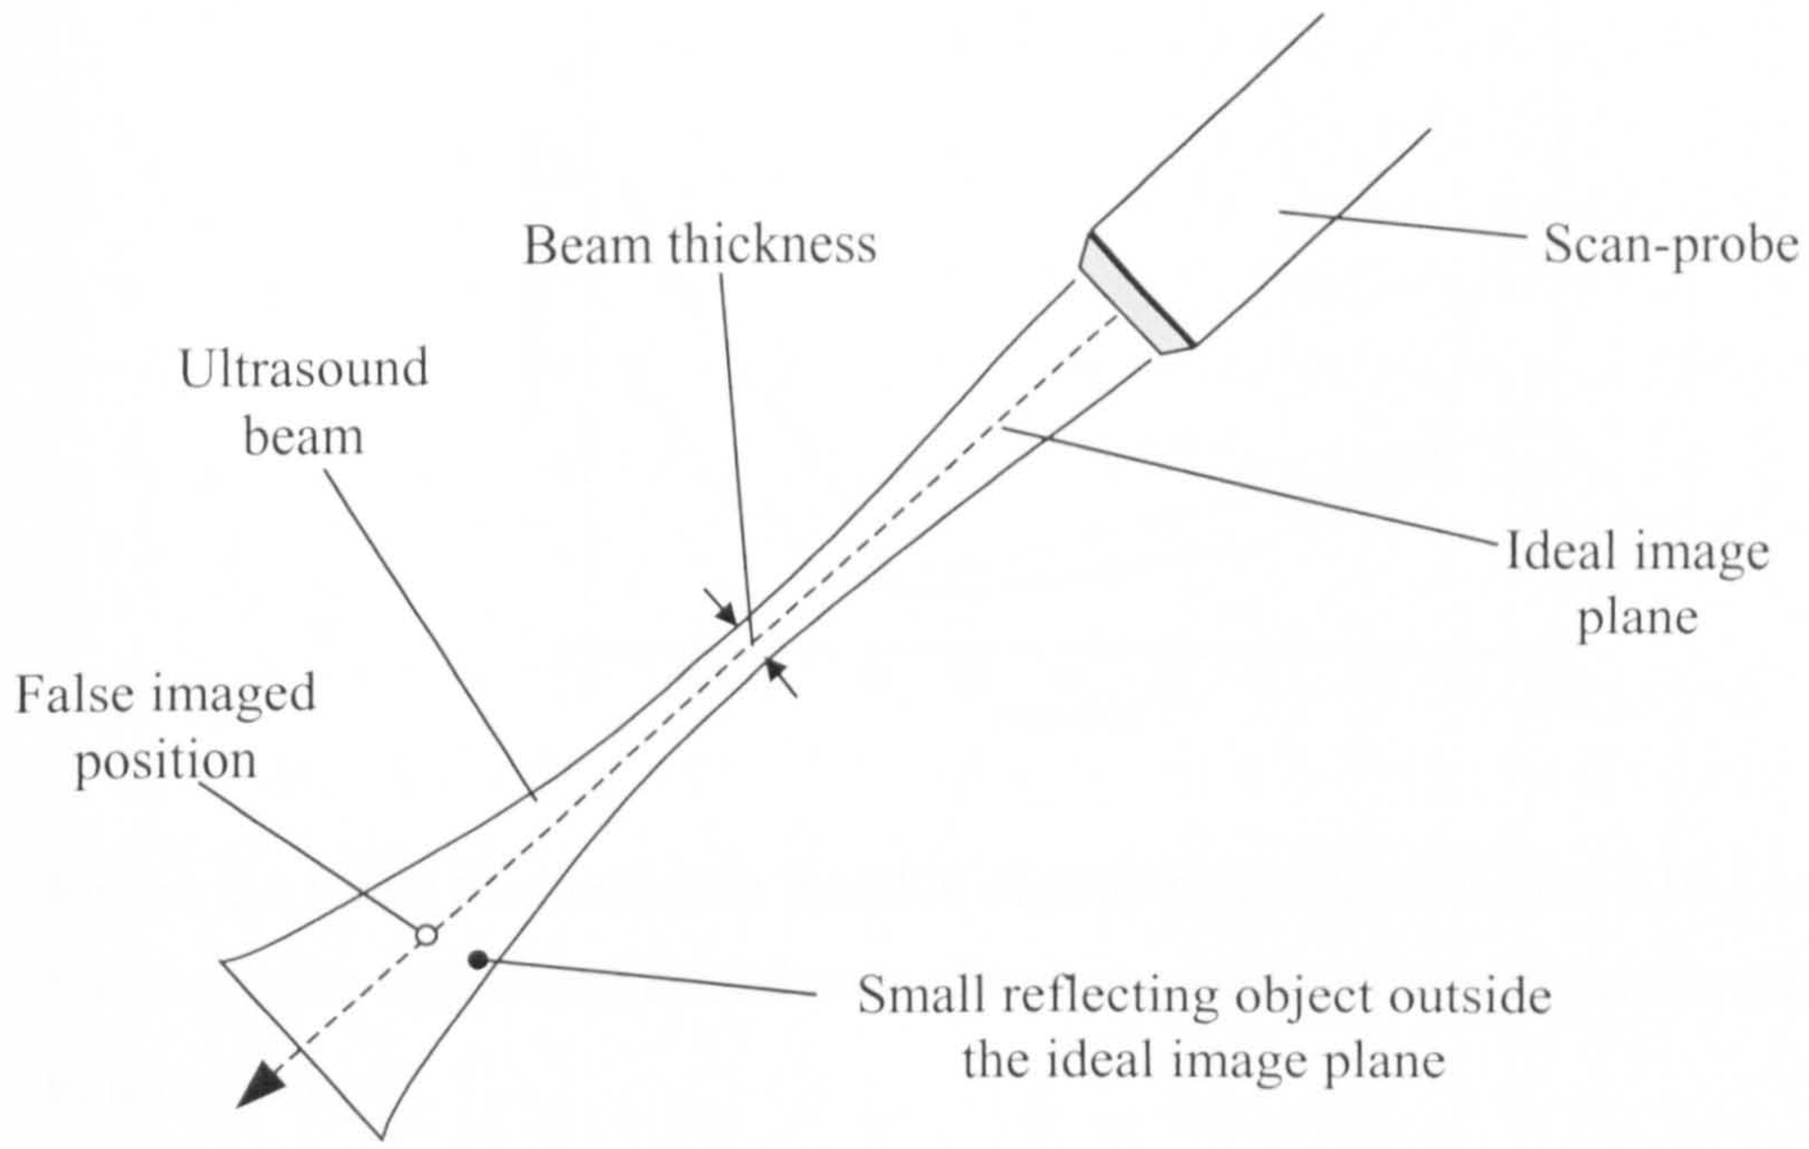
a speed of less than a few mm/s. Furthermore, at the high frame rates achievable with modern ultrasound scanners (> 50 Hz), errors due to temporal misregistration become negligible.

Unlike images captured from an analogue video output, the images captured using digital image capture are those currently being formed at the time the capture trigger is detected. This is not generally the case for a video-based system because the frame rate of the video output is usually fixed at 25 Hz, and will differ from the frame rate of the scanner. When the frame rate of the scanner exceeds 25 Hz, an additional delay is introduced and the image output from the video port may lag the image being currently formed by multiple imaging cycles. This effect, together with other possible delays introduced by the video electronics in the scanner, gives rise to significant uncertainty in the synchronisation of image capture with position and orientation measurements. The digital image capture method employed for this thesis, however, has the advantage that the amount of temporal misregistration is generally much lower than video-based systems.

5.1.3 Finite beam thickness

The elevational thickness of the beam produced by linear- and phased-array transducers is typically of the order of several millimetres, but the analysis used in the calibration of the 3DUS system assumed an infinitely thin beam. In practice, if a structure of size smaller than the beam thickness is imaged, there will be uncertainty in its measured position, as shown in Figure 5.3.

It can be seen from Figure 5.3 that the thickness of the ultrasound beam varies in the axial direction, leading to a non-uniform resolution in the elevational direction. The beam narrows near the focal zone, but is considerably wider at the top and bottom of the image. The beam thickness for the ATL L10-5 and L12-5 scan-probes, operating in B-mode, are shown in Figure 5.4. The uncertainty in the location of structures in the elevational direction gives rise to an approximate position error of less than $w/2$, where w is the beam thickness at a particular depth.



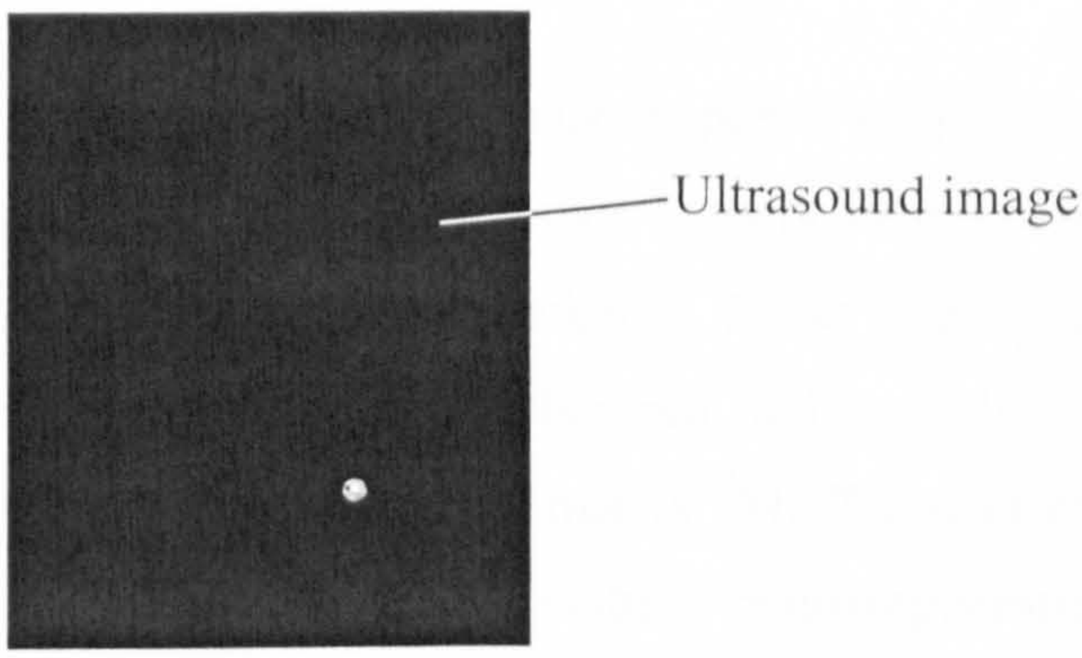


Figure 5.3 The effect of ultrasound beam thickness on locating a small object lying within the beam. When imaged, the object produces a reflection, but it does not lie within the ideal image plane (assumed to bisect the beam in the elevational direction). Measurement of the position of the object within the image reveals no information concerning its position perpendicular to the image. Therefore, there is uncertainty in the 3D location of the object due to the assumption that it lies on the ideal imaging plane.

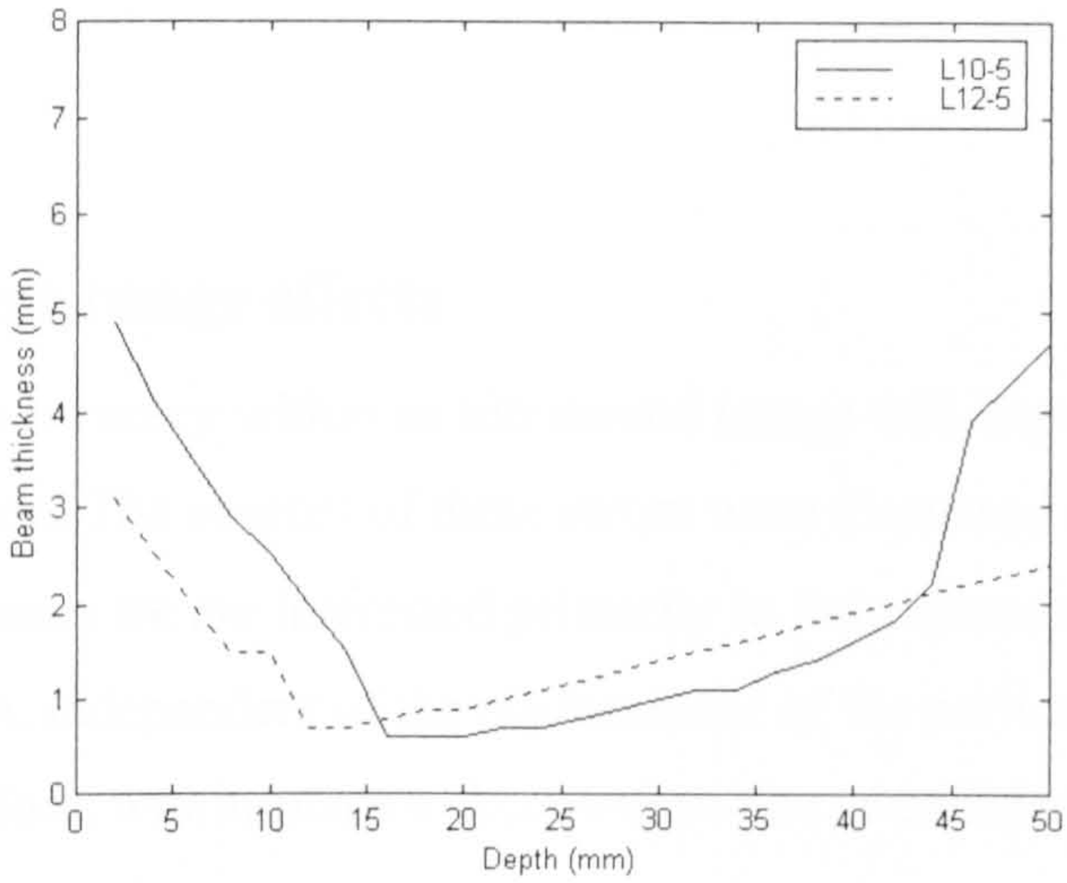


Figure 5.4 Beam thickness as a function of penetration depth for the L10-5 and L12-5 scan-probes with a single focal zone placed at a depth of 20 mm (source: ATL-Philips Ltd).

5.1.4 Imperfect calibration

The calibration procedure described in the previous chapter is subject to errors through two of the effects already highlighted above; namely, errors in *pcBIRD* measurements and the finite ultrasound beam thickness. These directly affect the accuracy of locating the pinhead target of the pin phantom, and result in an inevitable error in each of the parameters of the calibration matrix ${}^{R}\mathbf{M}_{I}$. These errors will propagate through to the reconstructed volume resulting in misregistration.

5.1.5 Motion artefacts

It has already been noted that tissue motion is a significant source of artefact for vascular 3DUS imaging (Chapter 3). An effective means of minimising misregistration due to tissue motion is to employ physiological gating. For the case of the carotid arteries, there are two sources of vessel motion: motion due to respiration and motion due to pulsatile blood flow through the vessel. Motion imposed by blood flow is much more significant than motion due to respiration during quiet breathing at rest [Chan 1993], Therefore, ECG-gating alone is effective in reducing misregistration errors if the patient breathes quietly during acquisition. Optimal implementation of ECG gating was discussed in the previous chapter. In certain patients, however, res-

piratory gating the acquisition may be desirable, although such a measure has not yet been implemented.

5.1.6 Within-plane image effects

The axial and lateral accuracy within an ultrasound image will impact upon the accuracy of 3D registration. The sources of these errors were discussed in Chapter 1, and in the present discussion we are interested primarily in the various sources of error for the 3DUS system, independent of the performance of the particular ultrasound scanner used. Therefore, within-plane effects will not be considered further.

4.2 Error propagation

It was shown in Chapter 4 that an image pixel with co-ordinates (m, n) is located within a reconstruction volume (V) by applying a series of matrix transformations:

$$\begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}_V = {}^{V}\mathbf{M}_T\,{}^{T}\mathbf{M}_R\,{}^{R}\mathbf{M}_I\,{}^{I}\mathbf{M}_J \begin{bmatrix} m \\ n \\ 0 \\ 1 \end{bmatrix}_J \tag{5.9}$$

This equation describes the 3D registration function applied to acquired images, and as such provides a convenient means of investigating the propagation of errors through the 3DUS system. If an error is introduced into one of the matrices on the right-hand side of Equation (5.9), then an estimate of the magnitude of the error in pixel position is found by calculating the quantity |e| as follows:

$$|\mathbf{e}| = |\mathbf{p}_V - (\mathbf{p}_V)_{\text{true}}| = \sqrt{(x - x_{\text{true}})^2 + (y - y_{\text{true}})^2 + (z - z_{\text{true}})^2}\,, \tag{5.10}$$

where (x_{true}, y_{true}, z_{true}) are the correct pixel co-ordinates specified by the vector $(\mathbf{p}_V)_{\text{true}}$. It was highlighted earlier that errors in *pcBIRD* position measurements are independent of the position and orientation of the receiver. Therefore, an error in the EPOM device position measurement results in the identical error in pixel location, and |e| is equal to the absolute position error for the EPOM device.

Errors in orientation measurements, on the other hand, are independent of the position of receiver, but dependent on its orientation. This means that identical errors in the measured rotation angles will result in a value of |e| which varies with the true rotation angles. In order to provide an estimate of the pixel registration error due to EPOM orientation errors, a computer simulation was carried out. The mean and maximum values of |e| were calculated using Equation (5.10) for 729 receiver orientations in steps of 45° about each axis of rotation. The predicted errors in locating the centre pixel in a 40 × 40 mm image are shown in Figure 5.5. In the simulation, the rotation error, in the range -1.5 – 1.5°, was assumed to be the same for all three rotation angles. The perpendicular distance between the surface of the ultrasound transducer and the origin of the receiver was set to 30 mm, with the receiver centred laterally on the scan-probe.

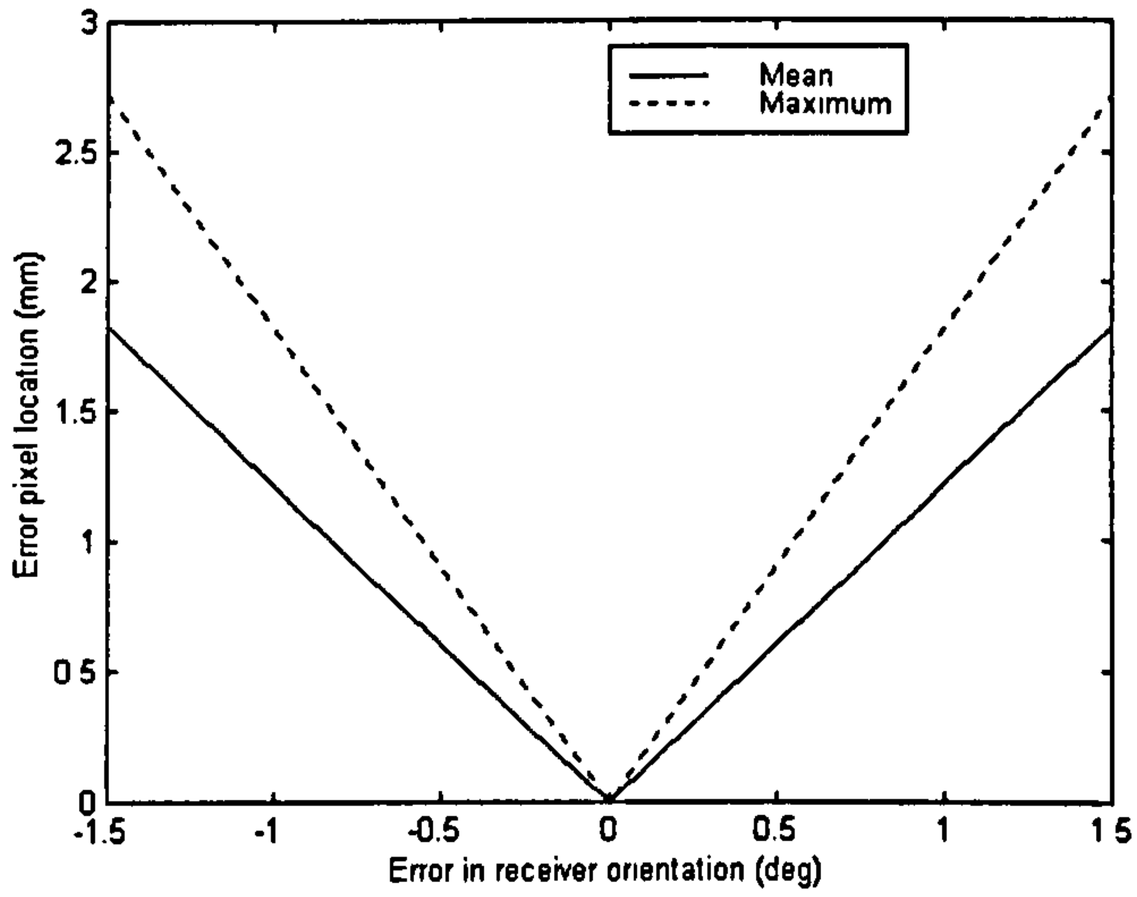


Figure 5.5 The effect of errors in *pcBIRD* receiver orientation measurement on image pixel registration. The predicted mean and maximum error in the registered location (|e|) of the centre pixel in a 40 × 40 mm image is plotted as a function of error in orientation. The orientation error shown on the horizontal axis was applied to all three rotation angles, and |e| was calculated over 729 different receiver orientations at regular intervals.

It can be seen from Figure 5.5 that there is a linear relationship between the predicted mean/maximum value of |e| and the magnitude of error in receiver orientation. For the case shown, the relationship between the mean value of |e| and the angular

error, δθ, was found to be:

$$\text{mean}(|\mathbf{e}|) = K\delta\theta \approx 1.2\delta\theta\,. \tag{5.11}$$

The multiplication factor, K, in this relationship is dependent on the perpendicular distance between the surface of the ultrasound transducer and the origin of the EPOM receiver. This relation was investigated and also found to be approximately linear, as shown in Figure 5.6. The significance of the lever-arm effect is illustrated here as K increases with the receiver-transducer distance. As already noted, an important ramification of the lever-arm effect is that mounting the EPOM receiver far away from the transducer surface will magnify the effect of errors in the measured orientation angles.

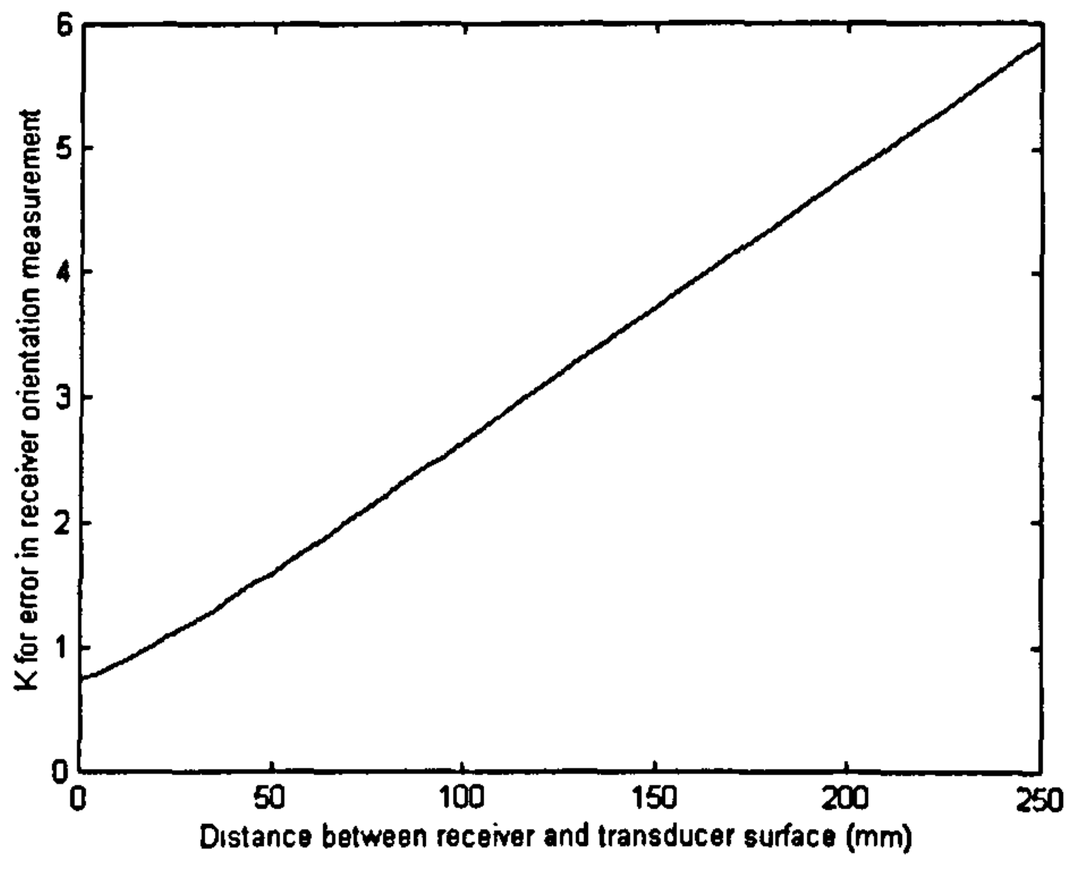


Figure 5.6 The predicted dependency of K on the perpendicular distance between surface of the ultrasound transducer and the origin of the EPOM receiver.

5.3 Experimental evaluation of the Ascension *pcBIRD*

5.3.1 Interference in *pcBIRD* measurements

It is well recognised that EPOM devices are susceptible to interference. In particular, the presence of some types of metal or extraneous electromagnetic fields may distort position and orientation measurements. When used in a free-hand 3DUS system, such interference will lead to inaccurate image registration in the reconstructed volume. The clinical ultrasound environment contains numerous potential sources of

interference, such as, the scanner, the scan-probe, the examination couch, CRT monitors, and other auxiliary electrical equipment.

There are two strategies for eliminating, or at least reducing, interference in an EPOM device: the first and most obvious is to move sources of significant interference. This strategy is effective up to a point, as it is often possible to rearrange the scanning environment. However, there are normally severe limitations on where essential equipment, in particular the ultrasound scanner, can be moved.

Another strategy is to attempt to optimise the parameters of the EPOM device to reduce its sensitivity to interference. The level of noise in measurements provides a useful indication of the influence of a particular operating environment. This will be affected by the measurement rate and filter settings of the device. The term 'noise' is usually reserved for the standard deviation of repeated measurements made under the same conditions. In the experiments below, an alternative measure suitable for vector measurements and referred to as the mean deviation was adopted. In order to avoid confusion, the mean deviation, Δ, for a series of vector measurements, $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N$, is given by:

$$\Delta = \frac{1}{N}\sum_{i=1}^{N}|\mathbf{x}_i - \bar{\mathbf{x}}|, \tag{5.12}$$

where $\bar{\mathbf{x}}$ is the mean vector. The mean deviation is a convenient indication of noise in position and orientation measurements made using the *pcBIRD*.

5.3.1.1 Reorganisation of the scanning environment

The layout of the room to be used for 3DUS imaging was reorganised as shown in Figure 5.7. As a precaution, the amount of metal in the immediate vicinity of the *pcBIRD* receiver and transmitter was minimised by replacing the metal examination couch with a wooden one. The *pcBIRD* transmitter was positioned at the head of the bed near to the scan-region appropriate for carotid imaging. The PC was placed on a moveable trolley positioned approximately 2 m away from the transmitter with the monitor easily visible to the ultrasonographer. All other equipment not part of the

3DUS acquisition system was moved more than 2 metres away from the scanning region.

The positions of the scanner and couch shown in Figure 5.7 are those used for routine ultrasound investigations of the carotid arteries. An important consideration in the reorganisation of the scanning environment was the preservation of this arrangement to ensure the straight-forward transition between routine and 3D ultrasound scanning. In a wider context, this is important because, to assist the widespread clinical acceptance of 3DUS, there should be minimal disruption to routine scanning technique.

5.3.1.2 Optimisation of *pcBIRD* parameters

In order to optimise the parameters of the *pcBIRD*, the effect of measurement rate and filter settings on repeated measurements was investigated. For position measurements the effect was quantified by calculating the Euclidean distance, ΔD, between the measured position vector, (x, y, z), and the mean measured position, $(\bar{x}, \bar{y}, \bar{z})$, where:

$$\Delta D = \sqrt{(x-\bar{x})^2 + (y-\bar{y})^2 + (z-\bar{z})^2}\ . \tag{5.13}$$

For repeated orientation measurements the *screw angle*, $\Delta\phi$, between a measured rotation matrix, $\mathbf{R}$, and the 'average' rotation matrix, $\overline{\mathbf{R}}$, was used as a measure of the influence of the environment. The screw angle was calculated using the following formula [Spoor 1980]:

$$\Delta\phi = \sin^{-1}[\sin(\Delta\phi)], \tag{5.14}$$

if $\sin(\Delta\phi) \geq \sqrt{2}/2$, or

$$\Delta\phi = \cos^{-1}\left[\frac{1}{2}(\mathbf{A}_{11} + \mathbf{A}_{22} + \mathbf{A}_{33} - 1)\right], \tag{5.15}$$

(*continued on page 151*)

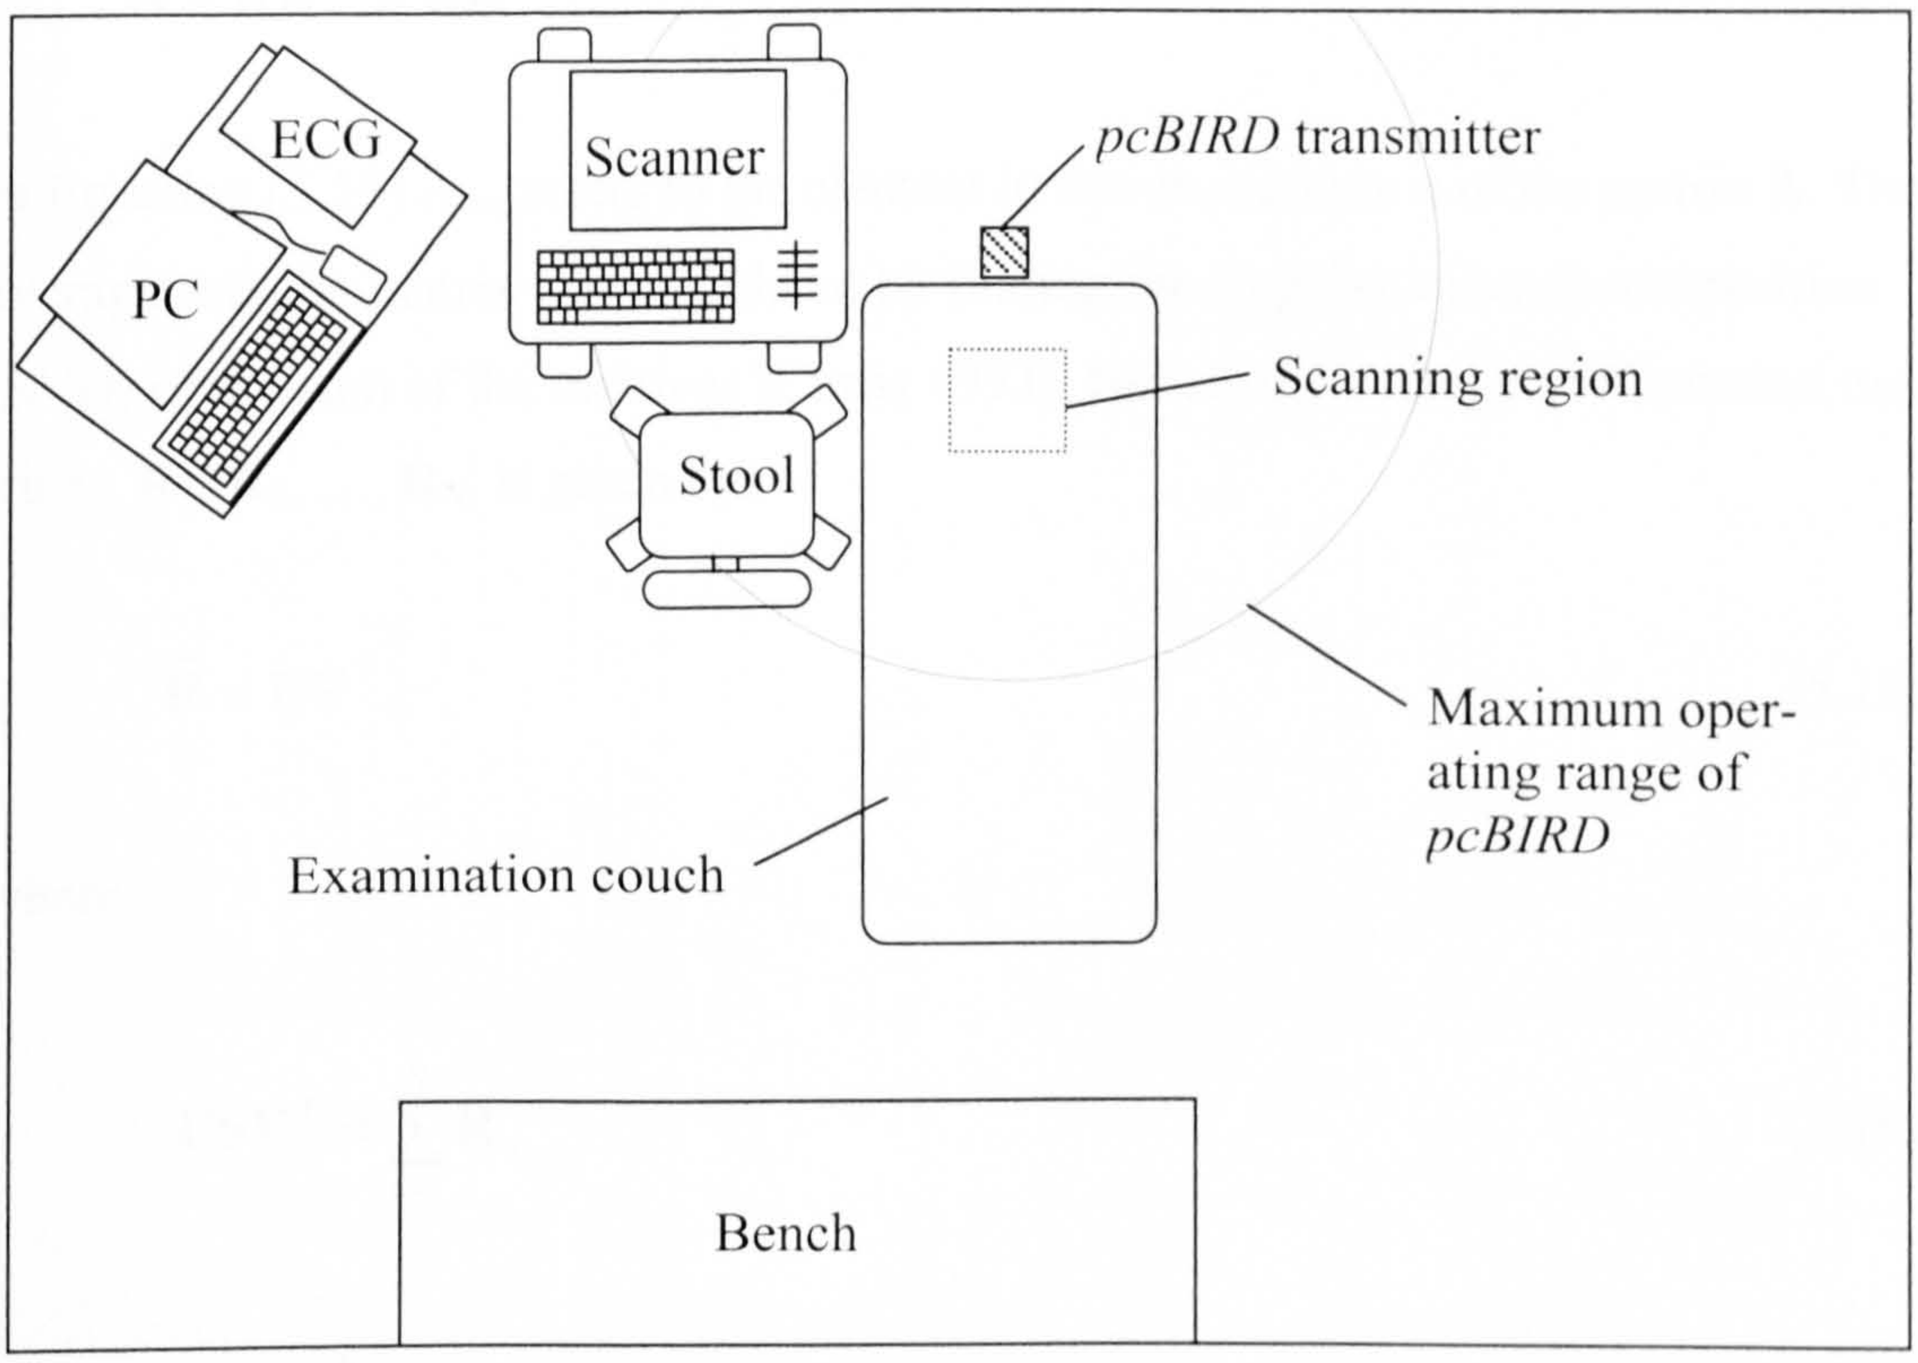


Figure 5.7 The floor plan for the vascular ultrasound laboratory used for 3DUS imaging of the carotid arteries.

if $\sin(\Delta\phi) < \sqrt{2}/2$, where

$$\sin(\Delta\phi) = \frac{1}{2}\sqrt{(\mathbf{A}_{13} - \mathbf{A}_{31})^2 + (\mathbf{A}_{21} - \mathbf{A}_{12})^2 + (\mathbf{A}_{32} - \mathbf{A}_{23})^2} \tag{5.16}$$

and

$$\mathbf{A} = \mathbf{R}\overline{\mathbf{R}}^{-1} = \mathbf{R}\overline{\mathbf{R}}^{\mathrm{T}}. \tag{5.17}$$

In Equation (5.16) $\mathbf{A}_{mn}$ refers to the element in row m, column n of the matrix **A**. The average rotation matrix was calculated by finding the singular value decomposition (SVD) of the sum of the matrices [Curtis 1993]. Hence, the average of *N* rotation matrices, $\mathbf{R}_1, \mathbf{R}_2, \ldots, \mathbf{R}_N$, is given by:

$$\overline{\mathbf{R}} = \mathbf{U}\mathbf{V}^{\mathrm{T}}, \tag{5.18}$$

where

$$\mathbf{U}\mathbf{S}\mathbf{V}^{\mathrm{T}} = \sum_{n=1}^{N}\mathbf{R}_n \tag{5.19}$$

is the SVD of $\sum_{n=1}^{N}\mathbf{R}_n$.

The mean values of ΔD and $\Delta\phi$ represent the mean deviations for repeated 3D position and orientation measurements, and provide two measures of the interference in the clinical environment. To test the effect of changing the measurement rate and filter settings, the following experiment was performed: the receiver was mounted on the scan-probe and fixed in place within the centre of the scanning region. The couch was placed as close to the scanner as possible with the centre of the *pcBIRD* transmitter approximately 350 mm from the side of the scanner. One of the six filter settings listed in Section 4.1.2 was selected, and 100 position and orientation measurements were recorded for each integral measurement rate from 70 to 144 Hz. The scan-probe was active throughout with the scanner in power Doppler mode. This procedure was then repeated for each of the possible filter combinations, and using

the four receiver/scan-probe orientations shown in Figure 5.8. The values of ΔD and $\Delta\phi$ were then calculated for each set of measurements using the Matlab mathematical software package (The Mathworks Inc., MA, USA).The experiment was repeated for the *pcBIRD* receiver alone, with the scan-probe removed, in order to compare the influence of the spatial relationship between the receiver and transmitter on measurements with and without the scan-probe present. In other words, are measurements significantly influenced when the scan-probe was positioned between the receiver and transmitter?

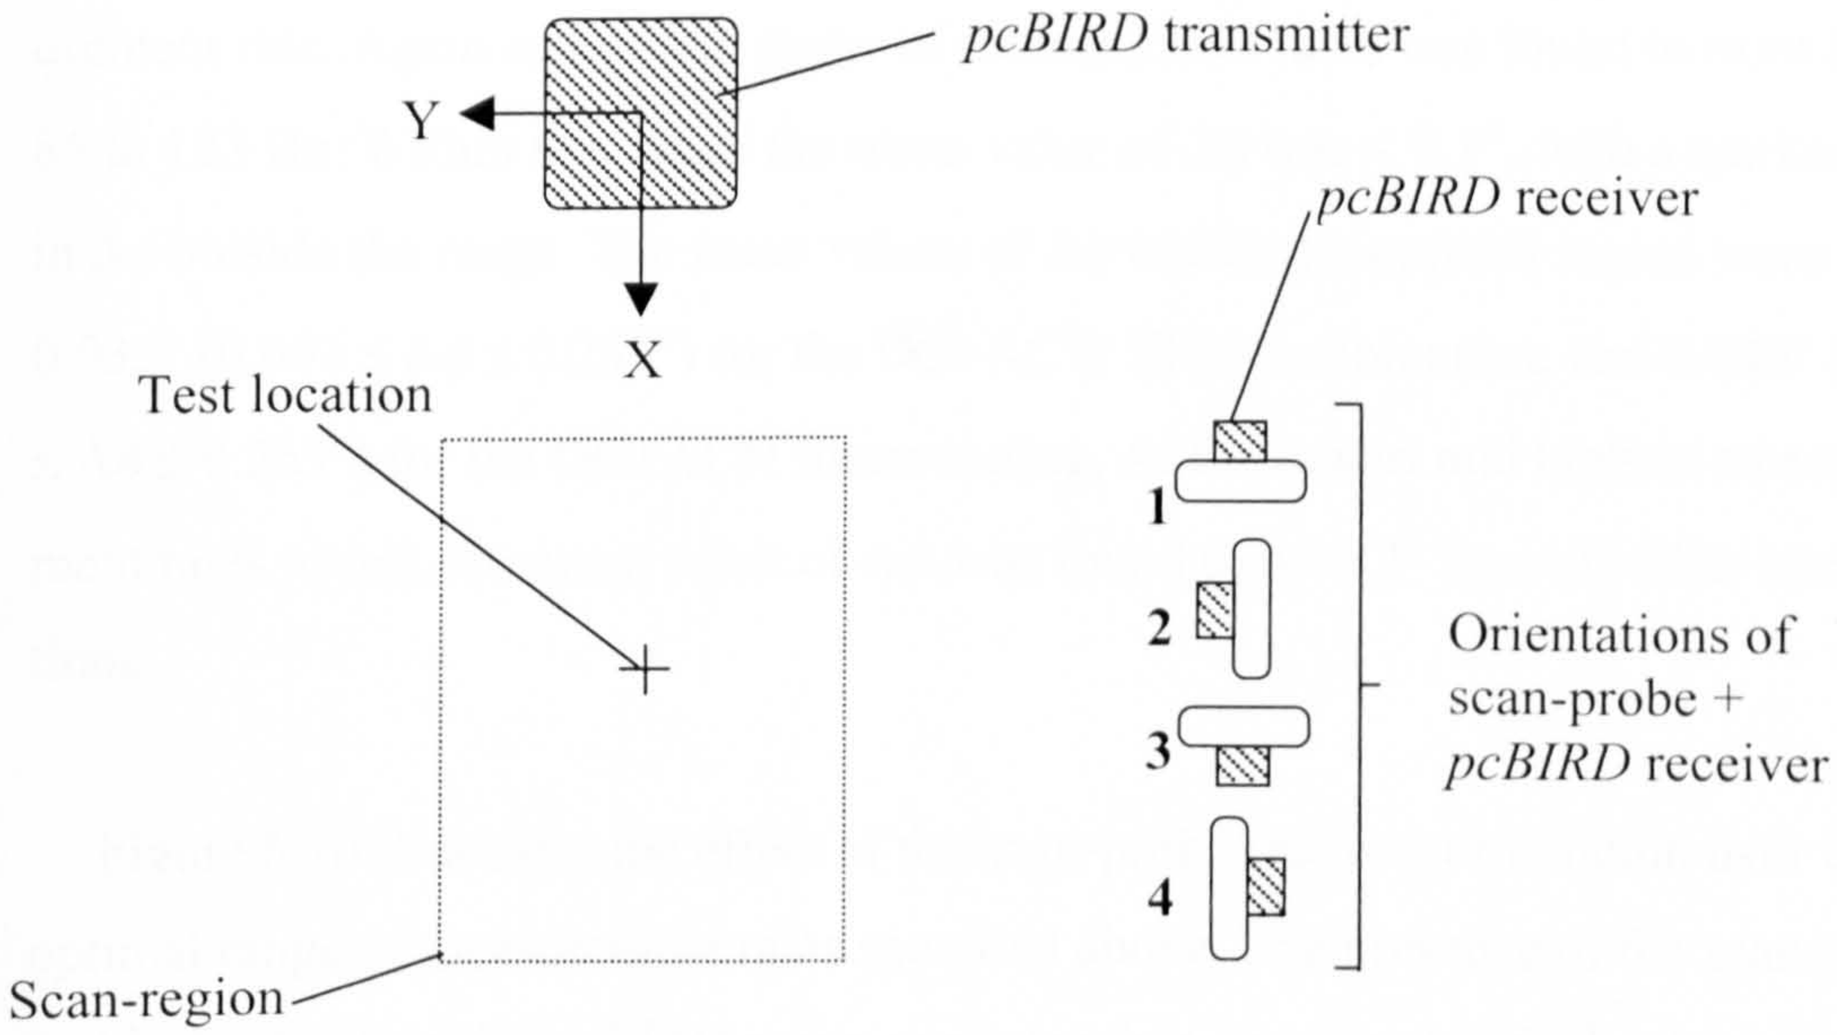


Figure 5.8 Four orientations of the scan-probe and receiver combination used during the measurement noise tests (viewed from above). In all cases the scan-probe was secured in an upright position.

Figure 5.9(a) shows the mean value of ΔD, averaged over the four receiver/scan-probe orientations in Figure 5.9, plotted versus measurement rate. The results are only presented for the two best filter settings – i.e. DC+ACW and DC+ACN – where the mean value of ΔD was less than the acceptable cut-off value of 0.5 mm for one or more measurement rate. The optimal range of measurement rate for these filter settings was 79 – 137 Hz. Over this range, the mean values of ΔD were 0.151 mm ($0.007 \leq \Delta D \leq 1.067$ mm) for the DC+ACW filter combination, and 0.197 mm ($0.006 \leq \Delta D \leq 1.557$ mm) for the DC+ACN filter combination.

The magnitude of ΔD was found to rise sharply at rates below 79 Hz and above 137 Hz for both filter combinations. At the lowest and highest measurement rates tested (70 and 144 Hz), ΔD was found to rise above 5 mm. All of the other filter settings tested did not meet the mean $\Delta D < 0.5$ mm criterion for any measurement rate.

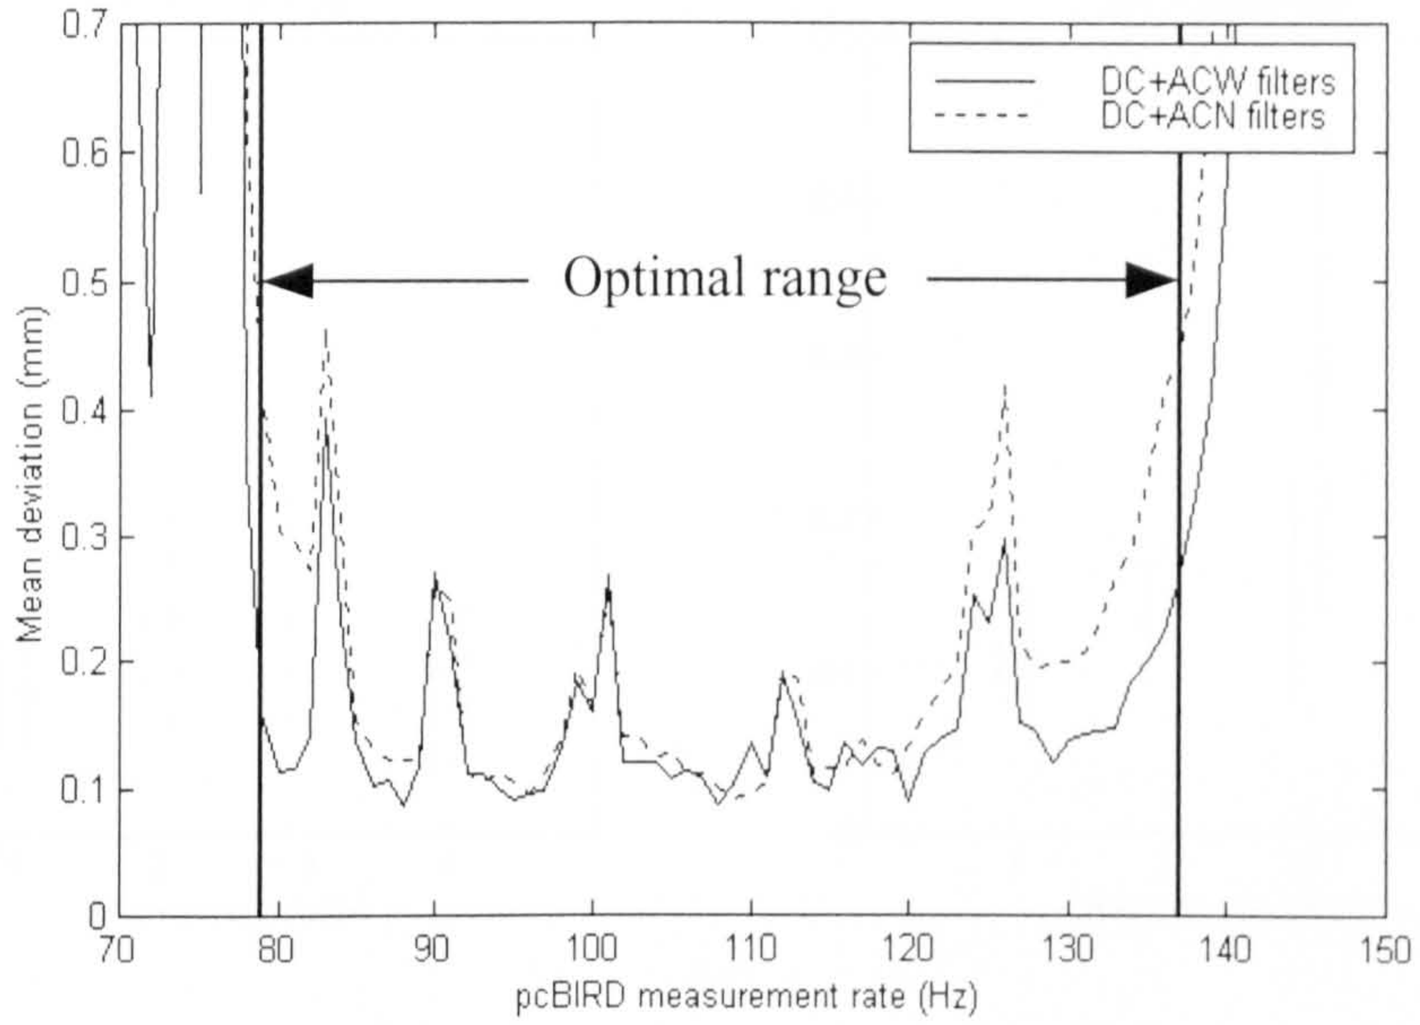
The DC+ACW and DC+ACN filter combinations were selected for the repeated orientation measurement tests on the basis of their favourable performance for position measurements. The results are shown in Fig. 5.9(b) with the mean value of $\Delta\phi$ (averaged over for the four receiver/scan-probe orientations) plotted against measurement rate. Again an optimal range of measurement rates was found to exist from 85 to 123 Hz. Within this range the mean value of $\Delta\phi$ was $< 0.1°$, with a marked rise in $\Delta\phi$ outside the range. The mean values of $\Delta\phi$ within the optimal region were $0.035°$ ($0.002 \leq \Delta\phi \leq 0.249°$) for the DC+ACW filter combination, and $0.040°$ ($0.002 \leq \Delta\phi \leq 0.269°$) for the DC+ACN filters setting. At the lowest and highest measurement rates tested, the mean value of $\Delta\phi$ was found to be $> 2°$ for both filter combinations.

Figure 5.10 illustrates the effect of the scan-probe/receiver orientation over the optimal range of measurement rates specified above. The presence of the scan-probe had little effect on the *pcBIRD* measurements when compared with the deviations observed for the receiver alone, except for orientation 3 where the path between the receiver and the transmitter was obstructed by the scan-probe. In this case, an appreciable rise in the mean and standard deviation for both ΔD and $\Delta\phi$ was observed.

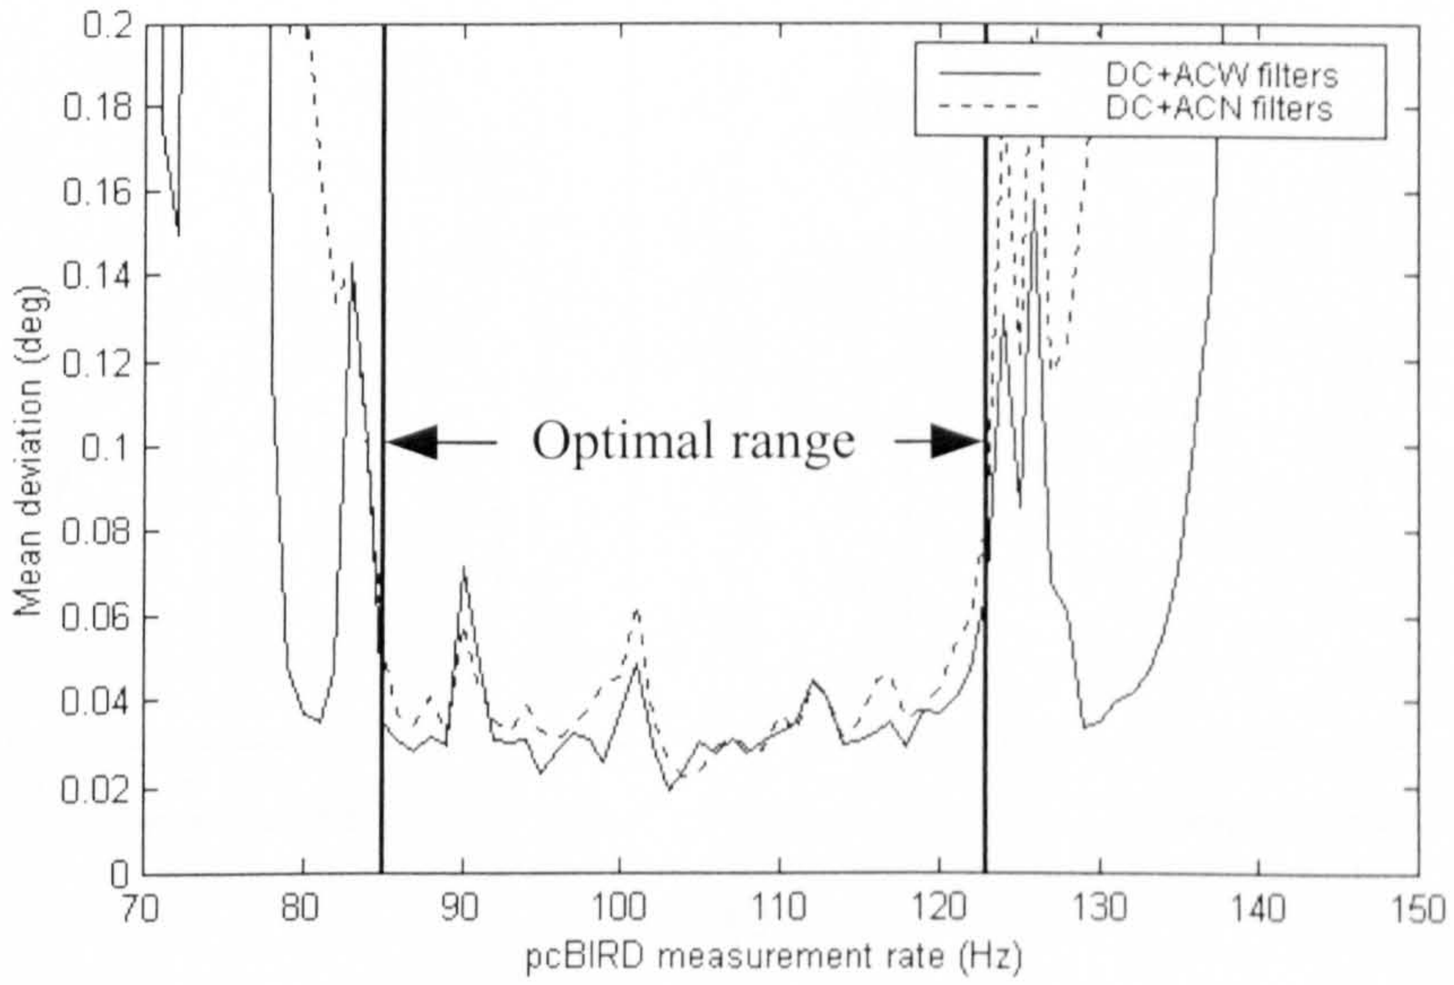
The optimised *pcBIRD* settings selected for the subsequent accuracy experiments were a measurement rate of 111 Hz in conjunction with the DC+ACN filter combination. This filter was chosen as it provided a good compromise between robustness to interference and dynamic performance, due to the lower time required to compute position and orientation measurements.

Although ΔD and $\Delta\phi$ provide useful measures of the deviation in repeated measurements, they yield no information about how accurate the device is in measuring

(*continued on page 156*)

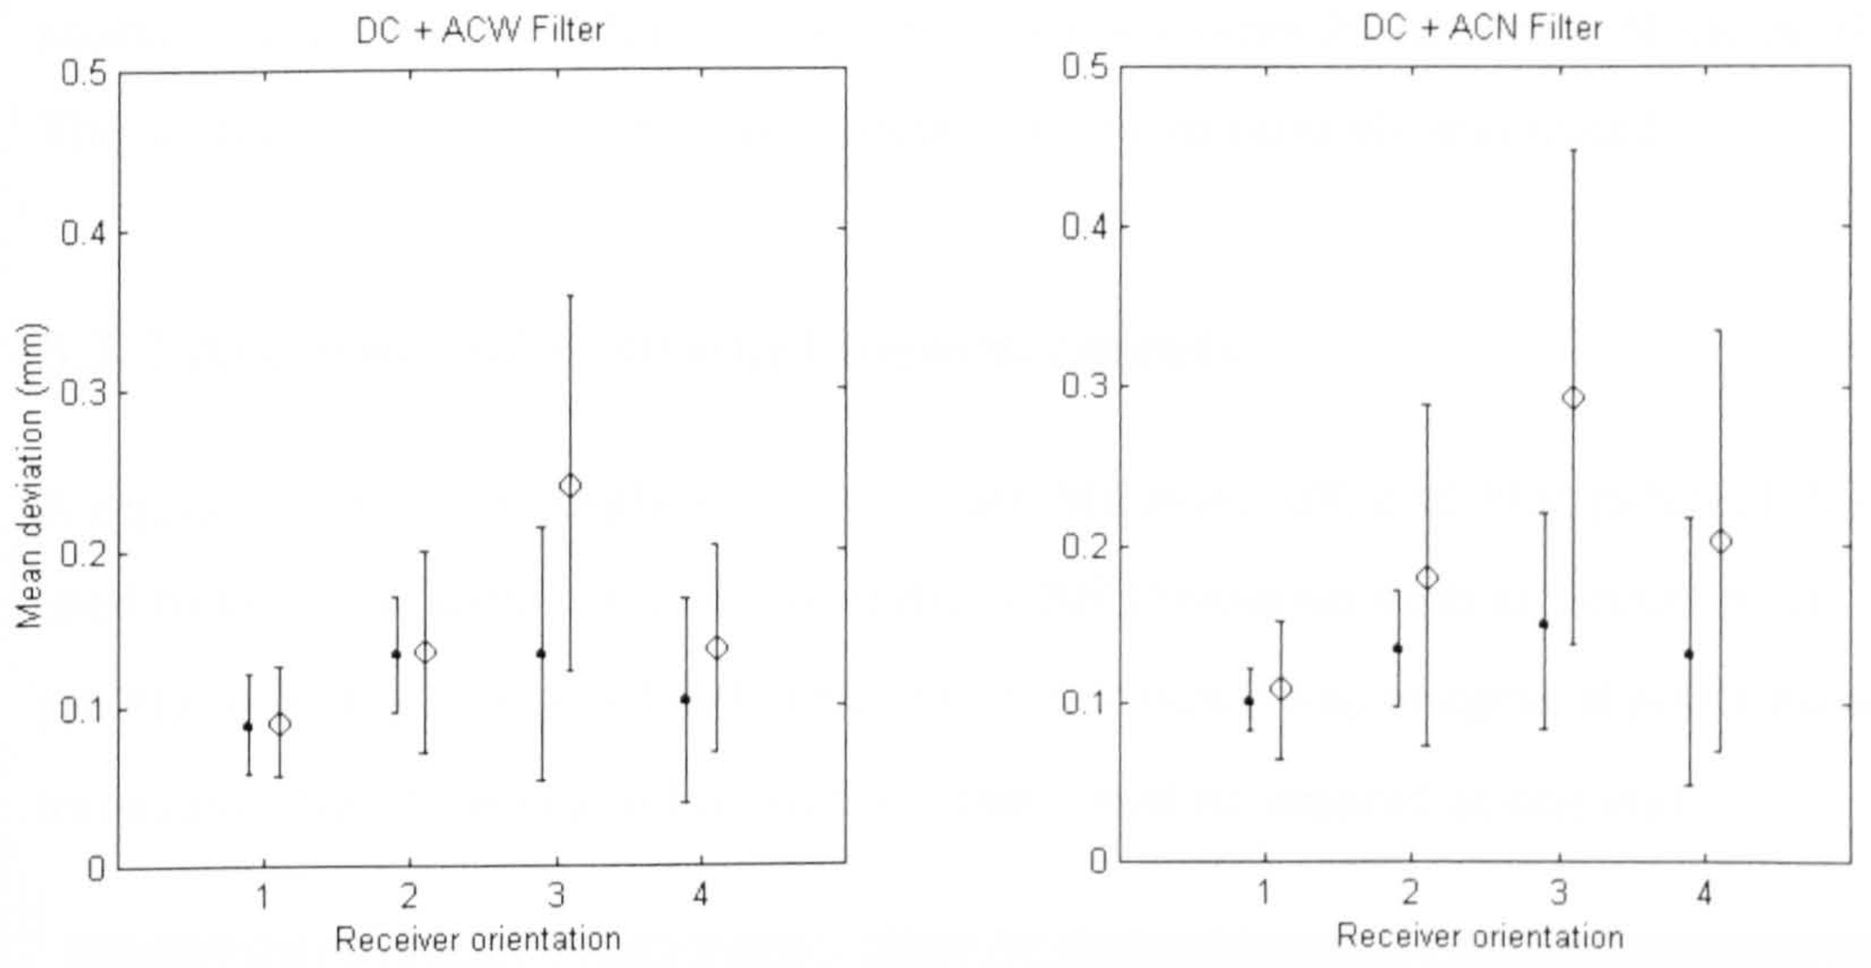


(a)

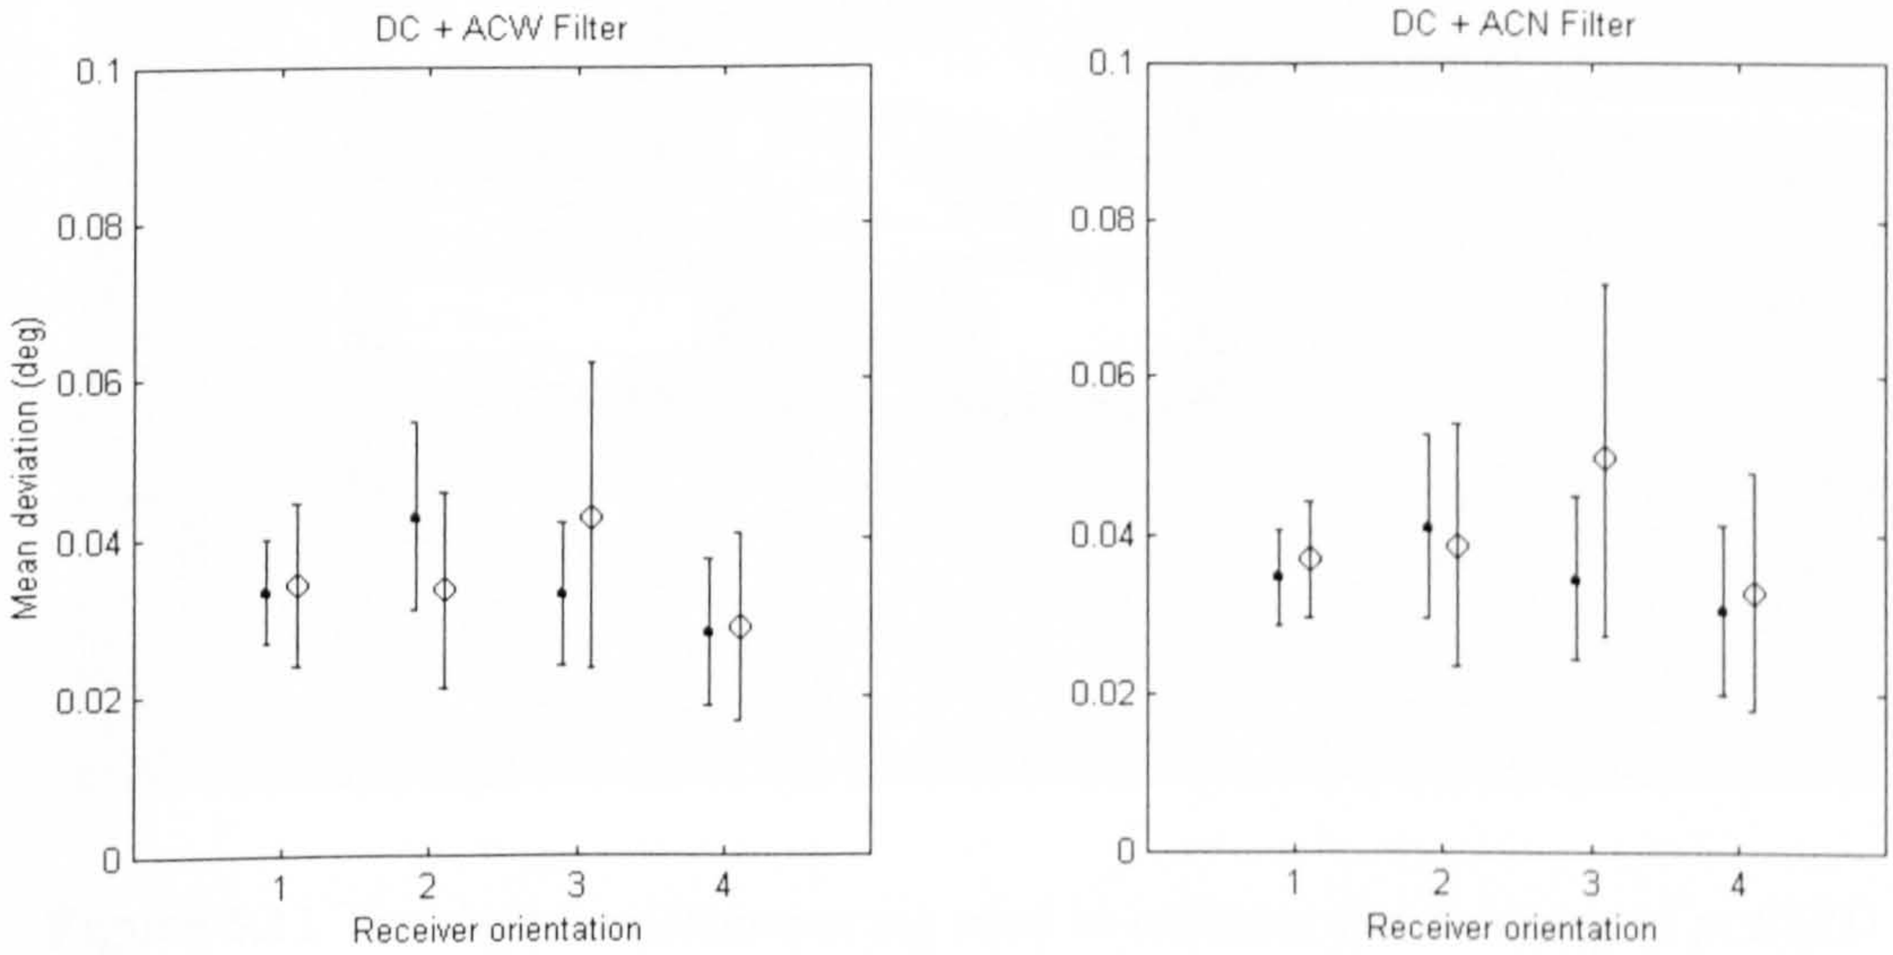


(b)

Figure 5.9 Mean measurement deviations (ΔD and $\Delta\phi$) plotted against measurement rate for the DC+ACW and DC+ACN filter settings. (a) shows the mean deviation for positional measurements, whereas (b) shows the mean deviation for rotational measurements. The values plotted were averaged over the 4 different receiver/scan-probe orientations tested (see Figure 5.8).



(a)



(b)

Figure 5.10 Mean ± SD deviations for *pcBIRD* position and orientation measurements: (a) position measurements (ΔD), and (b) orientation measurements ($\Delta\phi$) for the 4 different receiver/scan-probe orientations depicted in Figure 5.8. Key: • – measurements made by the *pcBIRD* receiver alone; – measurements made by the *pcBIRD* receiver when mounted on scan-probe. All values were over the appropriate optimal measurement rate ranges shown in Figure 5.9.

the relative distance between two points in space, or the relative orientations of two bodies. High accuracy is required for satisfactory 3D registration of acquired ultrasound images, and, therefore, it was necessary to assess the accuracy of the *pcBIRD*. The accuracy of both position and orientation measurements was tested.

5.3.2 Accuracy of positional measurements

A digital micrometer (Digimatic Scale Unit, Mitutoyo UK Ltd, Hampshire, UK) was used to measure linear translations of the *pcBIRD* receiver with an accuracy of ±0.005 mm over a range of 0-90 mm. The micrometer was integrated into a manual translation rig, as shown in Figure 5.11, that could be secured at one end.

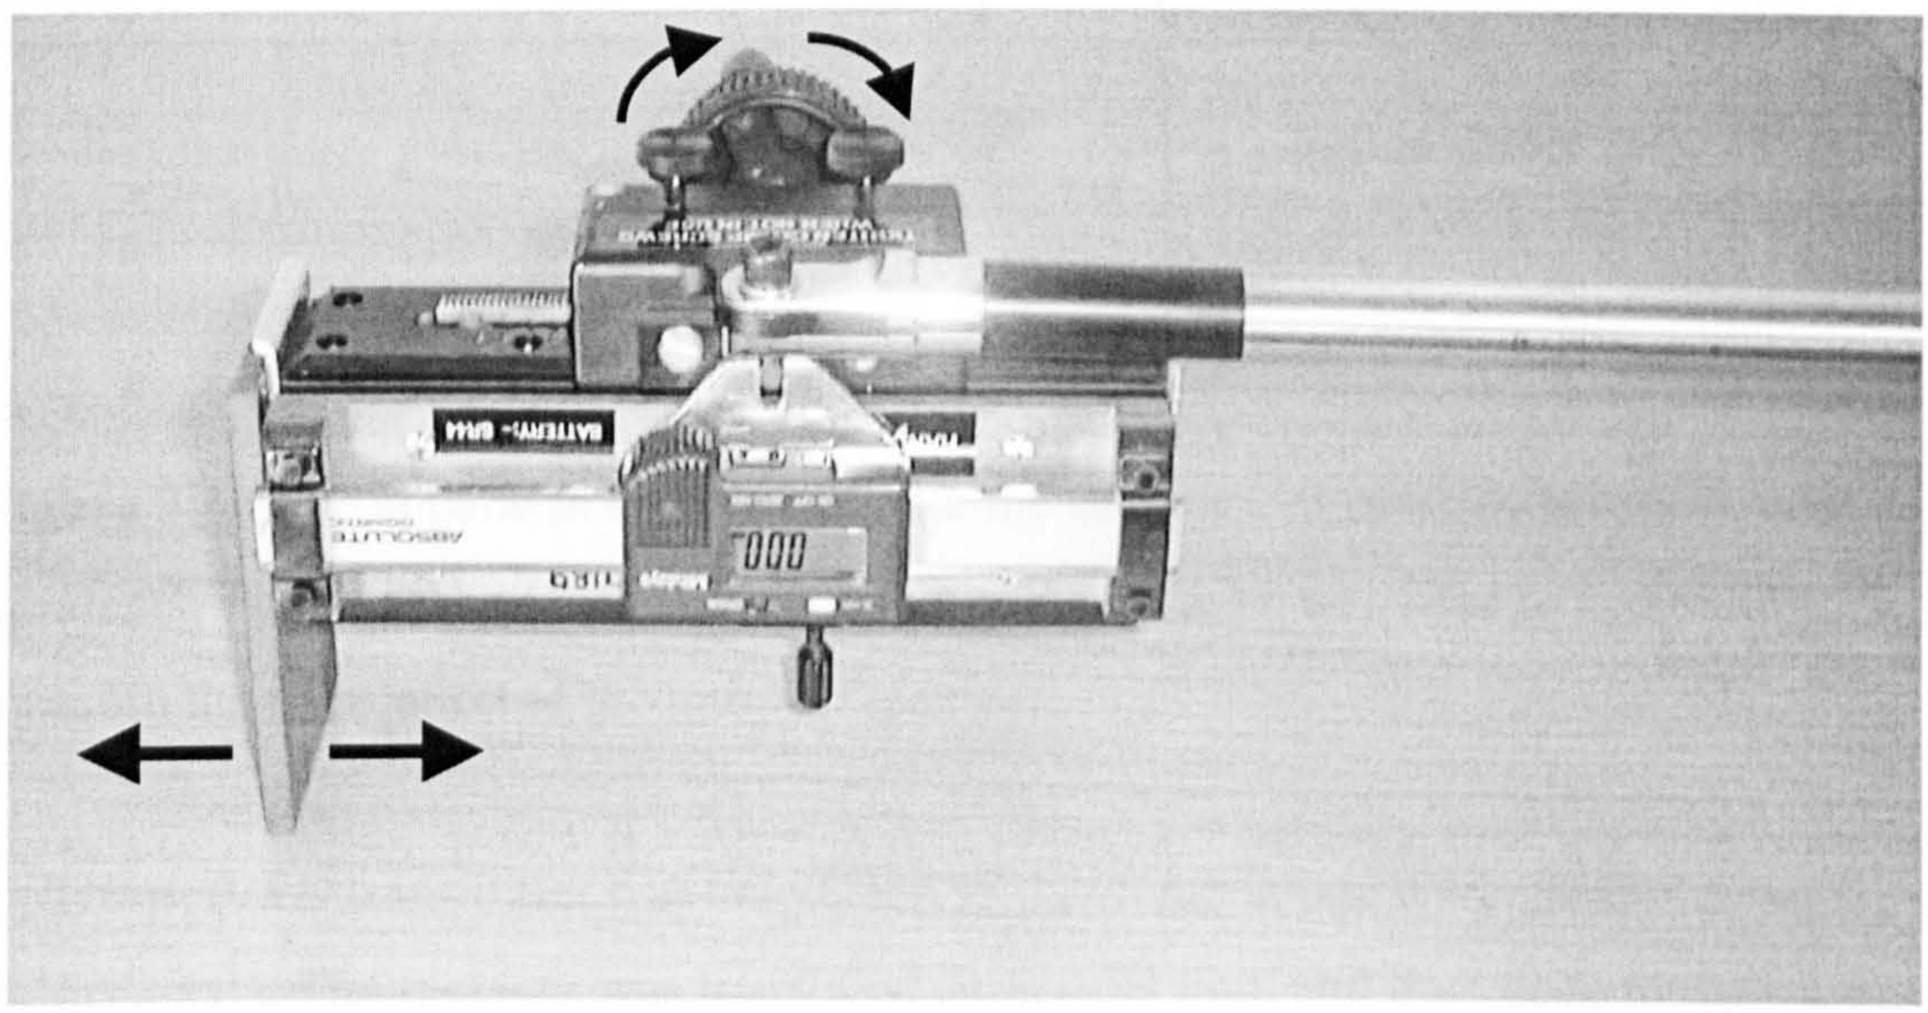


Figure 5.11 The digital micrometer rig used to estimate the accuracy of *pcBIRD* position measurements. The wheel visible at the top is rotated manually, and the free end of the rig translates linearly, as indicated by the arrows. The translation (in mm) is shown on the digital display.

The *pcBIRD* receiver (or receiver plus scan-probe combination) was attached to the end of a 1.2 m rigid wooden beam, as shown in Figure 5.12. This was necessary to prevent interference in the *pcBIRD* measurements due to metal within the micrometer rig. The other end of the beam was screwed to the micrometer rig, and the rig clamped to a retort stand fixed securely to the end of the examination couch. A Nylatron board with dimensions 400 × 250 × 10 mm was placed on the couch to provide a smooth, flat surface for the beam to move across. Two plastic brackets at-

tached to the board provided a guide for the beam, and Teflon gliders were glued to the brackets and to the underside of the wooden beam to minimise frictional resistance as the beam was translated.

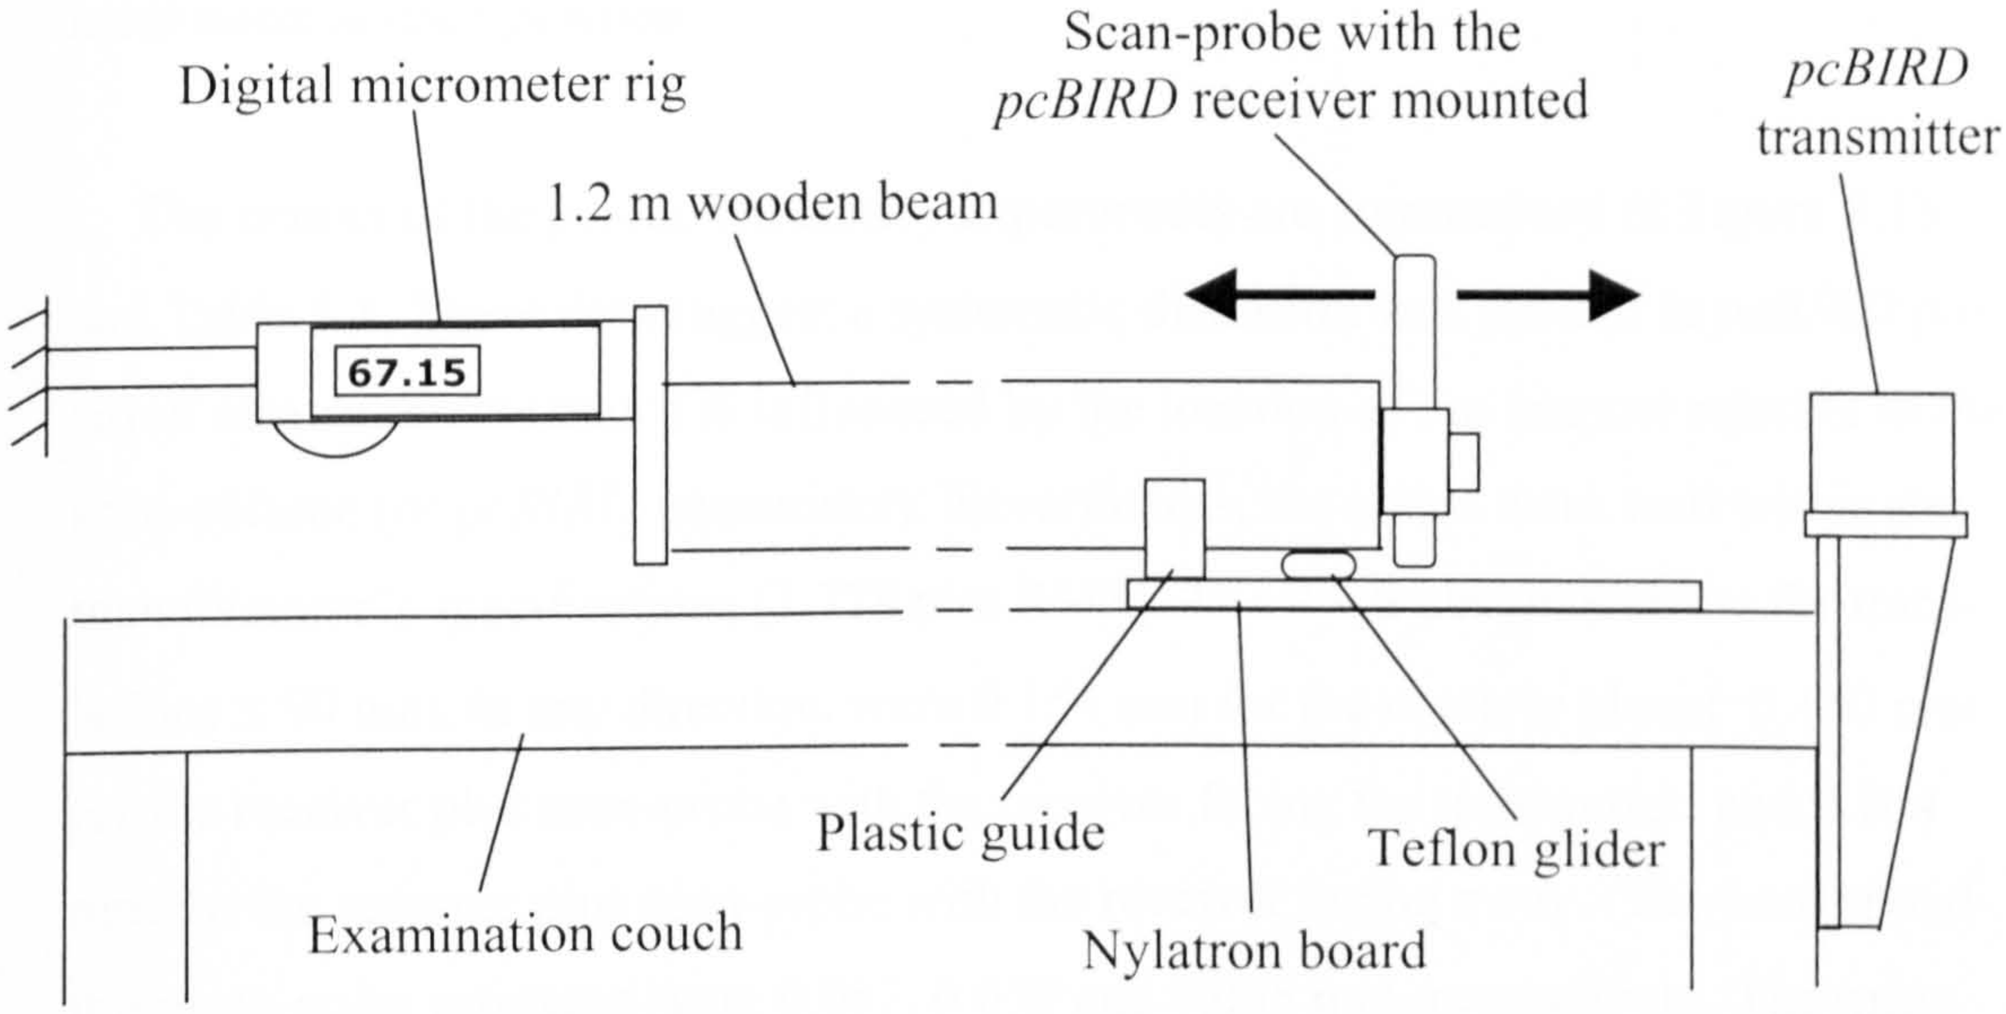


Figure 5.12 The apparatus used to estimate the accuracy of *pcBIRD* position measurements. The digital micrometer depicted in Figure 5.11 was used to accurately measure the translation of the *pcBIRD* receiver.

The *pcBIRD* transmitter was orientated so that translations were parallel to its x-, y-, or z-axis. The receiver was translated from 0-90 mm, and then back to the start position, in increments of 10 mm. For each incremental position, 20 *pcBIRD* measurements were recorded and averaged to give a single mean position. The distance, l, between the initial measured position, (x_0, y_0, z_0), and a subsequent measured position, (x_i, y_i, z_i), is given by:

$$l = \sqrt{(x_i - x_0)^2 + (y_i - y_0)^2 + (z_i - z_0)^2} \,. \tag{5.20}$$

The value of l was calculated for each *pcBIRD* position measurement and compared with the 'true' translation measured by the digital micrometer. This procedure was repeated, translating along each of the three transmitter axes, for the following configurations: i) the *pcBIRD* receiver alone; ii) the *pcBIRD* receiver mounted on the scan-probe with the receiver facing towards the transmitter; and iii) the *pcBIRD* re-

ceiver mounted on the scan-probe with the receiver facing away from the transmitter. The experiments were repeated with the scanner placed 350, 500, and 750 mm from the transmitter. For each set of 20 position measurements, the RMS value of ΔD defined in Equation (5.12), was calculated to give an estimate of the position measurement noise at each position.

The results of the position accuracy experiments are summarised in Figure 5.13 and Table 5.1. These data suggest a systematic distortion was present in *pcBIRD* position measurements which is influenced by the location of the scanner relative to the scan-volume (or *pcBIRD* transmitter). Nevertheless, the errors were well within the manufacturer's specifications (1.778 mm RMS). The RMS absolute errors for translations ≤ 90 mm, in any direction, were 0.154 mm for the receiver alone; 0.160 mm for the receiver plus scan-probe with the receiver facing the transmitter; and 0.204 mm for the receiver plus scan-probe with the receiver facing away. The corresponding mean noise estimates were 0.087, 0.059 and 0.065 mm, respectively. The mean separation between the origins of the transmitter and receiver co-ordinate systems during these experiments was approximately 390 mm.

	Scanner-transmitter separation (mm)		
	350	**500**	**750**
Receiver alone	0.215	0.121	0.137
Receiver + scan-probe combination 1*	0.207	0.086	0.164
Receiver + scan-probe combination 2**	0.256	0.132	0.204

Table 5.1 RMS absolute errors (mm) for positional measurements over translations of 0 to 90 mm. * Combination 1 refers to the configuration where the *pcBIRD* receiver faced towards the transmitter; and ** combination 2 refers to the configuration where the receiver faced away from the transmitter.

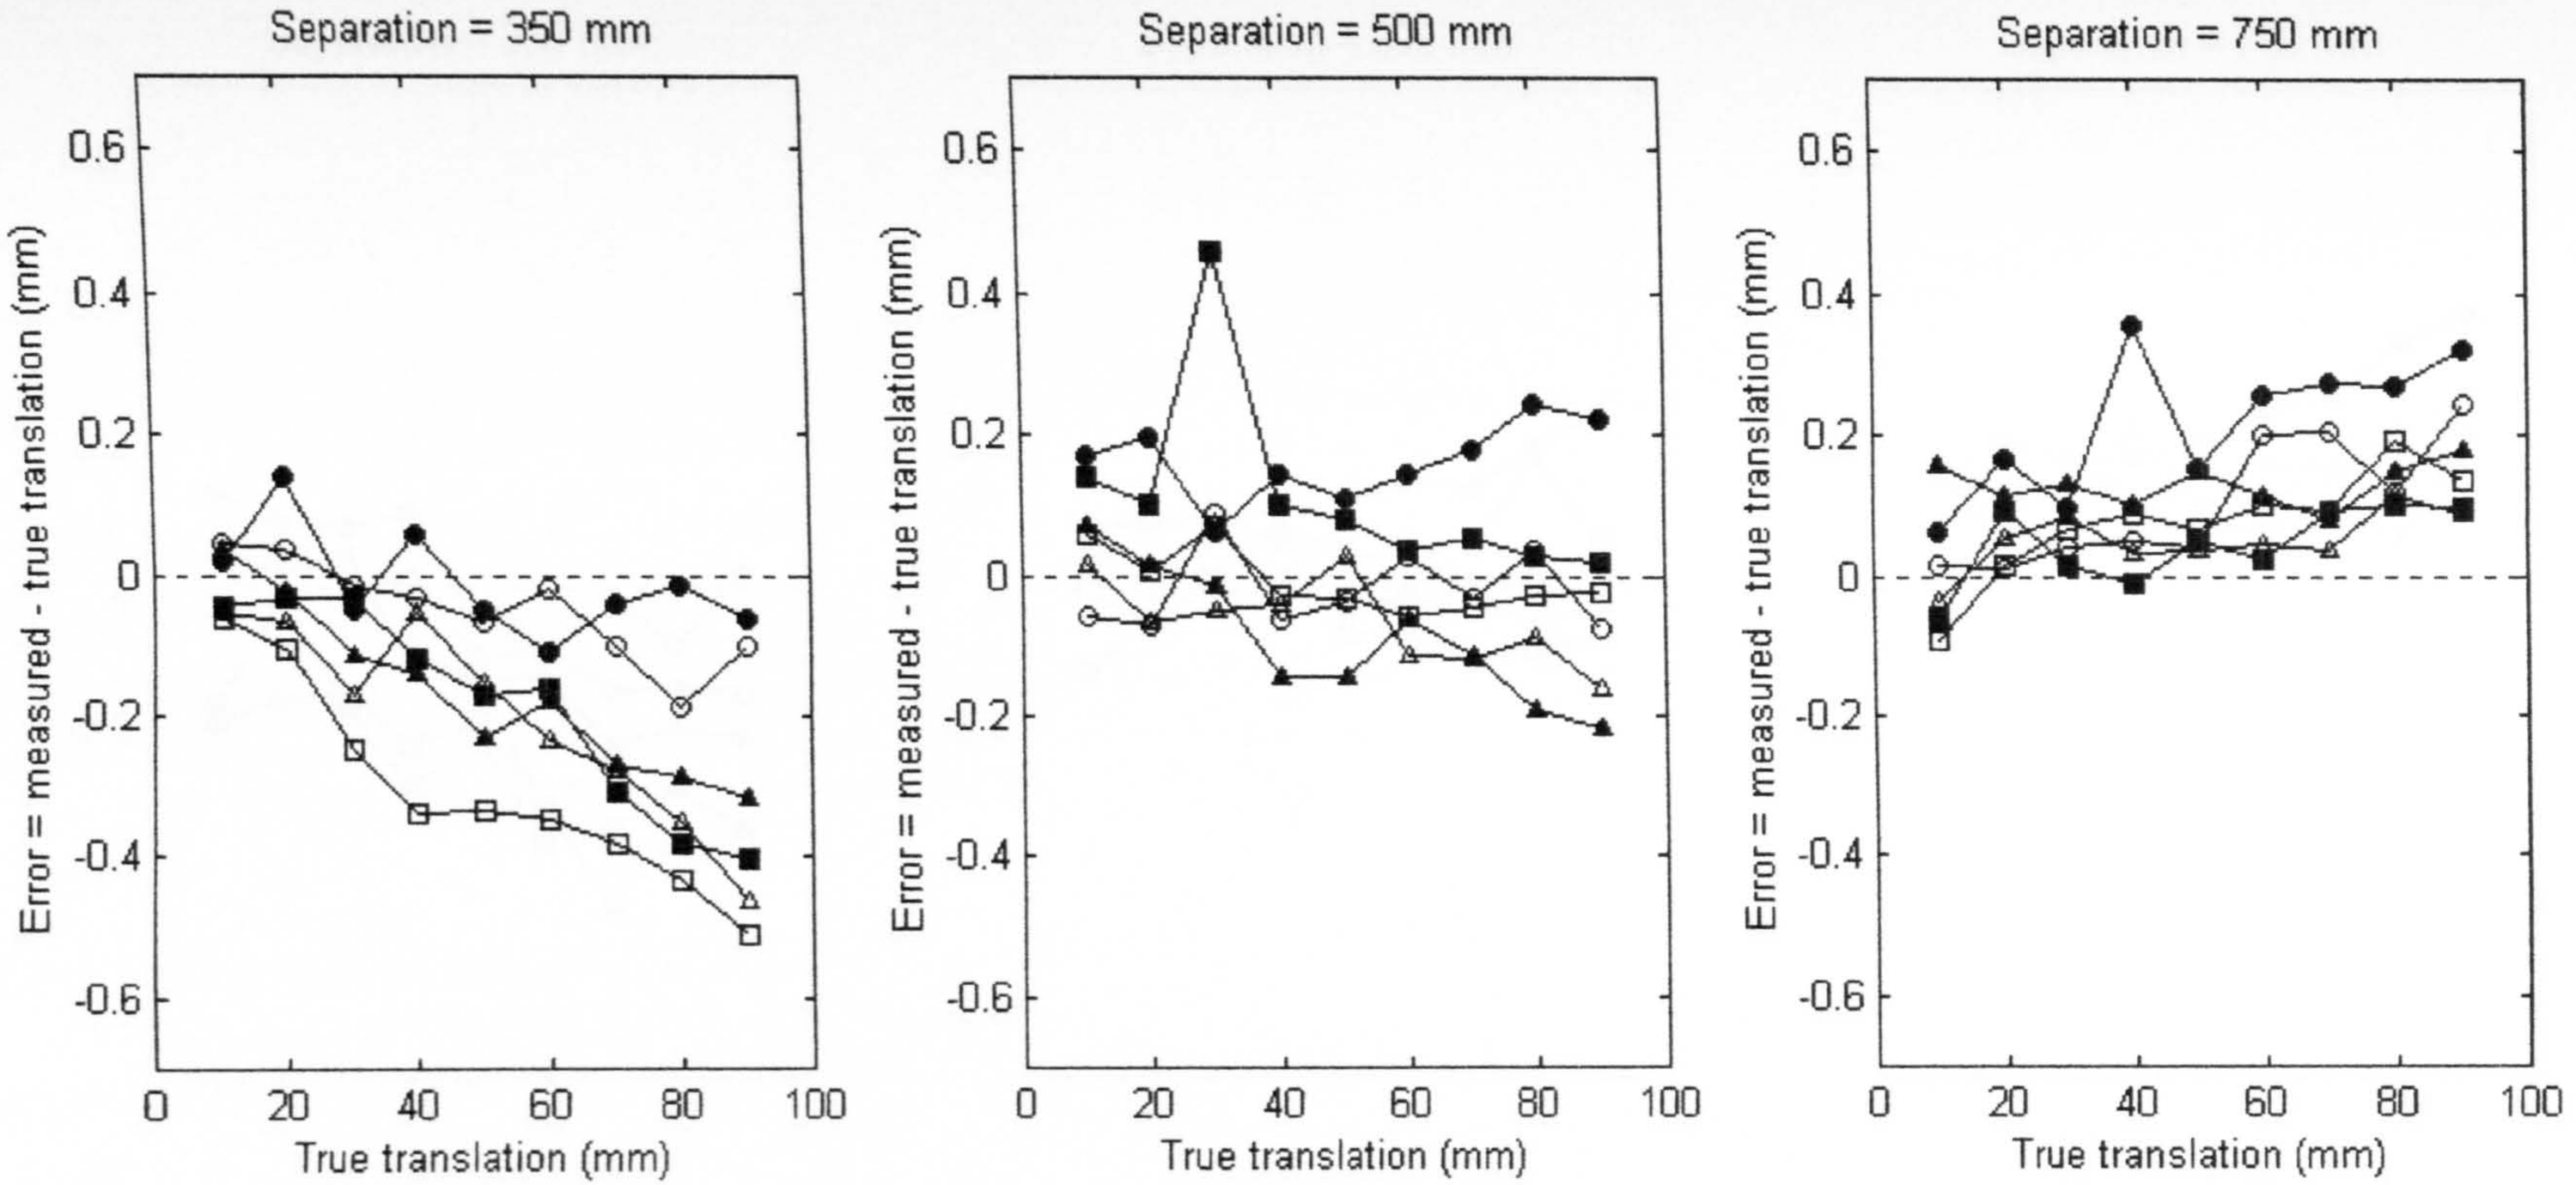


Figure 5.13 (a) The absolute errors in *pcBIRD* position measurements over 90 mm translations for the receiver alone.

Key: ■ – translations parallel to the x-axis moving towards transmitter; □ – translations parallel to the x-axis moving away from transmitter; ▲– translations parallel to the y-axis moving towards transmitter; Δ – translations parallel to the y-axis moving away from transmitter; ● – translations parallel to the z-axis moving towards transmitter; ○ – translations parallel to the z-axis moving away from transmitter.

(*continued overleaf*)

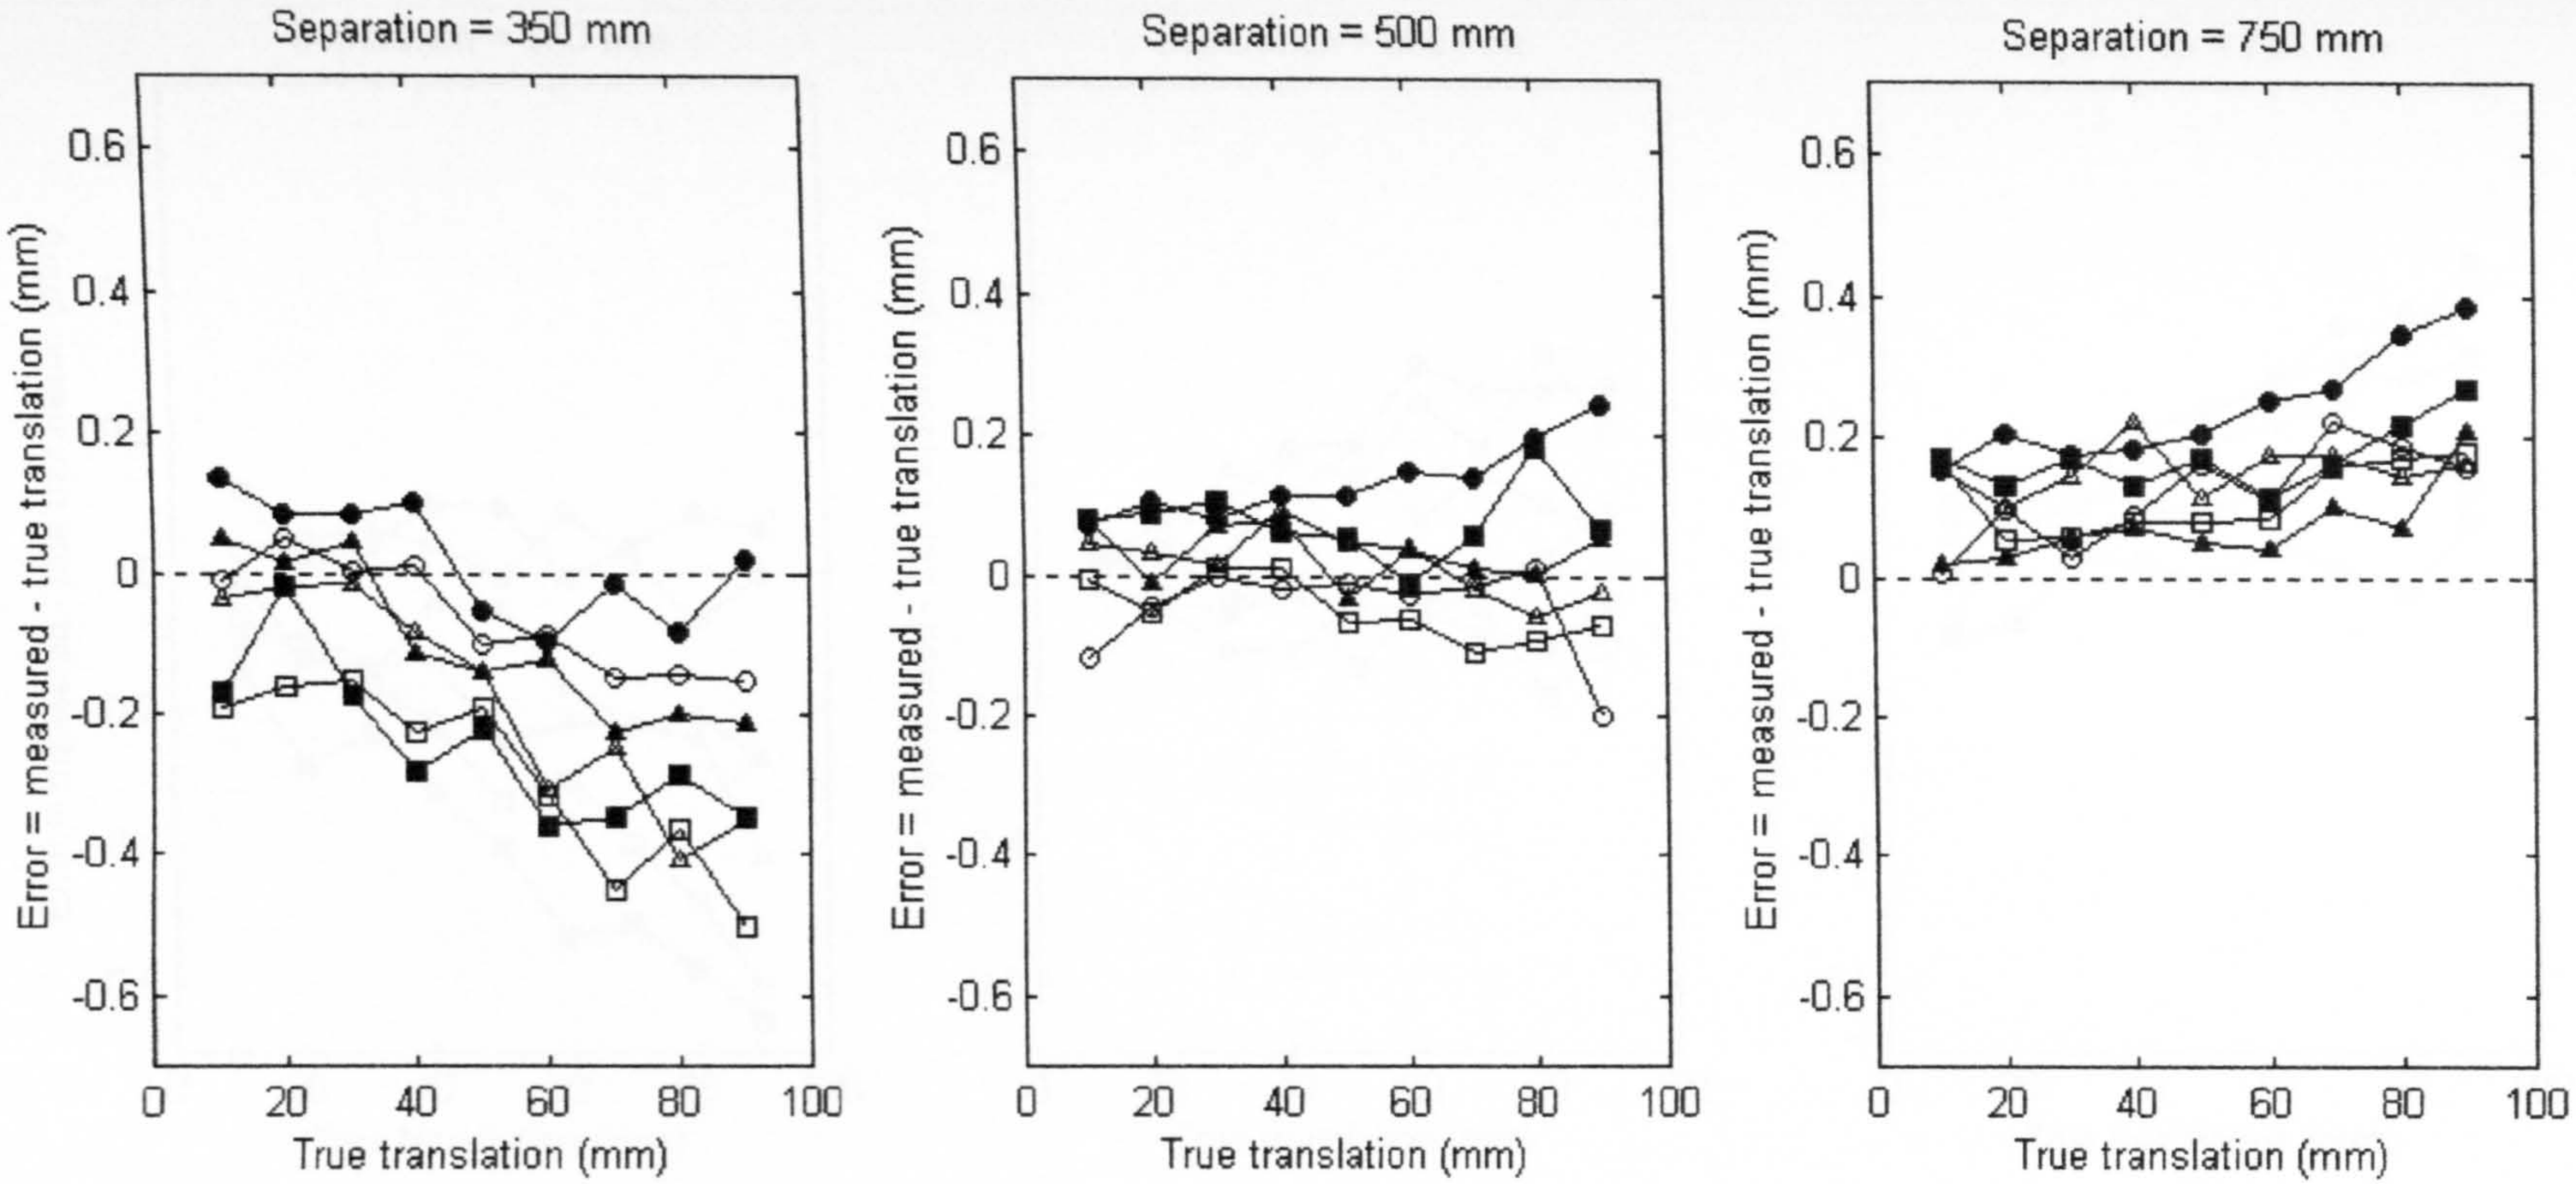


Figure 5.13 (b) The absolute errors in *pcBIRD* position measurements over 90 mm translations for the receiver mounted on the scan-probe and facing towards the transmitter. **Key:** ■ – translations parallel to the x-axis moving towards transmitter; □ – translations parallel to the x-axis moving away from transmitter; ▲– translations parallel to the y-axis moving towards transmitter; Δ – translations parallel to the y-axis moving away from transmitter; ● – translations parallel to the z-axis moving towards transmitter; ○ – translations parallel to the z-axis moving away from transmitter.

(*continued overleaf*)

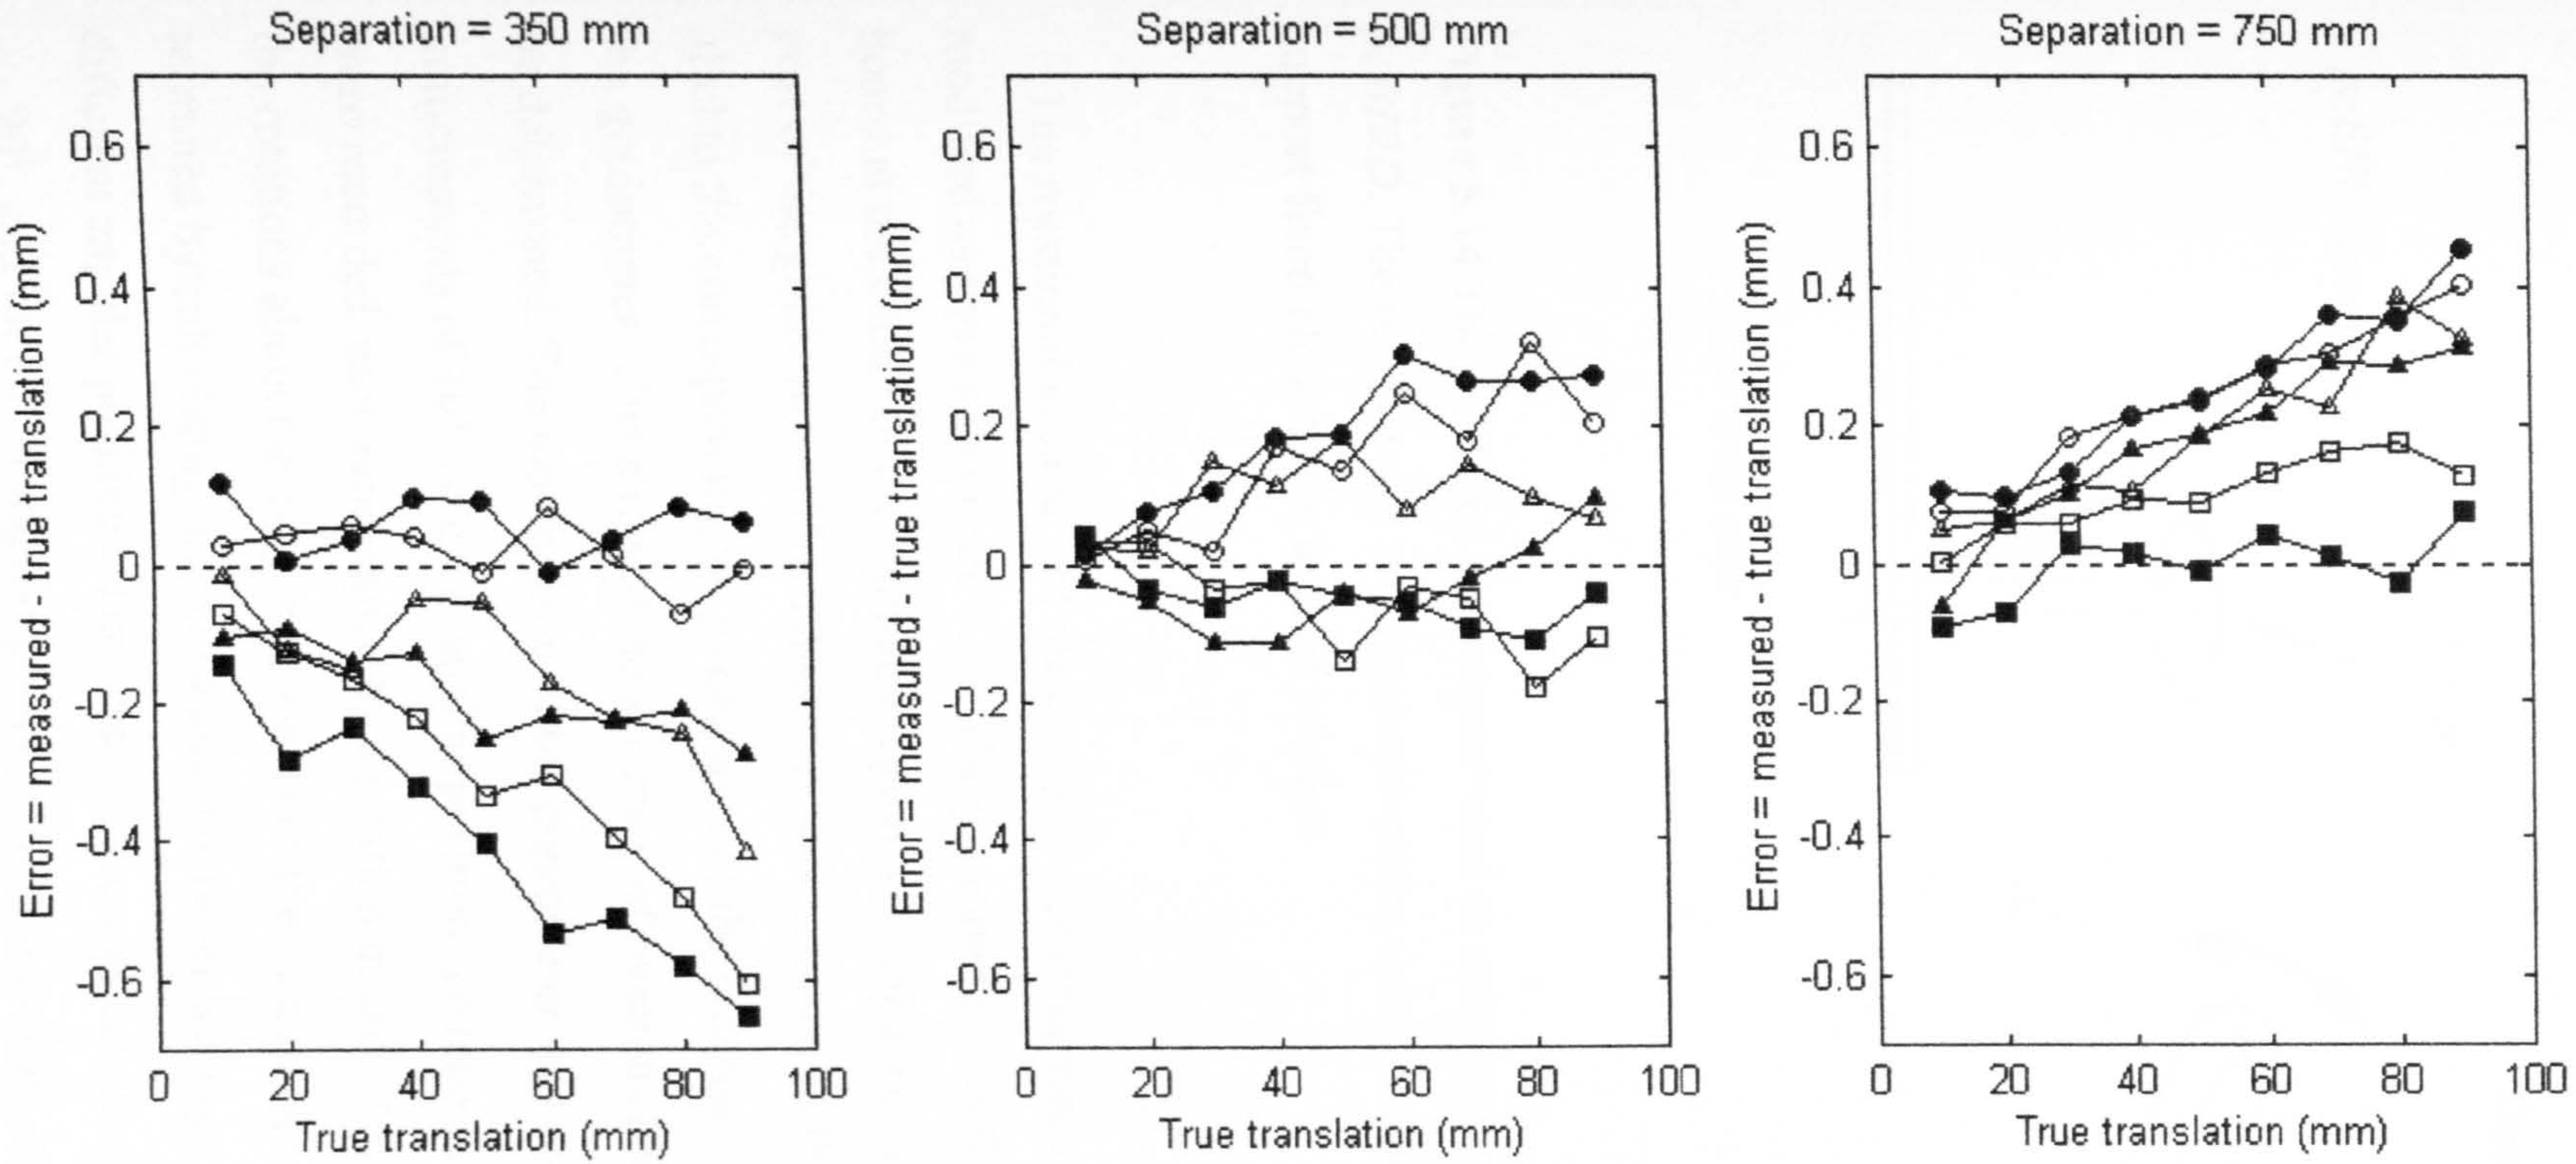


Figure 5.13 (c) The absolute errors in *pcBIRD* position measurements over 90 mm translations for the receiver mounted on the scan-probe and facing away from the transmitter. **Key:** ■ – translations parallel to the x-axis moving towards transmitter; □ – translations parallel to the x-axis moving away from transmitter; ▲– translations parallel to the y-axis moving towards transmitter; Δ – translations parallel to the y-axis moving away from transmitter; ● – translations parallel to the z-axis moving towards transmitter; ○ – translations parallel to the z-axis moving away from transmitter.

5.3.3 Accuracy of rotation measurements

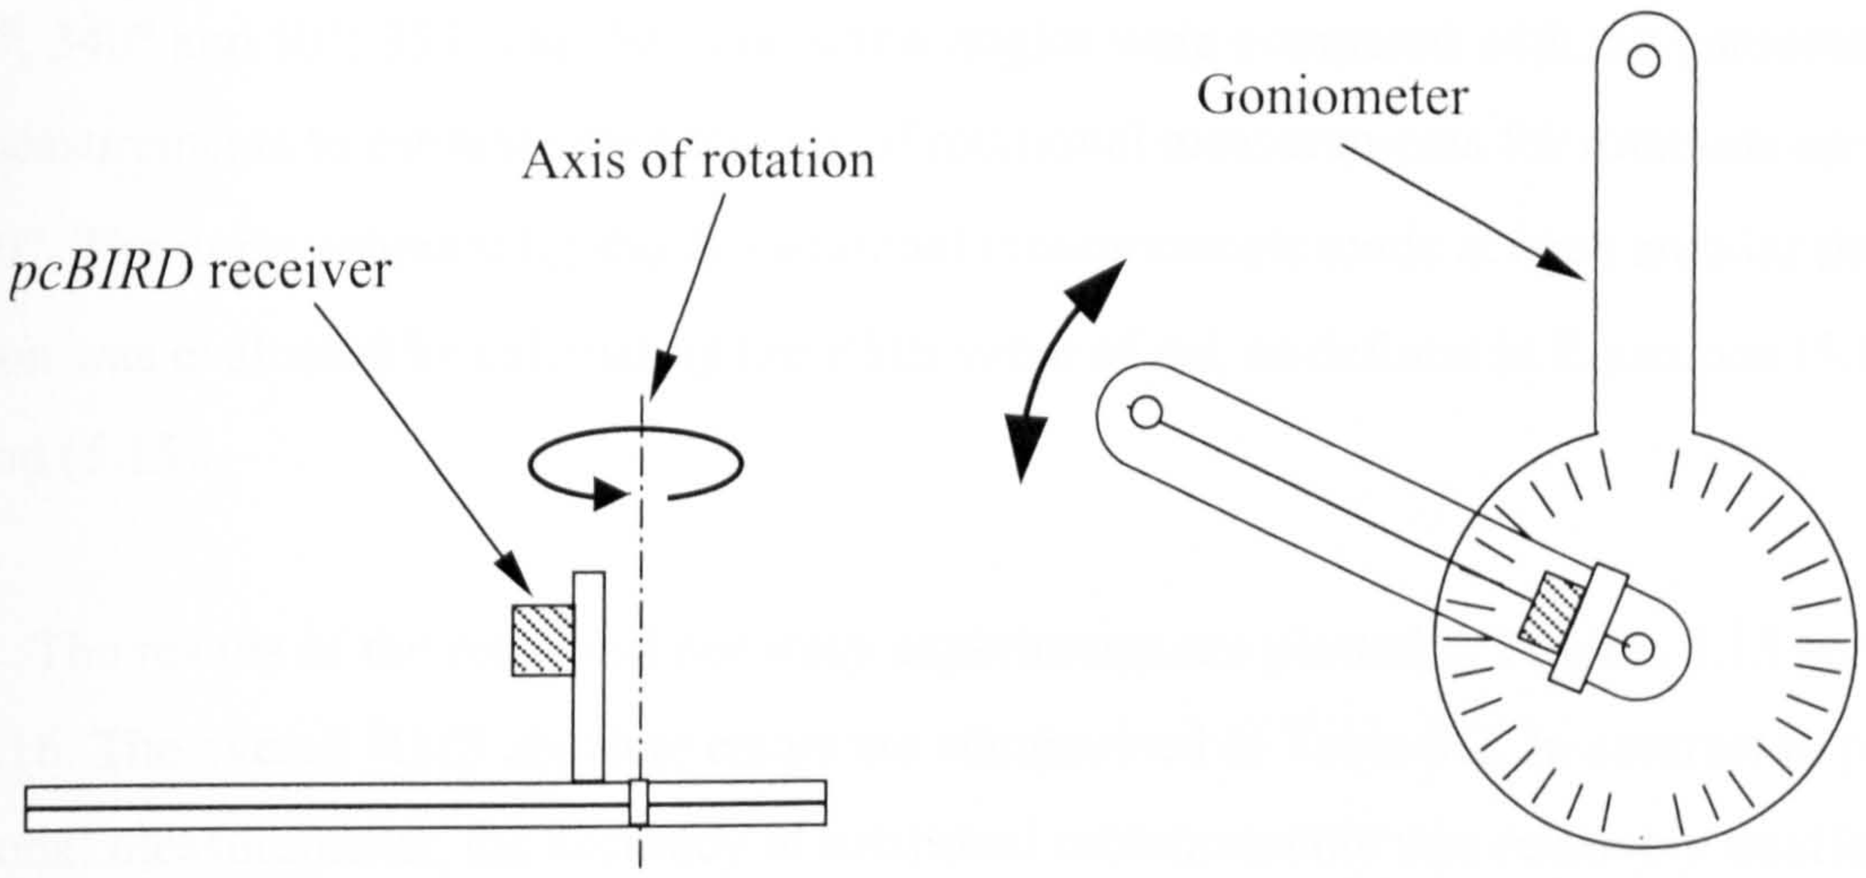


Figure 5.14 The modified goniometer used to estimate the rotational accuracy of the *pcBIRD*. The receiver (or receiver/scan-probe combination) was attached to the upright support fixed close to the axis of rotation.

The rotational accuracy of the *pcBIRD* was evaluated using a plastic goniometer modified as shown in Figure 5.14. The goniometer was fixed in position on the Nylatron board at the centre of the scanning region with the receiver attached to the upright support of the goniometer. The transmitter was oriented so that the x-, y-, or z-axis was parallel to the corresponding x-, y- or z-axis of the receiver, and to the axis of rotation of the goniometer. This allowed the accuracy of rotations about all three reference axes to be determined. The moving arm of the goniometer was rotated clockwise through 360° in increments of 10°. At each angular position, 20 receiver orientation measurements were recorded, each represented by a rotation matrix. This procedure was repeated for the rotations about the x-, y- and z-axis of the transmitter. The rotational error was determined by calculating the screw angle between average rotation matrices measured at different angular positions. The screw angle was calculated for the true rotations, 10, 20, …, 90°, and for all 36 angular positions from 0 to 360°. For example, for a true rotation

of 30°, the rotational error was calculated between rotation matrices recorded at the following 36 angular position pairs: 0° and 30°; 10° and 40°; ... ; 320° and 350°; 330° and 0°; 340° and 10°; 350° and 20°. The screw angles were compared with the goniometer measurements to estimate the accuracy of rotational measurements for rotations up to 90°. The noise estimate for the 20 rotational measurements made at each angular position was evaluated by calculating the RMS value of $\Delta\phi$, as defined in Equations (5.14) and (5.15).

The results of the rotational accuracy experiments are plotted in Figures 5.15 and 5.16. The overall RMS absolute errors are summarised in Table 5.2. In contrast to positional measurements, the accuracy of rotational measurements was relatively unaffected by either the scanner location or the presence of the scan-probe. There was, however, a tendency for the range of absolute error to increase with increasing true rotation angle up to 90°, as can be seen in Figure 5.15. This had the observed effect of an increasing RMS absolute error as the magnitude of the true rotation angle increased. The reason for this is unknown, although there may be a geometric explanation.

	Scanner-transmitter separation (mm)		
	350	**500**	**750**
Receiver alone	0.284	0.286	0.311
Receiver + scan-probe combination	0.288	0.292	0.258

Table 5.2 RMS absolute error (°) for rotational measurements for rotations up to 90°.

The RMS absolute rotation error with the receiver mounted on the scan-probe and all scanner locations was 0.280° (range: -0.842 to 0.868°). Again, this was well within the manufacturer's specification of 0.500° RMS error. The mean noise estimates for orientation measurements were 0.019° and 0.021° for the receiver alone and receiver plus scan-probe combination, respectively.

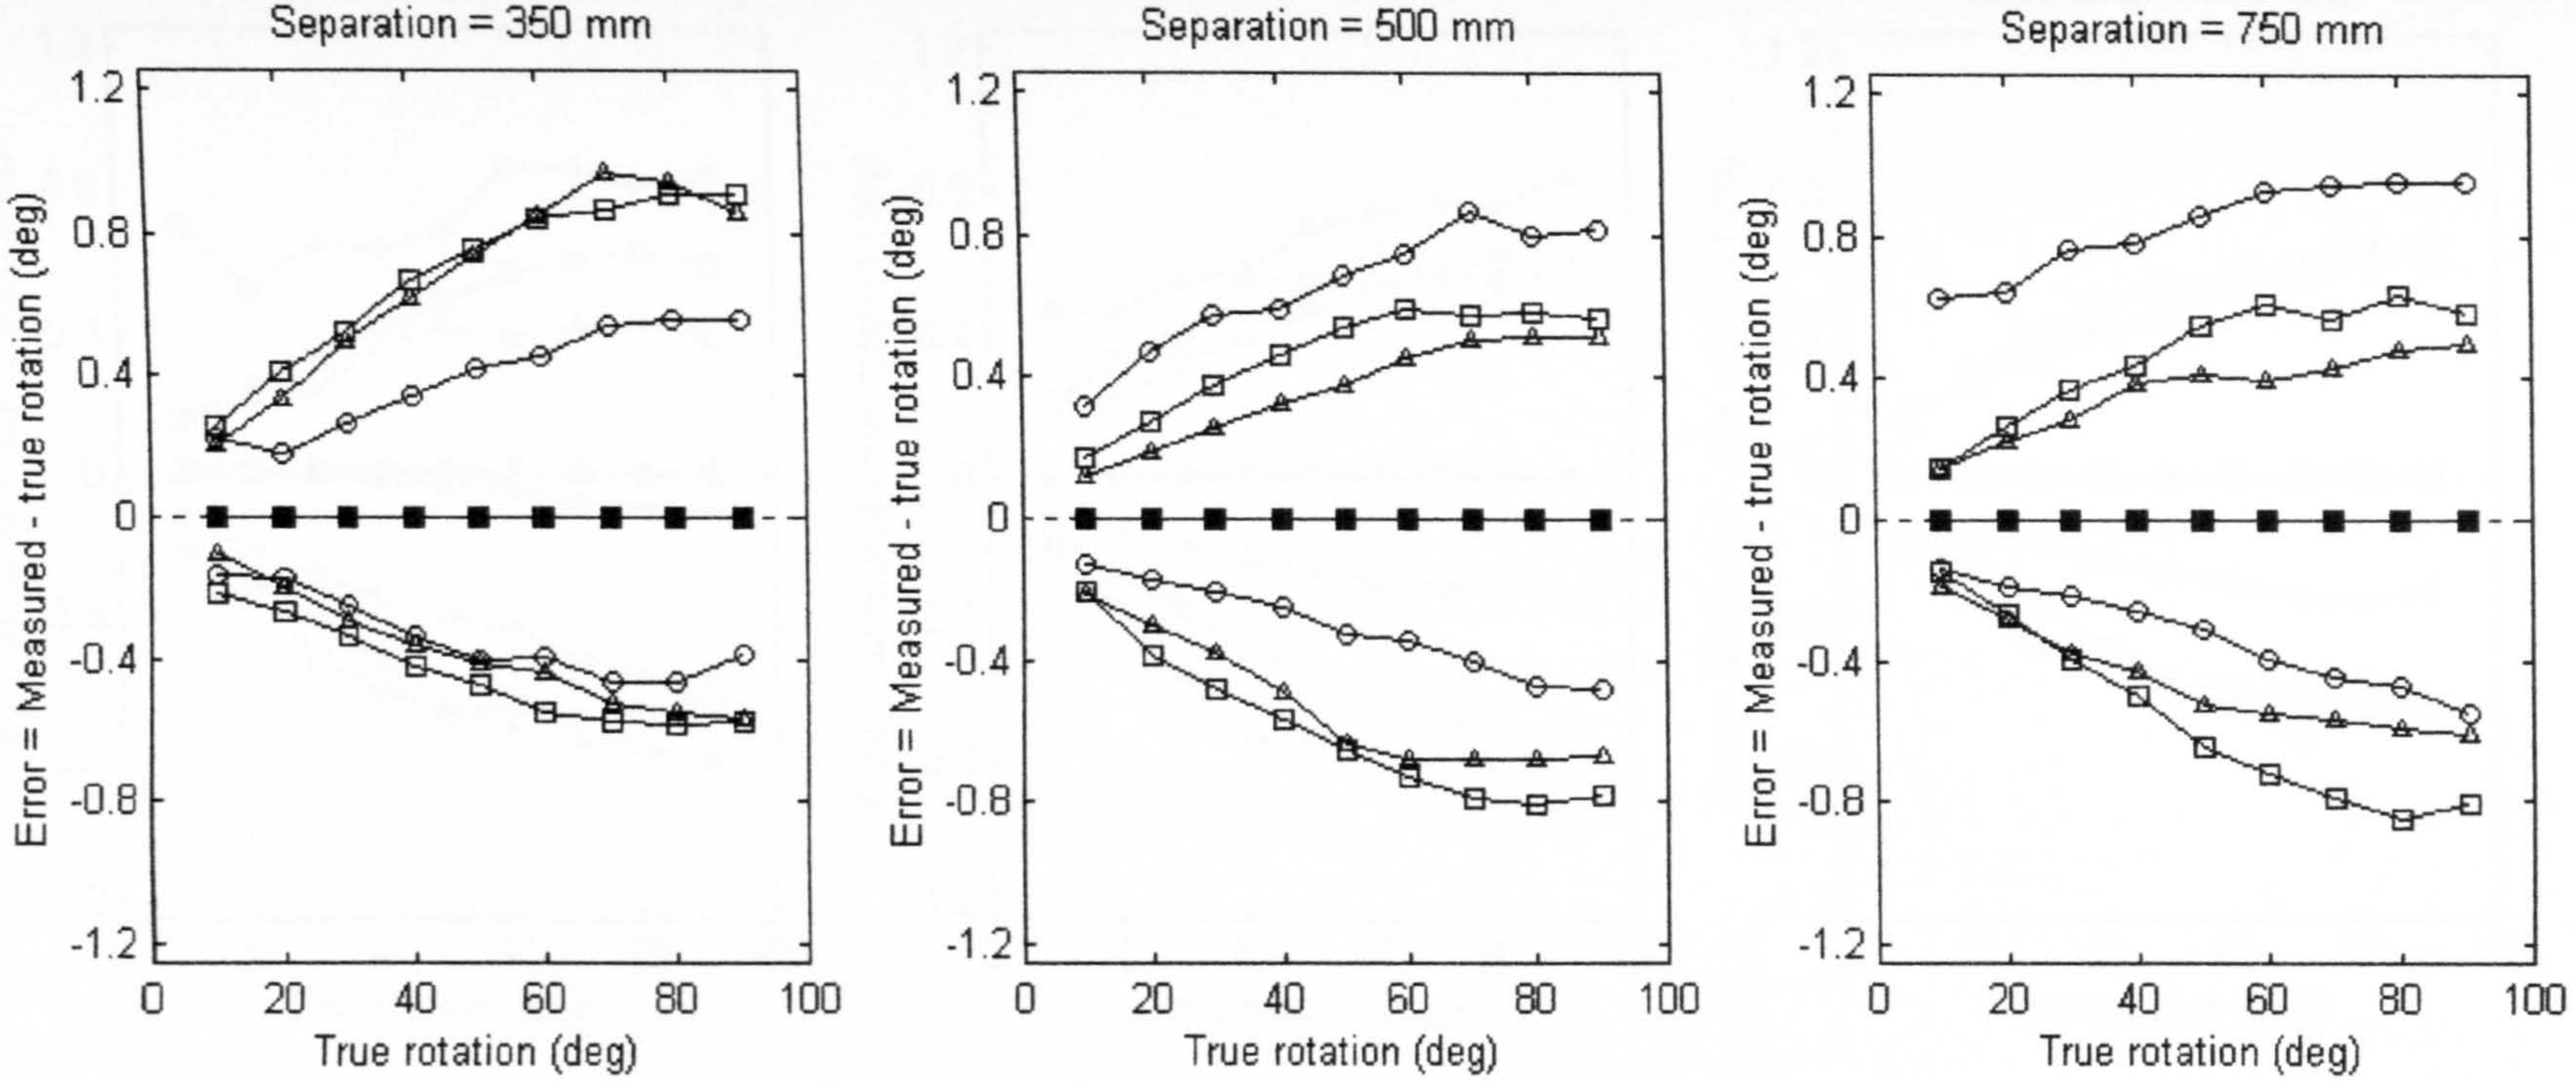


Figure 5.15 (a) The absolute errors in *pcBIRD* rotation measurements for the receiver alone and 3 different scanner-transmitter separations.

Key: ■ – mean error for rotations about the x-axis; □ – maximum/minimum error for rotations about the x-axis; ▲– mean error for rotations about the y-axis; Δ – maximum/ minimum error for rotations about the y-axis; ● – mean error for rotations about the z-axis; ○ – maximum/minimum error for rotations about the z-axis.

(*continued overleaf*)

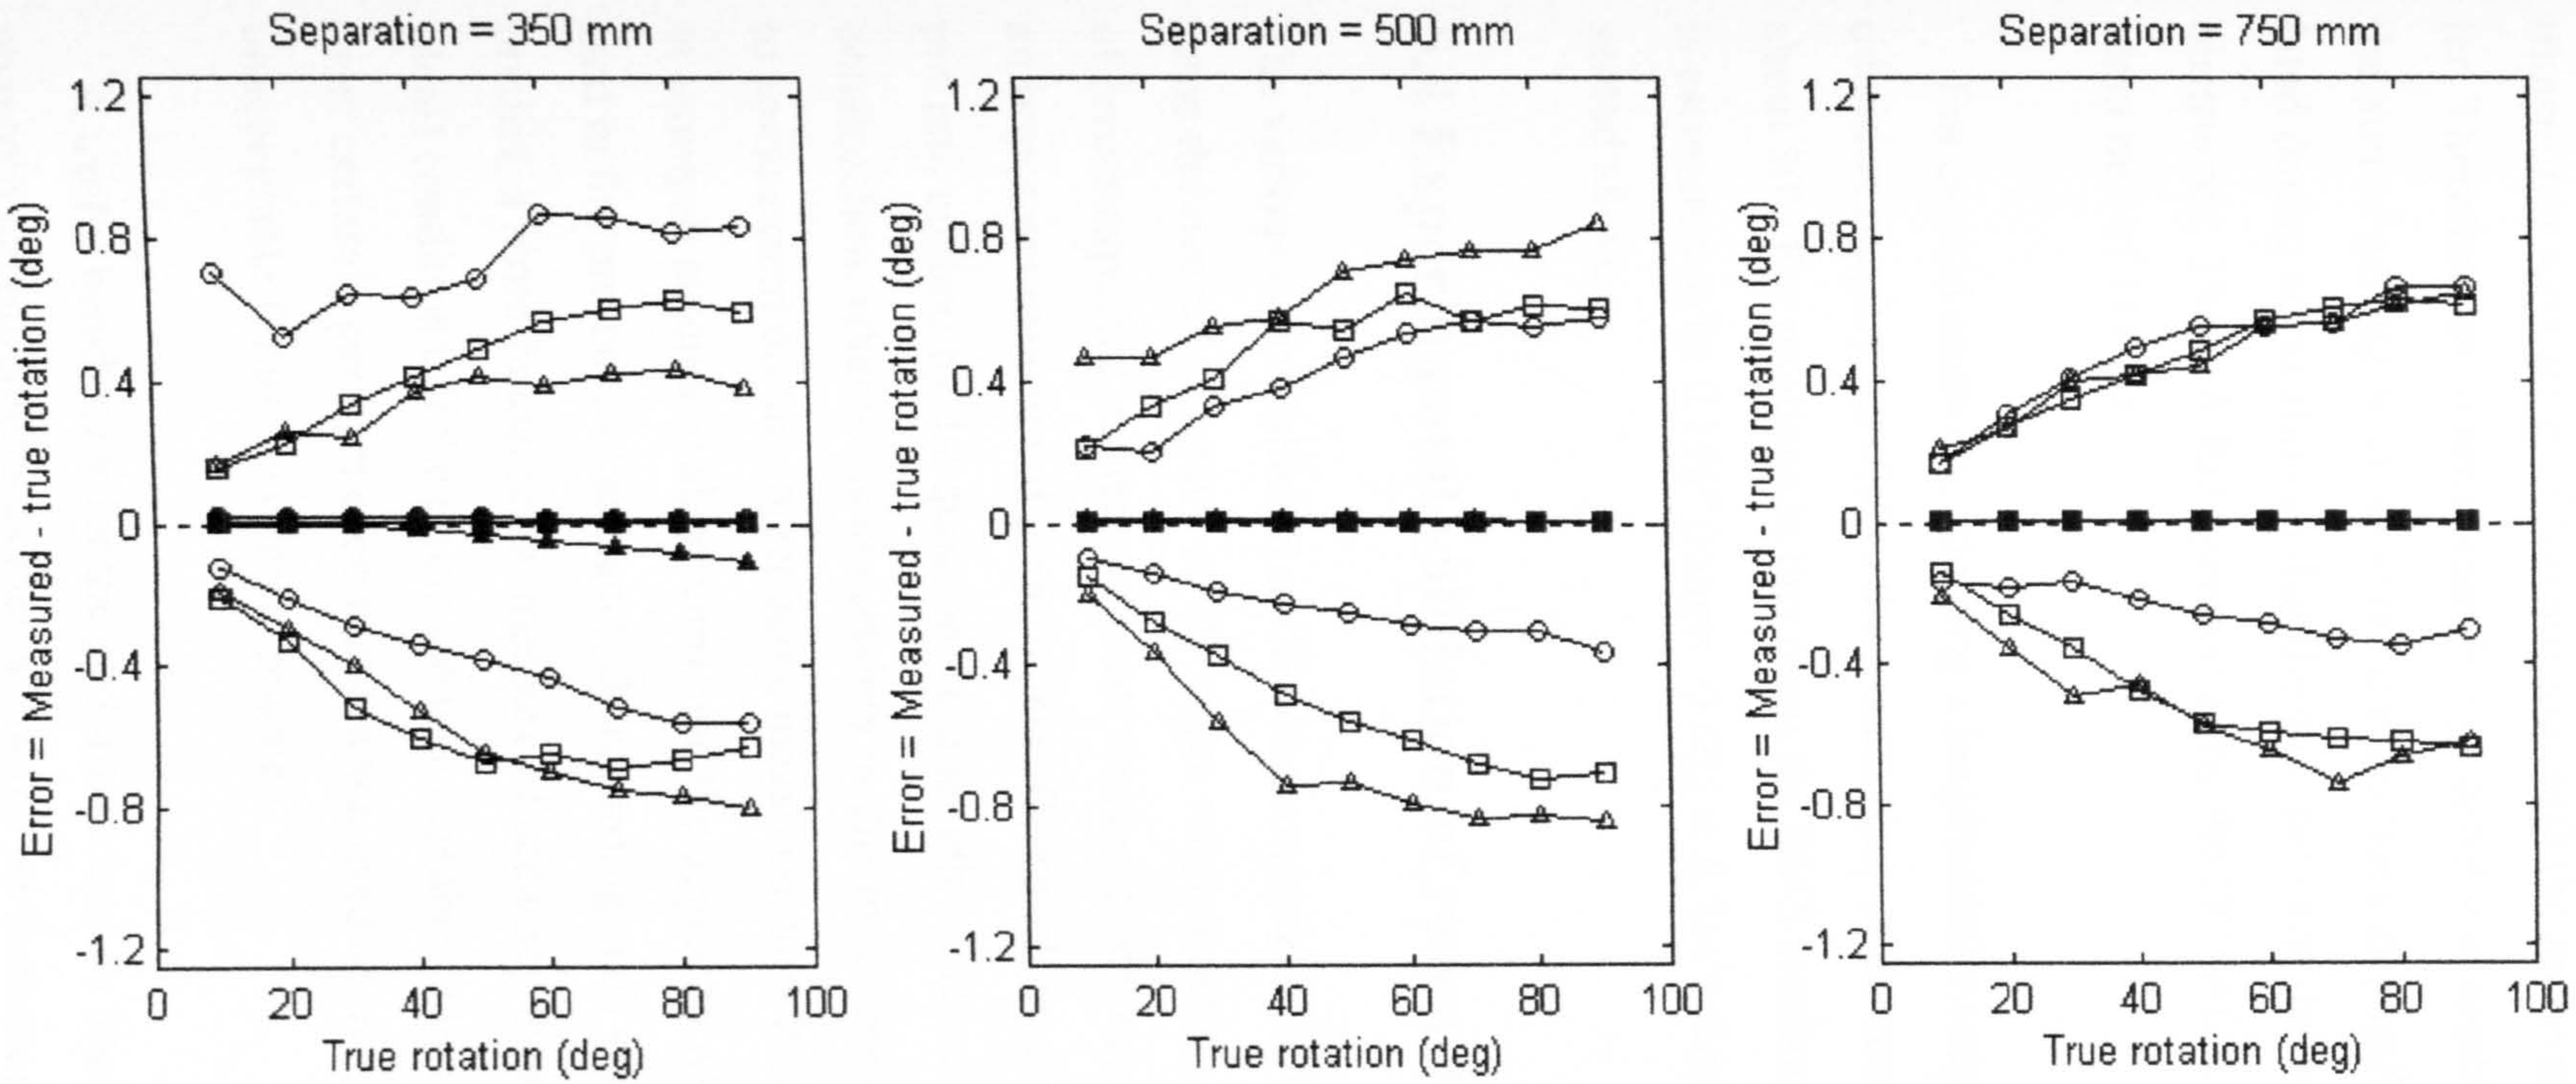


Figure 5.15 (b) The absolute errors in *pcBIRD* rotation measurements for the receiver mounted on the scan-probe and 3 different scanner-transmitter separations. **Key:** ■ – mean error for rotations about the x-axis; □ – maximum/minimum error for rotations about the x-axis; ▲– mean error for rotations about the y-axis; Δ – maximum/ minimum error for rotations about the y-axis; ● – mean error for rotations about the z-axis; ○ – maximum/minimum error for rotations about the z-axis.

5.3.4 Optimal protocol for carotid 3DUS scans

The results above suggest that the RMS position error was lowest for a scanner-transmitter separation of 500 mm and with the receiver facing towards the transmitter. Therefore, an optimal scanning protocol was established for carotid imaging with the aim of getting as close to these conditions as possible. The scan-probe was orientated during acquisition so that the path between the transmitter and receiver was not obstructed. The *pcBIRD* operated using the DC+ACN filter setting and a measurement rate of 111 Hz.

For carotid scanning, adopting a longitudinal sweep of the probe during a 3D acquisition roughly parallel to the x-axis, scan-probe rotations are limited to less than about 30°. For rotations in this range the RMS rotation error for all scanner locations is estimated to be 0.156° (range: -0.563 to 0.711°) from the results of the test presented above.

5.4 Experimental validation of reconstruction accuracy

The various errors inherent in the electromagnetic 3DUS system combine to determine the overall accuracy of 3D image reconstruction. Estimation of the magnitude of errors expected in measurements derived from reconstructed 3DUS data provides an important measure of performance. One such measurement is length, and may be put into context by the question: 'If a length is measured to be 20 mm in a 3DUS reconstruction, what is the uncertainty in this measurement?' The answer is dependent to some extent on the object being measured, how it is being measured (i.e. manually or using an automatic algorithm), and the degree to which the structure can be visualised in the presence of artifacts. Disregarding such variables, which are difficult to predict, a more fundamental measure of accuracy is provided under ideal (or near-ideal) conditions. If the errors are unacceptable in this case then the system will almost certainly perform even worse when used *in vivo*, and, therefore, errors will be unacceptable for use in clinical practice.

A useful benchmark for the performance of an imaging system is the accuracy of phantom measurements as they allow the performance of a system to be evaluated against some standardised criteria. This makes it possible to objectively assess per-

formance, and also to directly compare the performances of different systems. An essential aspect of the validation of the 3DUS system developed in this thesis was an experimental investigation of the accuracy of 3D reconstruction using a specialised ultrasound phantom, named the 'wedge phantom'.

5.4.1 Wedge phantom experiments

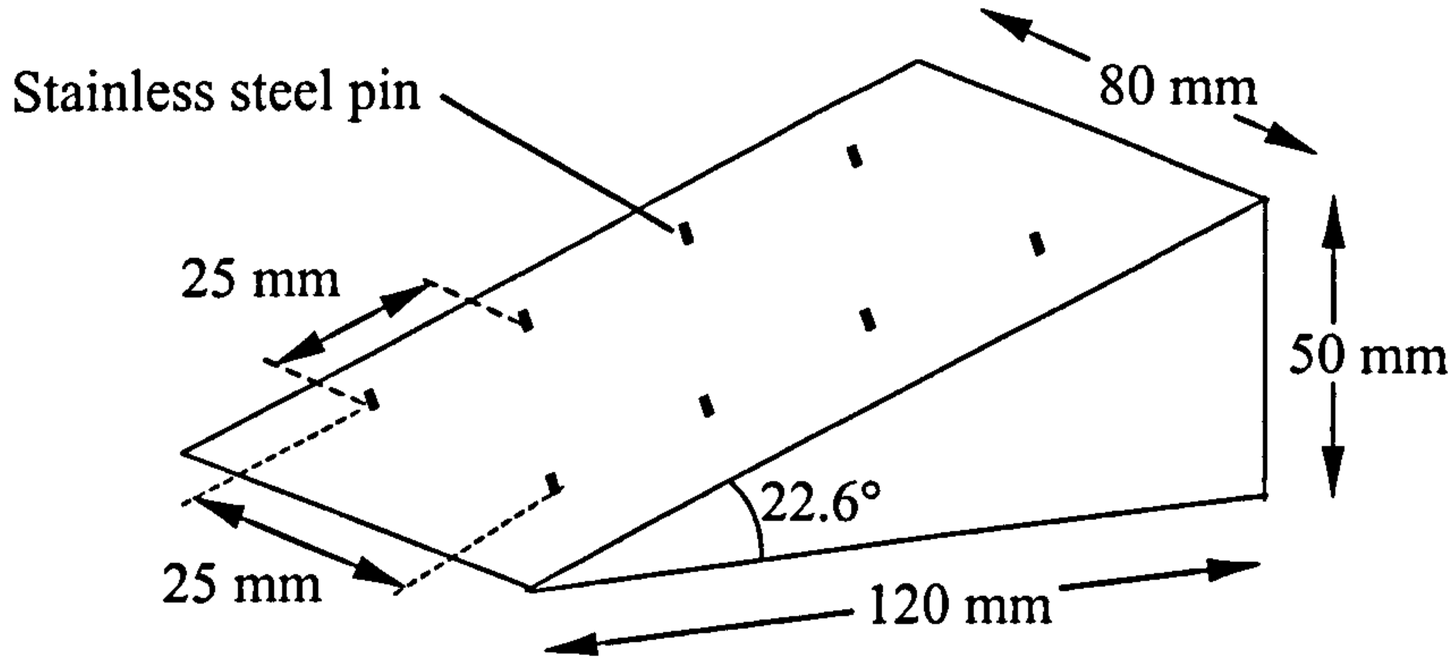


Figure 5.16 The 'wedge phantom' designed for assessing the reconstruction accuracy of the electromagnetic 3DUS system.

The wedge phantom, shown in Figure 5.16, was designed for the purpose of assessing the reconstruction accuracy of the system. The phantom was constructed from a tough plastic wedge with a 4×2 matrix of 0.5 mm stainless steel pins inserted perpendicularly into the inclined surface. The holes into which the pins were inserted were precision drilled 25 mm apart and to a depth so that 1 mm of each pin was exposed. The head of each pin forms a point-target for ultrasound imaging.

The reason for arranging the pins on an inclined surface is so that they appeared at a range of depths through the volume. Therefore, inter-pin measurements would be more representative of the real clinical imaging situation, where measurements are not necessarily made along the primary axes of the volume. The dimensions of the wedge phantom were chosen with carotid scanning in mind, but the phantom is of a size suitable for many other applications.

Two sets of experiments were carried out using the wedge phantom. During all experiments the L12-5 scan-probe was used. The wedge phantom was scanned in a

water bath of known temperature, and axial distance measurements were corrected in the same way as those of the pin phantom (see Section 4.3.3). For all experiments, the phantom was positioned on the examination couch within the region appropriate to carotid imaging.

The first experiment involved scanning the matrix of pins at different imaging depths with the wedge phantom tilted so that the plane of the pins was parallel to the bottom of the water bath, as depicted in Figure 5.17. The purpose of this experiment was to investigate whether the accuracy of inter-pin distance measurements, and/or the precision in reconstructed pinhead locations, was dependent on scan depth. Each of the 4 rows of 2 pins was imaged multiple times by slowly rotating the scan-probe through angles < 30° about each of the axes of the *pcBIRD* receiver. B-mode images, together with their associated 3D positions and orientations, were acquired regularly at 300 ms intervals. The pin matrix was scanned at 4 different depth ranges: 0-10 mm, 10-20 mm, 20-30 mm, and 30-40 mm.

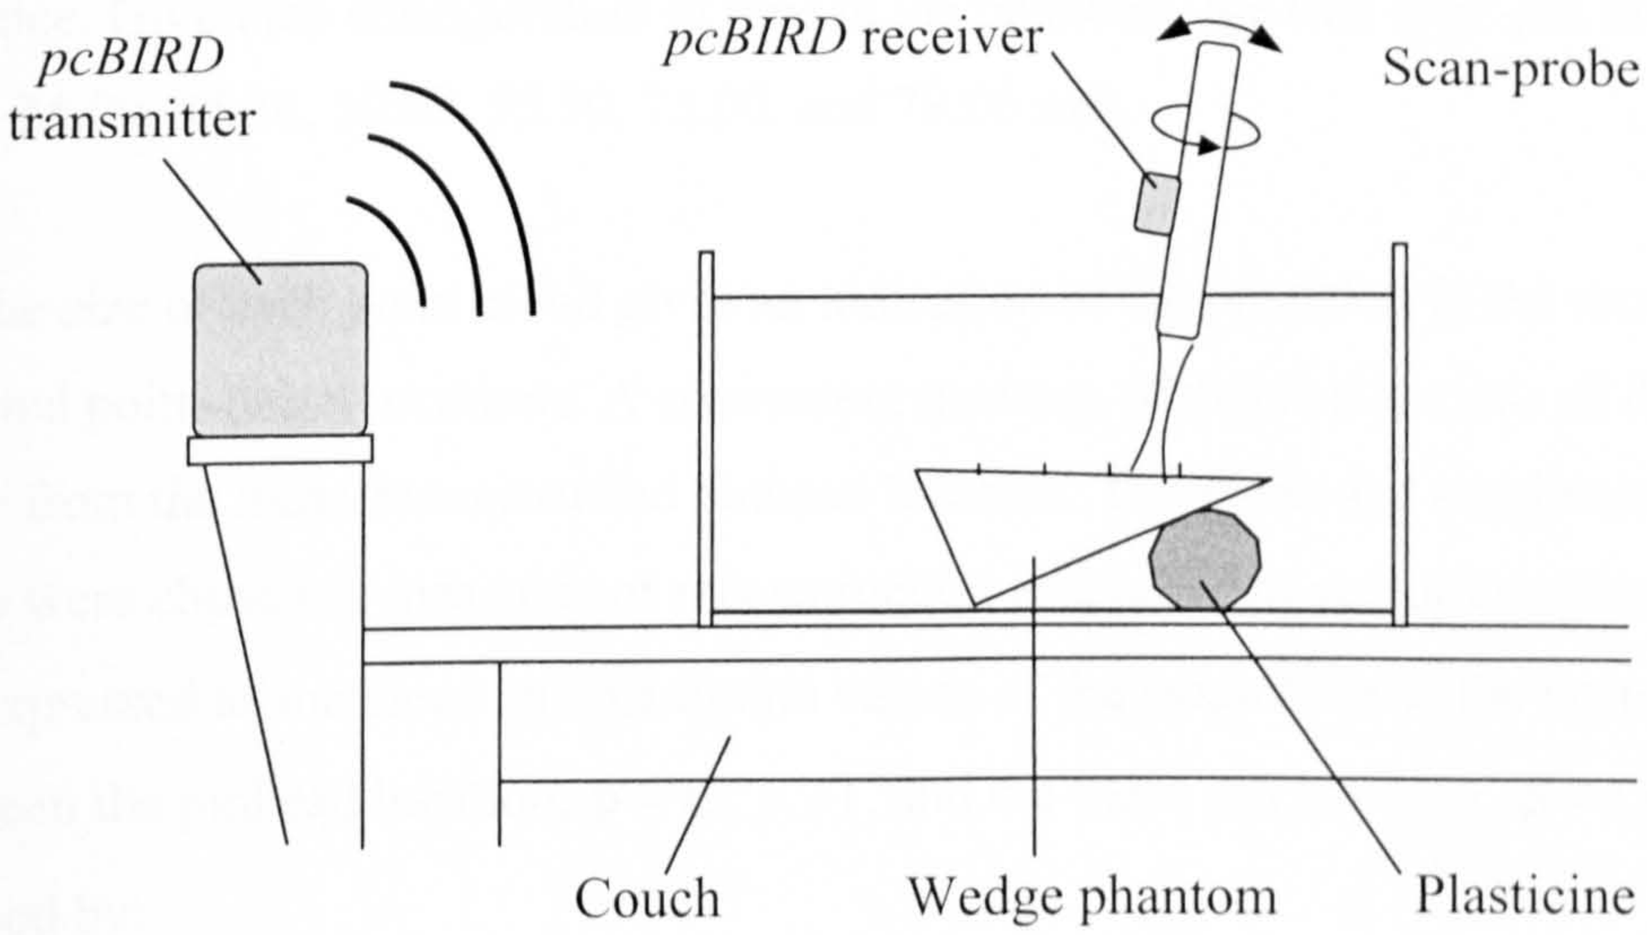


Figure 5.17 Imaging the wedge phantom at varying scan depths.

Figure 5.18 shows an example of a B-mode image of the wedge phantom. The within-image co-ordinates of the pinheads were found manually with the same software used in the calibration procedure (see Chapter 4). The locations of *all* imaged pinheads were identified, irrespective of their brightness in the image. This was considered to provide a more realistic indication of the errors encountered in clinical practice, since it takes into account the uncertainty of the pinhead location in the elevational direction of the ultrasound beam.

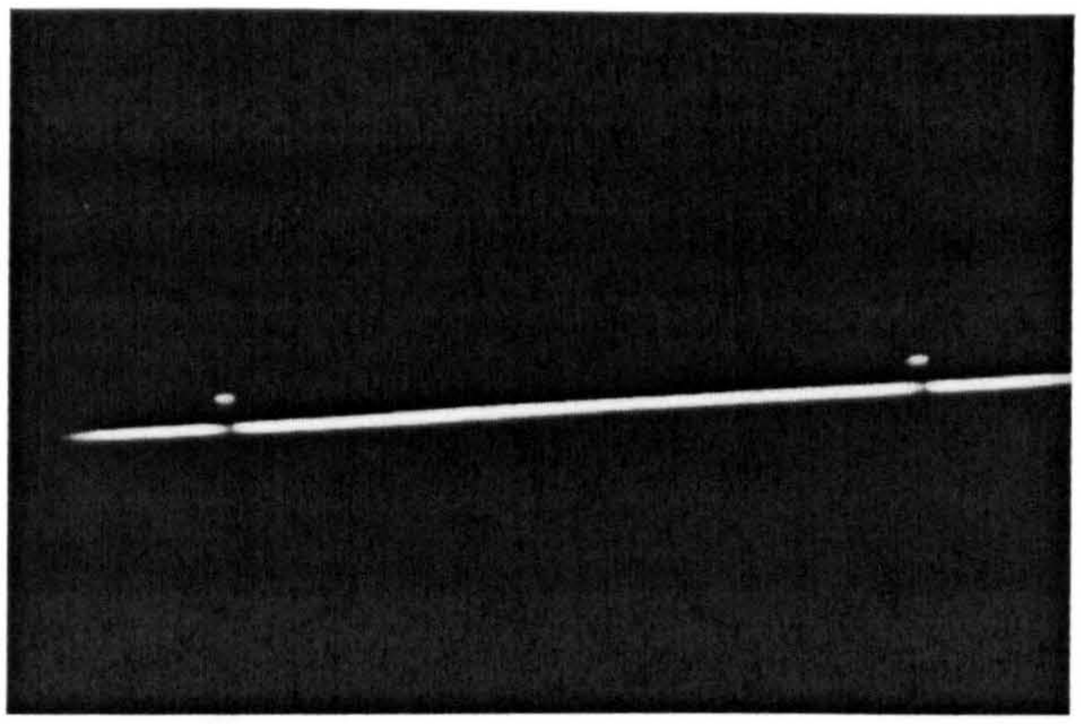


Figure 5.18 B-mode image of the wedge phantom

The identified pinhead co-ordinates were grouped and transformed into 3D co-ordinates. Thus, for each pinhead, there were multiple estimated 3D positions which can be thought of as forming a small 'cloud' of points around the true pin location, as shown in Figure 5.19. To determine the accuracy of inter-pin distance measurements, the distance between each pair of reconstructed pinhead locations drawn from two point clouds, in all possible combinations, was calculated and compared with the true distance. Given the configuration of pins in the phantom, the true inter-pin distances were 25.00, 35.36, 50.00, 55.90, 75.00, and 79.06 mm.

The size of each point cloud gives an indication of the precision in the reconstructed point-target locations. A convenient measure of this is a statistic of the distance from the mean reconstructed pinhead location. The mean and maximum distance were chosen as measures of reconstruction precision. In vector notation, these are expressed as the mean and maximum values of the magnitude of the vector, $\Delta\mathbf{p}$, between the pinhead location, $\mathbf{p} = (x, y, z)$, and the mean pin location, $\overline{\mathbf{p}} = (x, y, z)$, defined by:

$$\left|\Delta\mathbf{p}\right| = \sqrt{(x-\bar{x})^2 + (y-\bar{y})^2 + (z-\bar{z})^2}\ . \tag{5.21}$$

In addition, asymmetry in each point-cloud was revealed by examining the magnitudes of the x-, y-, and z-components of $\Delta\mathbf{p}$, given by:

$$\left|\Delta\mathbf{p}_x\right| = \left|x-\bar{x}\right|, \tag{5.22}$$

(*continued on page 171*)

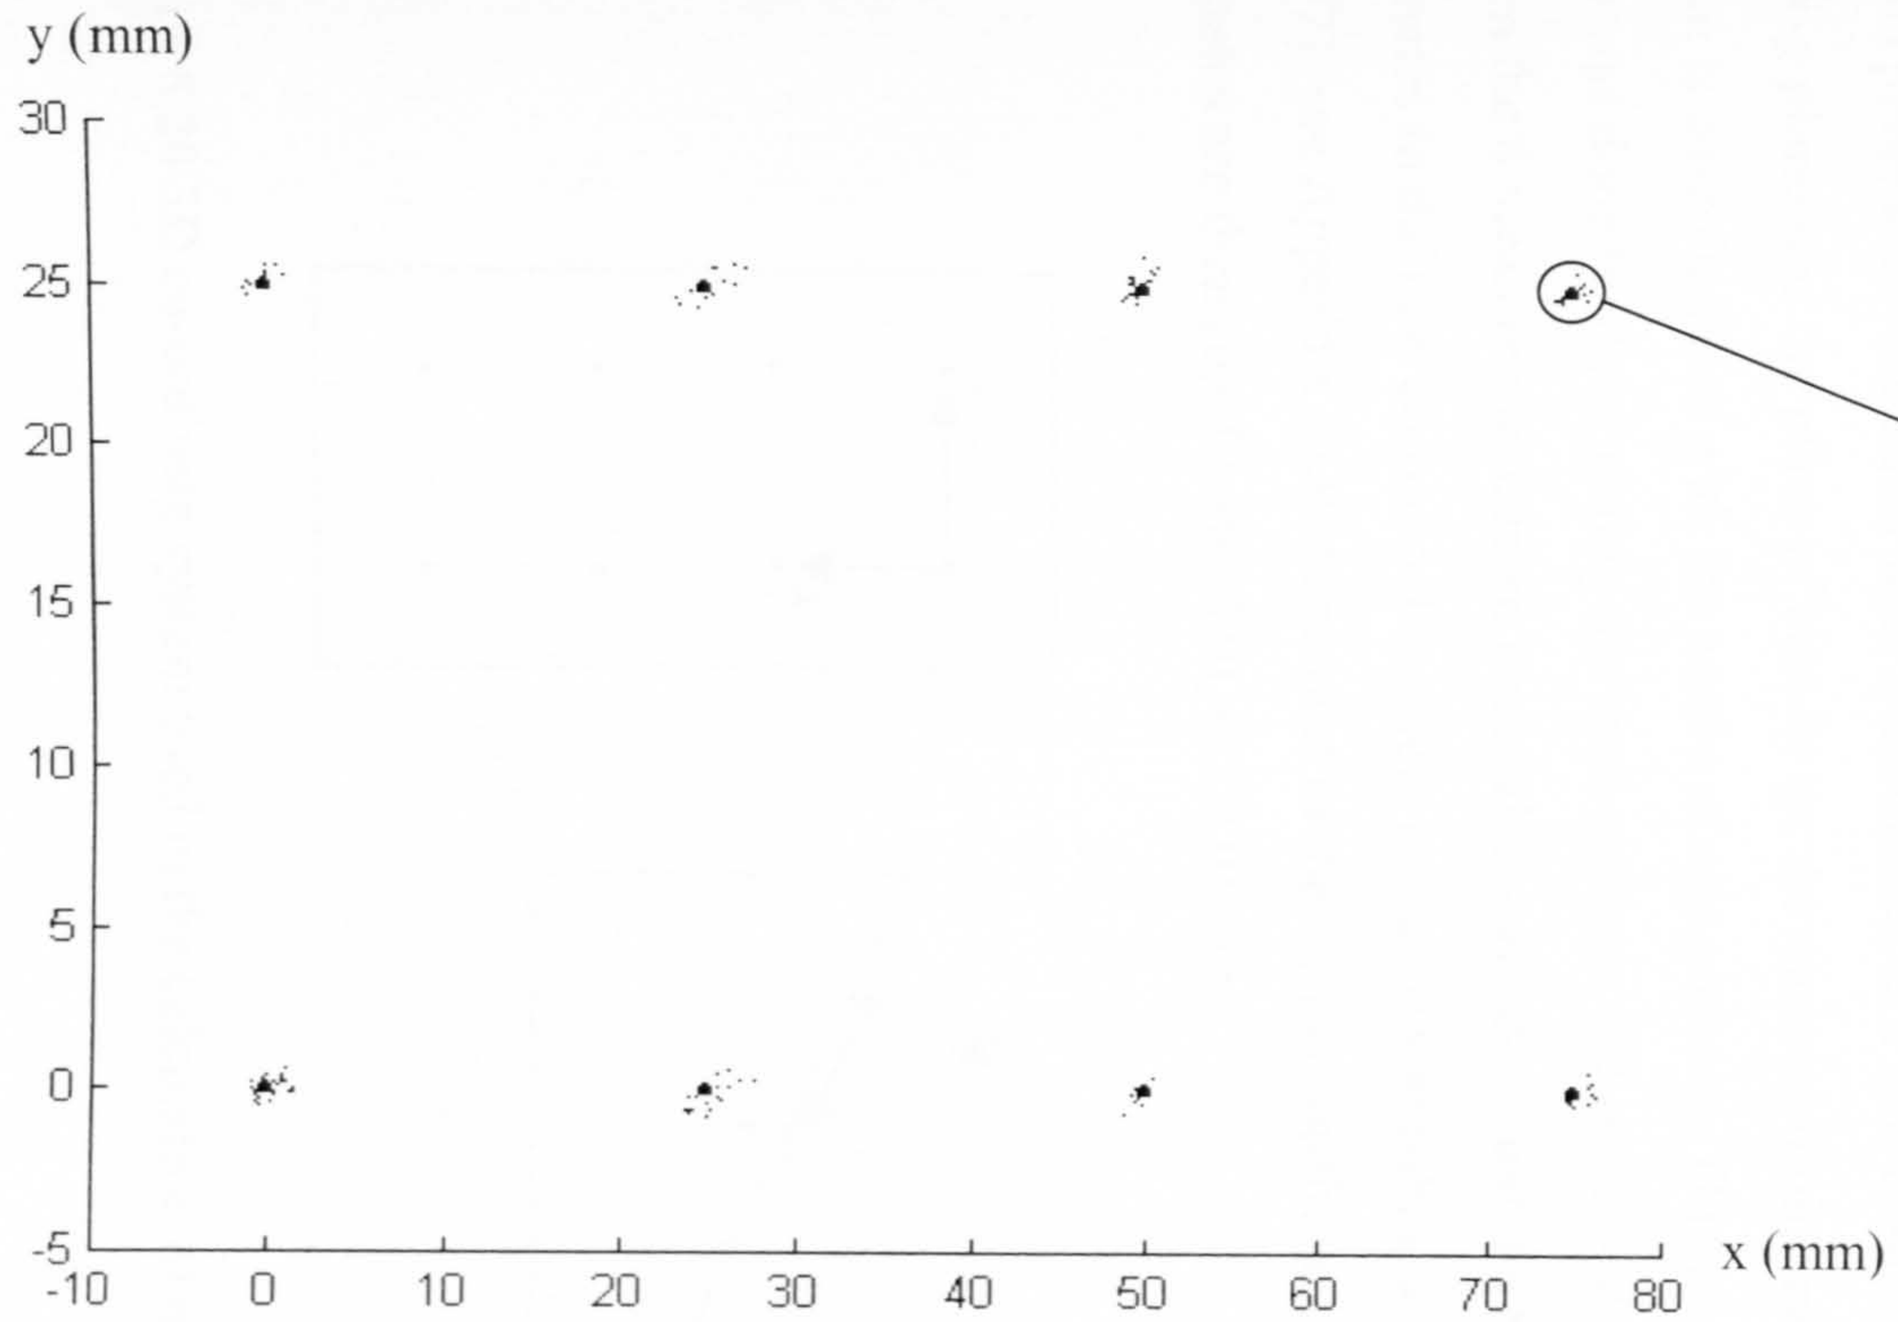


Figure 5.19 The reconstructed pinhead locations (•) form a cloud of points around the true location (shown on the right). The mean reconstructed location (denoted by '+') gives an estimate of the true location (denoted by 'o'), and the mean value of the magnitude of the vector $\Delta\mathbf{p}$ provides a measure of the reconstruction precision.

$$\left|\Delta \mathbf{p}_y\right| = \left|y - \bar{y}\right|, \tag{5.23}$$

and

$$\left|\Delta \mathbf{p}_z\right| = \left|z - \bar{z}\right|. \tag{5.24}$$

Such a component-wise analysis is useful to highlight any asymmetry in uncertainty which might arise as a consequence of the imaging process.

In order to provide a consistent frame of reference for each wedge phantom scan, all 3D pinhead co-ordinates were transformed to a new co-ordinate system wherein the x-y plane was the plane of best fit for the measured pinhead locations, and the origin is co-coincident with the centre of one pinhead. This is shown in Figure 5.20. This was done by determining the least-squares, 3D rigid-body transformation between the 8 mean reconstructed pinhead co-ordinates and the 8 'true' pinhead co-ordinates in the new co-ordinate system using an algorithm described by Arun et al. [1987] (see Appendix 2).With respect to the new co-ordinate system, the pinhead co-ordinates are then (0,0,0), (0,25,0), (25,0,0), etc.

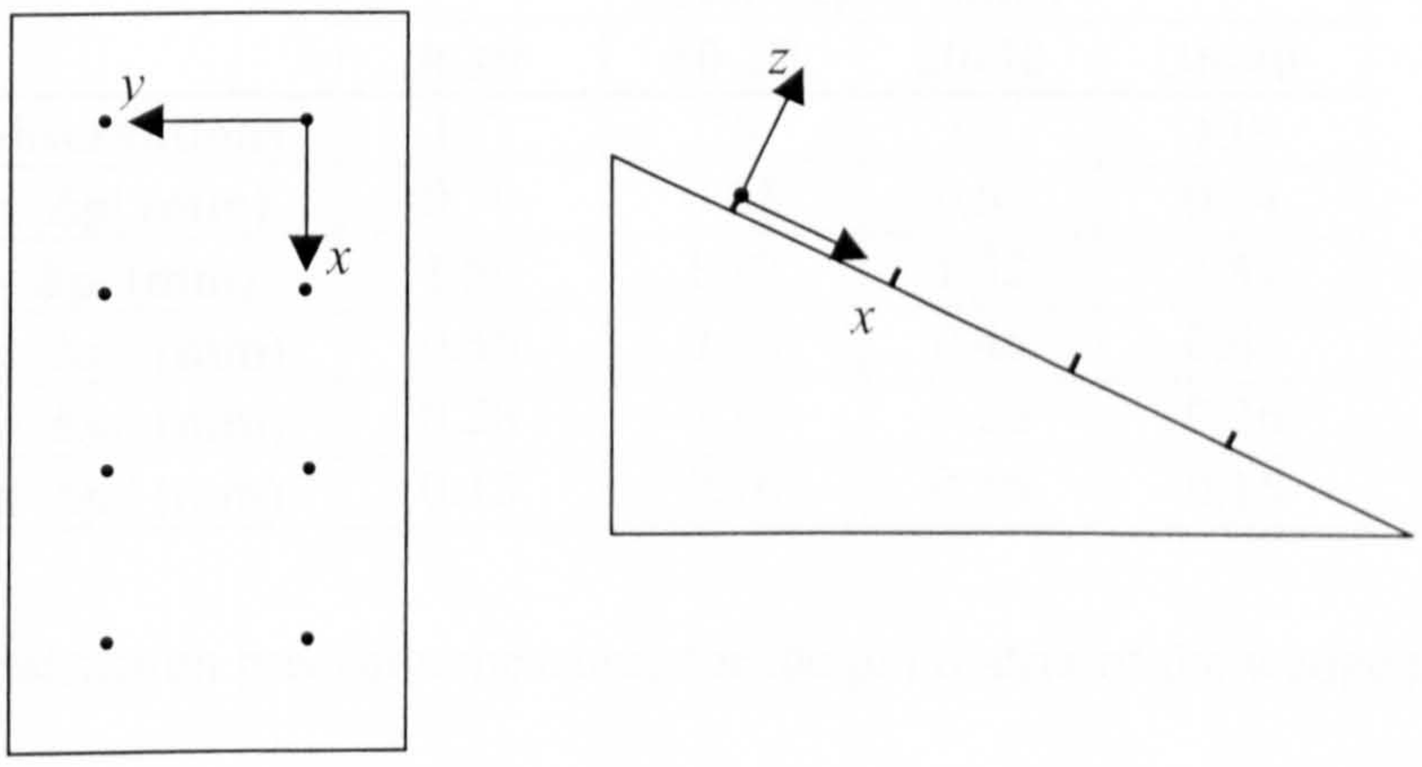


Figure 5.20 3D co-ordinate system used in the calculation of reconstruction precision.

In the second experiment, aimed at estimating the reconstruction accuracy under conditions closer to those found during clinical imaging, the wedge phantom was placed in position with its base flat, and its length parallel to that of the examination couch. A 3DUS scan of the pin matrix was performed using a slow scan-probe sweep

speed (< 2 mm/s), in the same manner as a scan of the carotid arteries. B-mode images and *pcBIRD* data were acquired at a rate of 1 per second, comparable to the rate for ECG-gated acquisition *in vivo*.

For each 3D scan, the within-image location of each point-target was specified, and the corresponding mean 3D co-ordinates calculated. The inter-pinhead distances were then calculated for each pair of pinheads, drawn from all reconstructed locations, and compared with the true distances. The experiment was repeated 6 times using 6 different phantom locations within the region appropriate to carotid scanning.

5.4.3 Wedge phantom results

The values of precision measures for variable-depth scans are given in Table 5.3 and shown graphically Figure 5.21. The mean point-target depth for each depth range was: 5.17, 15.16, 25.04, and 33.73 mm. It can be seen from Table 5.3 that there was only a small difference (< 0.13 mm) in the overall reconstruction precision ($|\Delta \mathbf{p}|$) at different scan-depths, although the precision appears to be poorer at the lowest and highest scan-depths.

	Scan-depth (mm)			
	0-10	**10-20**	**20-30**	**30-40**
No. observations	107	79	93	138
Mean $\|\Delta p\|$ (mm)	0.70	0.65	0.61	0.74
Max $\|\Delta p\|$ (mm)	1.56	1.29	1.32	2.57
Mean $\|\Delta p_x\|$ (mm)	0.56	0.48	0.46	0.61
Mean $\|\Delta p_y\|$ (mm)	0.28	0.32	0.26	0.26
Mean $\|\Delta p_z\|$ (mm)	0.13	0.16	0.16	0.16

Table 5.3 Reconstruction precision measures for the pin matrix of the wedge phantom.

The mean values of $|\Delta p_x|$, $|\Delta p_y|$, and $|\Delta p_z|$ reveal that there were significant differences between the uncertainty in point-target location for each of the principal directions. The precision was poorest in the x-direction which approximately corresponds to the elevational direction of the ultrasound image. The precision was best in the z-direction, corresponding to the axial direction of the image. There is also an observed

(*continued on page 174*)

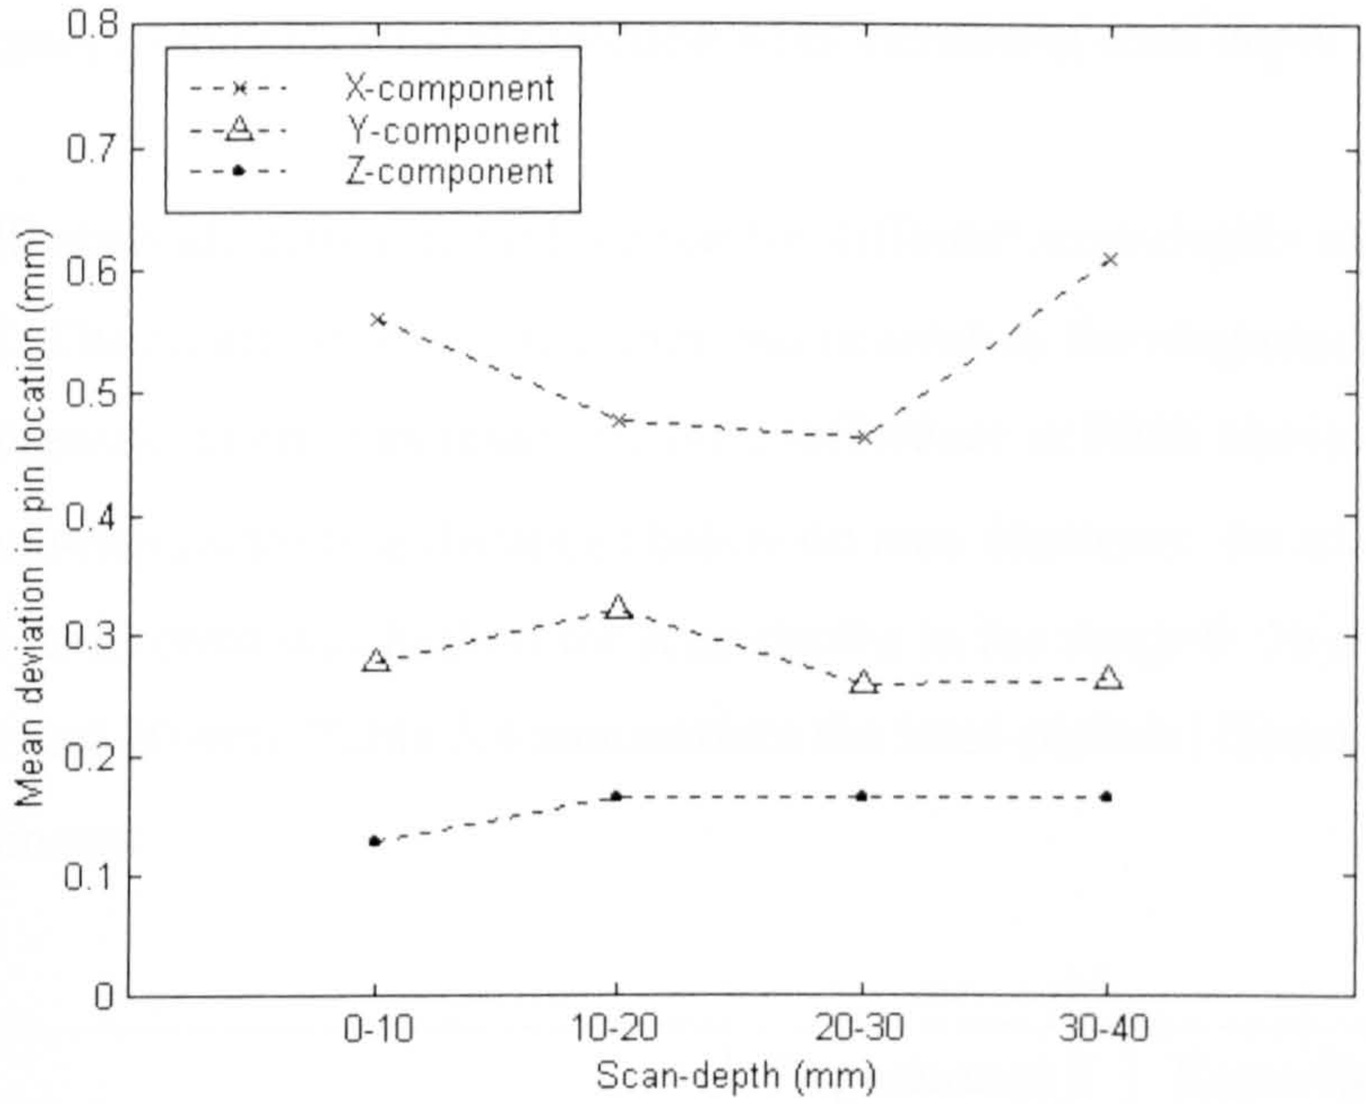


Figure 5.21 Reconstruction precision measures for the pin matrix.

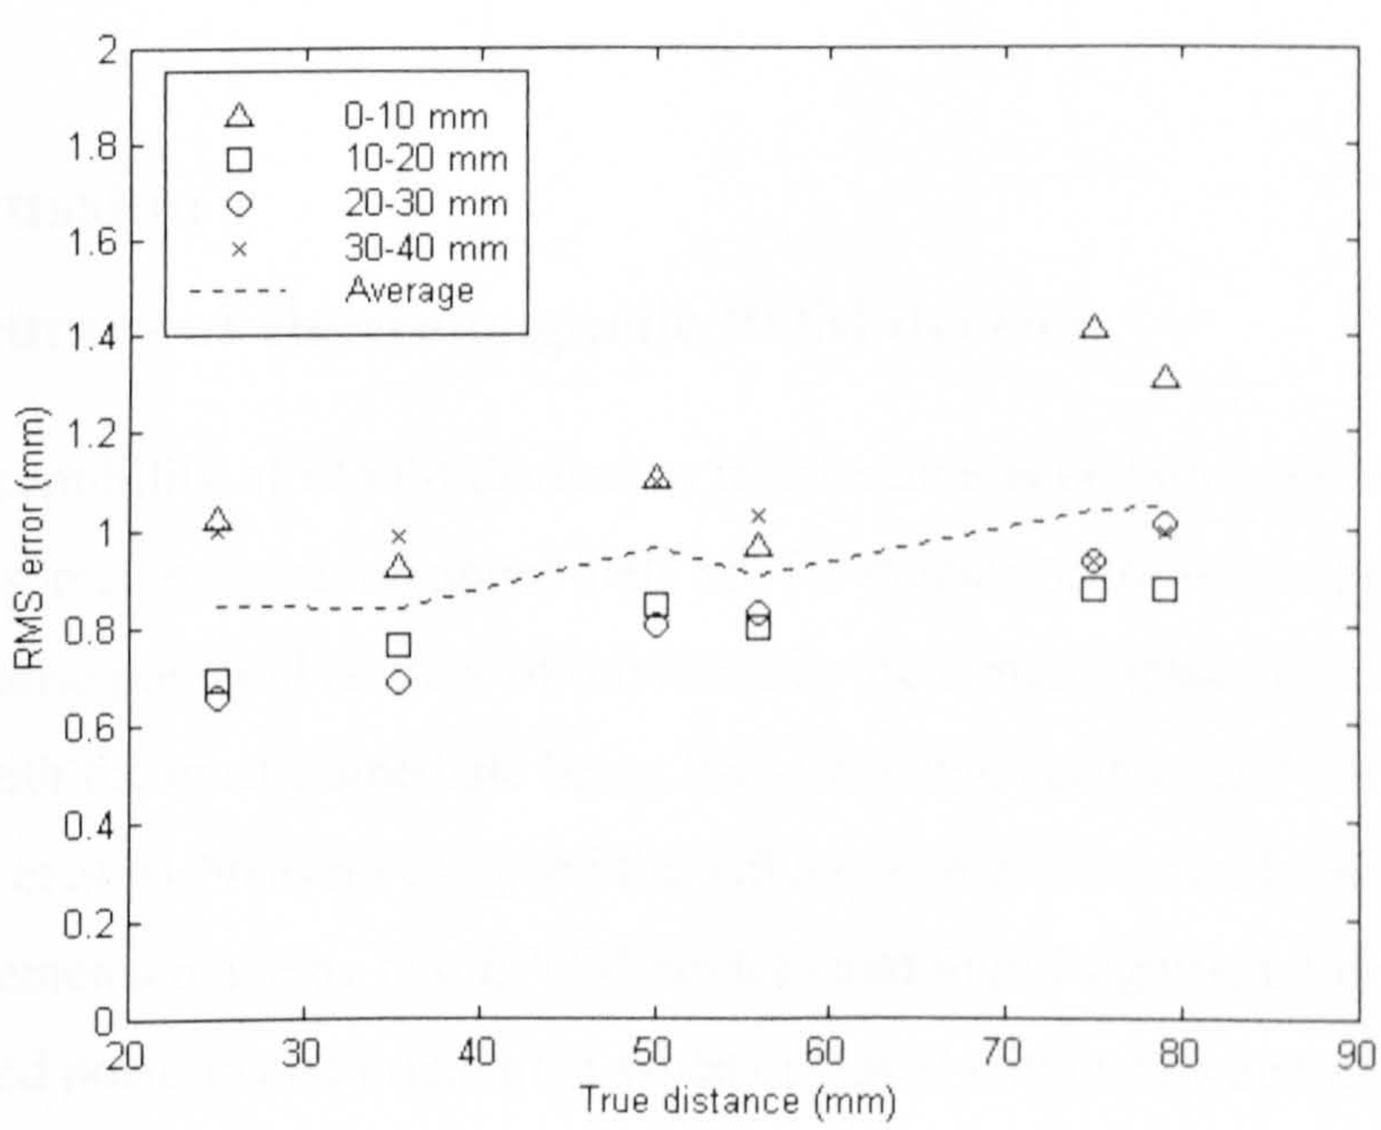


Figure 5.22 RMS absolute errors in pin location for 4 different scan depth ranges (indicated in the top left-hand corner).

variation in $|\Delta \mathbf{p}_x|$ with scan-depth, again with the highest mean values at the lowest and highest scan-depths. This is presumably due to the variation in the ultrasound beam thickness in the elevational direction with increasing scan depth.

The RMS absolute errors in pin location for different scan-depths are plotted in Figure 5.22. The results show a slight increase in error as the magnitude of the true distance increases. There was relatively little difference in RMS absolute errors at different scan-depths for true distances below 60 mm. However, for true distances ≥ 75 mm, the RMS error was highest for scan-depths in the range 0–10 mm, and lowest between 10 and 20 mm. Table 5.4 summarises the inter-pinhead distance errors for both experiments.

	Experiment 1	**Experiment 2**
No. observations	19596	1158
Mean ± SD absolute error (mm)	-0.08 ± 0.94	-0.45 ± 1.30
Mean ± SD relative error (%)	-0.13 ± 2.73	-0.53 ± 3.39

Table 5.4 Inter-pinhead distance errors for two wedge phantom experiments.

5.5 Discussion

5.5.1 Accuracy of electromagnetic POM devices

The susceptibility of EPOM devices to interference in certain environments has led to widespread reservation about their use for medical ultrasound applications. This is because potential sources of interference are commonplace in a clinical environment, with the most immediate being the ultrasound scanner and the scan-probe. Birkfellner et al. [1998a] investigated the influence of 3 ultrasound scan-probes on the measurements made by two EPOM devices and found significant distortions in the measured position and orientation when operated within 150 mm of the receiver. The authors concluded that the material used to shield scan-probes acts as the main source of distortion. Some researchers have chosen to mount the EPOM receiver away from the scan-probe electronics to avoid such interference [Barry 1997; Prager 1998]. However, as the earlier analysis of error propagation shows, the benefit of this measure is likely to be offset by increased errors due to the lever-arm effect.

The location of the scanner in relation to the operating region of an EPOM device is also an important consideration. It is a logical precaution to place the scanner as far away as possible from the scanning region to minimise the likelihood of distortion of EPOM device measurements. For most applications, the position of the scanner relative to the scanning region depends on the site of the anatomy under investigation, and on the position best suited to scanning; that is, the scanner must be placed so that controls are within reach of the ultrasonographer, and the display can be seen comfortably without contortion. Together these requirements place some limitation on the possible positions for the scanner. Since the EPOM receiver and transmitter necessarily have to be reasonably close to the scanning region, there are many situations where the scanner lies unavoidably within the operating region of the device when placed in the position most suitable for routine (2D) scanning. The location of the scanner is an important issue, not only because it has the potential to influence EPOM device measurements, but also because, in order for 3DUS to gain widespread clinical acceptance, there should be minimal disruption to routine scanning technique.

The majority of studies that describe electromagnetic 3DUS systems fail to specify adequately the location of scanner relative to the scan-volume. Notable exceptions are Detmer et al. [1994] and Leotta et al. [1997]: Detmer et al. state that experiments to validate their 3DUS system were performed with the scanner within 0.5 m of the scan-volume with "no observable noise and no loss of precision" (although no data is provided to demonstrate this). Leotta et al. report placing the scanner 1 m away from the EPOM transmitter. The results of the experiments carried out for this thesis provide further evidence that it is possible to obtain very high measurement accuracy using a commercial EPOM device with the scanner placed within 1 m of the scan region (and EPOM transmitter), and with the EPOM receiver attached to the scan-probe casing. Indeed, over a limited range of receiver motion, the estimated accuracy of both 3D position and orientation measurements was considerably higher than that specified by the manufacturers (the manufacturer's specification are generally calculated over the full operating range – 1.2 m in this case). This has important implications for 3D carotid ultrasound applications, since, in most cases, 3D acquisition involves sweeping the scan-probe along the neck with the longitudinal translation limited to less than 90 mm, and scan-probe rotations limited to < 30°. With these conditions satisfied, the RMS absolute error in position and rotation measurements is

estimated to be < 0.2 mm and 0.2°, respectively. The highest positional accuracy was obtained with the scan-probe oriented such that the EPOM receiver faces the transmitter throughout the scan, but, overall, the *pcBIRD* was influenced little by the presence of either the scan-probe or the scanner.

Milne et al. [1996] and Bull et al. [1998] tested independently the accuracy of a similar pulsed-DC EPOM device intended for use in tracking the motion of the knee-joint. Milne et al. [1996] found a mean error of 1.8% of the translation step size for steps of 25-152 mm along each of the transmitter axes, and a mean rotational error of 1.6% of the rotational increment for rotations from 1-20°. Using very similar experimental methods, Bull et al. [1998] found a mean positional error of 0.2% for step sizes of 25-200 mm and a mean rotational error of 2.0% for steps of 0.7-29.8°. In both studies a systematic error was observed which increased with the true distance or rotation between two measurement points. Systematic errors in position and orientation measurements were also observed from the experiments described in this chapter: in general, the magnitude of the absolute error was found to increase as the magnitude of the true distance/rotation increased. This effect was influenced by the location of the scanner, but was independent of the transmitter-receiver separation since the EPOM receiver was translated both towards and away from the transmitter and relative rotational measurements were made through a full 360°.

The magnitude of EPOM errors estimated in this thesis were small, so their systematic nature does not present a significant problem for carotid scanning. However, it may be a problem if much larger volumes are scanned, assuming the same pattern of increasing error with increasing true distance continues for translations greater than 90 mm. In this case, correction may be appropriate to reduce errors in the 3D image reconstruction. Mathematical correction of systematic errors is possible, but this requires a detailed description of x, y and z position errors throughout the working volume. Given the complex nature of the interaction between the magnetic field and the measurement environment, this can be a difficult and time-consuming task which required specialized apparatus. It has been attempted previously for surgical environments[Birkfellner 1998b].

A common method for evaluating the performance of EPOM devices in 3DUS systems is to determine the uncertainty in the 3D location of a point target from arbi-

trarily oriented receiver positions [Detmer 1994; Leotta 1997; Prager 1998]. Systematic errors in EPOM device measurements have been noted before in the context of 3DUS by Detmer et al. [1994] and Leotta et al. [1997]: both these studies observed an increase in measurement uncertainty as the receiver was moved further away from the transmitter over the full operating range. Systematic errors in position/orientation measurements of the type discussed in this chapter have received little attention previously, particularly in the context of 3DUS imaging. Such errors are likely to be encountered in practice whenever a 3D image is reconstructed from image slices acquired during a single sweep.

An interesting finding was the effect of changing the measurement rate of the *pcBIRD*. Experiments demonstrated that changing the measurement rate from the default setting can introduce very large measurement errors at low and high measurement rates. If a very high measurement rate is required then the parameters of the digital filters may also need to be tuned to reduce noise. Optimal ranges existed for the device tested, and these included the default rate (103 Hz). Either the default filter setting (DC and wide-notch AC combination), or the DC and narrow-notch AC filter combination were found to be optimal. However, in the accuracy tests, the latter was chosen on the basis of its superior dynamic performance

5.5.2 3D reconstruction accuracy

Leotta et al. [1997], Prager et al [1998], and Riccabona et al. [1995] have reported on the reconstruction accuracy of electromagnetic 3DUS systems. In the study by Leotta et al. [1997], a 4×2 arrangement of beads was used with the same dimensions as the pin matrix used in this thesis. The study reports a mean inter-target distance error of 0.06 ± 0.68 mm for scan-depths between 8.7 cm and 12.4 cm. Prager et al. [1998] also used a 4×2 matrix of point-targets (made from dressmaker's pins), with an inter-target separation of 100 mm. This study also reports high accuracy with a mean inter-target distance error of 0.04 ± 1.12 mm following calibration using a cross-wire phantom. Riccabona et al. [1995] report a mean error of $0.02 \pm 3.65\%$ for 3D distance measurements in the range 10-150 mm using electromagnetic 3DUS system. Similar experiments using a motor-driven scan-probe translation device for 3D acquisition, found a mean distance error of 1.0 ± 0.8 % in the plane of acquisition, and 1.0 ± 0.6 % for planes in other orientations, and distances between 10 and 100 mm

[Riccabona 1996b]. King et al. [1991] tested the performance of a 3DUS system, based on an acoustic POM device. In this study a more elaborate arrangement of pin-head point-targets was used, and the overall mean inter-target distance error was found to be 0.75% (true range: 43.13 to 110.38 mm) for an imaging depth of 6 cm.

The 3D reconstruction errors presented in this thesis are of the same order of magnitude as those reported in the studies mentioned above. An important aspect of the approach used in this study is the dynamic nature of the 3D acquisition – i.e. the scan-probe was translated and/or rotated slowly during data acquisition, simulating the motion for *in vivo* carotid scans. Furthermore, single sweep scans were performed, whereas two of the aforementioned studies obtained multiple images of each point-target using different scan-probe orientations. The wedge phantom was designed so that the plane of the point-target matrix was inclined, and hence they are located through a range of scan depths. The small target diameter had the effect of increasing the uncertainty in target location because of the low resolution of the ultrasound beam in the elevational direction, and highlighted a significantly lower precision in target location in this direction, as predicted by theoretical analysis.

Chapter 6

Reconstruction and quantification of the diseased carotid bifurcation

This chapter presents a method for the 3D reconstruction of the wall and lumen of the diseased carotid artery bifurcation from a series of spatially-located 2D ultrasound images, acquired using the system described in Chapter 4. User-intervention is required to define cross-sectional contours for each image, and a software tool for this task is described. Once reconstructed, the wall and lumen surfaces are represented by a set of longitudinally-oriented 3D curves. This representation is memory-efficient, and allows reslicing and quantitative measurement operations to be performed rapidly and automatically without further user-intervention. A number of quantitative measures are derived from the 3D reconstruction. In particular, methods are presented for quantifying plaque volume, volumetric stenosis index, extent of disease, and plaque surface area, as well as conventional stenosis indices from 'true transverse' cut-planes.

6.1 Introduction

As introduced in Chapter 3, 3DUS data can be reconstructed using either a volumetric (3D array of voxels) or surface representation. In the latter, the surfaces of objects are typically represented by a geometric model, such as a triangulated surface mesh, reconstructed from a series of cross-sectional contours. The use of such a geometric model has the advantage of being a relatively compact representation. The model may also be rendered quickly, allowing rapid visualization of an object. Such 3D surface rendered images are usually straight-forward to interpret, and quantities such as volume may be readily calculated from the underlying model [Barratt DC 2001].

3D surface reconstruction is currently limited by the lack of robust automatic segmentation algorithms for ultrasound images, The most reliable method of extracting tissue boundaries remains manual delineation of boundary curves, save for a few exceptional cases. An attempt to address this problem is not the focus of this thesis. Rather, the thesis concentrates on methods for the quantification of carotid artery disease, assuming the availability of a pre-segmented 3D data set. Accurate cross-sectional contours of the wall and lumen of the carotid arteries were obtained for this purpose using a manual procedure described in Section 6.3. The method requires a significant investment of time, with a typical data set taking up to 45 minutes to segment.

Once a series of transverse, cross-sectional contours – equivalent to a set of regularly sampled surface points – are available, there are many methods for reconstructing the surface [Bolle 1991]. A powerful technique, already discussed, is surface mesh generation. In this scheme, the surface is approximated by a large number of small, interconnecting polygonal facets (see Chapter 3). Triangulated meshes formed by connecting the vertices of adjacent contours have been employed for this purpose [Haber 2000; Barequet 1996].

Another approach, proposed by Lin et al. [1988], uses dynamic elastic interpolation surface reconstruction. Other researchers have adopted a generalised cylinder representation, where cross-sectional contours are 'mounted' onto a 3D central axis, or spine curve, as illustrated in Figure 6.1 [Burdin 1996; O'Donnell 1995; Parker 1999]. Intuitively, this representation may be thought of as 'sweeping out' the surface.

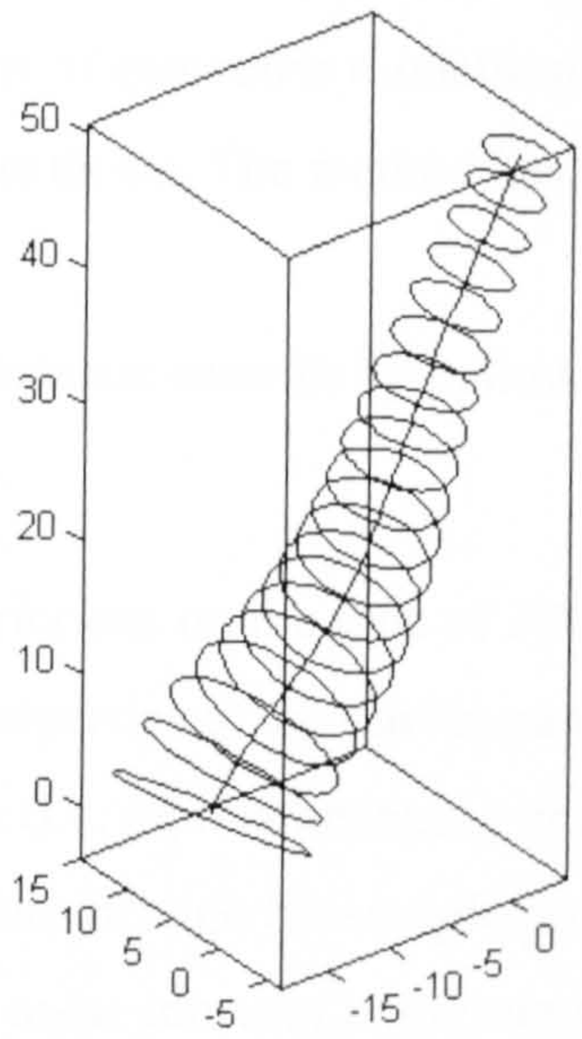


Figure 6.1 The generalised cylinder representation: a central axis curve (or spine) is defined onto which cross-sectional contours are mounted.

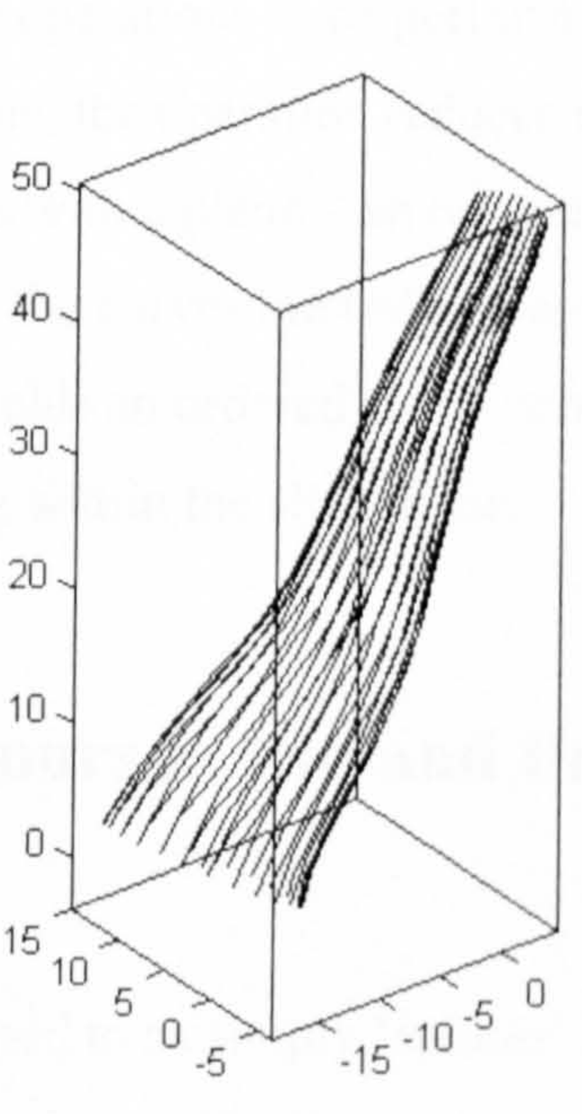


Figure 6.2 The surface shown in Figure 6.1 represented by a series of smooth, longitudinal curves.

A variation on the generalised cylinder approach is to represent a surface by a set of 3D curves oriented in the axial direction, as shown in Figure 6.2. This representation was adopted as the basis of geometric modelling of both the wall and lumen of the carotid bifurcation in this thesis. The method has the following attractions:

- First, since blood vessels are essentially hollow tubes, it is a simple and natural representation.

- There are also no restrictions on the type of 3D curve used to define a surface. This means that the properties of the curves can be varied depending on the surface characteristics (i.e. rapidly changing or smooth). Curves can be selected which do not pass through the sample points, and these are useful for the purposes of removing noise inherent in the sampling process.

- The freedom to choose from a variety of curve definitions also enables compact representation of a surface, since many curves can be stored in an efficient manner.

- The representation provides a geometric surface description which allows reslicing and quantitative operations to be performed rapidly. In the case of reslicing in a transverse plane, the operation reduces to finding the intersections of the longitudinal curves with a plane – an operation which is easily implemented. Further, since the curves are ordered around the vessel, transverse reslicing automatically yields an ordered set of vertices which define the cross-sectional contour lying within the slice plane.

6.2 Curves and contours in two and three dimensions

6.2.1 Spline curves

Spline curves – usually referred to as simply 'splines' – are formed from a number of sections, or pieces; each piece is specified by a polynomial function in t, where t is some parameter [de Boor 1978; Dierckx 1993]. The simplest form of spline is the univariate function in t, illustrated in Figure 6.3. Each piece of this spline is defined by:

$$f_i(t) = \sum_{j=0}^{n} c_{ij}(t - \tau_{i-1})^j , \qquad (6.1)$$

where $\tau_{i-1} \le t \le \tau_i$ and c_{ij} is the j^{th} scalar coefficient of the polynomial that defines the i^{th} piece. The values $t = \tau_i$ are known as breaks, and n is the degree of the spline: for example, n is 3 for a cubic spline.

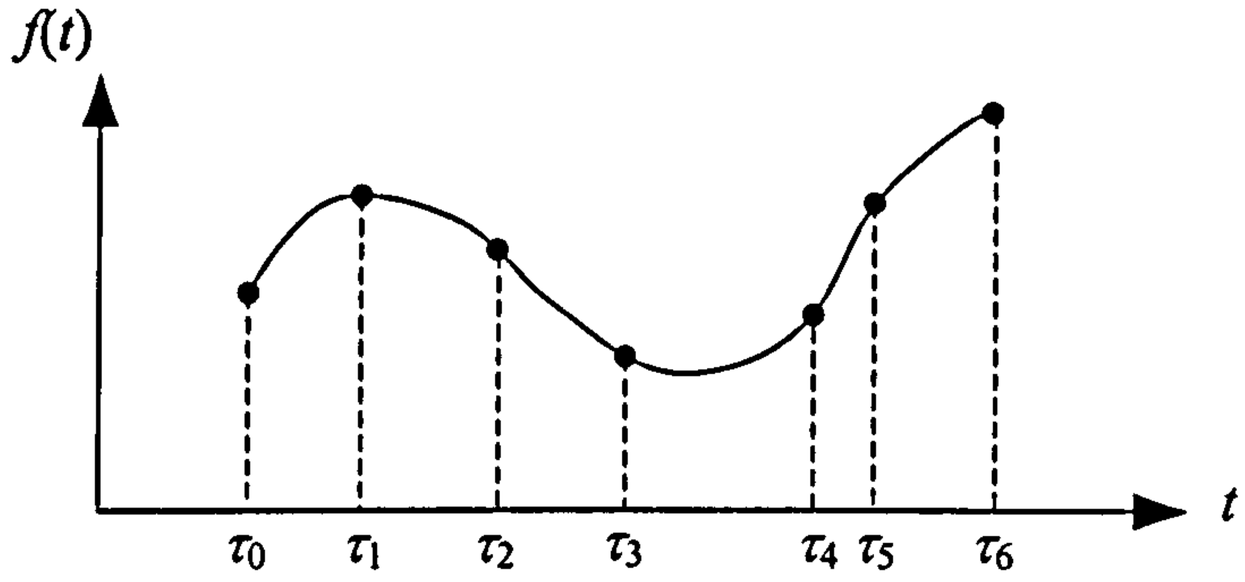


Figure 6.3 An example of a univariate spline function. The spline is formed from 6 polynomial pieces, each defined by a polynomial function over the (inclusive) interval $t = [\tau_{i-1}, \tau_i]$, $1 \le i \le 6$. The values of τ_i where two pieces are connected are known as breaks.

An important feature of a spline is continuity between pieces; that is, polynomial functions have continuous derivatives at breaks. In general, a spline is said to be CN-continuous when

$$\left(\frac{\mathrm{d}^N f_i}{\mathrm{d}t^N}\right)_{t=\tau_i} = \left(\frac{\mathrm{d}^N f_{i+1}}{\mathrm{d}t^N}\right)_{t=\tau_i} . \qquad (6.2)$$

CN-continuity ensures that splines are smooth. A popular choice are cubic splines which satisfy C0-, C1- and C2-contininutiy, and provide a good compromise between smoothness and the ability to approximate functions with complex shape. Intuitively, C0-continuity implies that pieces are joined at break points; C1-continuity implies that the tangents to adjoining pieces are equal at the break points; and C2-continuity implies that the curvature of adjoining pieces is equal at the break points.

Splines are easily extended to multiple dimensions. For a 3D spline, for instance, each polynomial piece is a vector function, say $\boldsymbol{f}_i$, with three components:

$$f_i(t) = \begin{bmatrix} x_i(t) \\ y_i(t) \\ z_i(t) \end{bmatrix}, \tag{6.3}$$

where $x(t)$, $y(t)$ and $z(t)$ are n^{th} degree polynomials of the form

$$c_n t^n + c_{n-1} t^{n-1} + \ldots + c_1 t + c_0, \tag{6.4}$$

with real coefficients, c_j, $0 \le j \le n$. With this definition, a 3D spline usefully describes a smooth space curve, and the function f_i defines 3D points on the spline over some interval in t. Combining Equations (6.3) and (6.4) leads to the alternative matrix form:

$$f_i(t) = \begin{bmatrix} c_n^x & c_{n-1}^x & \cdots & c_0^x \\ c_n^y & c_{n-1}^y & \cdots & c_0^y \\ c_n^z & c_{n-1}^z & \cdots & c_0^z \end{bmatrix}_i \begin{bmatrix} t^n \\ t^{n-1} \\ \vdots \\ 1 \end{bmatrix} = \mathbf{G}_i \mathbf{T}, \tag{6.5}$$

where the superscripts 'x', 'y', and 'z' of the coefficients c denote the x, y, or z component, respectively. $\mathbf{G}_i$ is termed the coefficient matrix, and $\mathbf{T}$ is the parameter matrix. This form has the advantage of being very concise, and the N^{th} derivative of a polynomial piece can be calculated easily, since:

$$\frac{\mathrm{d}^N f_i}{\mathrm{d}t^N} = \frac{\mathrm{d}^N (\mathbf{G}_i \mathbf{T})}{\mathrm{d}t^N} = \mathbf{G}_i \frac{\mathrm{d}^N \mathbf{T}}{\mathrm{d}t^N}, \tag{6.6}$$

for $N < n$.

When dealing with splines that represent space curves, it is common to use the arc-length, s, as a physically meaningful parameter, such that $0 \le s \le L$, where L is the total length of the curve. For a multi-dimensional spline, $f(t)$, with arbitrary parameter t, s and t are related by:

$$s(t) = \int_0^t \left| \frac{\mathrm{d}f}{\mathrm{d}t} \right| \mathrm{d}t \,. \tag{6.7}$$

A very large class of functions can be approximated to an arbitrary degree of accuracy using piecewise-polynomial (PP) splines. In addition, splines are an efficient way to describe smooth functions. These two properties make them attractive for representing 3D space curves and surfaces. A 3D cubic spline is fully specified by storing a 3×4 coefficient matrix ($\mathbf{G}_i$) for each piece, along with a list of the breaks. This representation is sometimes referred to as PP-form [de Boor 1978].

In the present application, the reason for fitting curves to sample points is to attempt to recover the shape of the arterial surface. A straight-forward method for fitting a spline to sample data points is to choose the parameters so that the curve passes through all sample points. This a called an interpolating spline, but does not give satisfactory results when noise is present in the sample points. A more appealing alternative is to use a smoothing spline [de Boor 1978; Dierckx 1993]. If noisy samples of a function have values y_i then a smoothing spline, $f(t)$, minimizes the expression:

$$p\sum_i w_i\left[y_i - f(t_i)\right]^2 + (1-p)\int \lambda(t)\left(\frac{\mathrm{d}^2 f}{\mathrm{d}t^2}\right)^2 \mathrm{d}t, \tag{6.8}$$

where $p \in [0,1]$ is a smoothing parameter, w_i is a weight which governs the effect of sample i, t_i is the value of the parameter t corresponding to the i^{th} sample, and $\lambda(t)$ is a weight function, usually set to $\lambda = 1$ for all t. The summation component of the expression in Equation (6.8) represents an measure of error, whereas the integral component is a measure of 'roughness'. A full account of the theory and implementation of smoothing splines may be found in [de Boor 1978]. For the purposes of this thesis, smoothing splines are considered to be merely a method for smoothing sample data for a 3D space curve by choosing appropriate values for p and w_i. In the work for this thesis, the Matlab Spline Toolbox (The Mathworks Inc., MA, USA) was used to generate smoothing splines using algorithms given in [de Boor 1978].

6.2.2 Contours

Contours may be defined as planar curves which represent physical boundaries, such as the vessel wall or the surface of a plaque. Although contours may be either

open or closed (i.e. the last and first point are the same), here we will be mainly interested in the closed cross-sectional contours which delineate the vessel wall and lumen. Open contours may be drawn to approximate a plaque surface, but these are later combined with the wall contour to form the closed contour of the lumen. The conventions shown in Figure 6.4 were adopted: contours which represent the wall are defined in an anti-clockwise sense, whereas contours which enclose the lumen are defined in a clockwise sense. The four cases depicted in Figure 6.4 represent a diseased-free artery and carotid plaque in various locations around the vessel wall.

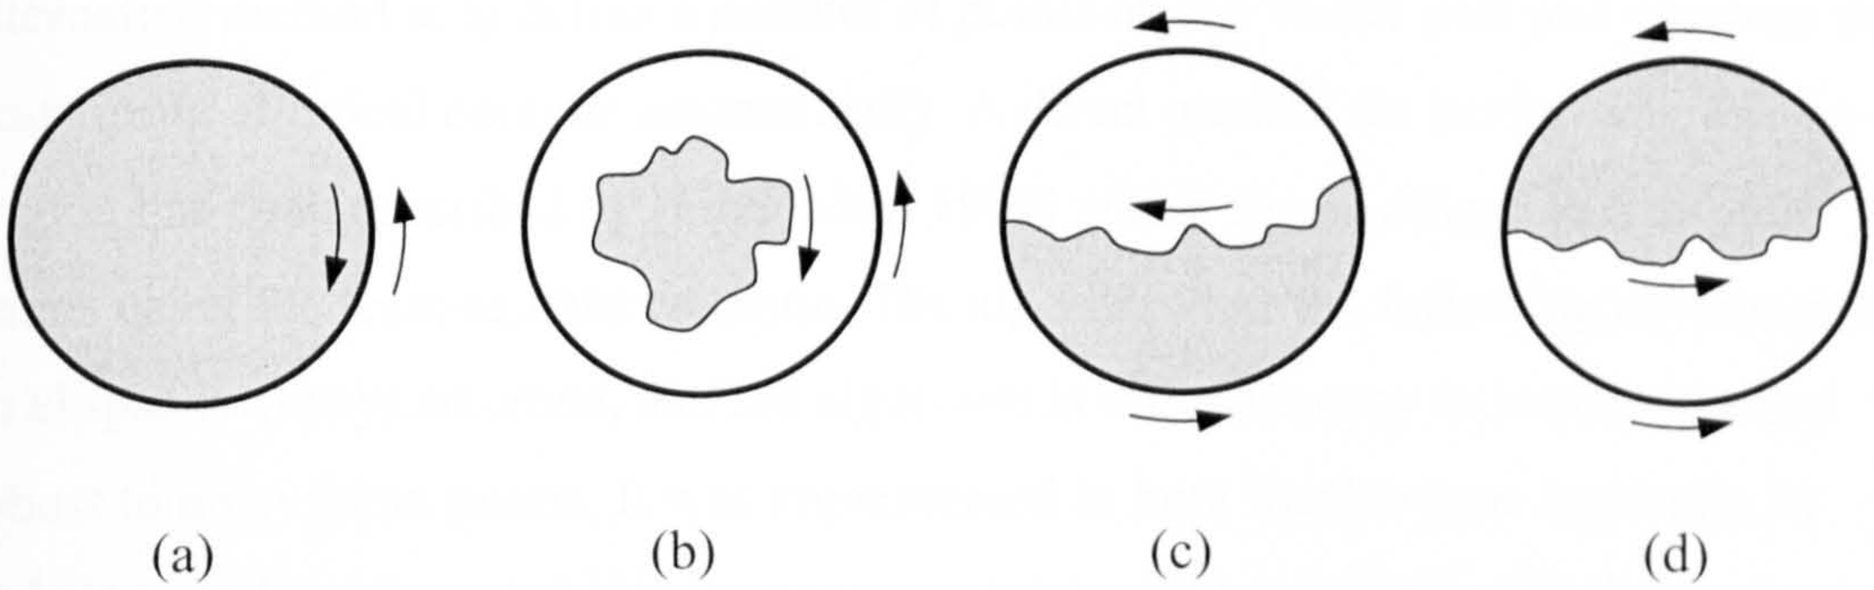


Figure 6.4 Conventions for the definition of wall and plaque contours (the lumen is shaded): (a) no disease is present, and the lumen and vessel wall contours are considered to be coincident, with the wall contour defined in the anticlockwise direction and the lumenal contour in the clockwise direction; (b) plaque is present all around the vessel wall, and the wall and lumen contours are defined in the anticlockwise and clockwise directions, respectively; (c) plaque is present on the anterior wall; (d) plaque is present on the posterior wall.

A contour may be approximated by an N-sided polygon or a smooth spline. In the former, which can in any case be derived by computing discrete points along a spline, the area is given by:

$$A = \frac{1}{2}\sum_{i=1}^{N} x_i y_j - x_j y_i, \quad j = \begin{cases} i+1, & i < N \\ 1, & i = N \end{cases} \tag{6.9}$$

where (x_1, y_1), (x_2, y_2),…, (x_N, y_N) are the co-ordinates of the polygon vertices. If the vertices are ordered in a clockwise direction then A is negative, indicating a cavity; otherwise, A is positive. Plaque area is calculated from the algebraic sum of A for the wall and lumen contours, as defined in Figure 6.6. When a closed contour is repre-

sented by a spline, it is important that CN-continuity is preserved at the beginning and end points; such a spline is called periodic.

Elliptical contours are generally a good approximation to the vessel wall in the transverse (or near-transverse) section away from the bifurcation. Many commercial ultrasound scanners provide a facility for manually defining an ellipse, with standard implementations allowing the user to set the lengths of the principal axes, and then rotate the ellipse until the 'best fit' to the wall is found by eye. This method is adequate, but obtaining the best fit is often subject to multiple manual refinements. An alternative method is to define a number of points on the vessel wall and compute the best-fitting elliptical contour automatically. A direct method for performing this operation has been described by [Fitzgibbon 1999] which fits an ellipse to 6 or more points using the least-squares criterion. The algorithm has the following advantages: an ellipse is always returned, and the algorithm is efficient, easy to implement and robust to noisy input points. It was implemented in both Matlab (see Appendix 2) and C++, and incorporated into the segmentation software described in the next section. The method is particularly suited to defining contours for transverse ultrasound images since not all sections of the vessel wall are usually visible because of weak reflections produced on the lateral sides; only a few points on each of the visible portions of the vessel are required, and the fitted ellipse may then be further refined by adding more points as necessary. Typically, 10 or more points are sufficient to provide a good fit. Each of the manually defined contour points is subject to uncertainty due to the difficulty of discriminating the media-adventitia boundary, particularly in advanced atherosclerosis. The algorithm is robust to such uncertainty, which may be considered to be equivalent to noise in the input data points.

6.3 Segmentation of acquired ultrasound images

As indicated at the start of the chapter, a prerequisite for surface reconstruction is the availability of a set of cross-sectional contours. This was achieved by segmenting transverse carotid ultrasound images, acquired by the 3DUS system, using custom-written software. The software enables the boundaries of the vessel wall and lumen to be delineated manually using a mouse-controlled cursor to define boundary points. This process was accelerated by assuming an elliptical vessel cross-section in regions away from the bifurcation, allowing the direct ellipse fitting algorithm to be em-

ployed. For carotid plaque and regions immediately proximal to the bifurcation, contours were drawn over the image by specifying a series of points and fitting a cubic interpolating spline. The software includes operations to modify contours by shifting, inserting, or removing points. The complete segmentation of a typical data set takes 10-30 minutes, depending on the complexity of the disease and the quality of visualisation. Further details of the segmentation software, together with an outline of the steps involved in segmenting images of the diseased carotid bifurcation are included in Appendix 3.

6.4 Surface reconstruction

Once a series of cross-sectional contours have been defined for the set of acquired slices, the next step is to reconstruct the surfaces which define both the walls and lumen of the carotid bifurcation. The approach adopted is to fit longitudinal curves to the contours following registration in some convenient reconstruction volume. Cubic smoothing splines were chosen to model the vessel wall due to their ability to suppress noise and approximate complex geometries.

Clearly, it is reasonable to assume that the surface of the vessel wall is very smooth, undergoing only gradual changes in area along its length. On the other hand, the lumenal surface in the diseased state is, in general, highly irregular. Therefore, unlike the vessel wall, no simplifying assumptions can be made about the nature of the lumenal surface. Splines are applicable to model such irregular surfaces, but there will still be an inevitable smoothing effect because such splines are smooth by virtue of continuity in the derivatives. This has the effect of limiting the accuracy of the representation through the removal of very small-scale surface features. This is not a significant drawback, however, since such features are very difficult to discern with any certainty from ultrasound images, and, with the exception of plaque surface area, their absence should not significantly affect the gross quantitative measures of carotid disease described later in this chapter.

6.4.1 Data pre-processing

6.4.1.1 Transformation to a new 3D co-ordinate system

Before reconstructing 3D surface curves, it is expedient to transform the co-ordinates

of each of the contour vertices to another 3D co-ordinate system (denoted by V in the analysis of Chapter 5). An appropriate co-ordinate system for the carotid bifurcation is one in which the axial direction of the vessel lies roughly parallel to the z-axis, with the origin placed at the proximal end of the CCA, and z increasing in the proximal to distal direction. As shown in Figure 6.5, this co-ordinate system was found by determining the vector between the centres of the first and last acquired ultrasound slices after registration in the co-ordinate system of the pcBIRD transmitter, T. This vector was then normalised to give the unit vector, $\hat{\mathbf{v}}_z$. The vectors defining the x- and y-axes – $\hat{\mathbf{v}}_x$ and $\hat{\mathbf{v}}_y$, respectively – were then found by computing the two orthogonal vectors to $\hat{\mathbf{v}}_z$. The direction of the x-axis was first approximated by finding the vector, $\mathbf{v}_x^{est}$, which lies within the first image slice and connects the centre pixel and the centre-top pixel, vertically above the centre pixel (see Figure 6.5). The unit vector defining the y-axis was then calculated using the formula:

$$\hat{\mathbf{v}}_y = \frac{\mathbf{v}_x^{est} \times \hat{\mathbf{v}}_z}{\left| \mathbf{v}_x^{est} \times \hat{\mathbf{v}}_z \right|}, \tag{6.10}$$

where $\times$ is the vector cross-product operator. The orthogonal unit vector that defines the x-axis is then given by:

$$\hat{\mathbf{v}}_x = \hat{\mathbf{v}}_z \times \hat{\mathbf{v}}_y . \tag{6.11}$$

The vectors $\hat{\mathbf{v}}_x$, $\hat{\mathbf{v}}_y$, and $\hat{\mathbf{v}}_z$ define the directions of x-, y-, and z-axes, respectively, for the new co-ordinate system, V, relative to T. If the origin of V is placed in the centre of the first slice, with co-ordinates (x_0, y_0, z_0) relative to the origin of T, then:

$${}^{T}\mathbf{M}_{V} = \begin{bmatrix} \begin{bmatrix} & & \\ \hat{\mathbf{v}}_x & \hat{\mathbf{v}}_y & \hat{\mathbf{v}}_z \\ & & \end{bmatrix} & \begin{matrix} x_0 \\ y_0 \\ z_0 \end{matrix} \\ \begin{matrix} 0 & 0 & 0 \end{matrix} & 1 \end{bmatrix}, \tag{6.12}$$

and, hence:

$$ {}^{V}\mathbf{M}_{T} = \begin{bmatrix} \begin{bmatrix} & & \\ \hat{\mathbf{v}}_x & \hat{\mathbf{v}}_y & \hat{\mathbf{v}}_z \\ & & \end{bmatrix} & \begin{matrix} x_0 \\ y_0 \\ z_0 \end{matrix} \\ \begin{matrix} 0 & 0 & 0 \end{matrix} & 1 \end{bmatrix}^{-1} \tag{6.13} $$

Using the notation introduced in Chapter 5, this matrix defines the transformation from $\boldsymbol{T}$ to $\boldsymbol{V}$.

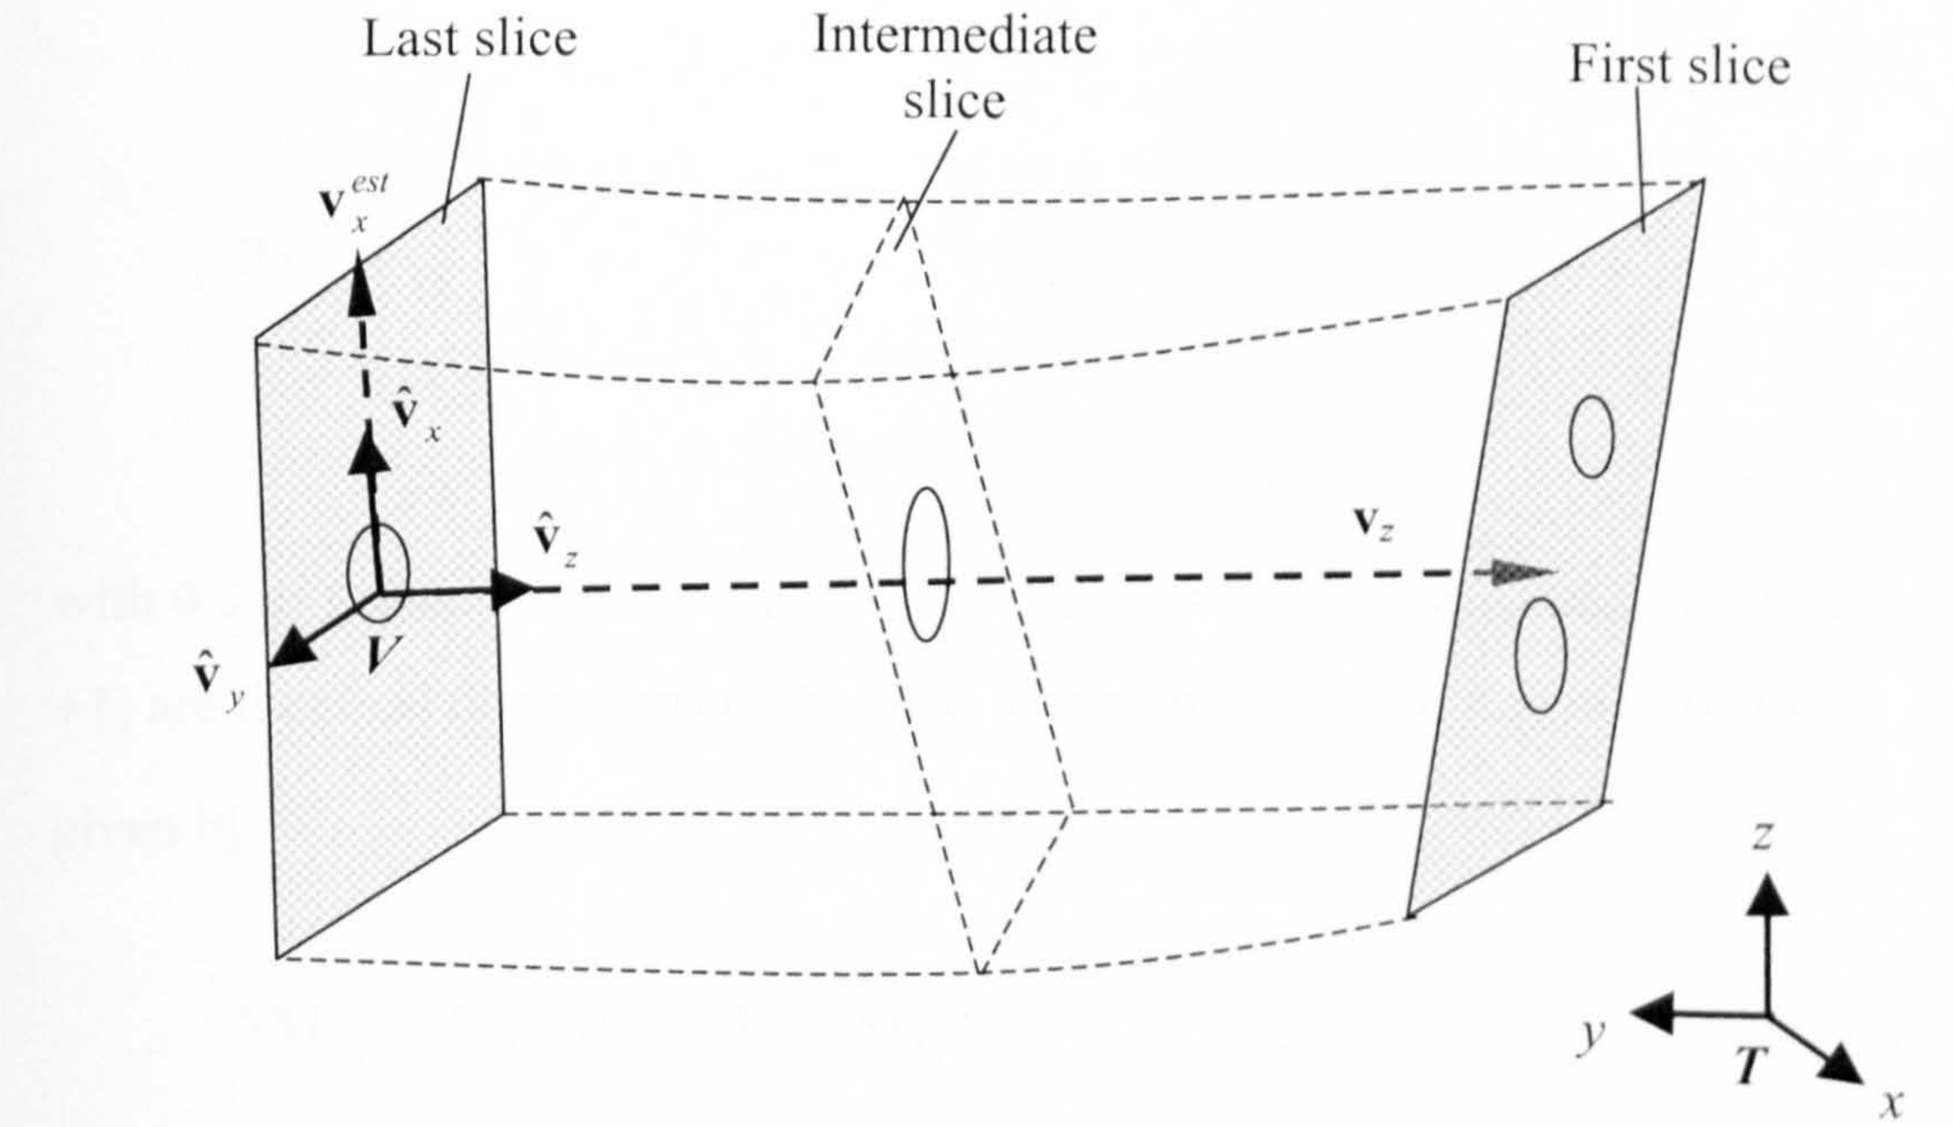


Figure 6.5 The local 3D co-ordinate system, $\boldsymbol{V}$, for 3D reconstruction of the carotid bifurcation. Initially, the first and last acquired slices are registered in $\boldsymbol{T}$, the co-ordinate system for the pcBIRD transmitter. Note that the conventional direction of the scan-probe sweep is in the distal to proximal direction.

6.4.1.2 Removal of overlapping slices

The algorithm developed to reconstruct the surfaces of the carotid bifurcation requires a set of non-overlapping cross-sectional contours. Overlapping contours are eliminated by detecting intersecting slices using the following method: for two slices with in-plane co-ordinate systems $\boldsymbol{J}_1$ and $\boldsymbol{J}_2$, points in each slice, specified by the position vectors $\mathbf{p}_{J_1}$ and $\mathbf{p}_{J_2}$, are related by:

$$\mathbf{p}_{J_1} = {}^{J_1}\mathbf{M}_{J_2}\mathbf{p}_{J_2}, \tag{6.15}$$

where

$$\mathbf{p}_{J_1} = \begin{bmatrix} u_1 \\ v_1 \\ w_1 \\ 1 \end{bmatrix}_{J_1} \tag{6.16}$$

and

$$\mathbf{p}_{J_2} = \begin{bmatrix} u_2 \\ v_2 \\ 0 \\ 1 \end{bmatrix}_{J_2}, \tag{6.17}$$

with $0 \le u_1 \le u_{\max}$, $0 \le u_2 \le u_{\max}$, $0 \le v_1 \le v_{\max}$, and $0 \le v_2 \le v_{\max}$; $(u_{\max}+1) \times (v_{\max}+1)$ are the pixel dimensions of the slice. The matrix ${}^{J_1}\mathbf{M}_{J_2}$ in Equation (6.15) is given by:

$${}^{J_1}\mathbf{M}_{J_2} = {}^{J_1}\mathbf{M}_T\,{}^{T}\mathbf{M}_{J_2} = {}^{T}\mathbf{M}_{J_1}^{-1}\,{}^{T}\mathbf{M}_{J_2}. \tag{6.18}$$

If slices 1 and 2 intersect, one or more co-ordinate pairs, (u_2, v_2), will exist for which w_1 equals zero. Thus, just considering the z-component of $\mathbf{p}_{J_1}$ in Equation (6.15) yields a linear equation of the form:

$$w_1 = au_2 + bv_2 + c = 0. \tag{6.19}$$

A straight-forward method for testing whether Equation (6.17) holds for some (u_2, v_2) is to use interval arithmetic, where u_2 and v_2 are treated as the *intervals* $\tilde{u}_2$ and $\tilde{v}_2$. The notation adopted for intervals here is: $\tilde{x} = [L,U] = \{x: L \le x \le U\}$, where L and U are real numbers that specify the (closed) lower and upper bounds on the interval $\tilde{x}$, respectively. Rearranging Equation (6.19) and expressing in terms of intervals gives:

$$\tilde{v}_2 = \frac{-c - a\tilde{u}_2}{b}, \tag{6.18}$$

where $\tilde{u}_2 = [0, u_{\max}]$. Now, if the intervals $\tilde{v}_2$ and $[0, v_{\max}]$ intersect then so do the slices; that is, there exists at least one pair values, (u_2, v_2), that are co-ordinates within the image, and $w_2 = 0$ and $0 \le u_1 \le u_{\max}$ and $0 \le v_1 \le v_{\max}$. Therefore, Equation (6.18) provides a condition for intersection of two rectangular slices.

This interval-based method above provides a simple and efficient method for detecting overlapping slices, and requires the computation of a single, linear interval expression following matrix multiplications. Interval computations were performed using the *b4m* toolbox for Matlab [Zemke 2001]. If two slices are found to intersect, it does not necessarily follow that contours defined within those slices will intersect, resulting in the removal of contours which in fact do not intersect with any others. This is not a significant problem in practice since there are a relatively large number of slices acquired in a typical sweep. The more involved analysis necessary to test for the more specific case of whether two 3D polygons intersect is also avoided. The useful property of the test is that if two slices do not intersect then contours defined within them are guaranteed not to intersect too. Thus, if overlapping slices are removed before segmentation, it is guaranteed that no user-defined contours will intersect when transformed into the reconstruction volume. If overlapping contours are identified prior to segmentation, there is also the added advantage that the user can be prevented from defining contours, which would be later removed because they violate the no-overlap condition.

6.4.1.3 Centroid registration of wall contours

Figure 6.6(a) shows a series of cross-sectional contours for the wall of the carotid bifurcation after they have been naively registered with respect to a reconstruction co-ordinate system. It can be seen from Figure 6.6(a) that some of these contours (indicated by arrows) are clearly misregistered. This presents a problem for surface reconstruction, since the sampled surface points effectively contain unwanted noise. An effective method correcting for this type of contour misregistration is to perform

a centroid registration operation, where the 3D curve which connects the centroids of the contours is smoothed. The contours are then shifted within their slice planes so that their centroids coincide with the smoothed curve. Moore et al. [1998] used this technique successfully to reduce errors in magnetic resonance reconstructions of a straight tube. The method assumes that the misregistration is confined to within the plane of the acquired image slices – i.e. misregistration of a contour to a position outside the slice plane is not accounted for. This is not an unreasonable assumption, however, since contour misregistration is often the result of motion artefacts in a direction approximately parallel to the slice plane. Irregularities and inconsistencies in the shape of contours on consecutive slices due to the segmentation phase are not corrected by centroid smoothing, but these are relatively small, and are corrected during reconstruction by fitting smoothing splines to the contour vertices, as explained in Section 6.4.2.

A 3D version of the smoothing spline, introduced in Section 6.2.1, was employed to fit a smooth curve to the centroids of registered wall contours. The simplest implementation uses equal weighting for all calculated centroid points (i.e. $w_i = 1$ for all i in Equation (6.8)). This scheme, however, was not found to work well because all points have equal influence to the fitted curve. It is evident from inspection of Figure 6.6(i) that, within a typical 3D ultrasound acquisition, the degree of misregistration will vary considerably along the vessel. Therefore, the ideal curve fitting scheme would assign a low degree of influence to those centroid points which exhibit the greatest misregistration. This observation suggests that using a non-uniform weighting for the smoothing spline will give more satisfactory results. The optimal weighting scheme, chosen to minimise the influence of the most misregistered centroid spline, is based on assessing the influence of each centroid point on a uniformly-weighted smoothing spline fitted to the centroid points. This is achieved by iteratively removing each point in turn, fitting a smoothing spline to the remaining points, and calculating the distance from the removed point to the fitted curve. When this distance is large, the influence of the point is high. Conversely, when the point is close to the fitted curve, it's influence is low. The following iterative, weighted smoothing spline (IWSS) algorithm was employed:

(*continued on page 195*)

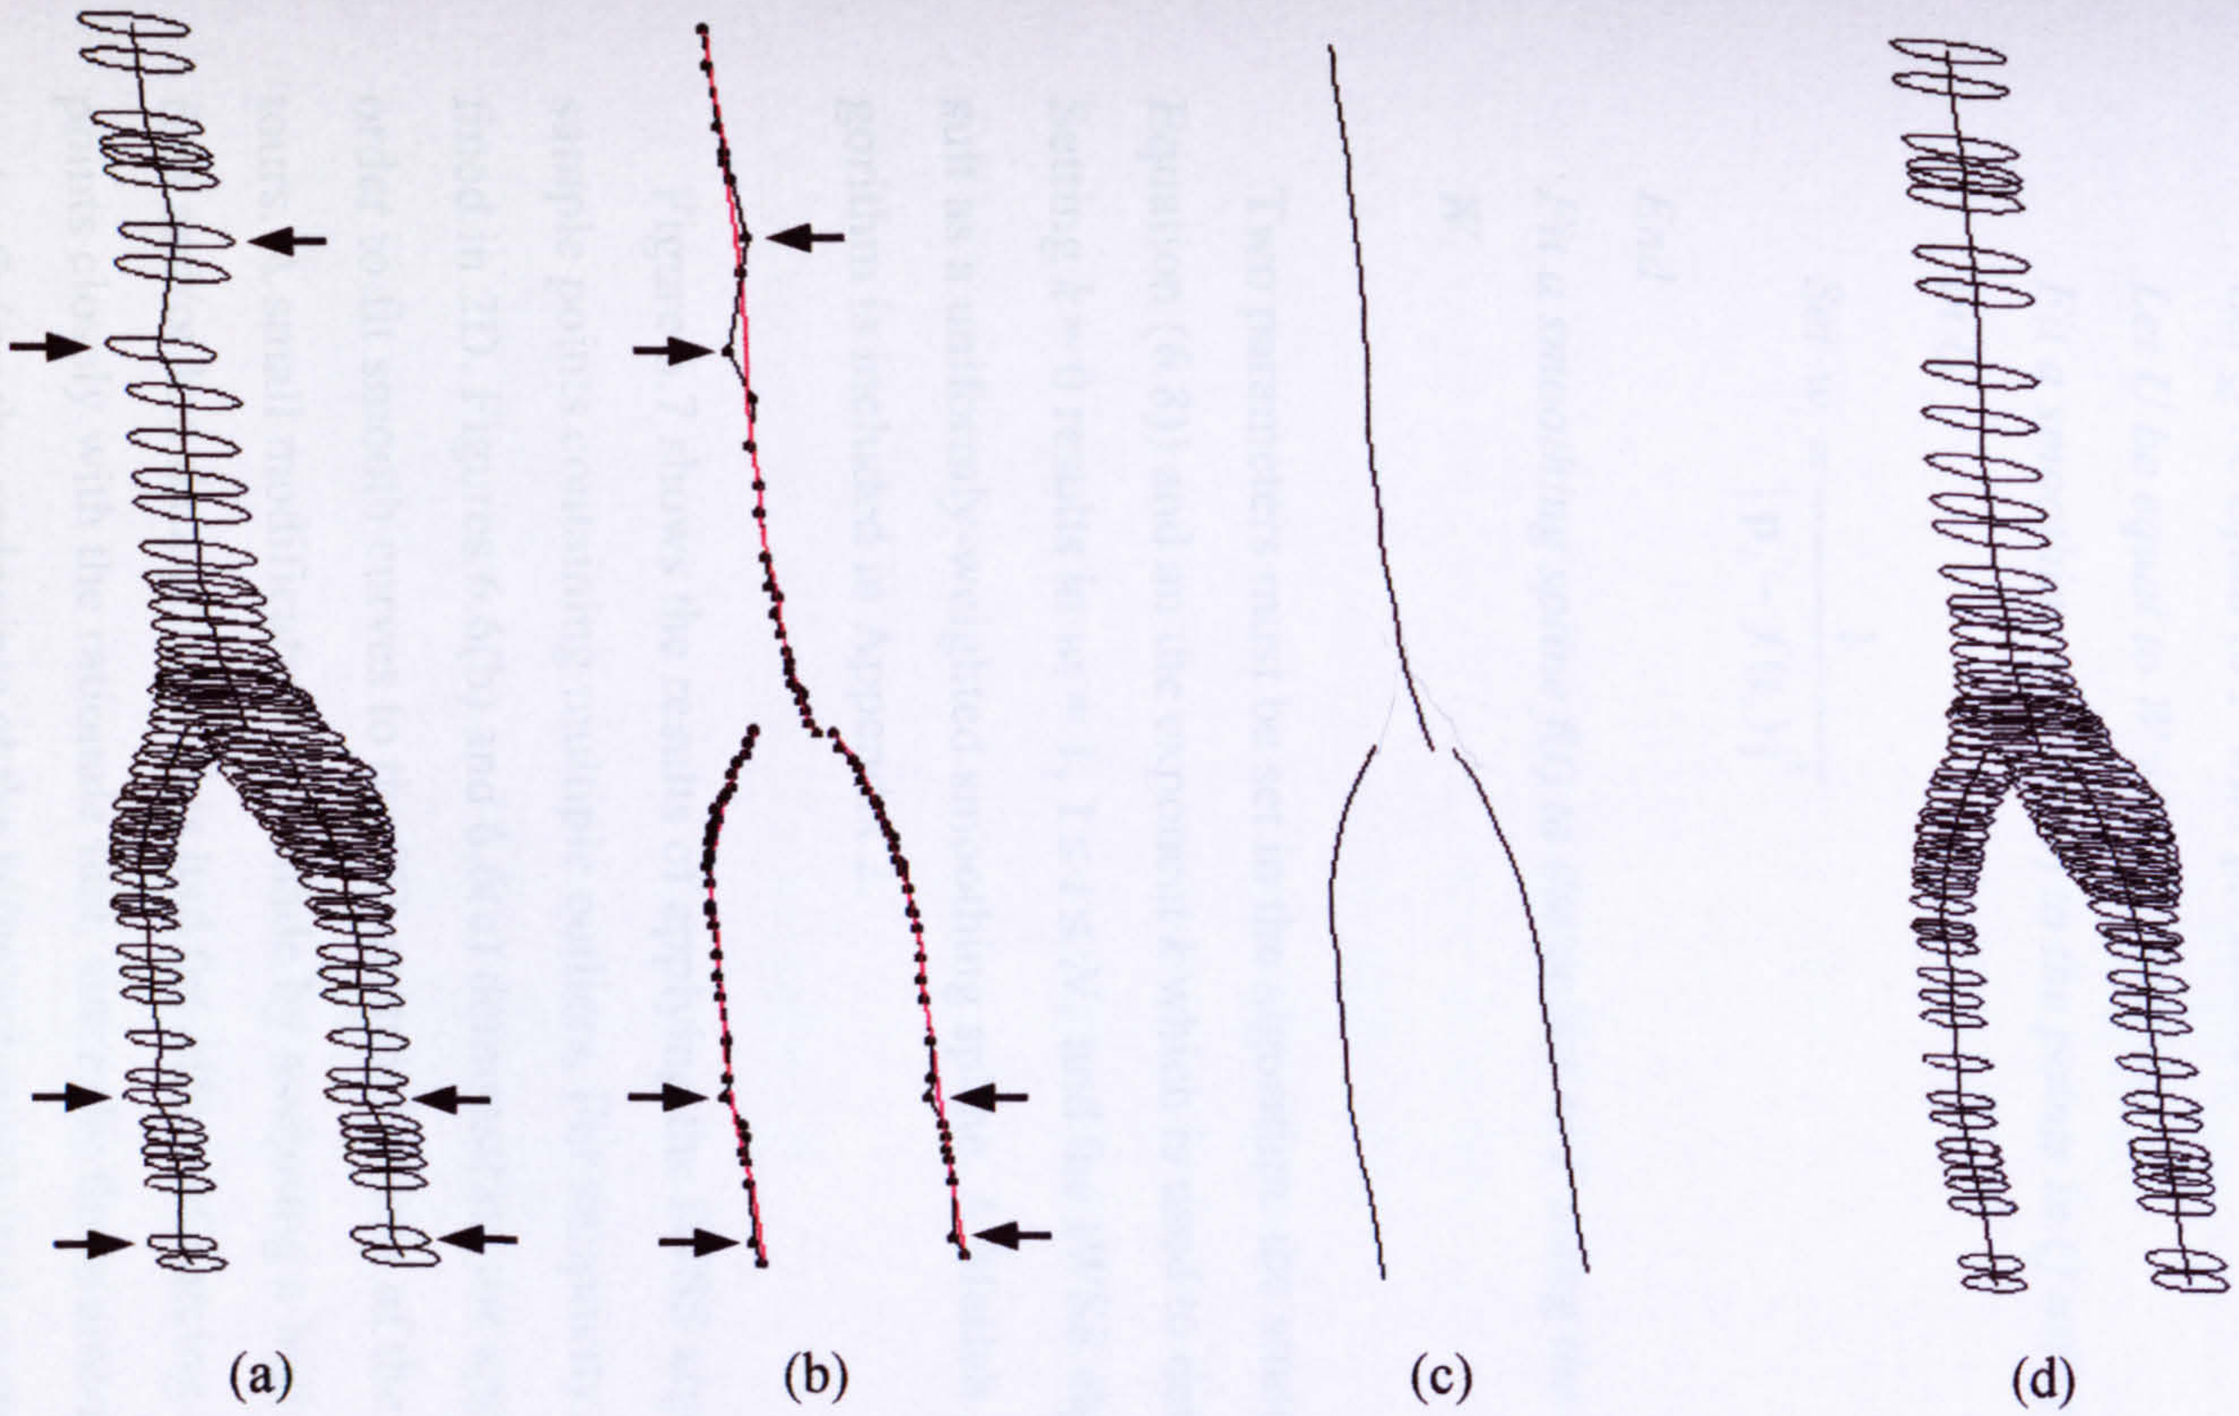


Figure 6.6 Centroid registration of cross-sectional wall contours. (a) The 3D-registered wall contours for a normal carotid bifurcation. The arrows highlight contours which are visibly misregistered. (b) The 3D curve connecting the centroids of the contours. The dots represent the locations of the original centroids, and the red curves represents the smooth curve fitted to these points which define the central axes for each vessel. (c) The smooth 3D central axis curves fitted to the centroid points. (d) The vessel wall contours following centroid registration. Note that the contours highlighted in (a) are now registered correctly.

Let $P = \{\mathbf{p}_1, \mathbf{p}_2, ..., \mathbf{p}_N\}$ *be the position vectors of the N noisy sample points, such that* $\mathbf{p}_i$ *specifies the 3D point* (x_i, y_i, z_i), $1 \le i \le N$.

Define the parameter set, $T = \{t_1, t_2, ..., t_N\}$, where $t_i = \dfrac{i-1}{N-1}$, $1 \le i \le N$.

Define the initial weights, $W = \{w_1, w_2, ..., w_N\}$, where $w_i = 1$, $1 \le i \le N$.

For i=1 to N

Let Q be equal to P with $\mathbf{p}_i$ *removed*

Let U be equal to W with w_i *removed*

Fit a smoothing spline f(t) to the points in Q using the corresponding weights in U

Set $w_i = \dfrac{1}{|\mathbf{p}_i - f(t_i)|^k}$

End

Fit a smoothing spline f(t) to the points in P using the corresponding weights in W.

Two parameters must be set in the algorithm: the smoothing parameter p (see Equation (6.8)) and an the exponent k which is used to determine w_i. Typically, $k < 4$. Setting $k = 0$ results in $w_i = 1$, $1 \le i \le N$, and the IWSS algorithm gives the same result as a uniformly-weighted smoothing spline. A Matlab implementation of the algorithm is included in Appendix 2.

Figure 6.7 shows the results of applying the IWSS algorithm to a set of noisy sample points containing multiple outliers. For simplicity, the sample points are defined in 2D. Figures 6.6(b) and 6.6(c) demonstrate the application of the algorithm in order to fit smooth curves to the 3D centroid points of the CCA, ICA, and ECA contours. A small modification was made by assigning a high weight (say, 100) to the first and/or last data points. This had the effect of forcing the spline to fit the end points closely with the rationale that, since the three arteries are treated independently, fixing the endpoints at the bifurcation ensured continuity between the smoothed contours. The effect of reregistering cross-sectional contours by smoothing the centroid curve can be seen in Figure 6.6 (d).

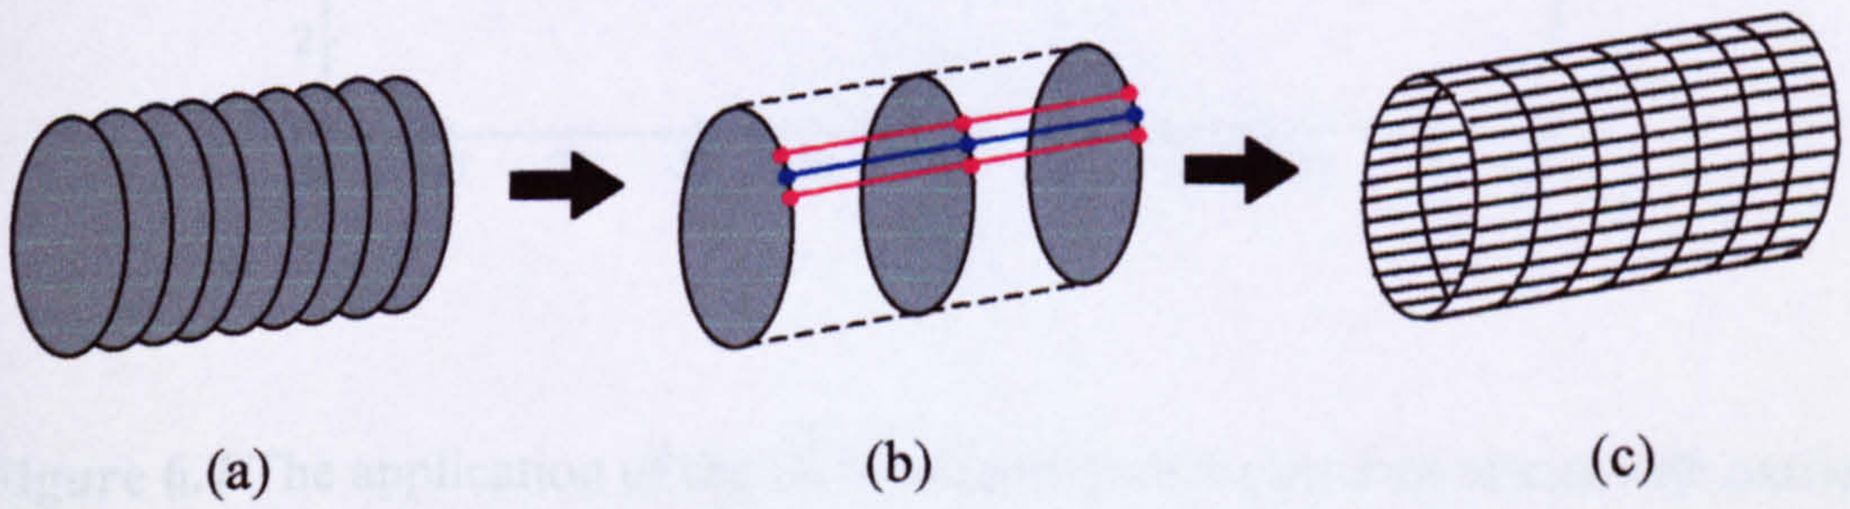
6.4.1.4 Contour resampling

Another prerequisite for the surface reconstruction algorithm is that the number of vertices within consecutive cross-sectional contours must be the same. This was achieved by resampling each contour to yield the desired number of equally-spaced vertices (typically, 100). Fitting a cubic spline to the planar contour vertices, using the arc-length parameter, s, provided an efficient means of resampling. Given N original contour vertices, $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_N$, the approximate arc-length parameter, s_i, corresponding to the i^{th} vertex ($1 \le i \le N+1$, where $\mathbf{p}_{N+1} = \mathbf{p}_1$), is given by:

$$s_i = \begin{cases} 0, & i = 1 \\ \sum_{j=1}^{i-1} \left| \mathbf{p}_{j+1} - \mathbf{p}_j \right|, & 1 < i \le N+1 \end{cases} \tag{6.19}$$

Provided N is reasonably large, Equation (6.19) will give a good approximation to the arc-length for smooth fitted spline. A periodic, cubic interpolating spline, $f(s)$, was then fitted to the contour points. Resampling the spline was achieved by dividing the range of s (= $[0, s_{N+1}]$) into M equal sections and evaluating the spline at these values of s to give $M+1$ new vertices. Since the first and last vertices will be the identical, the last is discarded, leaving M remaining vertices which are approximately equispaced.

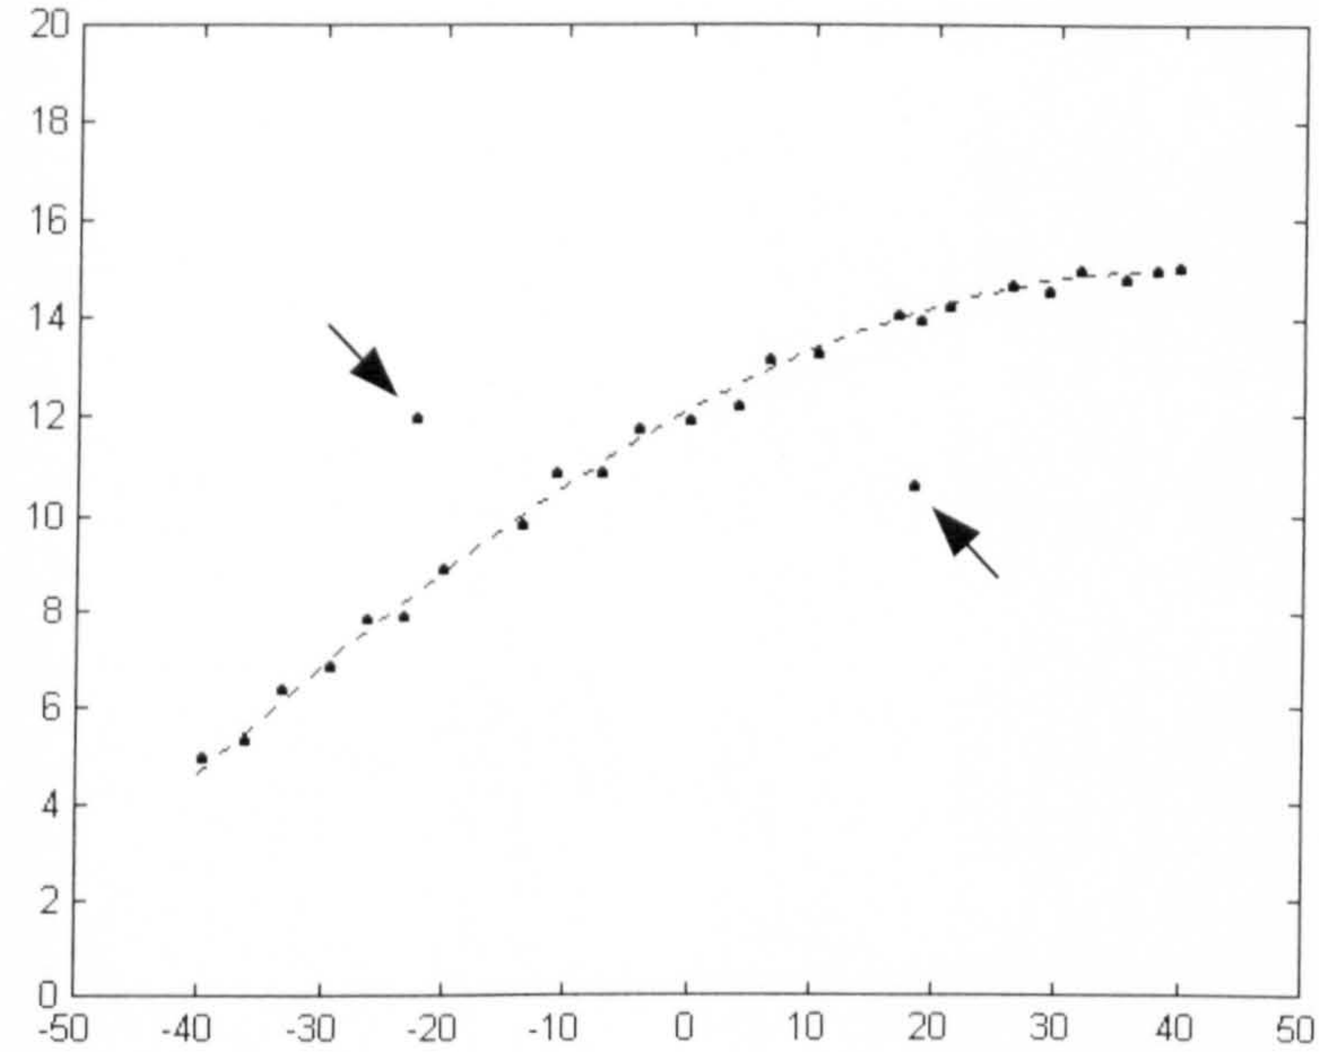
6.4.2 Reconstruction of the vessel wall



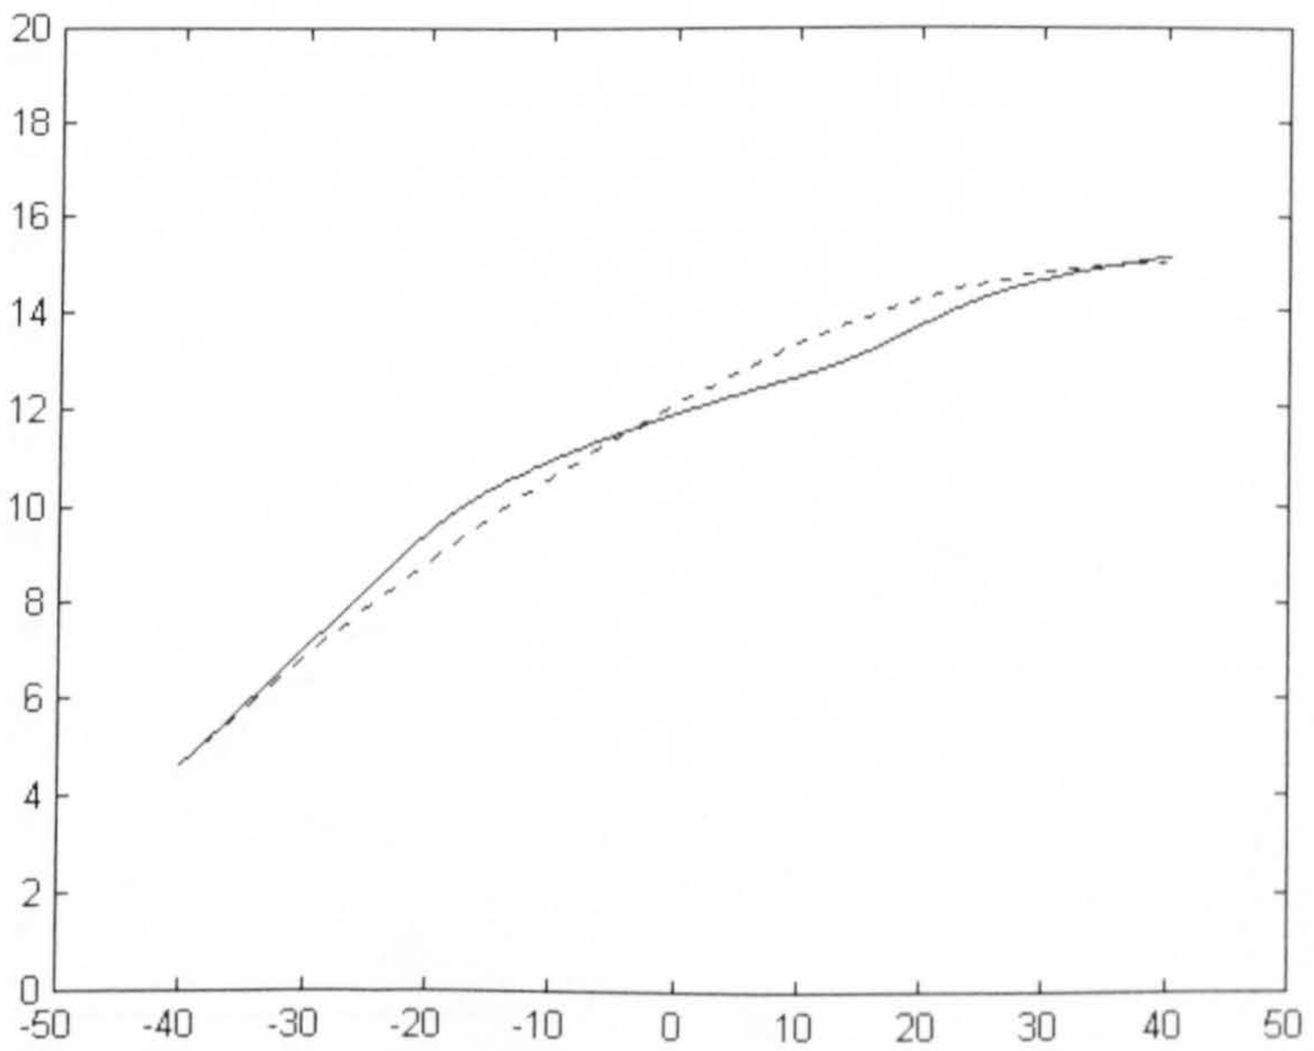
(a) (b) (c)

Figure 6.8 A schematic of the reconstruction of the 3D surface of a straight cylinder from cross-sectional contours: (a) the cross-sectional contours; (b) longitudinal curves are fitted to matching vertices on consecutive contours (matching vertices are similarly coloured); (c) the final reconstructed surface.

(*continued on page 200*)



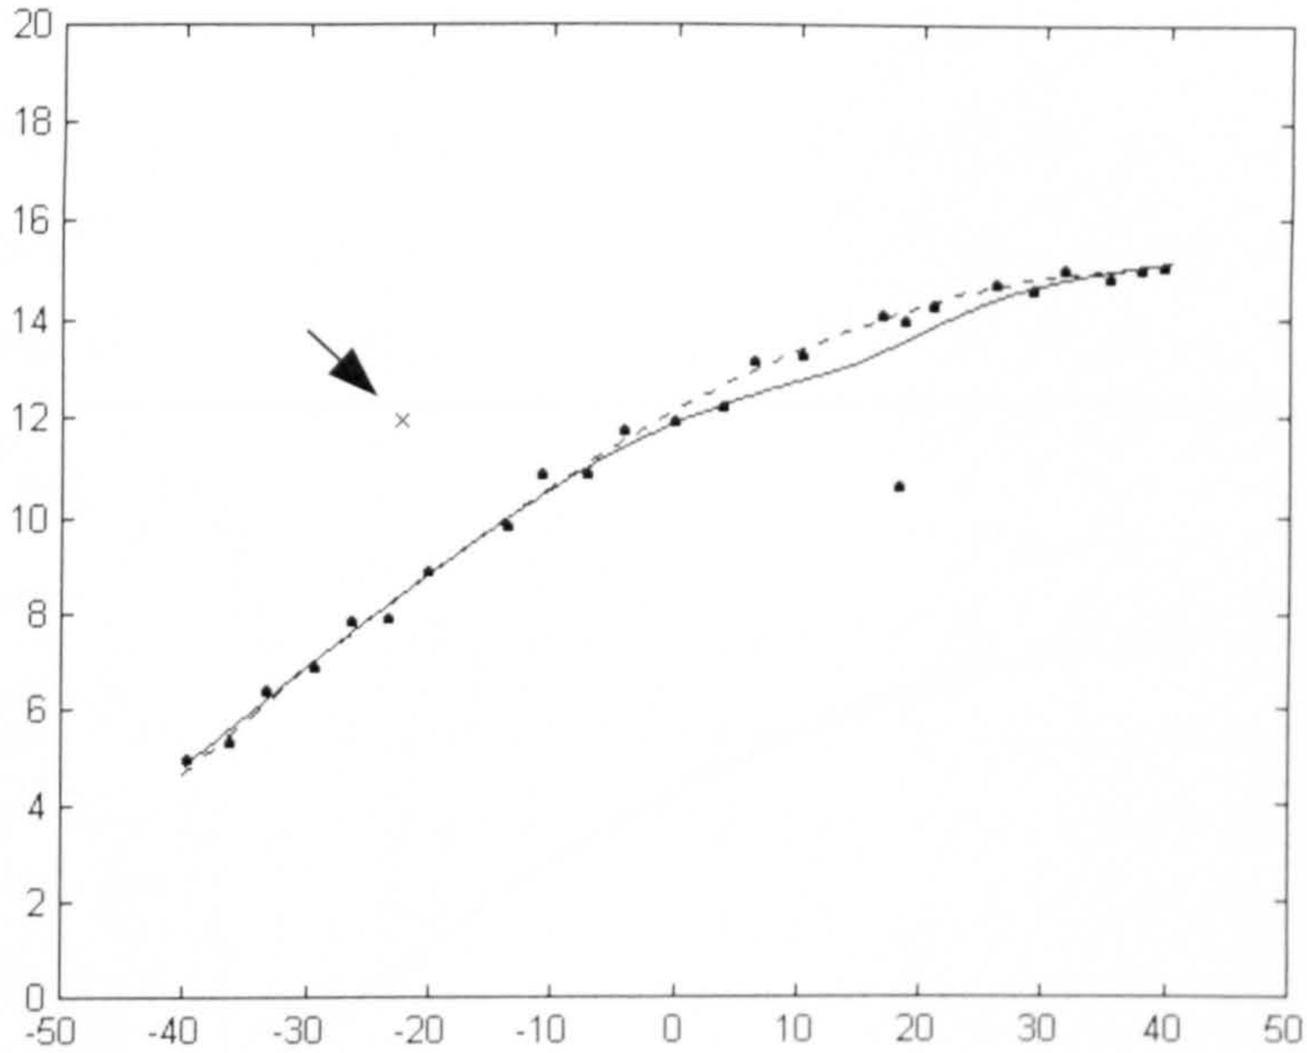
(a)



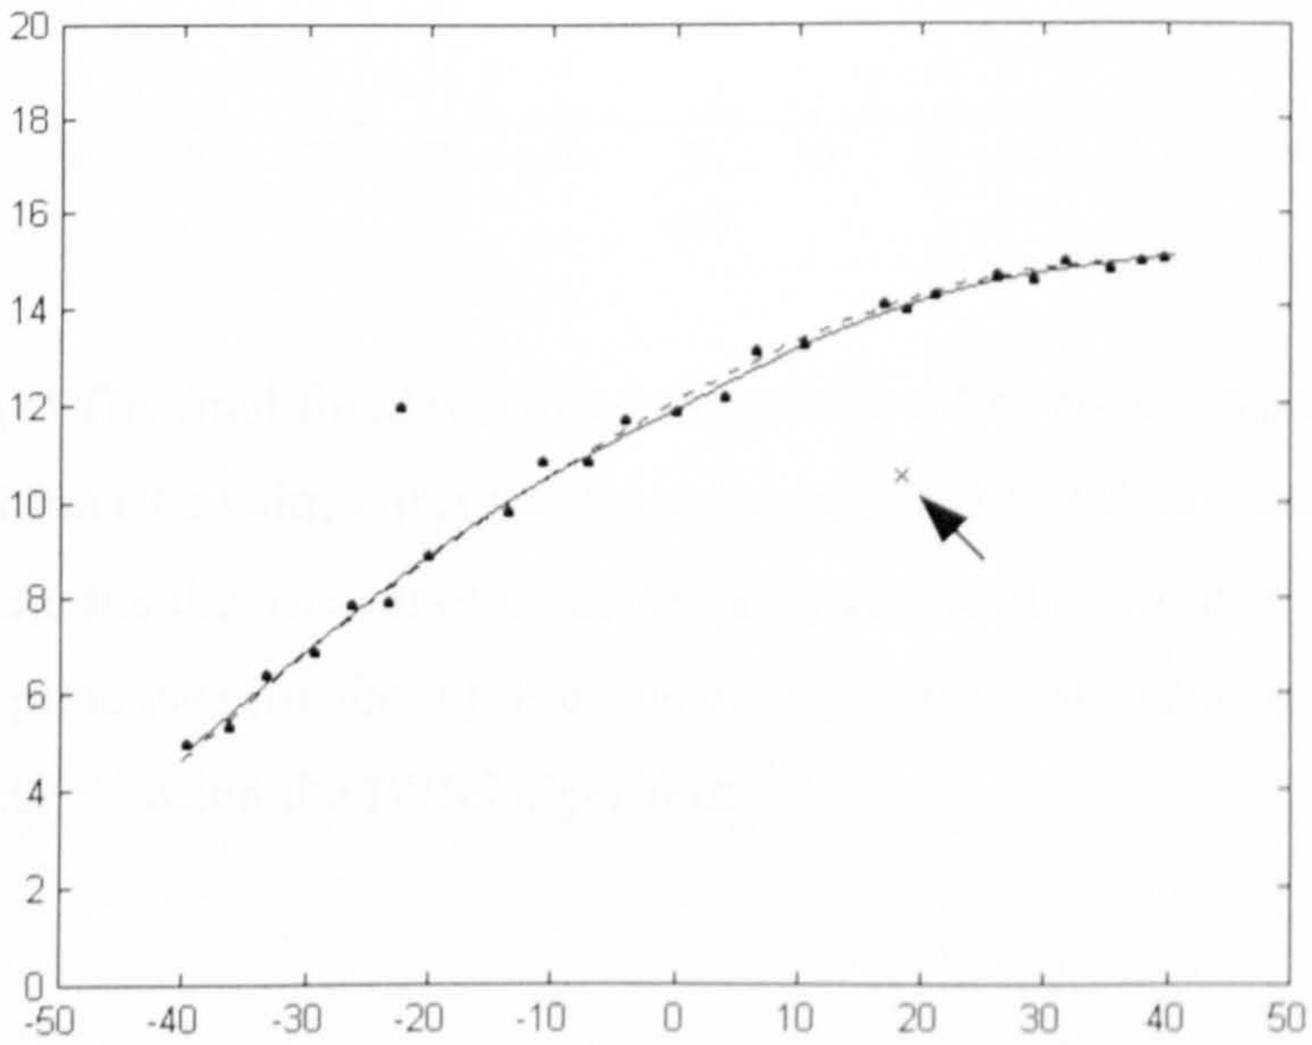
(b)

Figure 6.7 The application of the IWSS algorithm to noisy data points with outliers. (a) Noisy sample points with the true curve (shown as a dashed line). The arrows indicate outliers in the sample data. (b) the result of fitting a uniformly weighted smoothing spline to the points ($w_i = 1$ for all i, $p = 0.999$). Note how the outliers distort the fitted curve resulting in a poor fit to the true curve near these points.

(*continued overleaf*)



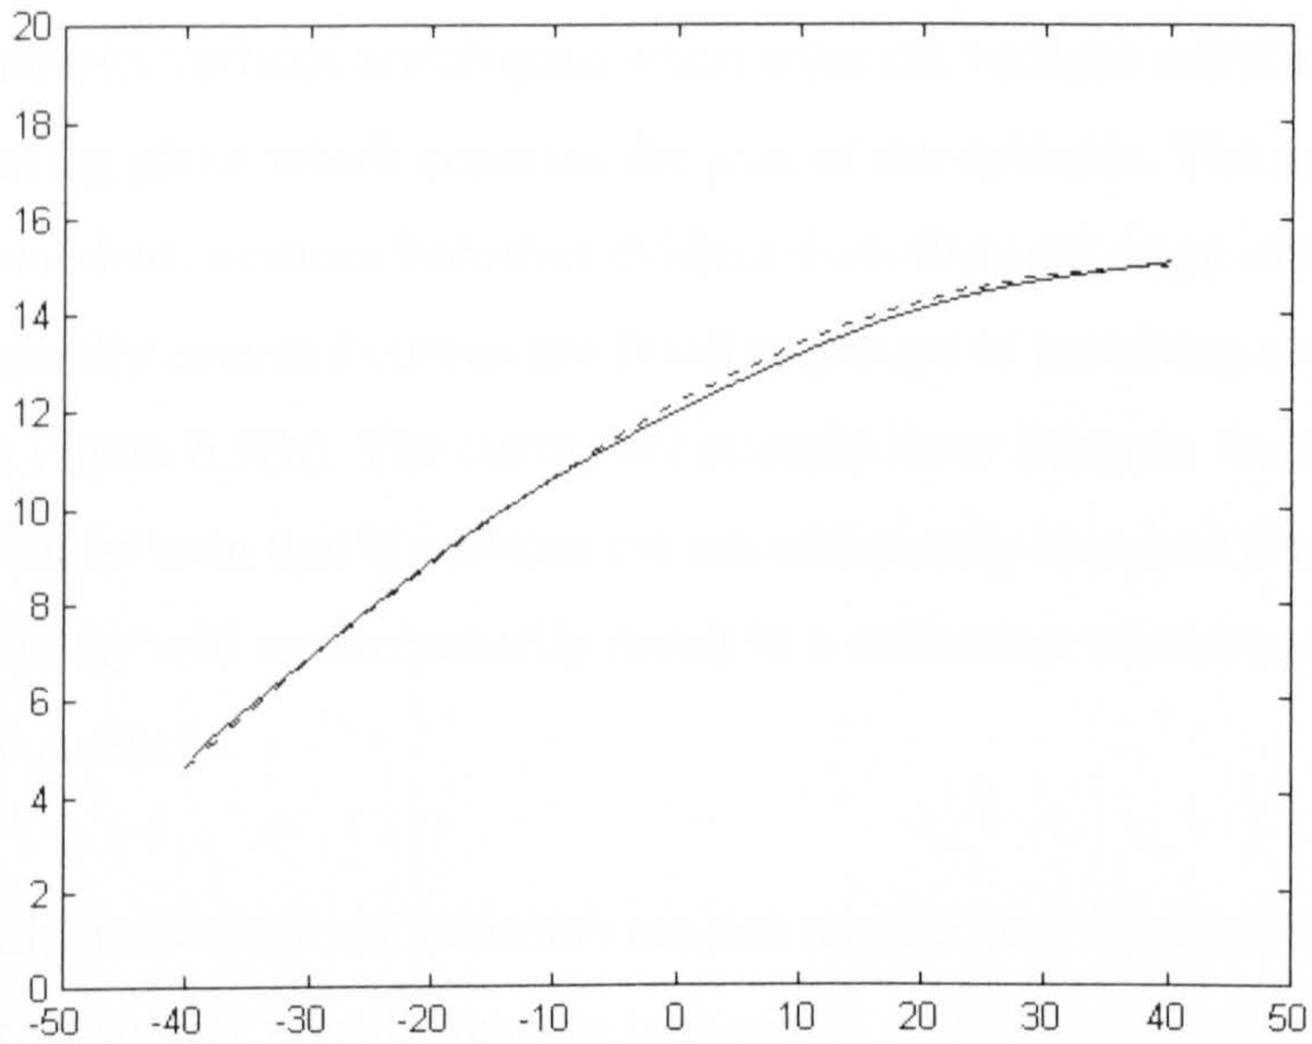
(c)



(d)

Figure 6.7 (c) The smoothing spline fitted during the 8th iteration of the algorithm (solid line). The outlying point, indicated by the arrow, has been removed and a weighted smoothing spline is fitted to the remaining points. (d) As for (c) but the spline fitted during the 17th iteration is shown.

(*continued overleaf*)



(e)

Figure 6.7 (e) The final fitted weighted smoothing spline following all 26 iterations of the algorithm (the values of p and k used were 0.999 and 1, respectively). Note that this spline fits the true curve much more closely than a uniformly weighted smoothing spline (see (b) above). For this example, the RMS error was reduced from 3.15 % to 1.05 % using the IWSS algorithm.

For clarity, the basic idea of the surface reconstruction method developed in this thesis is illustrated in Figure 6.8 using a cylindrical section as a simple example. Figure 6.8(a) shows a series of cross-sectional contours for the cylinder. It is assumed that the contours are non-overlapping, that each contour has the same number of vertices, and that the vertices of consecutive contours are aligned – i.e. the cumulative distance between adjacent pairs of vertices, one from each contour, is minimised. In this idealised example, vertices are aligned when adjacent vertices subtend the same angle to a bisecting plane which contains the axis of the cylinder. The necessity for aligned, or matched, vertices becomes evident from the next stage of the reconstruction, where axially oriented curves are fitted to groups of matching contours. This is illustrated in Figure 6.8(b). The curves are straight lines lying on the surface of the cylinder. It can be seen that if vertices are not adequately matched then fitting axial curves in this way will not necessarily result in a collection of curves that are coincident with the surface.

In general, cross-sectional contours are not regular, and fitted curves are splines which require a strictly non-decreasing parameter. Although it is possible to satisfy this condition by some reordering of the vertices from overlapping contours, this procedure was avoided by imposing the condition of non-overlapping contours using the method detailed in Section 6.4.1.2. In practice, very few slices actually overlap due the manner in which images are acquired within a single sweep of the scan-probe. Therefore, the data loss is minimal.

Figure 6.8(c) depicts the axial curves and the cross-sectional contours. Together, these form a mesh which approximates the surface of the cylinder. This method of surface reconstruction is applicable to semi-parallel contours of virtually any shape, such as those obtained during freehand 3DUS imaging. A minimum total Euclidean distance criterion was used to match the vertices of consecutive contours: in this scheme, if $\mathbf{p}_i^{C_j}$ is the 3D position vector of the i^{th} vertex ($0 \le i \le N-1$) in the j^{th} contour (see Figure 6.9), the vertices are matched by finding the value of k which minimises:

$$D = \sum_{i=0}^{N-1} \left| \mathbf{p}_i^{C_j} - \mathbf{p}_{(i+k) \bmod N}^{C_{j+1}} \right|. \tag{6.20}$$

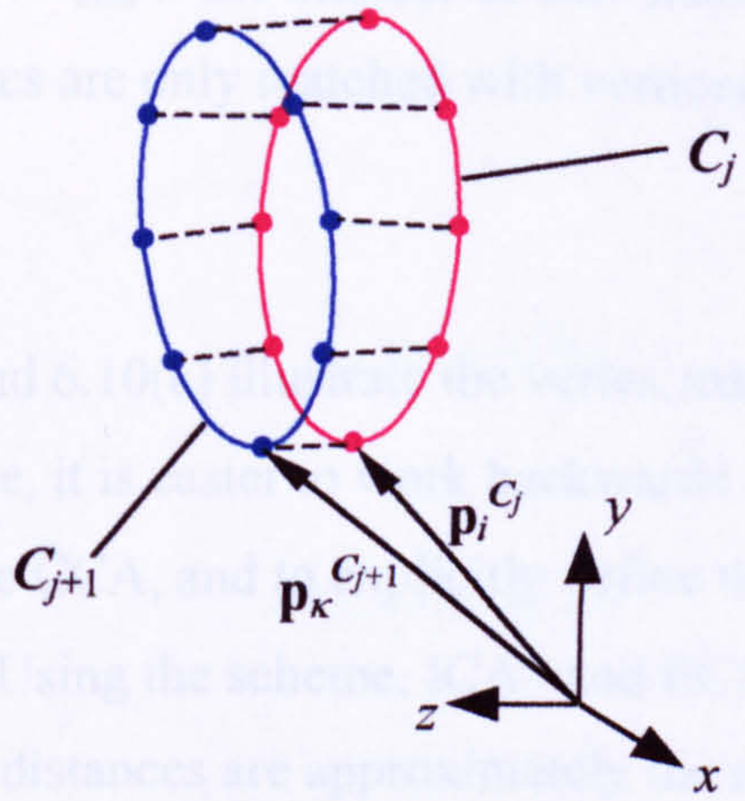


Figure 6.9 Matching vertices in consecutive contours C_j and C_{j+1}. Matched vertex pairs are connected by a dashed line, such that the quantity D, defined in Equation (6.20), is minimised.

Intuitively, each pair of matched vertices may be thought of as being connected by an elastic cord. Minimising D in Equation (6.20) then minimises the total tension – or, equivalently, elastic potential energy – in the cords. This simple vertex matching algorithm was found to work well for straight sections of vessel, as the shapes of contours do not differ considerably between consecutive slices. As it stands, however, this method cannot handle vessel bifurcations.

To deal with a bifurcating vessel the follow adaptations were made: the contour of the CCA immediately proximal to the apex of the bifurcation, denoted by CCA* for the purposes of this discussion, was assumed to be identical to the contour enclosing the union of two overlapping contours, ICA* and ECA*, as shown in Figure 6.10(a). Importantly, the CCA* contour shares all of its vertices with the ICA* and ECA* contours. Hence, non-shared vertices completely define the ICA* and ECA* contours where they intersect. The number of vertices in the ICA* and ECA* contours must also correspond to the number in ICA and ECA contours, respectively. In order to ensure a consistent number of vertices, the following condition must be satisfied:

$$N_{\mathrm{CCA}} + N_{\mathrm{extra}} = N_{\mathrm{ICA}} + N_{\mathrm{ECA}}\,, \tag{6.21}$$

where N_{CCA}, N_{ICA}, and N_{ECA} are the number of vertices in the CCA, ICA, and ECA

contours, respectively. N_{extra} is the number of non-shared vertices. It should be noted that these extra vertices are only matched with vertices of the subsequent ICA and ECA contours.

Figures 6.10(b) and 6.10(c) illustrate the vertex matching scheme for the carotid bifurcation. In practice, it is easier to work backwards from the distal end of the daughter vessels to the CCA, and to explicitly define the ICA* and ECA* contours during segmentation. Using the scheme, ICA* and ECA* are resampled first so that the mean inter-vertex distances are approximately the same for both contours. These contours are then merged to find the enclosing contour which determines CCA* and, hence, N_{CCA}.

In the vertex matching procedure described above, the apex of the bifurcation is not defined explicitly, but implicitly. Indeed, it is not usually possible to locate the apex of the bifurcation precisely from acquired ultrasound slices. This is due partly to low resolution in the elevational direction, but also to the fact that the acquisition process involves sampling the vessel along its length. Therefore, even with slow acquisition sweeps, images are typically captured proximal and distal to the apex of the bifurcation. The procedure can be applied to both wall or lumen contours, provided the shape of contours changes reasonably gradually from slice to slice. For each artery, information for vertex matching is stored most efficiently by storing a value of k (Equation (6.20)) for each contour. Each matched vertex can then be found and the vertices reordered appropriately. Once vertices have been successfully matched, splines are fitted longitudinally to the matched vertices. Again, smoothing splines were chosen for this purpose by virtue of their ability to remove noise in the sample points. Because the IWSS algorithm relatively expensive in terms of computation time, and has marginal benefits following centroid registration of the original wall contours, a uniformly weighted smoothing spline was used instead to provide an efficient means of approximating wall surfaces. Figure 6.11 shows the results of fitting longitudinal splines to the contours of the example shown previously in Figure 6.6.

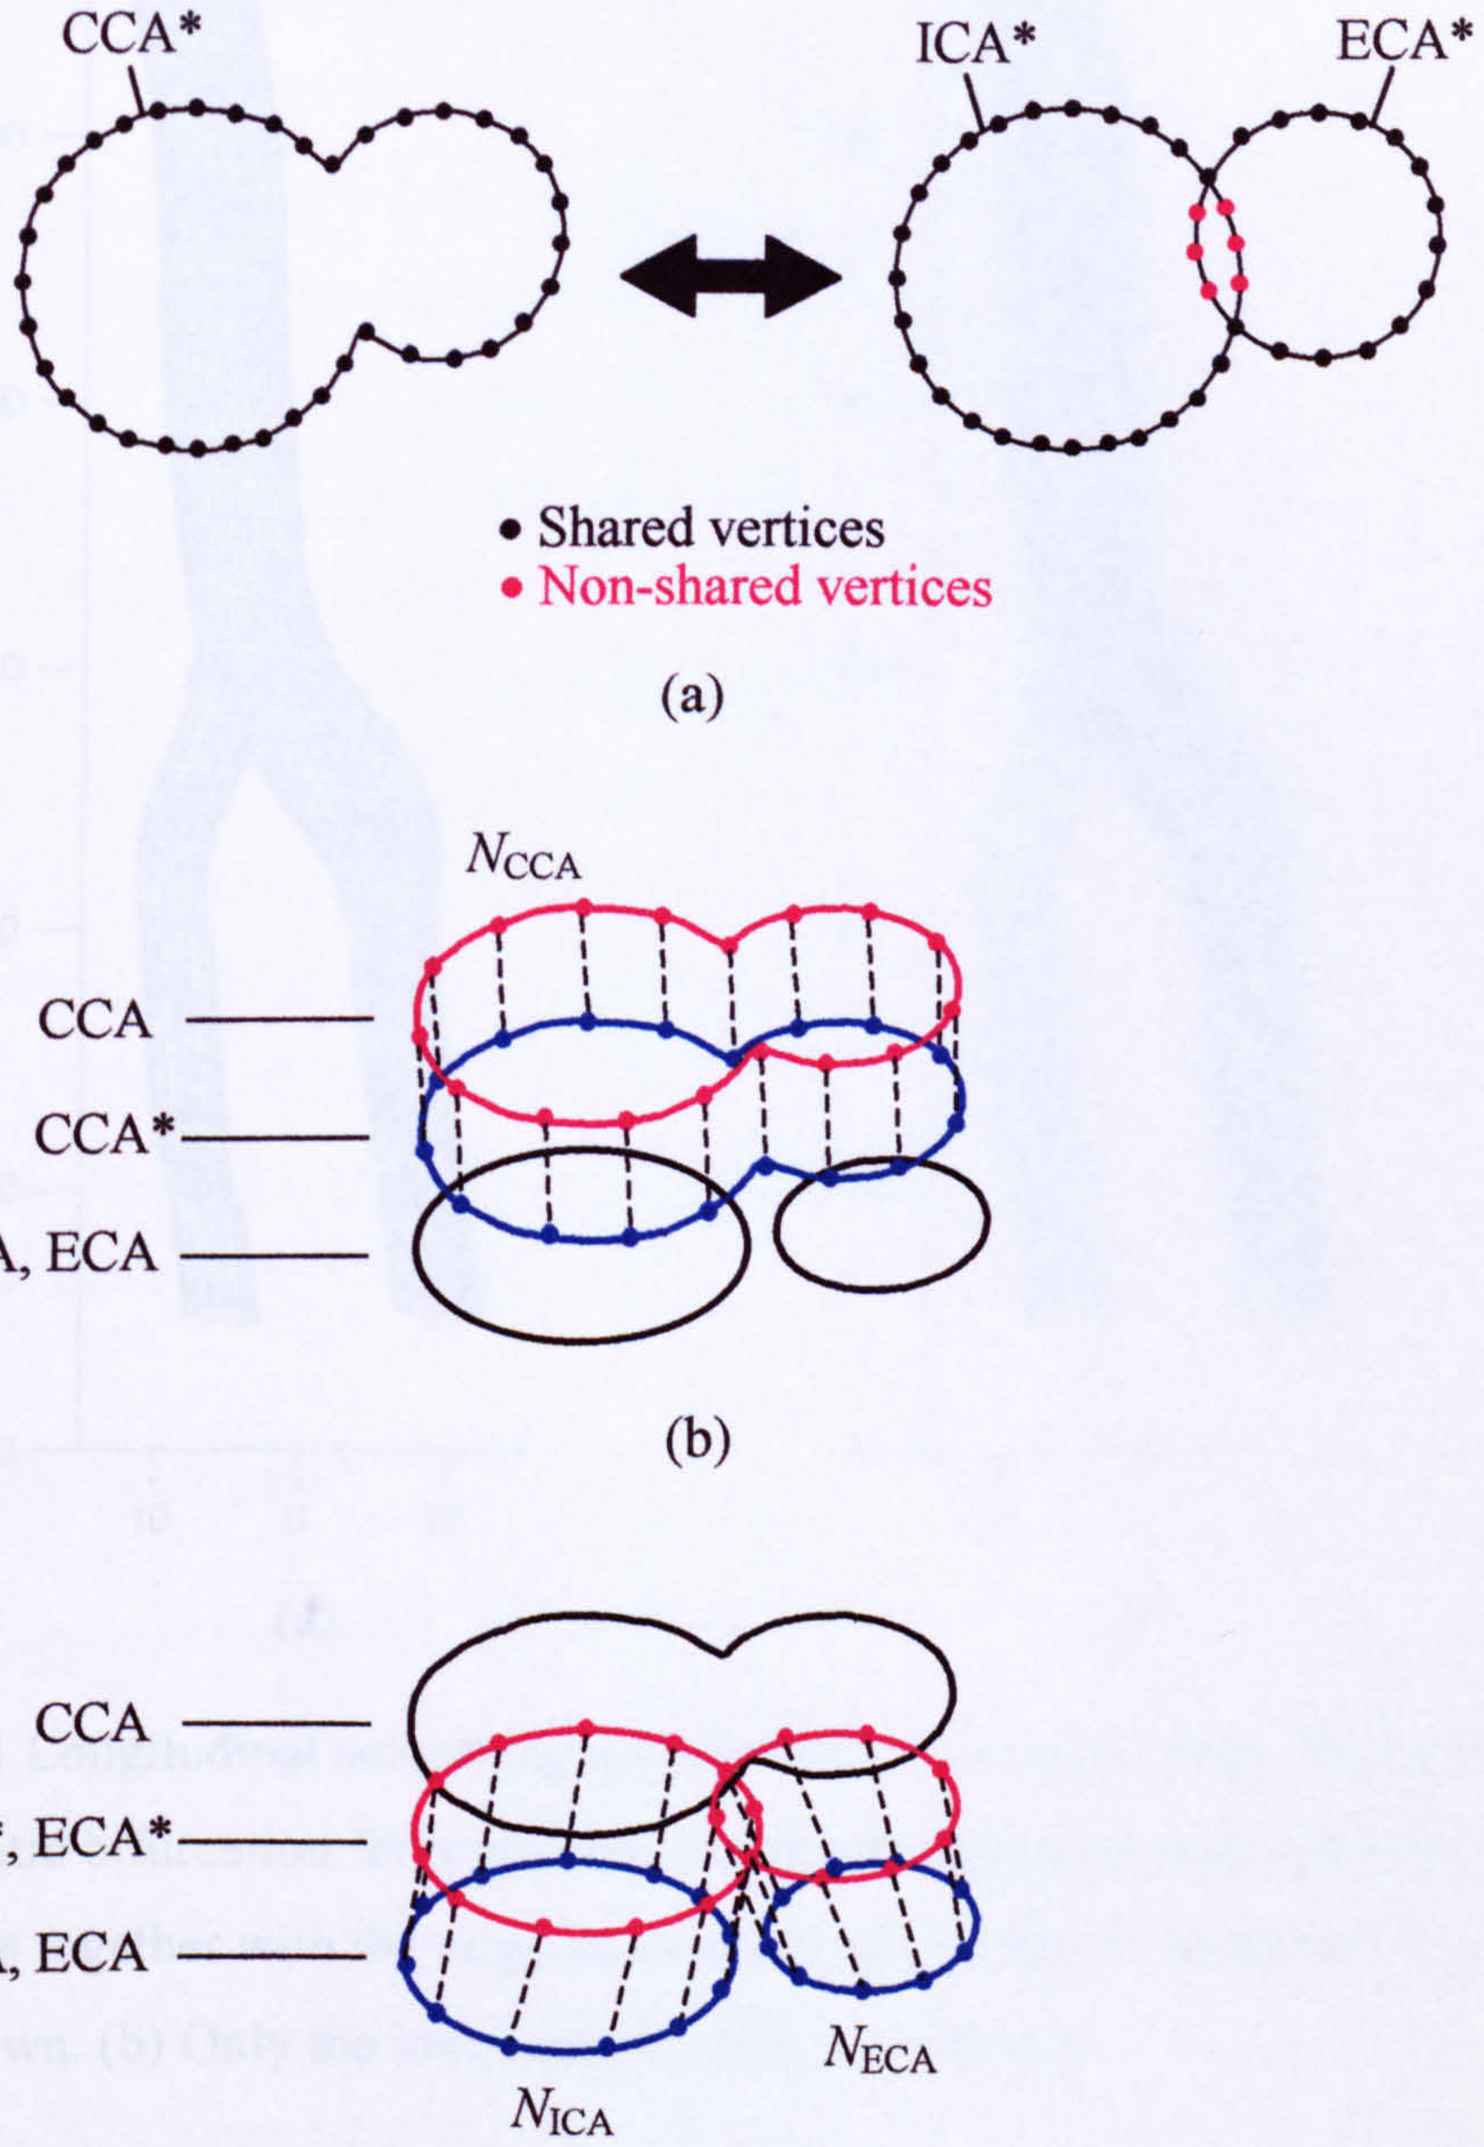


Figure 6.10 Vertex matching scheme at a bifurcation. (a) The CCA contour immediately proximal to the bifurcation (CCA*) encloses the union of the two intersecting contours, ICA* and ECA*. (b) Vertex matching scheme proximal to the bifurcation. N_{CCA} denotes the number of vertices in the CCA contours. (c) Vertex matching scheme immediately distal to the bifurcation. N_{ICA} and N_{ECA} denote the number of vertices in the ICA and ECA contours, respectively.

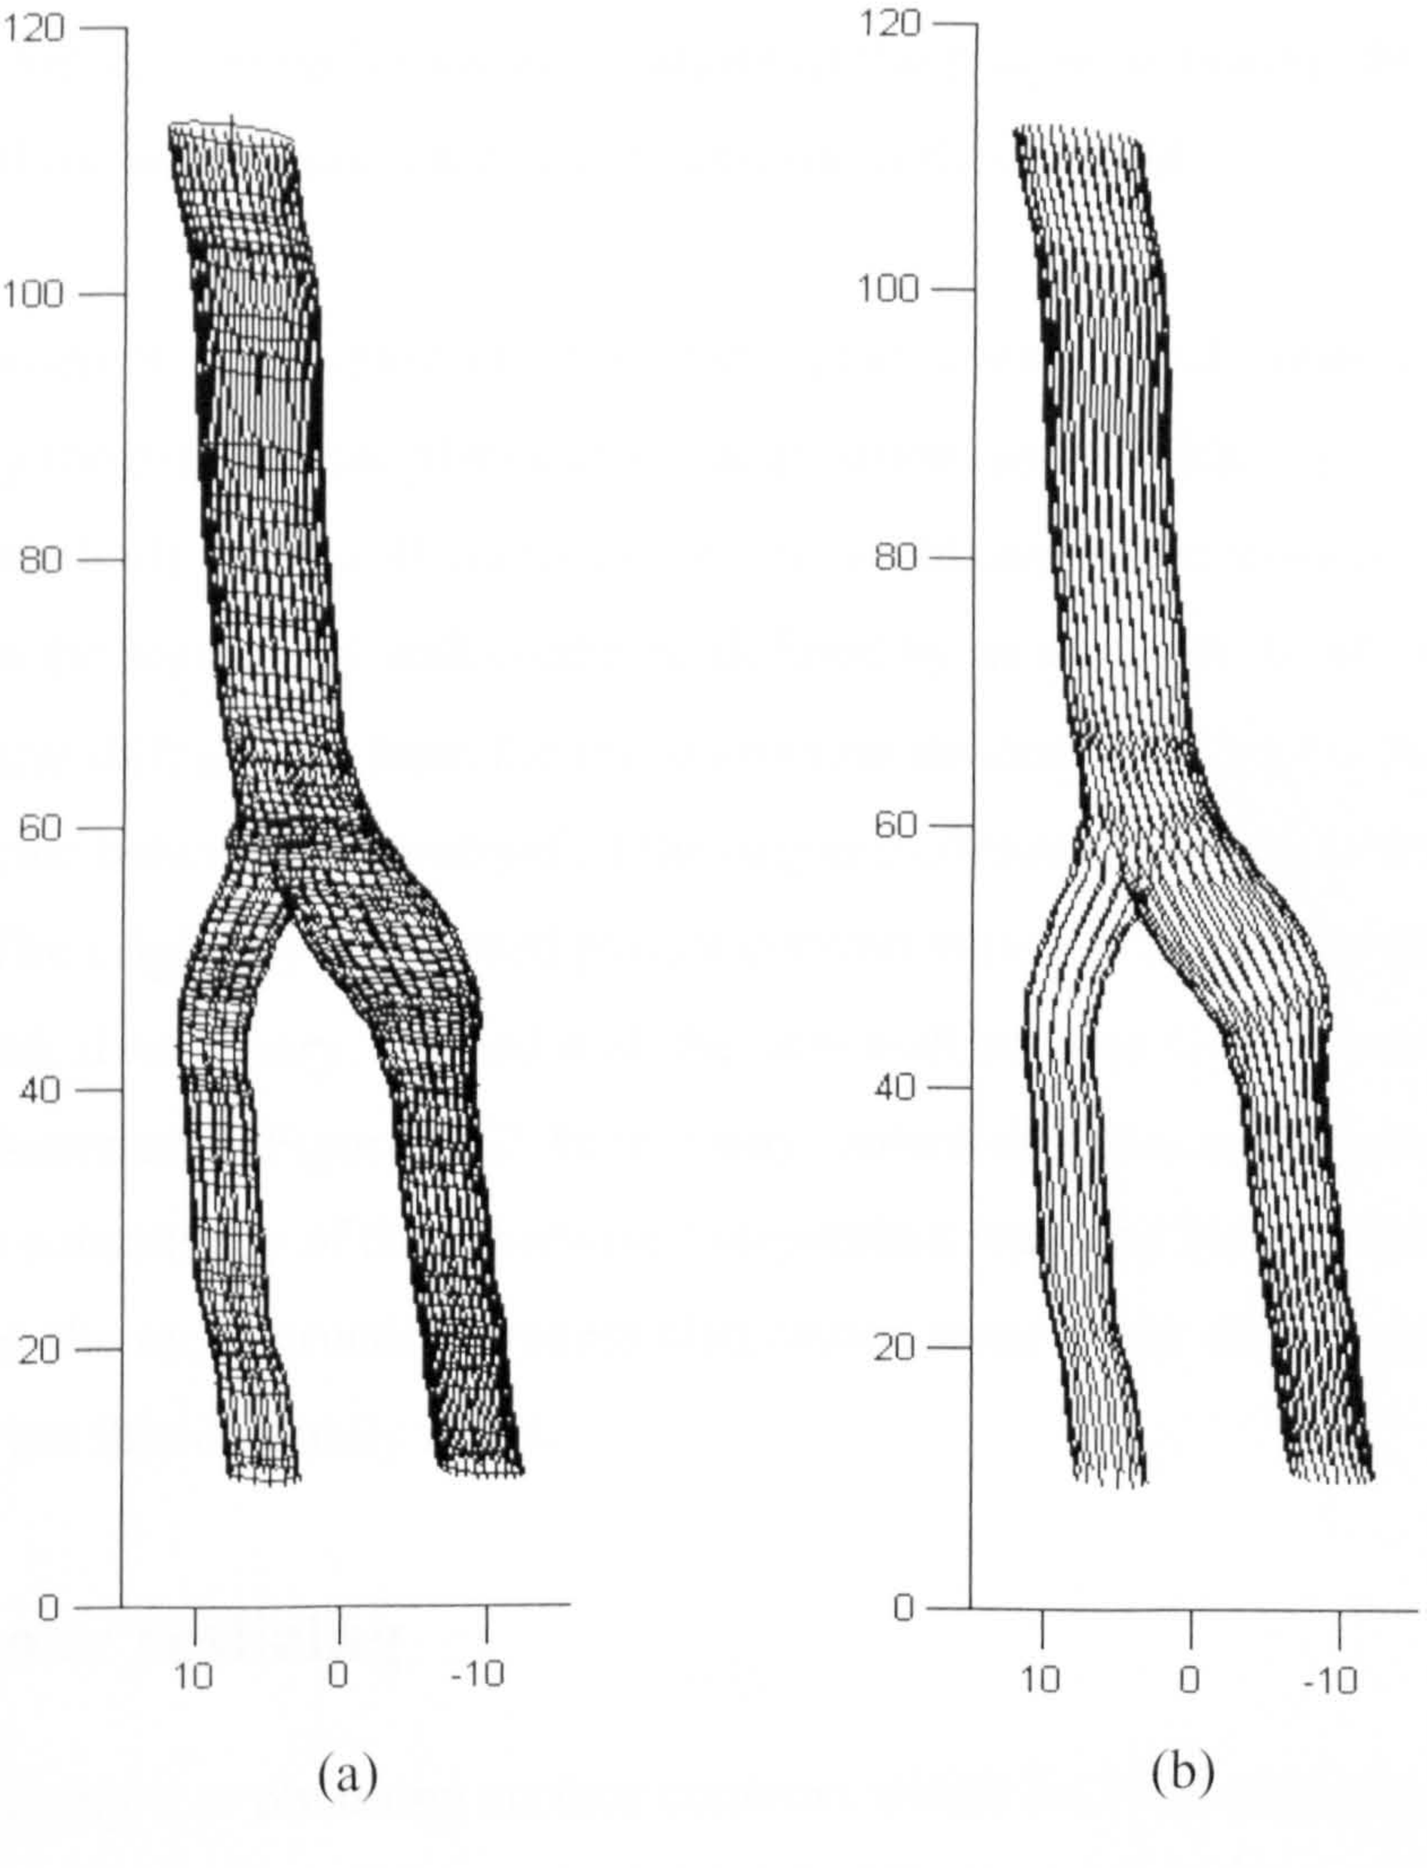


Figure 6.11 Longitudinal smoothing splines fitted to cross-sectional contours of a normal carotid bifurcation following the vertex matching procedure. (a) The longitudinal splines together with the original contours after centroid smoothed (see Figure 6.6) are shown. (b) Only the longitudinal splines are shown.

6.4.3 Reconstruction of the vessel lumen

The vessel lumen can be reconstructed using exactly the same technique as for the vessel wall. Due to the irregular nature of carotid plaque, however, centroid registration is not appropriate for lumen contours. To understand the effects of this, consider a disease-free vessel whose wall and lumen are defined by the same, single set of cross-sectional contours. Independent reconstruction of the wall and lumen, with centroid registration being applied to the wall but not the lumen contours, results in two surfaces which do not coincide exactly. A further complication arises in the diseased artery where it may be desirable to apply different spline smoothing parameters when reconstructing the vessel wall and lumen. This may be necessary, for ex-

ample, to preserve irregular surface features of the plaque. However, the disease-free portions of the artery may not coincide exactly as they should.

In an attempt to minimise these problems, the reconstructed vessel wall was resliced using the original cut-planes of the acquisition (see Section 6.5 for details of the reslicing method). New wall contours were then obtained from these slices and compared with the segmented wall contours, defined by an observer, to determine the within-plane shift arising from the reconstruction process. This shift was quantified by the vector between the centroid of the original contours and that of the resliced contour. The originally segmented plaque contour was then shifted by the same amount and, if necessary, merged with the new wall contour to form a lumen contour, as illustrated in Figure 6.12. In this way, lumen contours are registered to increase the consistency of the otherwise independent wall and lumen reconstructions. In practice, the reconstruction process also causes some shape distortion, in addition to a shift, but this is usually small.

6.5 Planar reslicing

Planar reslicing – or obtaining surface contours which lie in an arbitrary plane – is an important operation for the interrogation of the 3D reconstructed carotid bifurcation. Reslicing is most meaningful in transverse planes along the vessel. Using a central axis as a reference (defined in Section 6.6), it is possible to reslice the vessel along its length using cut-planes which are accurately perpendicular to the central axis. This, in turn, enables automatic cross-sectional area and/or diameter measurements to be made throughout the vessel without the need for user interaction. In general, such multiple measurements are likely to increase the accuracy and reliability of conventional measures of disease severity; in particular, the selection of a single maximum diameter or area stenosis measurement. In addition, more information is provided on the distribution of the disease in individual patients. For the longitudinal spline representation, the reslicing problem is essentially one of finding the intersections of all of the splines with a cut-plane. Once found, the intersection points for a contour definine a section through the vessel. With reference to Figure 6.13, a plane is uniquely defined in 3D-space by a point, $\mathbf{p}_0$, which lies in the plane, and two orthogonal unit vectors, denoted by $\hat{\mathbf{v}}_x$ and $\hat{\mathbf{v}}_y$, which also lie within the plane.

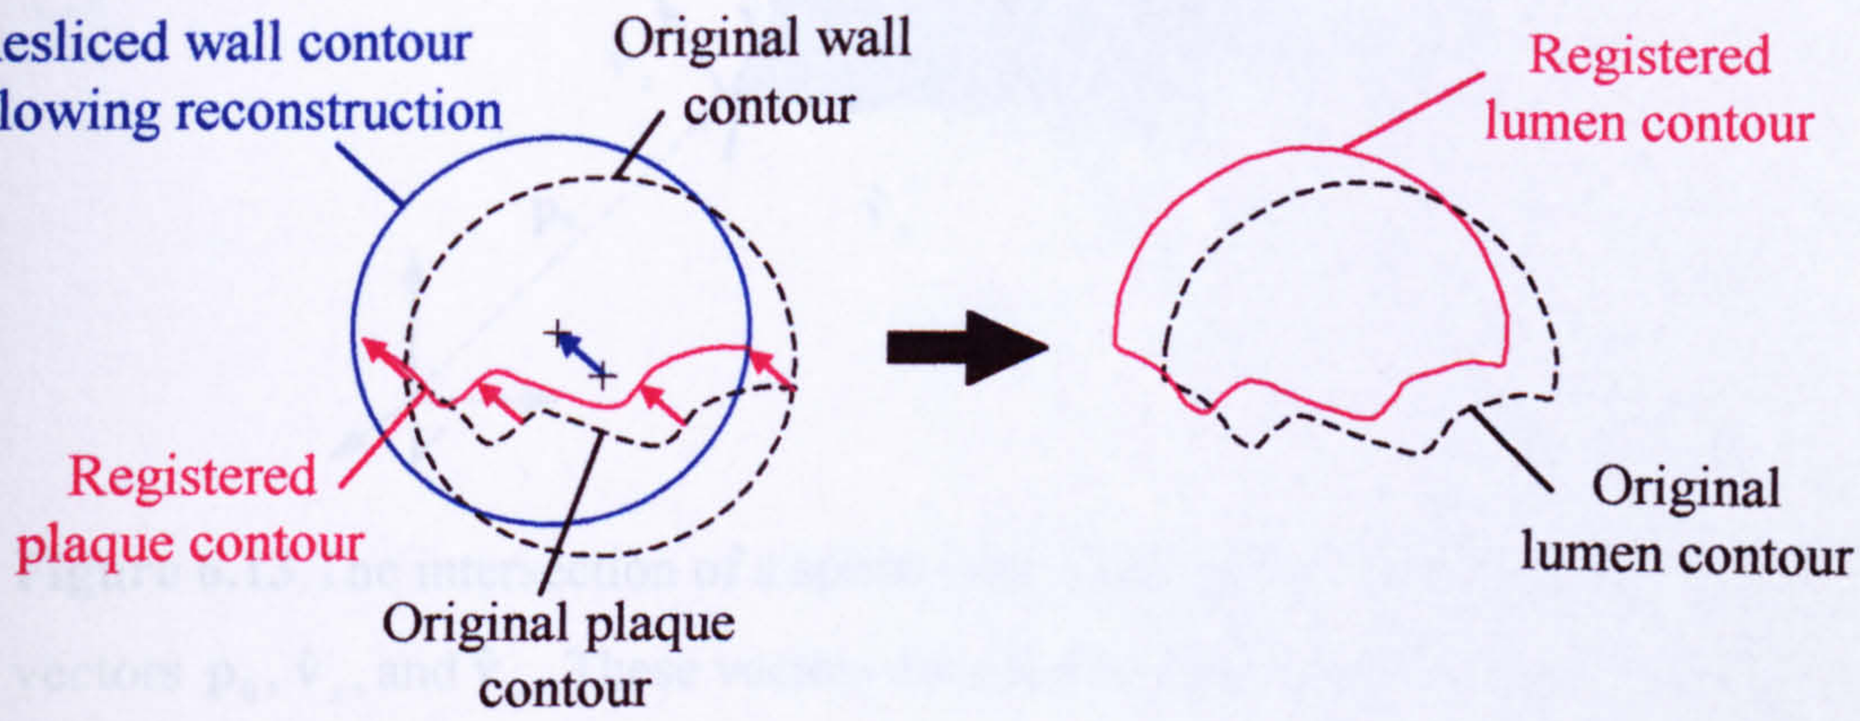


Figure 6.12 Within-plane registration of plaque contours for (a) an open plaque contour, and (b) a closed plaque contour. In both cases, the original, unregistered contours are indicated by dashed a line and registered contours are coloured. The arrows indicate the registration shift applied to plaque contours. This shift is defined by the vector between the centroids of the original wall contours and the contour obtained by reslicing the 3D reconstruction using a cut-plane corresponding to the plane of the original, acquired ultrasound image.

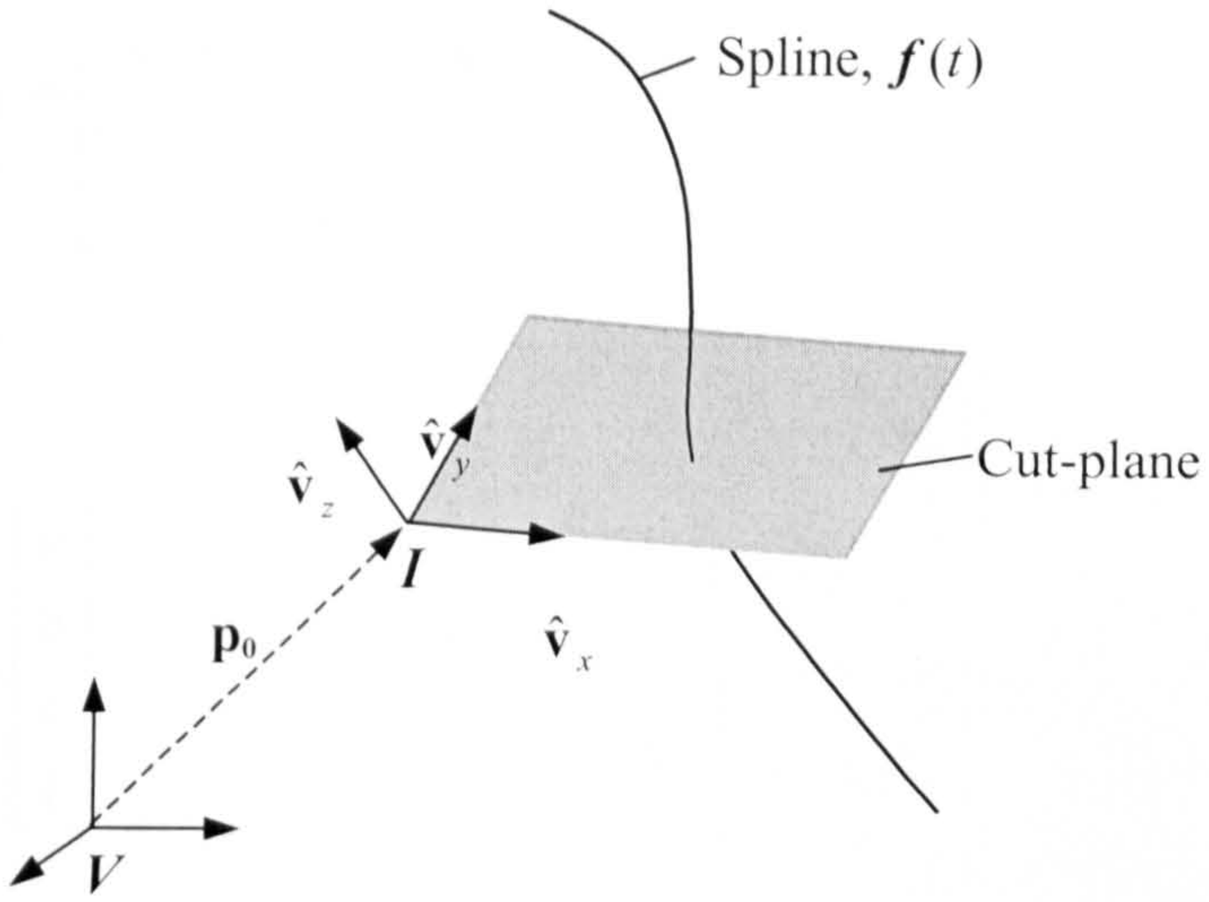


Figure 6.13 The intersection of a spline with a cut-plane, uniquely defined by the vectors $\mathbf{p}_0$, $\hat{\mathbf{v}}_x$, and $\hat{\mathbf{v}}_y$. These vectors are used to define a 3D co-ordinate system, $\boldsymbol{I}$, for the plane, and the point of intersection is found by transforming the spline to $\boldsymbol{I}$, and determining the point with a zero z-component.

With respect to the co-ordinate system $\boldsymbol{V}$, any point in the plane is specified by:

$$(\mathbf{p}_{\text{plane}})_V = \mathbf{p}_0 + a\hat{\mathbf{v}}_x + b\hat{\mathbf{v}}_y, \tag{6.22}$$

where a and b are scalars. Since $\hat{\mathbf{v}}_x$ and $\hat{\mathbf{v}}_y$ are orthogonal, they can be used to define a 3D co-ordinate system, $\boldsymbol{I}$, with x-, y-, and z-axes defined by the vectors $\hat{\mathbf{v}}_x$, $\hat{\mathbf{v}}_y$, and $\hat{\mathbf{v}}_z$, respectively, where

$$\hat{\mathbf{v}}_z = \hat{\mathbf{v}}_x \times \hat{\mathbf{v}}_y. \tag{6.23}$$

The position vectors which specify a single 3D point with respect to $\boldsymbol{V}$ and $\boldsymbol{I}$ are then related by:

$$\mathbf{p}_V = \mathbf{p}_0 + a\hat{\mathbf{v}}_x + b\hat{\mathbf{v}}_y + c\hat{\mathbf{v}}_z = {}^V\mathbf{M}_I\mathbf{p}_I, \tag{6.24}$$

where

$$ {}^{V}\mathbf{M}_{I} = \begin{bmatrix} \begin{bmatrix} \hat{\mathbf{v}}_x \\ 0 \end{bmatrix} & \begin{bmatrix} \hat{\mathbf{v}}_y \\ 0 \end{bmatrix} & \begin{bmatrix} \hat{\mathbf{v}}_z \\ 0 \end{bmatrix} & \begin{bmatrix} \mathbf{p}_0 \\ 1 \end{bmatrix} \end{bmatrix} \tag{6.25} $$

and

$$ \mathbf{p}_I = \begin{bmatrix} a \\ b \\ c \\ 1 \end{bmatrix}. \tag{6.26} $$

Applying Equation (6.24) to a 3D spline, $f(t)$, we have:

$$ f_V = {}^{V}\mathbf{M}_I f_I . \tag{6.27} $$

Rearranging and substituting Equation (6.5) gives:

$$ f_I = {}^{V}\mathbf{M}_I^{-1}\mathbf{GT} = \mathbf{HT} . \tag{6.28} $$

Now, following transformation into the co-ordinate system I, the z-component of f_I is zero at the intersection of the spline with the plane. Therefore, Equation (6.28) yields the following condition for intersection:

$$ \left(\mathbf{H}_i\mathbf{T}\right)_z = Q_i(t) = 0 , \tag{6.29} $$

where $Q_i(t)$ is a polynomial in t defined on the interval $[0, t_{max}]$. The intersection problem then reduces to finding the real roots of a polynomial, $Q_i(t)$, which lie within the interval $[0, t_{max}]$.

Since a spline is formed from multiple polynomial pieces, a brute force approach to finding the intersection of a spline with a plane would involve computing the roots for all polynomial pieces and determining whether an intersection exists for each. This method is somewhat inefficient, however, particularly when a spline has many pieces. A more efficient strategy is to pre-select the piece which is known to intersect the plane, and then find the roots of $Q_i(t)$ in order to determine the exact point of in-

tersection. This can be accomplished by first evaluating $Q_i(t)$ at the breaks, $t = \tau_j$. Then, a break-point lies below the plane if $Q_i(\tau_j) < 0$; above the plane if $Q_i(\tau_j) > 0$; or on the plane when $Q_i(\tau_j) = 0$, and Equation (6.29) is satisfied. Hence, a polynomial piece will intersect the plane when: (i) its first point is below the plane and its last point is above the plane; or (ii) its first point is above the plane and its last point is below the plane; or (iii) either its first or last point is in the plane. Mathematically, the condition for intersection then becomes:

$$Q(\tau_j)Q(\tau_{j+1}) \leq 0. \tag{6.30}$$

Applying this criterion assumes that the cut-plane is intersected only once by the spline. Even with this constraint, in the special case where a break-point lies on the plane, two polynomial pieces will satisfy the criterion, since break points are shared by consecutive pieces. In this situation, the pre-selection algorithm is terminated as soon as the criterion has been satisfied for one piece. Pre-selecting the intersecting piece of a spline in this way leads to a rapid implementation of the planar reslicing method.

A special case of planar reslicing uses cut-planes which are parallel to, say, the *x-y* plane of the reconstruction co-ordinate system. Because of the intentional orientation of the vessel surface within the reconstruction volume, with the longitudinal axis aligned roughly with the *z*-axis, parallel reslicing is carried out in the *z*-direction, with cut-planes defined by the equation $z = n.\delta$, where n is an integer and δ is a suitable interslice separation. Parallel reslicing is useful because rendering and quantification operations are simplified significantly when a series of parallel cross-sectional contours is available. Also, because virtually all other 3D imaging techniques acquire parallel slices, many algorithms are available to perform these operations.

6.6 Determination of the vessel central axis

The central axis for a segment of artery may be defined informally as the 3D curve which passes through the centroids of cross-sections along the length of the vessel. However, the geometry of the central axis using this definition will be sensitive to the orientation of the cross-sections in question. Another potential problem occurs

around bifurcations. Figure 6.6(c) depicts three central axis curves for the carotid arteries. The presence of a bifurcation causes a discontinuity between the end of the CCA central axis and the start of the ICA central axis. This arises because of the change in cross-sectional shape of the artery proximal to the bifurcation; in the region near to the bifurcation, the centroid ceases to become very helpful. This situation does not present too much of a problem when calculating 'true' transverse cut-planes, since the three carotid arteries can be considered separately. The discontinuity between the CCA and ICA axes is still not completely satisfactory, however, because a true transverse plane here may intersect all three carotid arteries at once. For this reason, attention is confined to regions where a single vessel is intersected.

It is desirable to have a continuous frame of reference, which, for the purposes of longitudinal measurements, spans the whole length of the reconstructed vessel. Furthermore, since carotid plaques frequently extend from the CCA into the ICA or ECA, it seems appropriate to construct a continuous central axis from the CCA to the either of the daughter vessels. One solution, which has been adopted by researchers interested in computational blood flow modelling, is to fit a smooth curve between the centroids of circular (or near circular) cross-sections of the distal CCA and the centroids of cross-sections of the of the proximal ICA/ECA. In this way, a continuous central axis is generated. The difficulty with this approach lies in the classification of 'circularity'. To minimise the variability inherent in this process, a dimensionless shape factor was chosen to provide an objective measure of circularity: for a cross-sectional contour with area A and circumference C, the shape factor, F, is defined as:

$$F = \frac{4\pi A}{C^2}. \tag{6.31}$$

When the contour is circular, $F = 1$, otherwise, $F > 1$.

The method developed to determine the central axis curves for 3D reconstructed vessels is a 3D extension of the method described by Smith et al. [1996], and is illustrated in Figure 6.14. The algorithm iteratively computes the local direction of the

(*continued on page 212*)

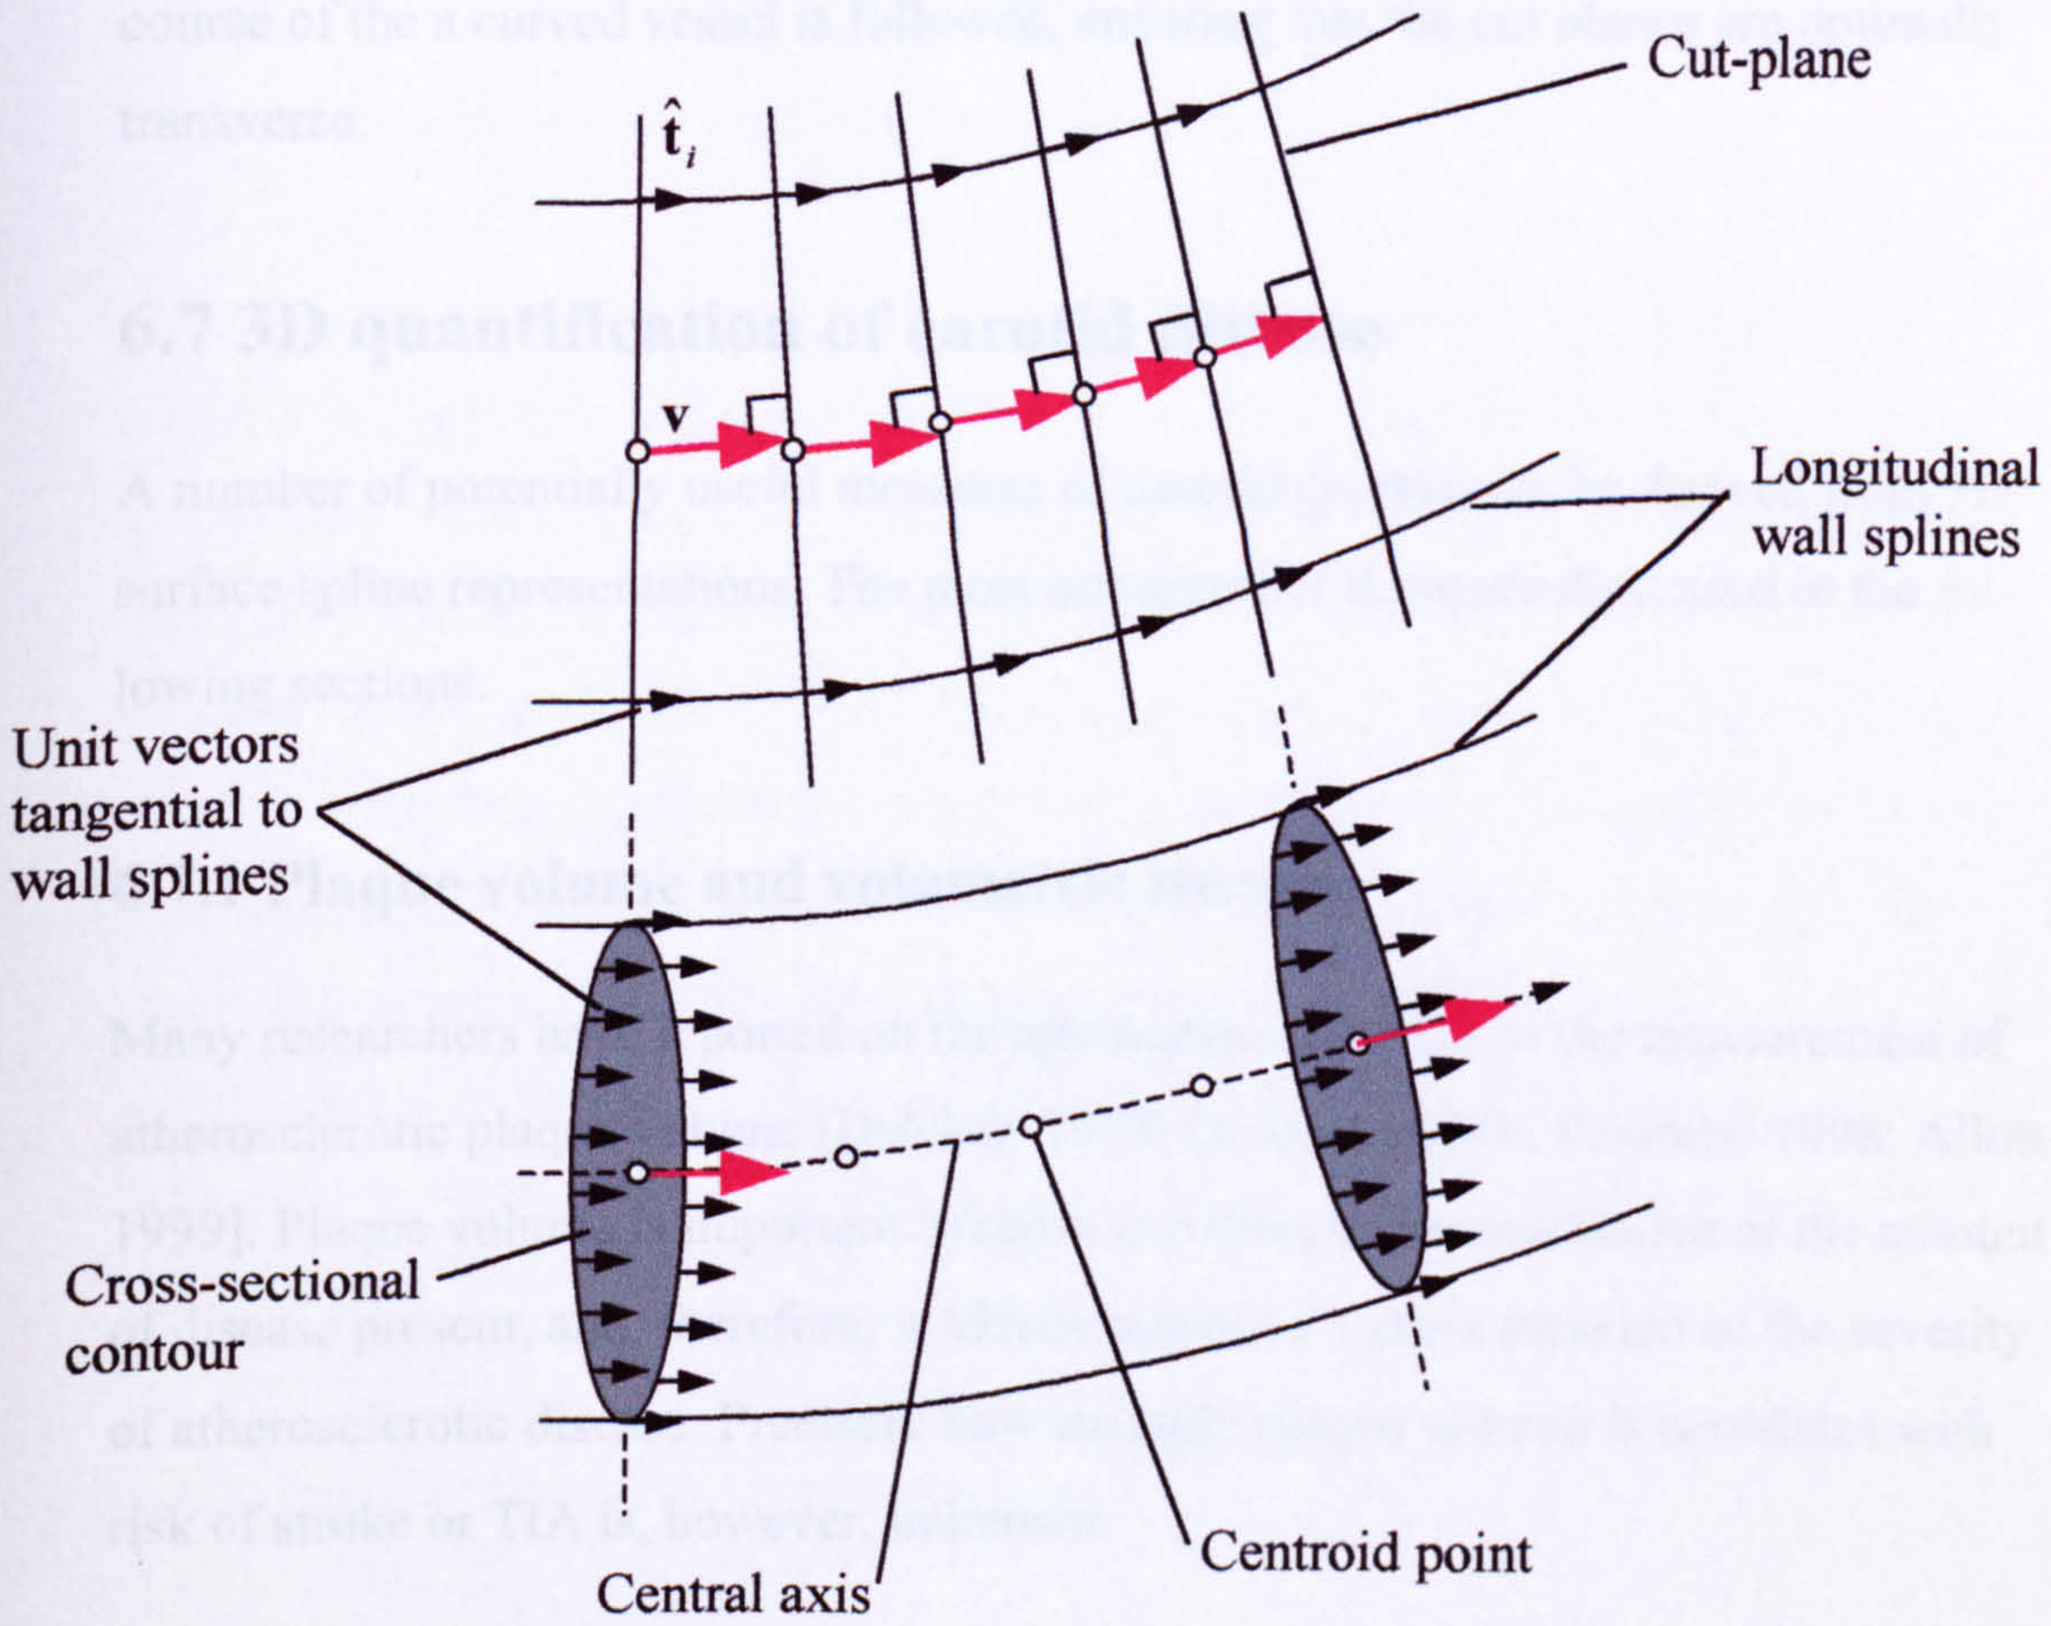


Figure 6.14 Determination of the vessel central axis from a 3D surface reconstruction using the longitudinal spline representation. A vector **v** is generated for each slice which has the same direction as the average of the unit tangential vectors ($\mathbf{t}_i$) evaluated around the cross-section contour. Subsequent cut-planes are then defined which are perpendicular to **v** for the previous slice as shown. The central axis is defined as the smooth curve which passes through all centroid points.

vessel based on the vector, **v**, obtained by averaging the unit tangent vectors at points on the wall spline intersected by the current cut-plane. The resultant vector is used to choose the cut-plane for the next iteration, as shown in Figure 6.14. An initial cut-plane is chosen which is approximately perpendicular to the vessel walls, and successive cut-planes are perpendicular to **v**. This algorithm has the advantage that the course of the a curved vessel is followed, ensuring that the cut planes are optimally transverse.

6.7 3D quantification of carotid disease

A number of potentially useful measures of carotid disease can be derived from 3D surface spline representations. The most important of these are discussed in the following sections.

6.7.1 Plaque volume and volumetric stenosis

Many researchers have reported on the application of 3DUS to the measurement of atherosclerotic plaque volume [Delcker 1998; Delcker 1994b; Palombo 1998; Allott 1999]. Plaque volume is important because it is directly representative of the amount of disease present, and, therefore, is ideally suited as a gross measure of the severity of atherosclerotic disease. Precisely how strongly plaque volume is correlates with risk of stroke or TIA is, however, unknown.

Plaque volume conveys no information on the distribution of disease within a vessel, and it is feasible that two arteries may contain plaque of identical volume but have very different patterns of disease. A further limitation of plaque volume is that, alone, it has not been normalised with respect to vessel size. This is important because carotid artery diameters vary significantly between individuals, particularly when vascular remodelling has taken place in response to atherosclerosis. Therefore, in order to account for inter-patient variations in vessel size, the volumetric stenosis index, proposed in Chapter 3, may be more useful.

A straight-forward, planimetric method was used to calculated volume from parallel cross-sectional contours, obtained by reslicing the 3D surface reconstruction of the carotid bifurcation. In particular, volume is given by the formula:

$$V = \frac{h}{2}\sum_{i=1}^{N-1}(A_i + A_{i+1}), \tag{6.32}$$

where A_i is the area of the i^{th} contour ($1 \le i \le N$), and h is the interslice separation. Another method, described by [Watanabe 1982], permits volume to calculated from non-parallel contours. This method was not adopted because parallel contours could be found rapidly and automatically from the vessel reconstructions.

The value of the volumetric stenosis index is dependent on the length of the reconstructed segment over which it is measured. Extending the length of the segment to include portions of the artery which are relatively unaffected by disease would cause the value of the index to decrease. Furthermore, *if the whole bifurcation is* considered, the geometry of the ECA will also influence the measurement. It seems sensible, therefore, to restrict the segment of measurement to the region where disease is present. This is simple when the disease is focal, but may not be possible when disease extends all the way along the length of the reconstruction. In this case, it is important to standardise serial measurements by quantifying the same segment from each reconstruction. This may be achieved by using the central axis as a frame of reference.

6.7.2 True transverse measurements

As explained in Chapter 3, planar stenosis indices are a well-established means of quantifying carotid disease. The freehand nature of conventional, 2D ultrasound means that the 'optimal' planes within which measurements are made are chosen subjectively, based on the experience of the ultrasonographer. This procedure is subject to significant error. One of the advantages of 3D reconstruction is that 'true' transverse planes can be found automatically, thus increasing the accuracy and precision of planar measurements.

A convenient definition for a true transverse plane is the plane which is perpendicular to the vessel central axis. In principle, therefore, a transverse plane is derived for all positions on the central axis. As discussed in Section 6.6.3, two central axis points are available immediately distal to the bifurcation of the central axis curve, and the transverse plane for both is likely to be of little use. Elsewhere, within the

CCA, ICA, or ECA, transverse planes are well-defined. Because of its clinical importance, the ICA is of primary interest, and, although the true transverse stenosis indices described below can be calculated all of the carotid vessels, they will only considered in the ICA.

The ability to quickly compute transverse planes and reslice the reconstructed vessel enables very many more measurements to be made than would be possible by hand. Thus, a more comprehensive, objective quantification of the severity of disease becomes feasible, with a greater chance of correctly identifying the true maximum diameter stenosis. A further advantage of multiple true transverse measurements over single measurements is that a spatial profile of disease along the vessel can be constructed. This provides information about the 3D distribution of disease, and may prove a valuable method of further characterising and classifying carotid disease.

6.7.2.1 Diameter-based measures

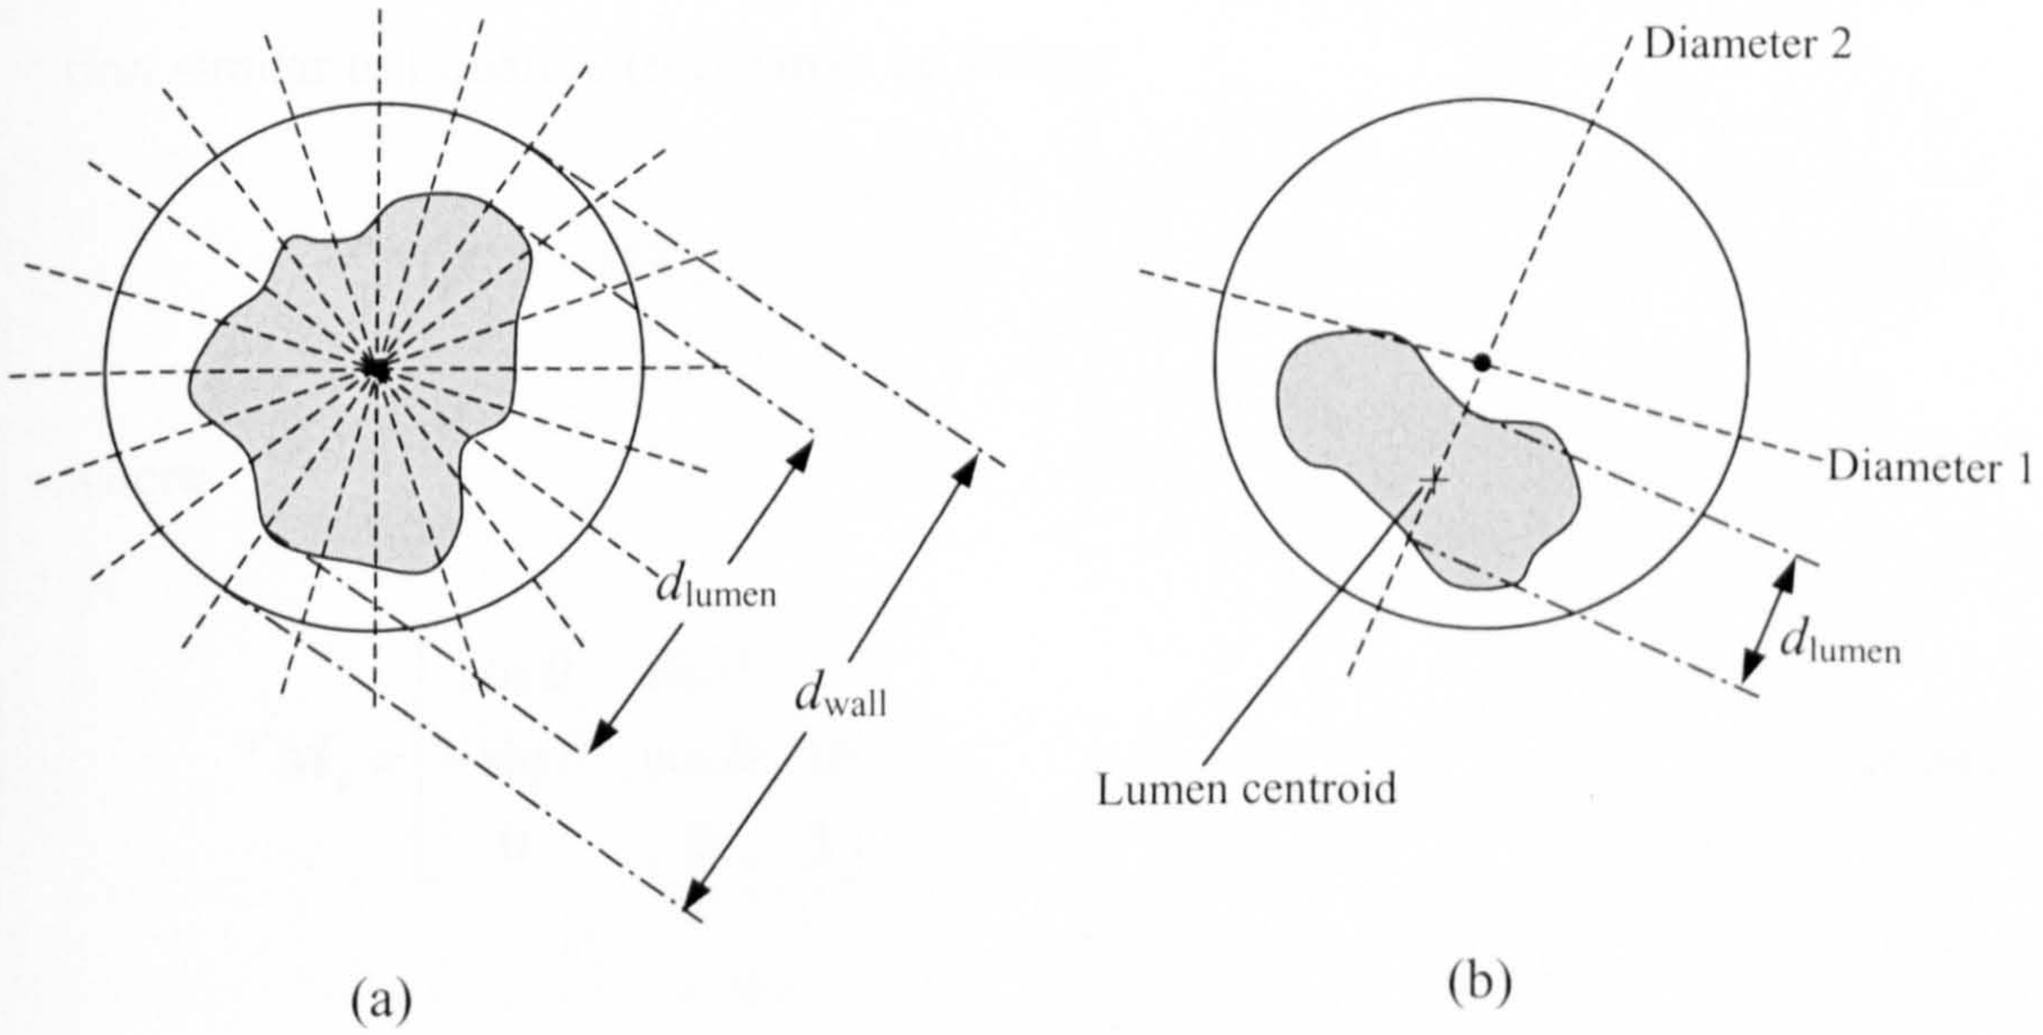


Figure 6.15 Calculation of lumen diameters from a transverse image. (a) The lumen contour (shaded) encloses the centroid of the wall contour. In this case, multiple diameter lines can be drawn which pass through the wall centroid, as indicated. The lumen diameter (d_{lumen}) can then be measured along any of these lines. (b) Where the wall centroid is not enclosed by the lumen contour, randomly oriented diameter lines result in lumen diameter measurements which may be arbitrarily small, as for Diameter 1. Therefore, a diameter line is chosen which passes through the wall and lumen centroids (Diameter 2).

After reslicing a vessel using a true transverse plane, diameter lines can be chosen at any angle, with each passing through the centroid of the wall contour. Figure 6.15 illustrates the method used to calculate lumenal diameter measurements. If the wall centroid is enclosed by the lumen contour (see Figure 6.15(a)) then multiple diameter lines are possible which pass through the wall centroid, and hence the centre of the vessel. If, however, the wall centroid lies outside the lumen, as shown in Figure 6.15(b), selecting an arbitrary diameter line will not necessarily lead to a sensible diameter measurements. Indeed, it is possible that some lumen diameters could be very close to zero if the diameter line intersects the lumen contour near the edge. This is illustrated in Figure 6.15(b) and will result in artificially high stenosis indices. In order to avoid this, the diameter line which passes through the both wall and lumen centroids was chosen, and diameter measurements made along this single line.

It is assumed that contours are intersected by a diameter line exactly twice. When both contours are represented by a spline, intersections are found using a very similar method as finding the intersection of a plane with a spline. For the 2D case, an equation similar to Equation (6.28) may be written:

$$f_{I'} = {}^{I'}\mathbf{M}_{I}\mathbf{G}\mathbf{T} = \mathbf{H}\mathbf{T}, \tag{6.33}$$

where

$${}^{I'}\mathbf{M}_{I} = \begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{6.34}$$

and I and I' are the 2D co-ordinate systems shown in Figure 6.16. Examining the x-component in Equation (6.34) yields a polynomial in t for the i^{th} polynomial piece of the spline, which is zero at the intersections with a diameter line. Thus:

$$\left(\mathbf{H}_{i}\mathbf{T}\right)_{x} = Q_{i}(t) = 0. \tag{6.35}$$

Like the algorithm outlined in Section 6.5, an initial search for the intersecting polynomial pieces of the lumen and wall splines reduces the computation time considerably.

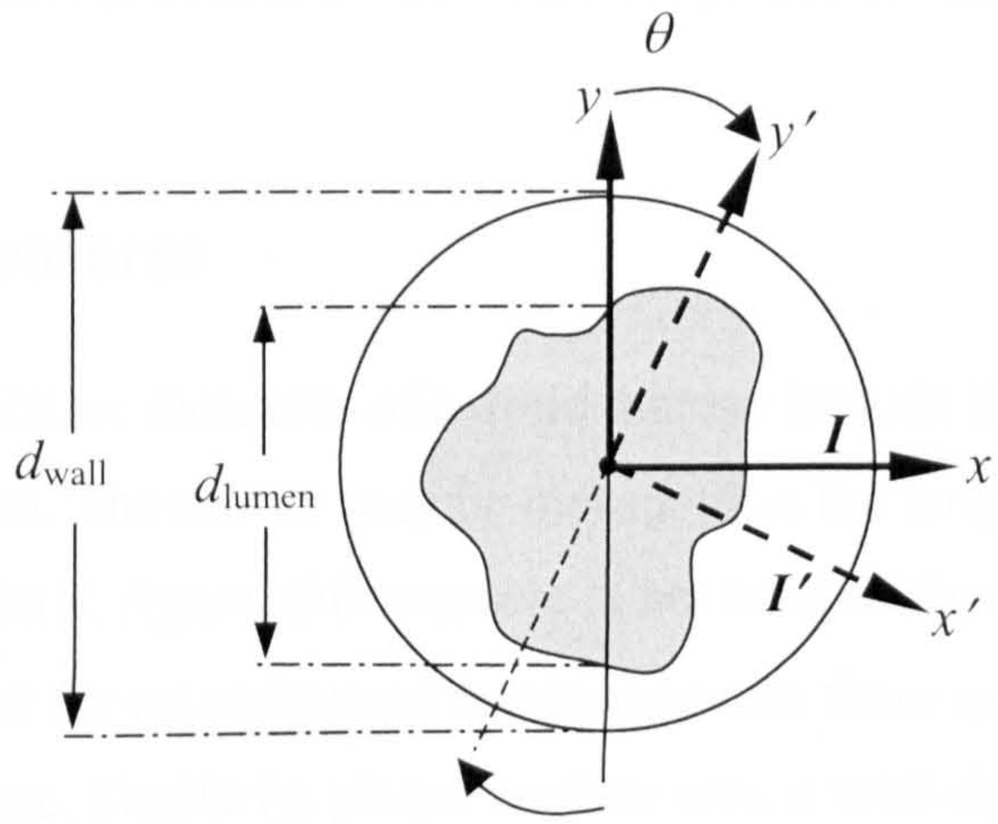


Figure 6.16 Co-ordinate system, ***I′***, aligned with a diameter line oriented at angle θ to the fixed co-ordinate system, ***I***. The wall and lumen diameters – d_{wall} and d_{lumen}, respectively – are calculated by the intersections with the wall and lumen contours.

Because intersections can be determined efficiently using the method above, it is a simple matter to compute diameter measurements for a large number of true transverse diameter lines oriented between 0 and 180° for lumen contours which enclose the wall centroid. The maximum diameter stenosis is then given using the formula:

$$S_{\max} = \max_{l,\theta}\left[100\times\left(1-\frac{d_{\text{lumen}}(l,\theta)}{d_{\text{wall}}(l.\theta)}\right)\right], \tag{6.36}$$

where the parameter l is a distance which indicates the position of the measurement along the central axis. It is also possible to investigate, for instance, the maximum or mean diameter stenosis as a function of l.

6.7.2.2 Area-based measures

True transverse area measurements can be easily derived. The areas of the lumen and vessel wall – A_{lumen} and A_{wall}, respectively – are estimated using the formula given in Equation (6.9). The degree of area stenosis index (as a function of l) is then given by:

$$S(l) = 100 \times \left(1 - \frac{A_{\text{lumen}}(l)}{A_{\text{wall}}(l)} \right). \tag{6.37}$$

Plotting $S(l)$ against l produces an area stenosis profile for the vessel.

6.7.3 Other measures

Additional quantitative measures of carotid disease include the extent of disease and plaque surface area. The former may be measured as the length of plaque measured along the central axis. Again, this measure is not adequately defined when disease is present throughout the reconstructed section because there is no independent start- or endpoint for disease. Similarly, plaque surface area is well-defined for focal disease, but dependent on the length of the reconstructed vessel otherwise.

Surface area is determined from N parallel cross-sectional contours in the same way as volume, using the formula:

$$A = \frac{h}{2} \sum_{i=1}^{N-1} (C_i + C_{i+1}), \tag{6.38}$$

where C_i is the total arc-length of the i^{th} contour, $1 \le i \le N$, and h is the interslice separation. Since relatively small changes in surface shape can lead to large changes in surface area, the surface smoothing inherent in both the segmentation and reconstruction processes are likely to significantly influence the accuracy of plaque surface area measurements. In particular, it is usually impossible to discern small-scale irregularities in plaque surface with any certainty using ultrasound, and this undermines confidence in the usefulness of surface area as a reliable measure. A gross measure of the area of a (smoothed) plaque or lumen surface might be useful, however, and, together with volumetric information, may possibly provide some insight into the mechanical aspects of plaque stability.

Chapter 7

Studies on phantom and patient carotid arteries

Chapter 7 focuses on the validation and clinical application of the techniques described in the previous chapter. This includes a validation of the accuracy of reconstruction and quantification of the carotid bifurcation using anatomically realistic ultrasound phantoms, as well as an assessment of the clinical feasibility of the 3D quantification techniques developed in this thesis.

7.1 Reconstruction accuracy

7.1.1 Anthropomorphic carotid phantoms

Three anthropomorphic ultrasound phantoms (R.G. Shelley Ltd, Ontario, Canada) were used to validate the accuracy of 3D reconstruction and quantification of the carotid bifurcation using the 3DUS acquisition system and algorithms described in earlier chapters. The phantoms each contain one anatomically realistic carotid bifurcation constructed from moulded silicon tubing and encased in agar gel. The geometry of all phantom vessels was fabricated to an exact specification using a low-melting-point metal to cast the lumen [Smith 2001]. This technique enables tight stenoses to be constructed which would otherwise be impossible to fabricate without sectioning the phantom to remove the cast. Each vessel geometry is specified by a mathematical model from which a cast was made using a computer-controlled milling machine. The phantoms can be used either as a static phantom, or as a flow phantom by pumping a blood mimicking fluid through the bifurcation They are, therefore, suitable for both ultrasound and MRA imaging.

Figure 7.1 shows the lumen geometries for the three carotid phantoms. One normal geometry and two stenosed geometries were used. All geometries are specified by two overlapping tubes with varying circular cross-sectional diameter as specified by Smith et al. [1996]. Since the geometries are circular in cross-section, with the central axes lying in a single plane, they are somewhat idealised. Nevertheless, vessel diameters are based on averaged dimensions determined from x-ray angiograms of patients with and without atherosclerotic disease. In this sense, they may be considered to be anthropomorphic. In the case of the diseased geometries (Figures 7.1(b) and 7.1(c)), the degree of stenosis was determined using the NASCET stenosis index (see Chapter 3). The two stenosed phantoms that were used in the experiments outlined in the next section had *eccentric* stenoses of 30% and 70%.

7.1.2 Carotid phantom experiments

The three phantoms were placed one at a time on the examination couch within the centre of the region appropriate for carotid scanning, and with an anatomical

(*continued on page 221*)

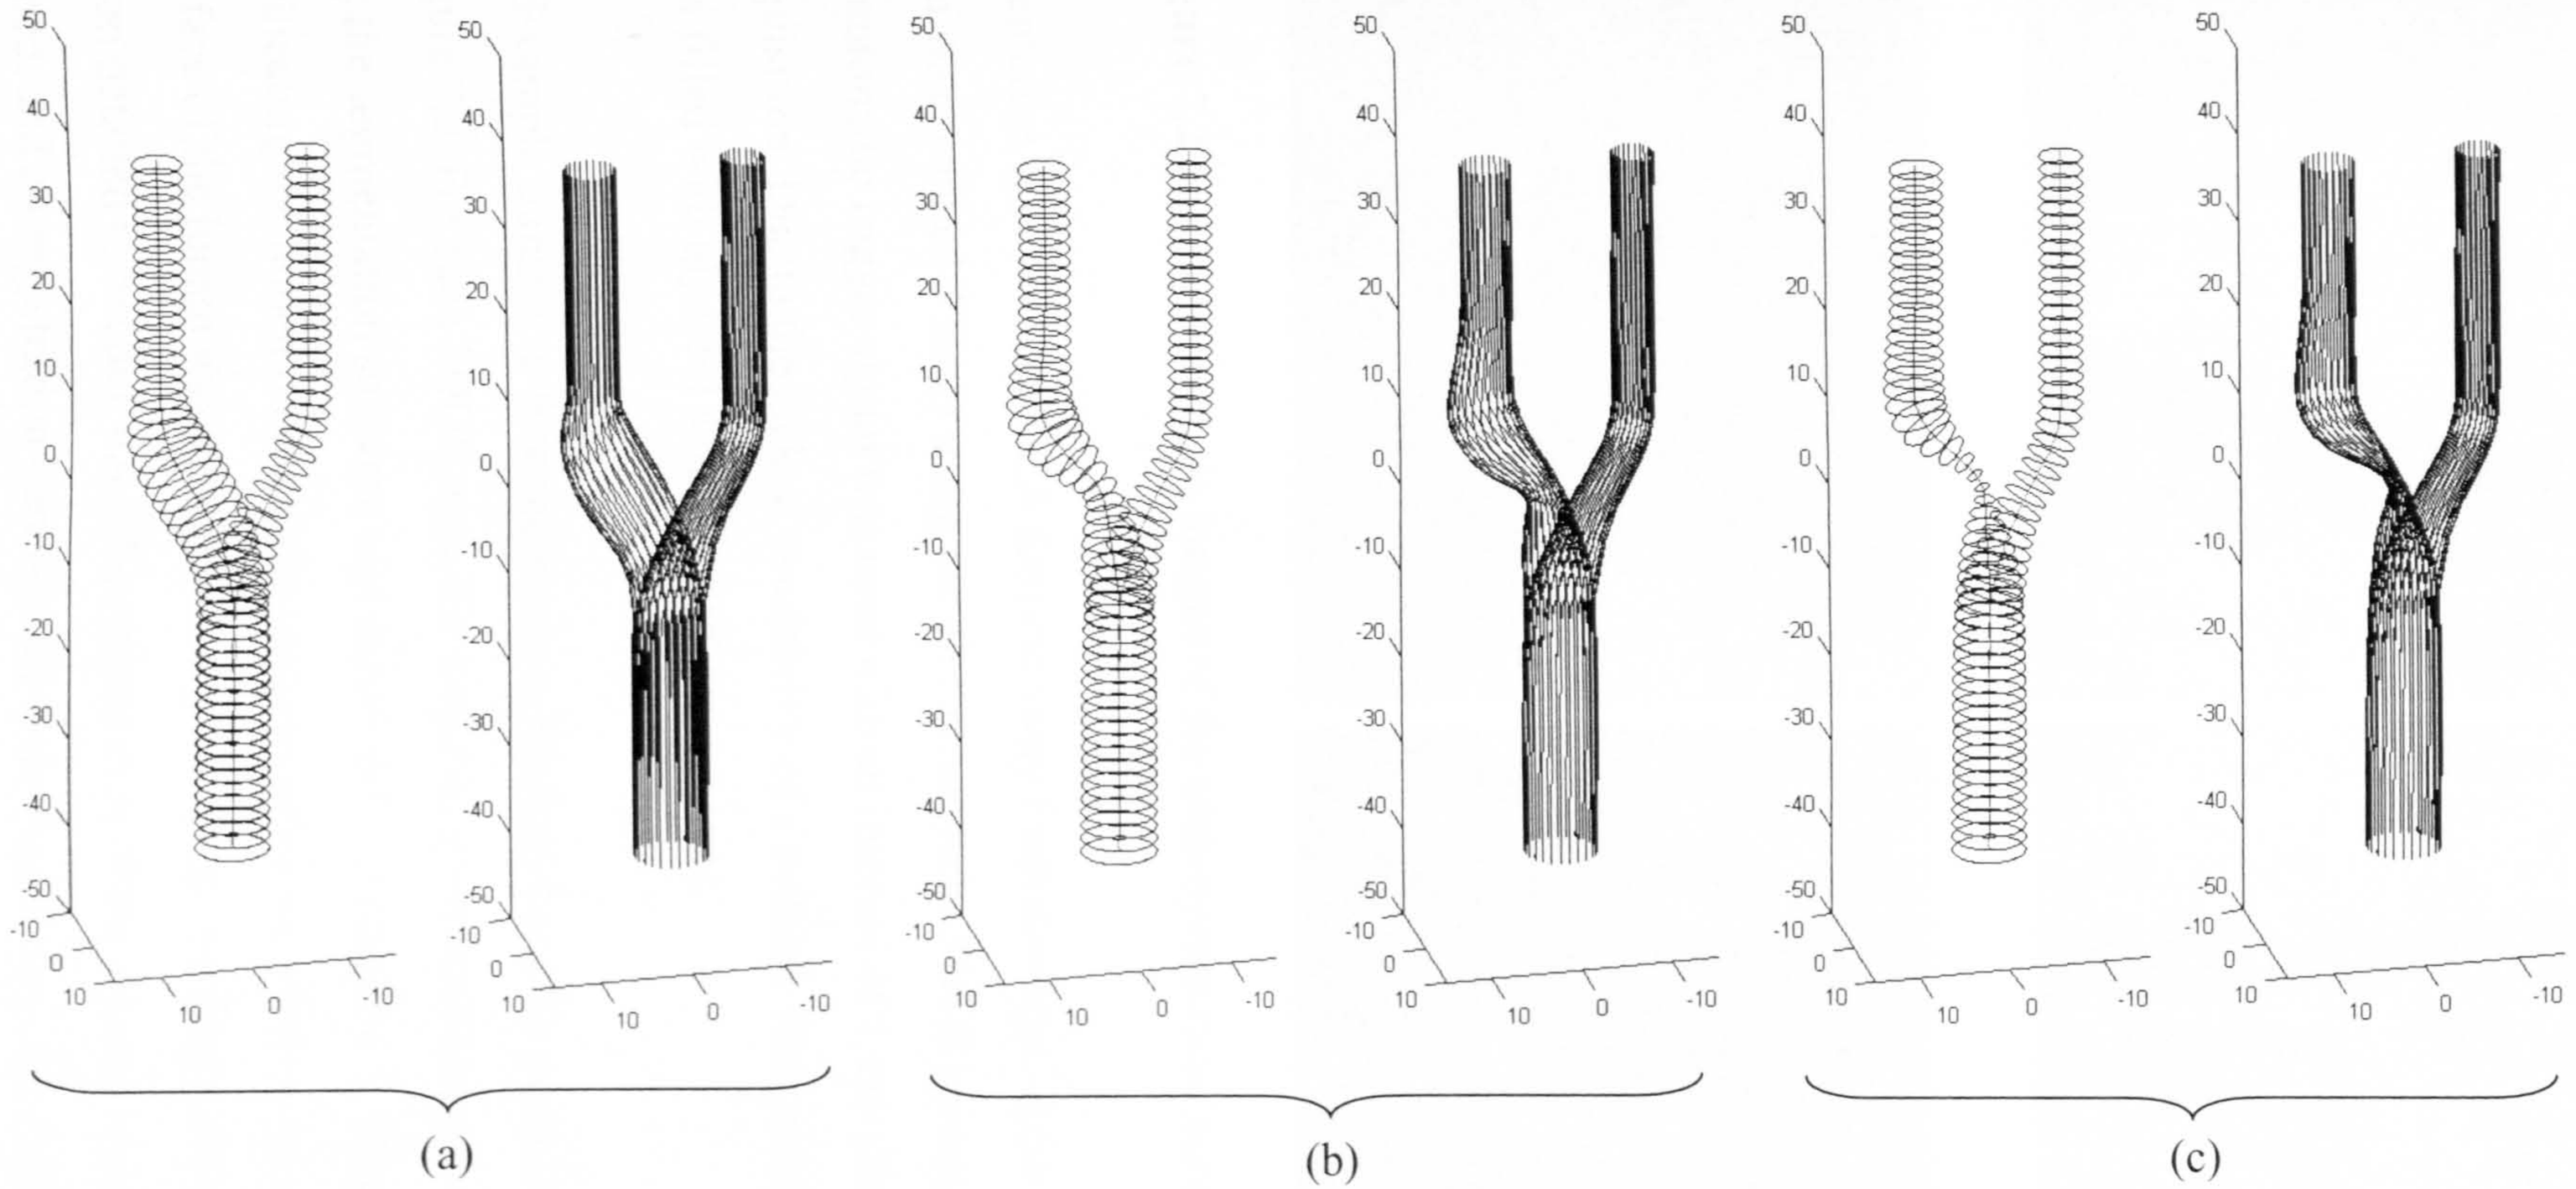


Figure 7.1 Lumen geometries of 3 anthropomorphic carotid ultrasound phantoms. (a) normal; (b) 30% eccentric stenosis; (c) 70% eccentric stenosis. For each geometry, perpendicular cross-sectional contours and fitted longitudinal splines are shown (on the left and right, respectively).

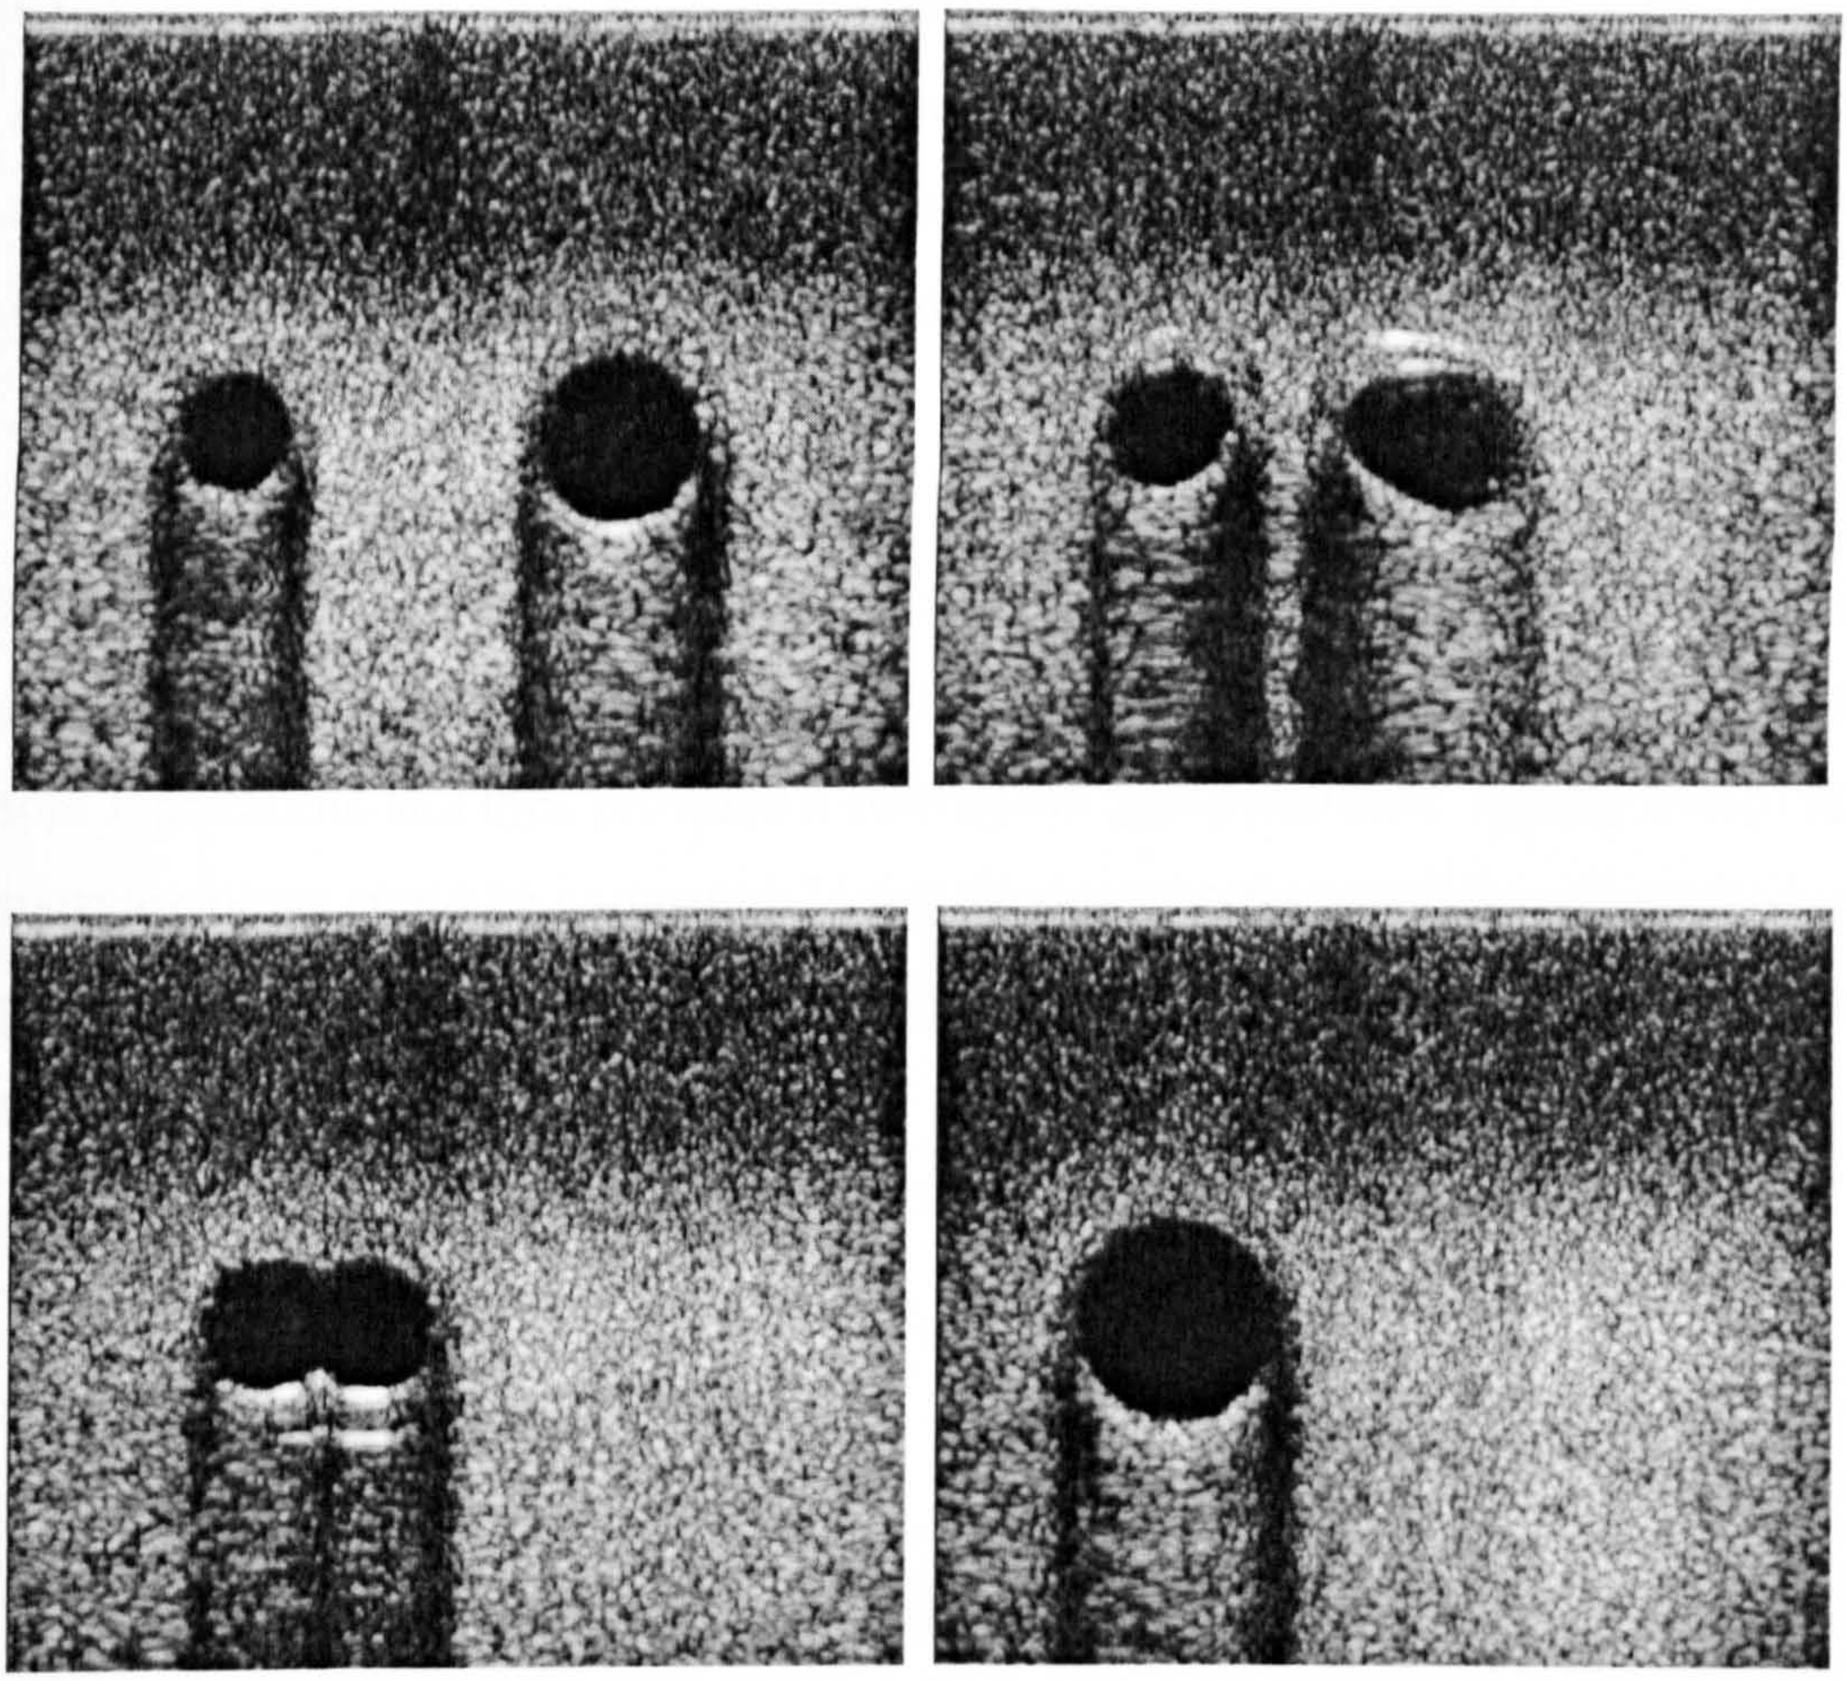


Figure 7.2 Example B-mode images of the anthropomorphic carotid phantoms.

orientation (i.e. with the CCA furthest away from the head of the couch). Each phantom was scanned 3 times using the electromagnetic 3DUS system. For convenience, an automatic trigger signal was used with an interval of 1 s to simulate ECG-gated acquisition. Due to the lack of availability of a suitable flow pump, each phantom was filled with glycerol solution (9% by volume).

Example transverse B-mode images obtained from the phantoms are shown in Figure 7.2. For each 3DUS scan, the acquired transverse images were processed using the segmentation software tool described in Appendix 3. This involved manually delineating the lumen on each slice to generate a set of cross-sectional contours. The surface of the lumen was then reconstructed using the longitudinal spline fitting algorithm detailed in Section 6.4. The central axes of the carotid vessels were determined, and co-registration between the 3 reconstructions for a single phantom was achieved by transforming the longitudinal splines to a new co-ordinate system where the z-axis was the best fit (using a least-squares criterion) to the central axis of the

CCA, and the bifurcation point of the central axes was coincident with the origin. Each reconstructed bifurcation was resliced in the z-direction to yield a series of parallel cross-sectional contours. The ICA was then also resliced perpendicularly to the central axis in order to compute true transverse measurements.

The above procedure was carried out for all of the phantom 3D scans in addition to 3 reference reconstructions derived from geometrical data provided by the phantom manufacturer. These data were used to specify the lumen geometry during phantom fabrication, and thus provided a convenient 'gold standard' against which 3DUS measurements could be compared. Using the quantitative algorithms described in the previous chapter, the following lumen measurements were derived for each phantom reconstruction:

- cross-sectional area as a function of length along the z-axis;
- cross-sectional circumference as a function of length along the z-axis;
- lumen volume;
- lumen surface area;
- distance between the centroids of the ECA and ICA contours as a function of length along the z-axis;
- volumetric stenosis index;
- true transverse lumen area as a function of length along the ICA central axis;
- true transverse mean lumen diameter as a function of length along the ICA central axis;
- minimum true transverse lumen area in the ICA;
- minimum true transverse lumen diameter in the ICA;
- true transverse area stenosis index at the point of minimum ICA lumen area;
- true transverse diameter stenosis index at the point of minimum ICA lumen diameter;
- NASCET stenosis index at the point of minimum ICA lumen diameter.

All measurements where made along the z-axis, over a distance of ± 30 mm with respect to the apex of the geometric bifurcation (This was determined as the position at which the CCA splits into two separate contours). The interslice separation was 0.5 mm. In the case of minimum vessel diameter, two different measurements were calculated: the minimum of the minimum within-slice diameter, and the minimum of

the mean within-slice diameter. The minimum and mean within-slice diameters were calculated from 50 diameter measurements per slice.

All stenosis indices were calculated using the following formula:

$$S = 100 \times \left(1 - \frac{u}{(u_{\text{Normal}})_{\text{REF}}} \right), \tag{7.1}$$

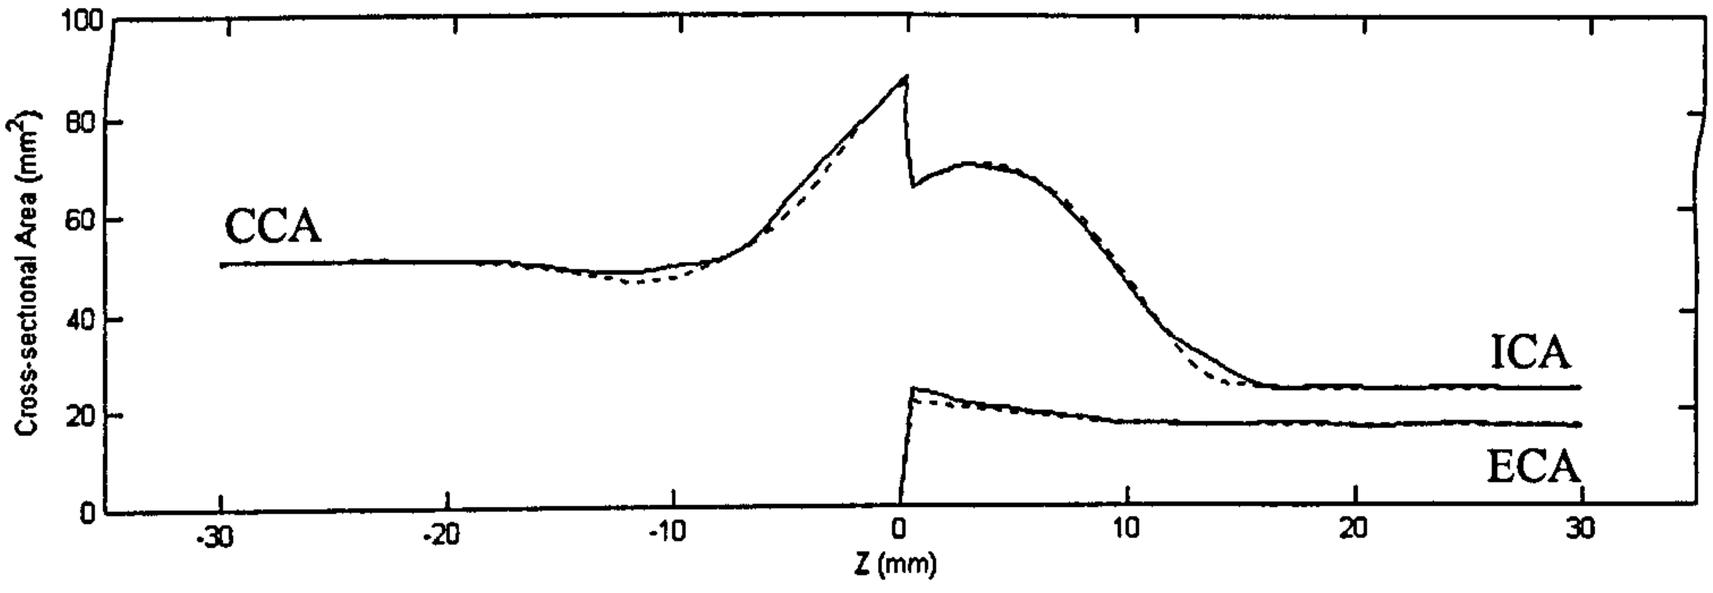
where u is an area, diameter, or volume measurement, and $(u_{\text{Normal}})_{\text{REF}}$ is the corresponding reference measurement determined from the true normal geometry: for true transverse indices, $(u_{\text{Normal}})_{\text{REF}}$ is the true area/diameter for the normal geometry determined for the transverse slice which has the minimum true area/diameter for the stenosed geometries. For the NASCET index, $(u_{\text{Normal}})_{\text{REF}}$ is the true distal ICA diameter for the normal geometry (5.552 mm).

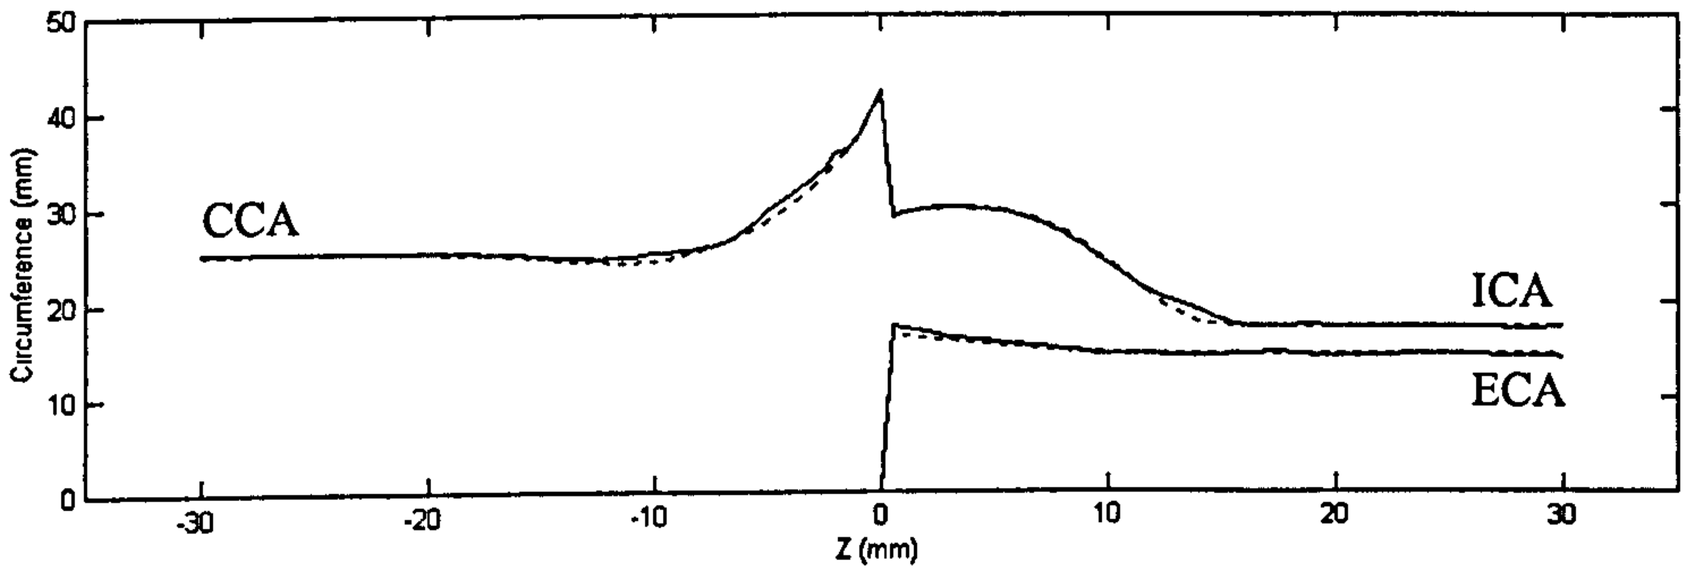
7.1.3 Results

Figure 7.3 shows examples of the cross-sectional area and circumference profiles determined from reconstructions for each of the 3 carotid phantoms. The numerical errors calculated for each of the 3D reconstruction measurements are summarised in Tables 7.1 and 7.2. Overall, errors were small. Errors averaged over all 9 phantom scans were: -0.25 ± 1.28 mm^2 (-0.99 ± 5.59 %) for cross-sectional area; -0.06 ± 0.44 mm (-0.43 ± 2.86 %) for cross-sectional circumference; -22.19 mm^3 (-1.75 %) for lumen volume; -5.71 mm^2 (-0.88 %) for lumen surface area; and 0.26 ± 0.28 mm for the ICA-ECA centroid separation. Although the absolute errors in cross-sectional area and circumference were reasonably consistent for each of the phantom geometries, relative errors were significantly higher for 70 % stenosis due to the smaller lumen cross-section in the proximal portion of the ICA.

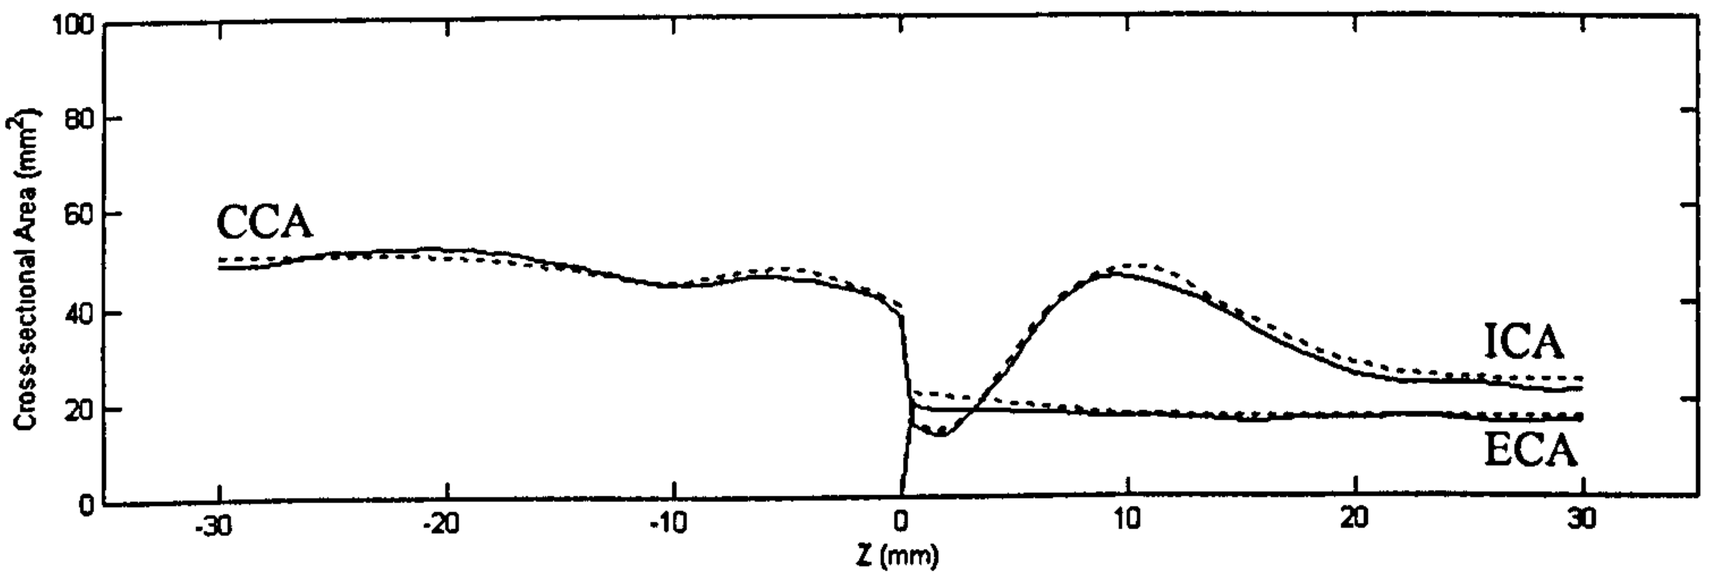
Inspection of Table 7.2 reveals that all of the ICA stenosis indices calculated for the phantom reconstructions overestimated the true stenosis. Indices based on diameter measurements produced larger errors when the minimum within-slice diameter was used as opposed to the mean. Furthermore, for both stenosed phantoms, the

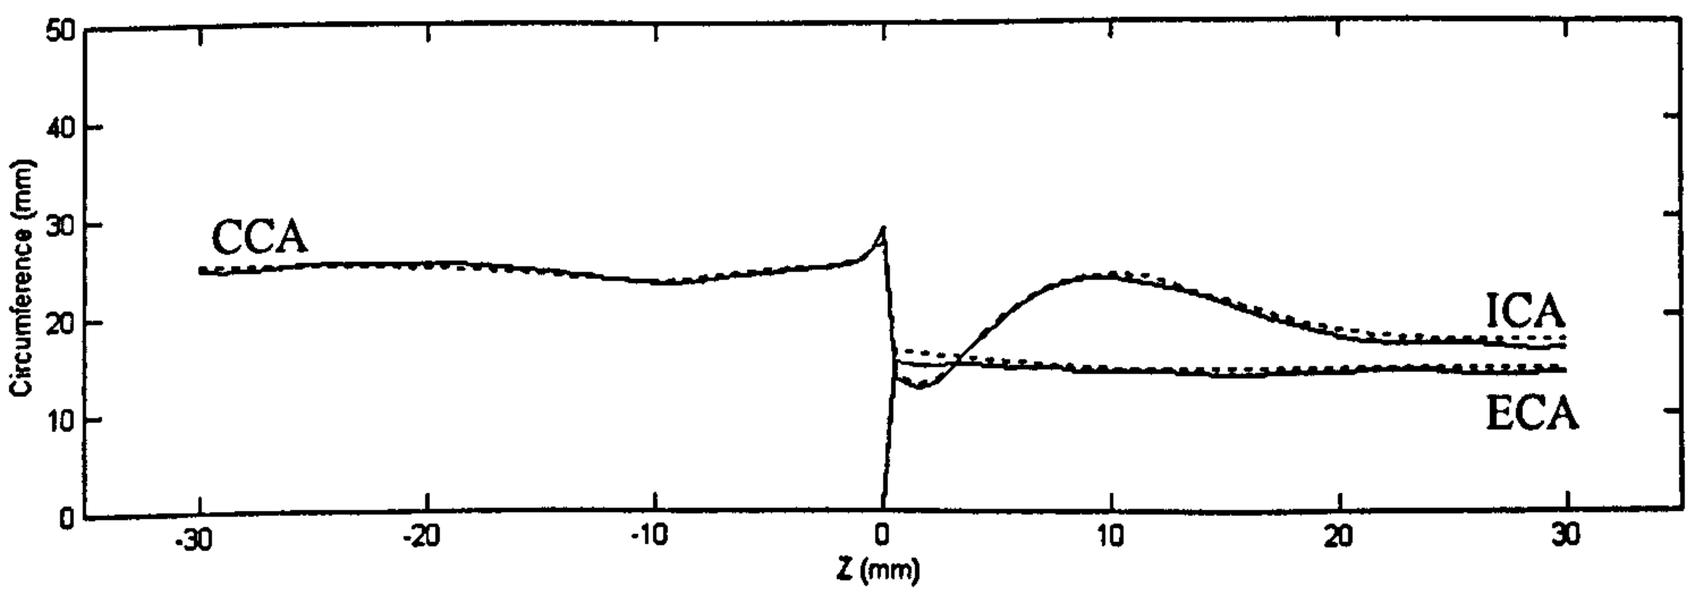
(continued on page 225)





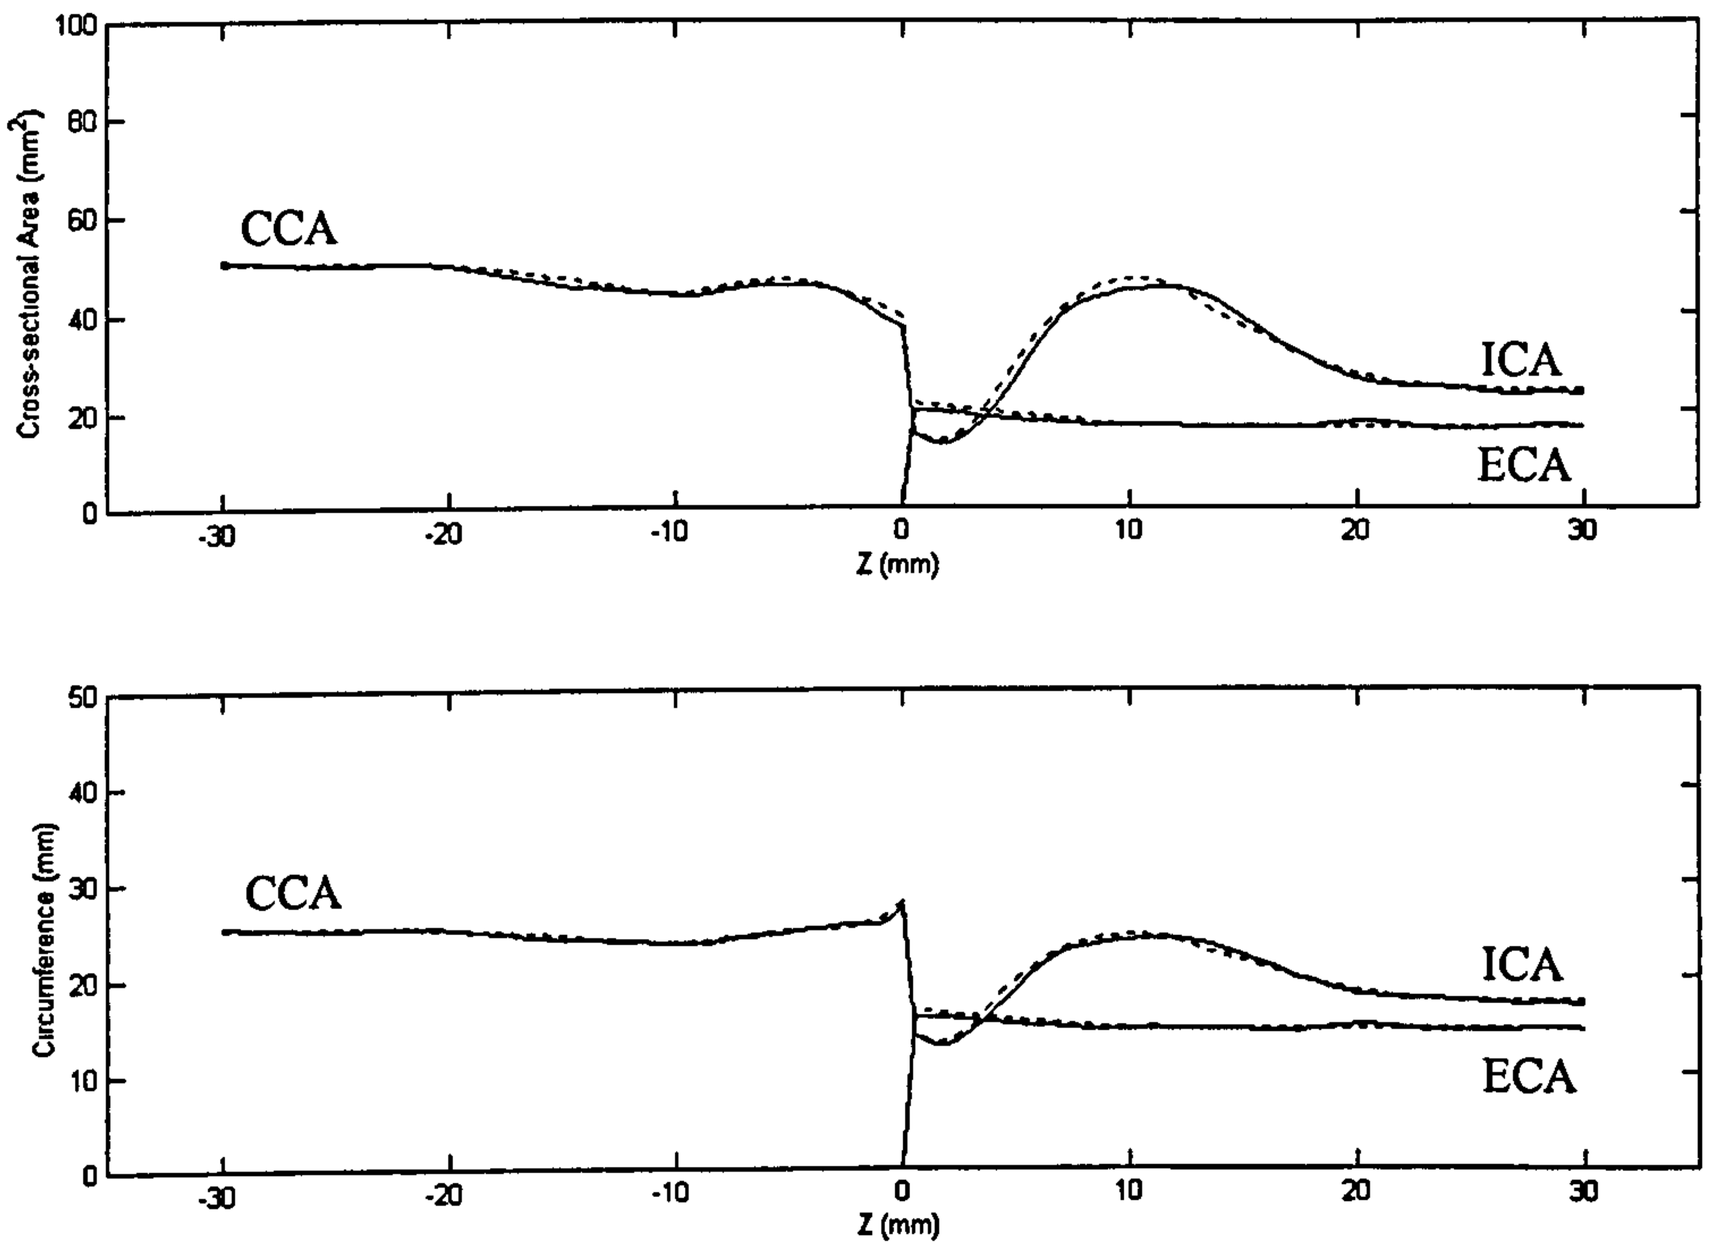
(a)





(b)

Figure 7.3 Sample cross-sectional area and circumference profiles for (a) the normal phantom geometry; (b) 30 % stenosis; and (c) 70 % stenosis (overleaf). The true profile is shown as a dashed line.



(c)

Figure 7.3 (continued)

largest absolute errors occurred for the established NASCET criteria, even though the value of the true NASCET stenosis was numerically lower than the true transverse indices (28 % vs. 55 and 80 %, and 69 % vs. 81 and 96 %, as indicated in Table 7.2). Therefore, it appears that the true transverse diameter stenosis index – which is in fact equivalent to the ECST index – may be more reliable than the NASCET index.

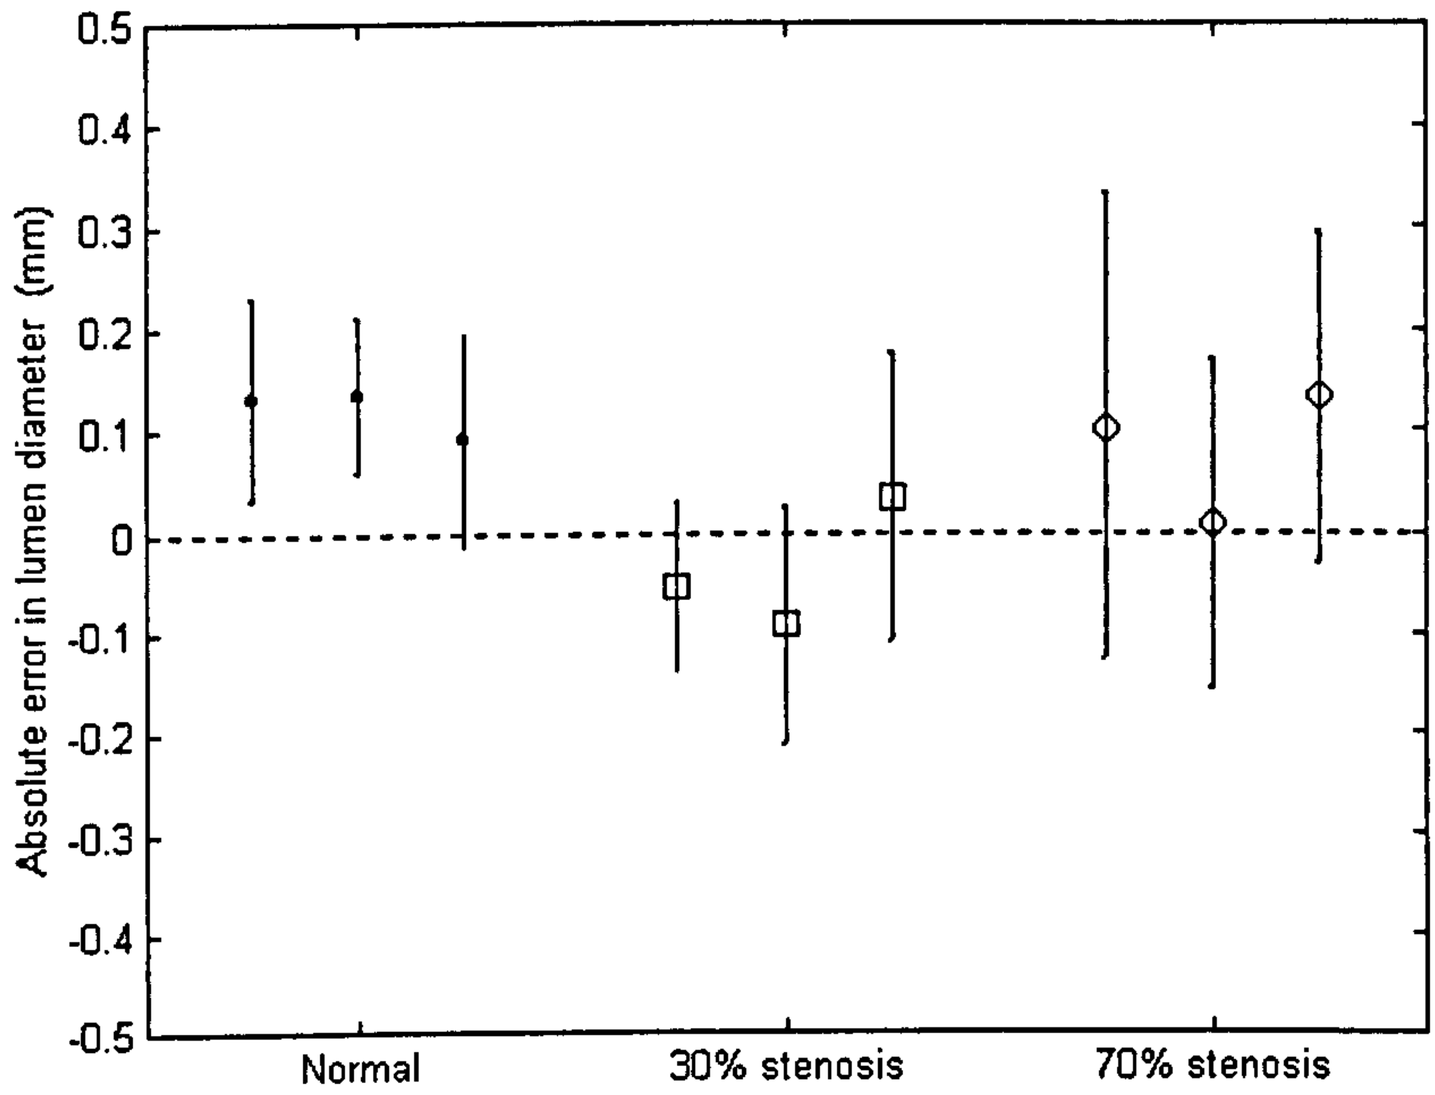
When considering the relative errors for all quantitative measurements, the results were poorest for the 70 % stenosis phantom. This is largely attributed to the very small ICA lumen for this phantom, since there is relatively little difference between the absolute errors in cross-sectional diameter and area measurements, as illustrated in Figure 7.4. It can be seen that diameter measurements were found to be less accurate for the 70% stenosis geometry, however. This is most likely a result of poorer visualisation of the narrow ICA lumen in this phantom, which gives rise to greater uncertainty in the accuracy of the segmented cross-sectional contours.

		Mean ± SD absolute and relative errors								
Phantom	**Scan #**	**Cross-sectional area (mm^2)**		**Cross-sectional circumference (mm)**		**Lumen volume (mm^3)**		**Lumen surface area (mm^2)**		**ICA-ECA centroid separation (mm)**
Normal	**1**	0.33 ± 0.96	(0.94 ± 2.92 %)	0.14 ± 0.32	(0.63 ± 1.53 %)	30.04	(0.90 %)	12.59	(0.67 %)	0.16 ± 0.25
Normal	**2**	-0.13 ± 1.83	(0.12 ± 4.76 %)	-0.01 ± 0.54	(0.08 ± 2.42 %)	-11.81	(-0.35 %)	-1.07	(-0.06 %)	-0.07 ± 0.14
Normal	**3**	-0.57 ± 1.171	(-0.78 ± 2.72 %)	-0.12 ± 0.33	(-0.35 ± 1.44 %)	-50.84	(-1.53 %)	-10.94	(-0.58 %)	0.17 ± 0.21
30% stenosis	**1**	-0.94 ± 1.13	(-3.99 ± 3.90 %)	-0.29 ± 0.39	(-1.77 ± 2.09 %)	-83.81	(-2.92 %)	-25.70	(-1.45 %)	0.45 ± 0.47
30% stenosis	**2**	-1.12 ± 1.46	(-4.41 ± 4.53 %)	-0.36 ± 0.47	(-2.09 ± 2.40 %)	-100.99	(-3.52 %)	-32.72	(-1.85 %)	0.06 ± 0.29
30% stenosis	**3**	-0.53 ± 0.88	(-1.86 ± 3.01 %)	-0.14 ± 0.28	(-0.74 ± 1.51 %)	-47.87	(-1.67 %)	-12.17	(-0.69 %)	0.17 ± 0.41
70% stenosis	**1**	0.70 ± 1.71	(2.43 ± 14.43 %)	0.22 ± 0.76	(1.03 ± 7.21 %)	61.91	(2.42 %)	19.78	(1.20 %)	0.67 ± 0.38
70% stenosis	**2**	-0.05 ± 1.09	(-1.73 ± 7.85 %)	-0.07 ± 0.42	(-0.94 ± 4.17 %)	-5.27	(-0.21 %)	-7.41	(-0.39 %)	0.33 ± 0.19
70% stenosis	**3**	0.59 ± 1.36	(1.69 ± 10.10 %)	0.19 ± 0.58	(0.80 ± 4.94 %)	52.93	(2.07 %)	17.13	(1.04 %)	0.56 ± 0.11

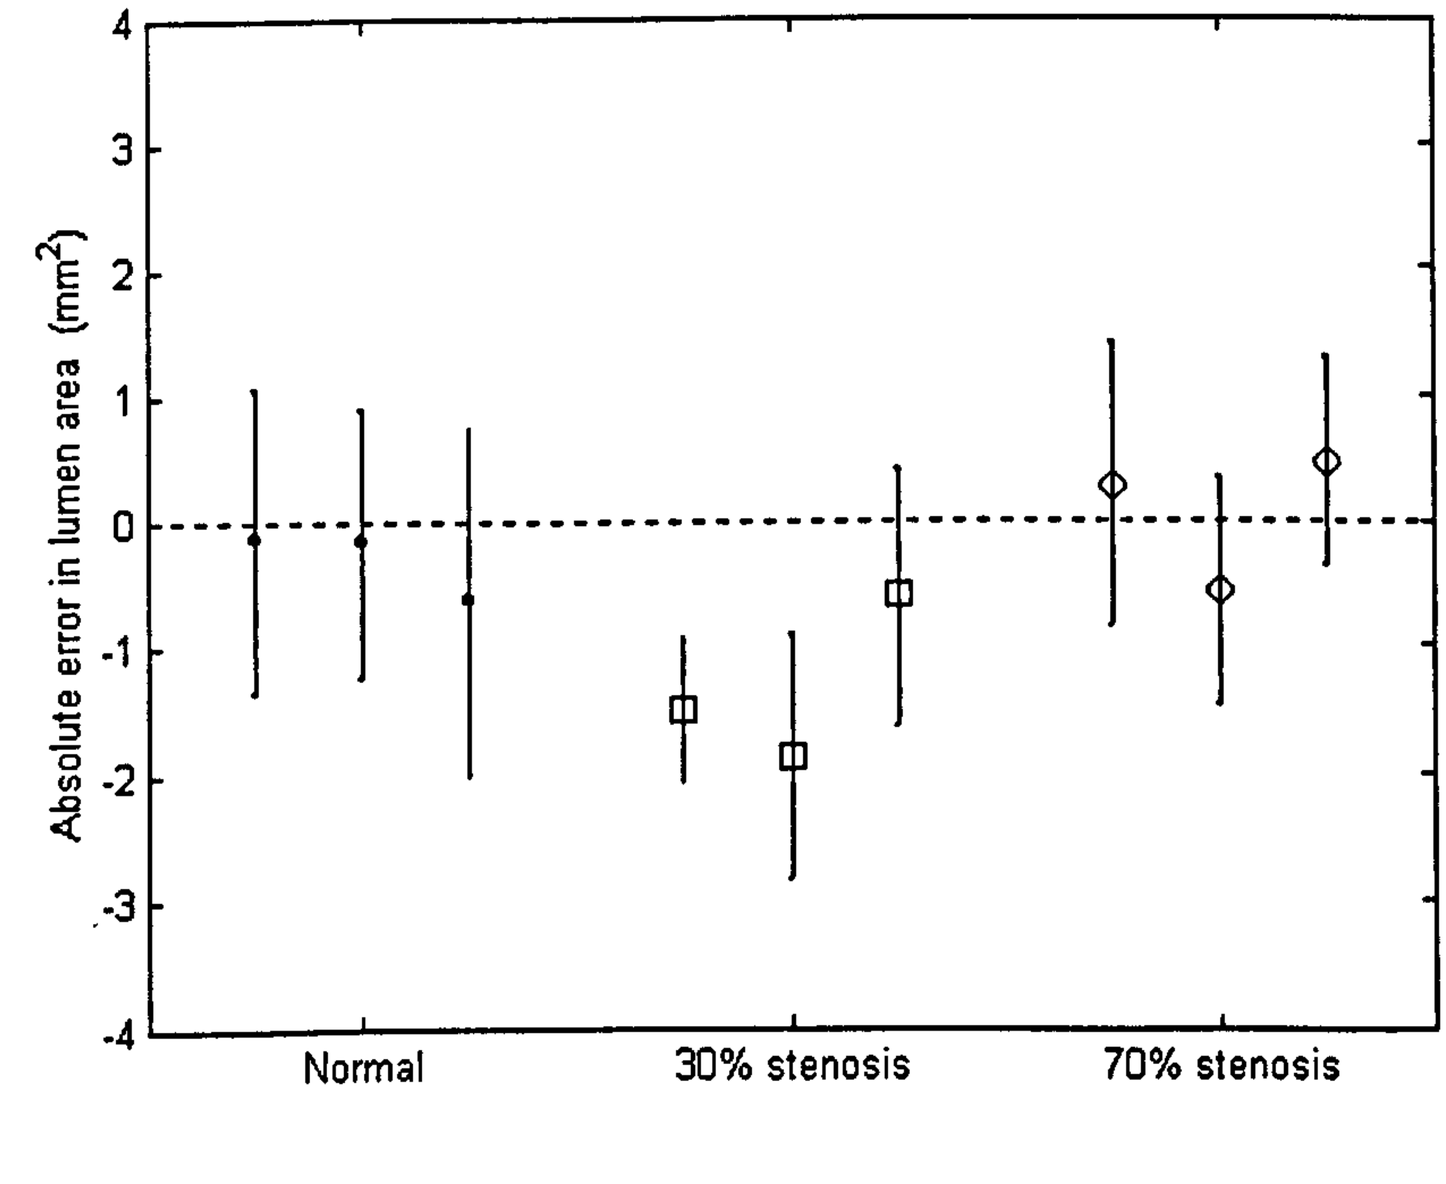
Table 7.1 Errors in phantom 3D reconstruction measurements. Where multiple measurements are made for a single reconstruction, errors are quoted as a mean ± SD. Relative errors, expressed as a percentage, are given in parentheses.

		Absolute and relative errors in stenosis indices (%) .											
Phantom	**Scan #**	**True transverse ICA diameter stenosis index (based on minimum within-slice diameter)**		**True transverse ICA diameter stenosis index (based on mean within-slice diameter)**		**NASCET stenosis index (based on minimum within-slice diameter)**		**NASCET stenosis index (based on mean within-slice diameter)**		**True transverse ICA area stenosis index**		**Volumetric stenosis index**	
30% stenosis		**True value = 54.77**		**True value = 54.77**		**True value = 27.86**		**True value = 27.86**		**True value = 79.55**		**True value = 13.85**	
30% stenosis	**1**	3.66	(6.68 %)	2.52	(4.60 %)	5.83	(20.93 %)	4.02	(14.43 %)	2.96	(3.72 %)	2.52	(18.19 %)
30% stenosis	**2**	2.93	(5.53 %)	2.15	(3.93 %)	4.67	(17.76 %)	3.42	(12.28 %)	2.64	(3.32 %)	3.03	(21.88 %)
30% stenosis	**3**	3.67	(6.70 %)	2.40	(4.38 %)	5.86	(21.03 %)	3.84	(13.78 %)	2.83	(3.56 %)	1.44	(10.40 %)
70% stenosis		**True value = 81.02**		**True value = 81.02**		**True value = 69.45**		**True value = 69.45**		**True value = 97.40**		**True value = 23.25**	
70% stenosis	**1**	8.82	(10.89 %)	7.55	(8.08 %)	14.20	(20.45 %)	10.54	(15.18 %)	2.09	(2.17 %)	-1.86	(-8.00 %)
70% stenosis	**2**	5.89	(7.27 %)	3.69	(4.55 %)	9.48	(13.65 %)	5.95	(8.57 %)	1.32	(1.37 %)	0.16	(0.69 %)
70% stenosis	**3**	4.87	(6.01 %)	2.65	(3.27 %)	7.84	(11.29 %)	4.26	(6.13 %)	0.99	(1.03 %)	-1.59	(-6.84 %)

Table 7.2 Absolute and relative errors in stenosis indices calculated from 3D phantom reconstructions. In each case, the relative error, expressed as a percentage, is given in parentheses.



(a)



(b)

Figure 7.4 Absolute errors in (a) lumen diameter and (b) lumen area for the 3 carotid phantoms, each reconstructed 3 times (the error bars represent Mean ± SD).

7.2 Patient studies

7.2.1 Methods

The clinical feasibility of the 3DUS plaque quantification techniques was assessed by scanning a series of patients with known carotid artery disease. In each session, both the left and right carotid arteries were investigated, but a 3D acquisition was only performed on an artery if initial 2D images revealed evidence of atherosclerotic plaque. In this case, images and Doppler US velocity measurements appropriate to a standard clinical scan were obtained prior to a 3DUS scan. The optimal protocol established in Chapter 5 was used for all 3DUS scans, with ECG-gated power Doppler enabled, and a slow scan-probe sweep in the distal to proximal direction. A single operator (DCB) scanned all patients in the group and was unaware of the results of previous clinical investigations. This remained the case during the subsequent 3D analysis of patient data sets, which was carried out by the same operator over 1 month after the initial scanning sessions.

During the analysis phase, 10 patients were selected randomly from all those who had undergone 3DUS investigations. These were then critically assessed for their suitability for 3D reconstruction and quantification using the methods previously described. This led to the rejection of those patient sets in which the vessel wall and/or atherosclerotic disease could not be visualised adequately. Specifically, the criteria used to evaluate acquired image slices were:

1. Adequate visualisation of the vessel wall and lumen boundaries in diseased portions of the artery.

2. No evidence of significant patient motion during acquisition, such as excessive neck movement, swallowing, or exaggerated respiratory motion.

3. A vessel geometry for which slices cut the vessels more or less transversely along its length (this may not be the case for a very tortuous or widely diverging vessel geometries).

Four sets out of 10 satisfied all of the above criteria. Of those rejected, 4 had inadequate visualisation of the vessel and disease mainly due to acoustic shadowing or echolucent plaque which was insufficiently delineated by colour information in the transverse view. The depth of the carotid bifurcation on one scan also precluded ade-

quate visualisation. One patient set was rejected due to a divergent bifurcation with an angle of bifurcation greater than 90°. A second patient was excluded because of evidence of pronounced respiratory motion.

Wall and lumenal cross-sectional contours were manually drawn on a subset of the acquired slices for each of the four remaining data sets. Contours were not always drawn on all slices in a set because they are typically closely spaced and, for much of the vessel, there is little change in contour geometry between consecutive slices. The vessel walls and lumen surfaces were reconstructed using the longitudinal spline fitting algorithm and surface rendered using a custom-written Matlab program. The central axes were then determined for each vessel reconstruction. Figure 7.5 summarises diagrammatically the steps involved in the 3D analysis. The following quantitative measures were computed:

- Plaque cross-sectional area as a function of length along the z-axis.
- Extent of disease with respect to the z-axis.
- Plaque volume.
- Lumen surface area.
- Volumetric stenosis.
- Maximum true transverse diameter stenosis index as a function of length along the central axis of the ICA.
- True transverse area stenosis index as a function of length along the vessel central axis of the ICA.
- Maximum true transverse diameter stenosis index for the ICA.
- Maximum true transverse area stenosis index for the ICA.

In the list above, 'z-axis' refers to the direction in which parallel reslicing was performed.

7.2.2 Results

The results of 3D reconstruction and quantitative analysis for each of the 4 example diseased carotid bifurcations are presented below. In each case, surface rendered 3D images of the reconstructed vessel wall and lumen are included to aid visualisation of 3D anatomy. Longitudinal power Doppler images are also included for comparison. The numerical results for all patients are collected in Table 7.3.

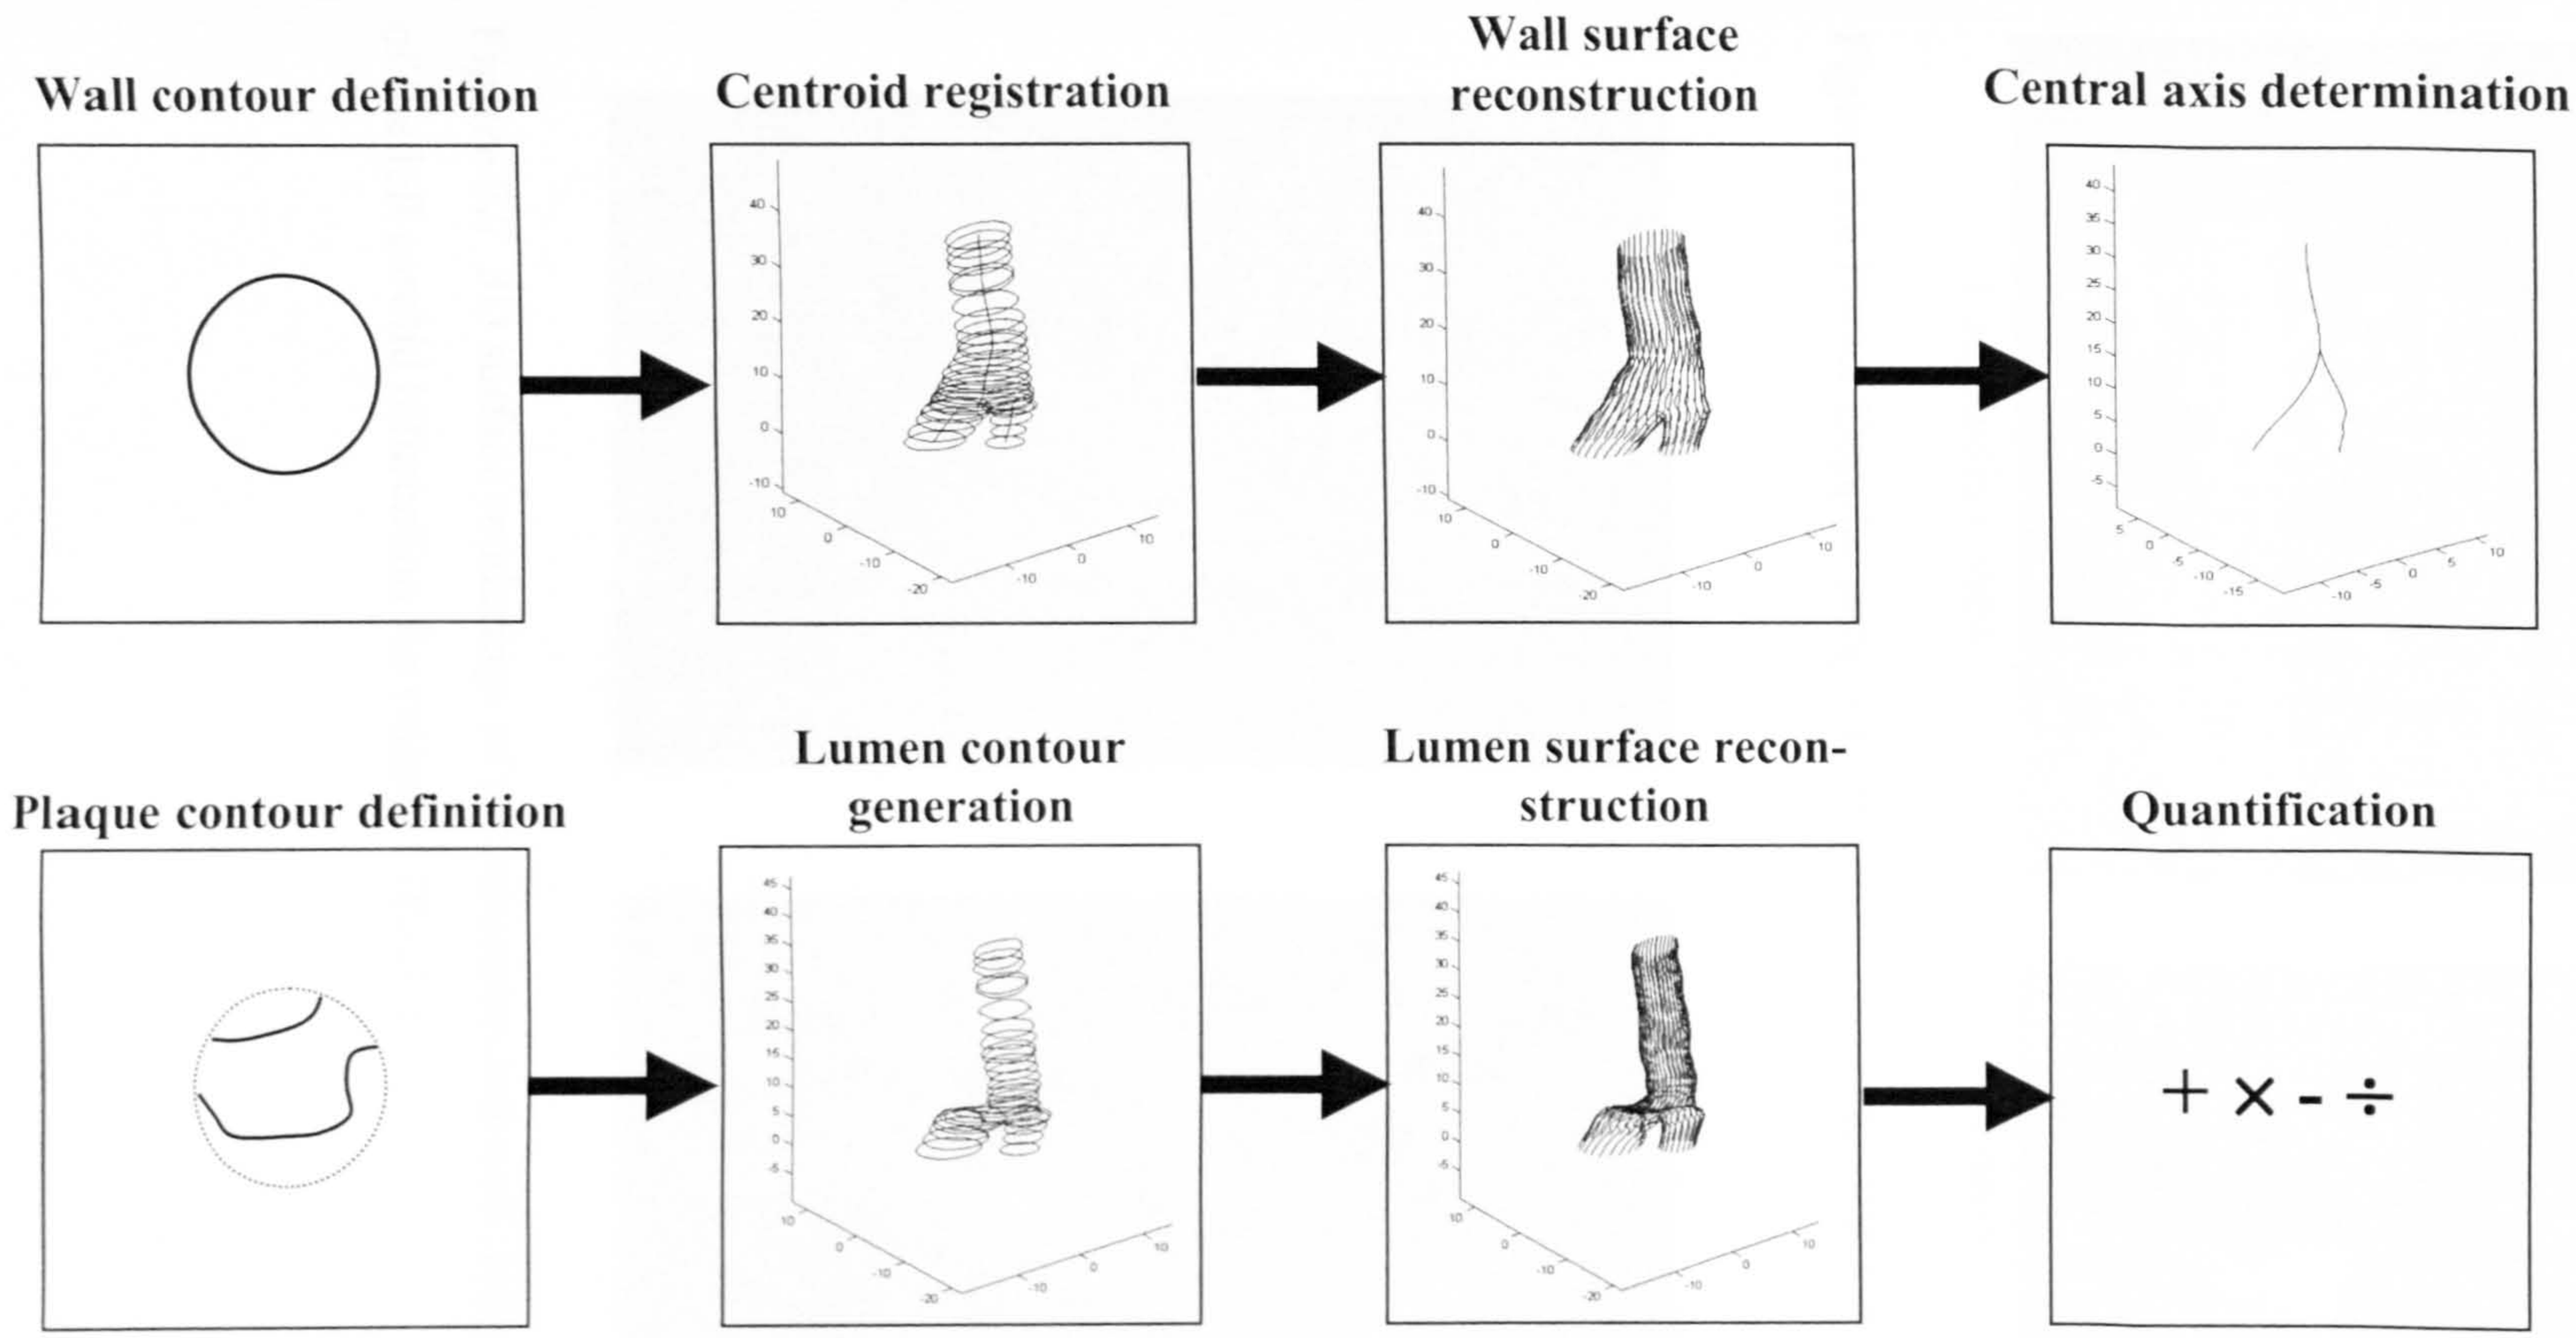
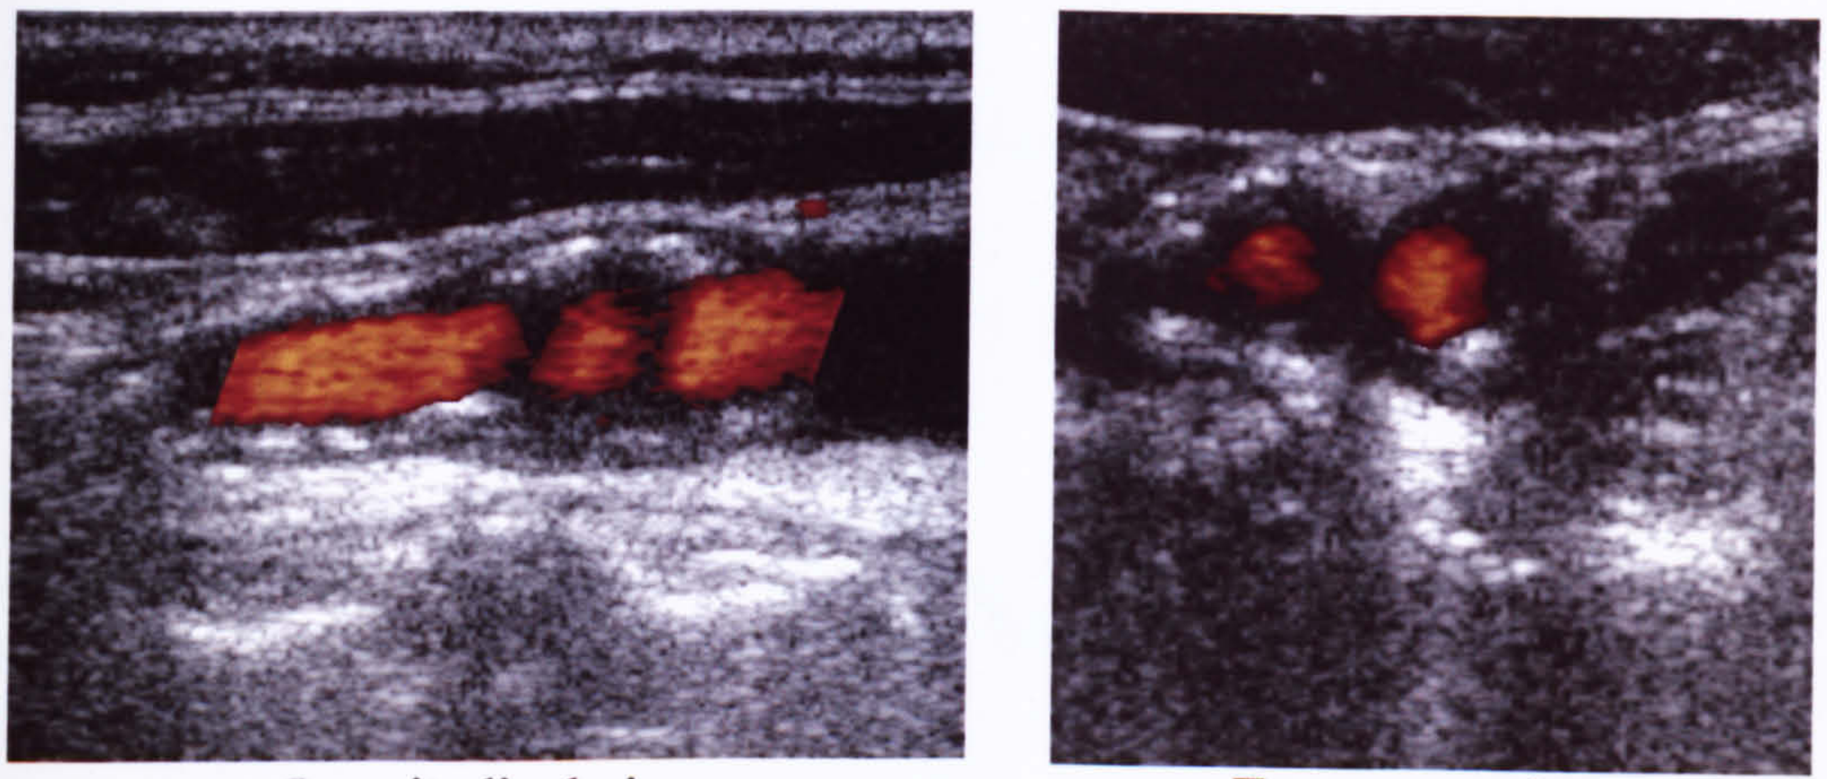


Figure 7.5 The steps involved in analysing a 3DUS data set. Details for each of these steps are given in Chapter 6.

7.2.2.1 Patient 1 (PA-L)



Longitudinal view Transverse view

Figure 7.6 Power Doppler images of the left ICA for patient 1 (PA-L).

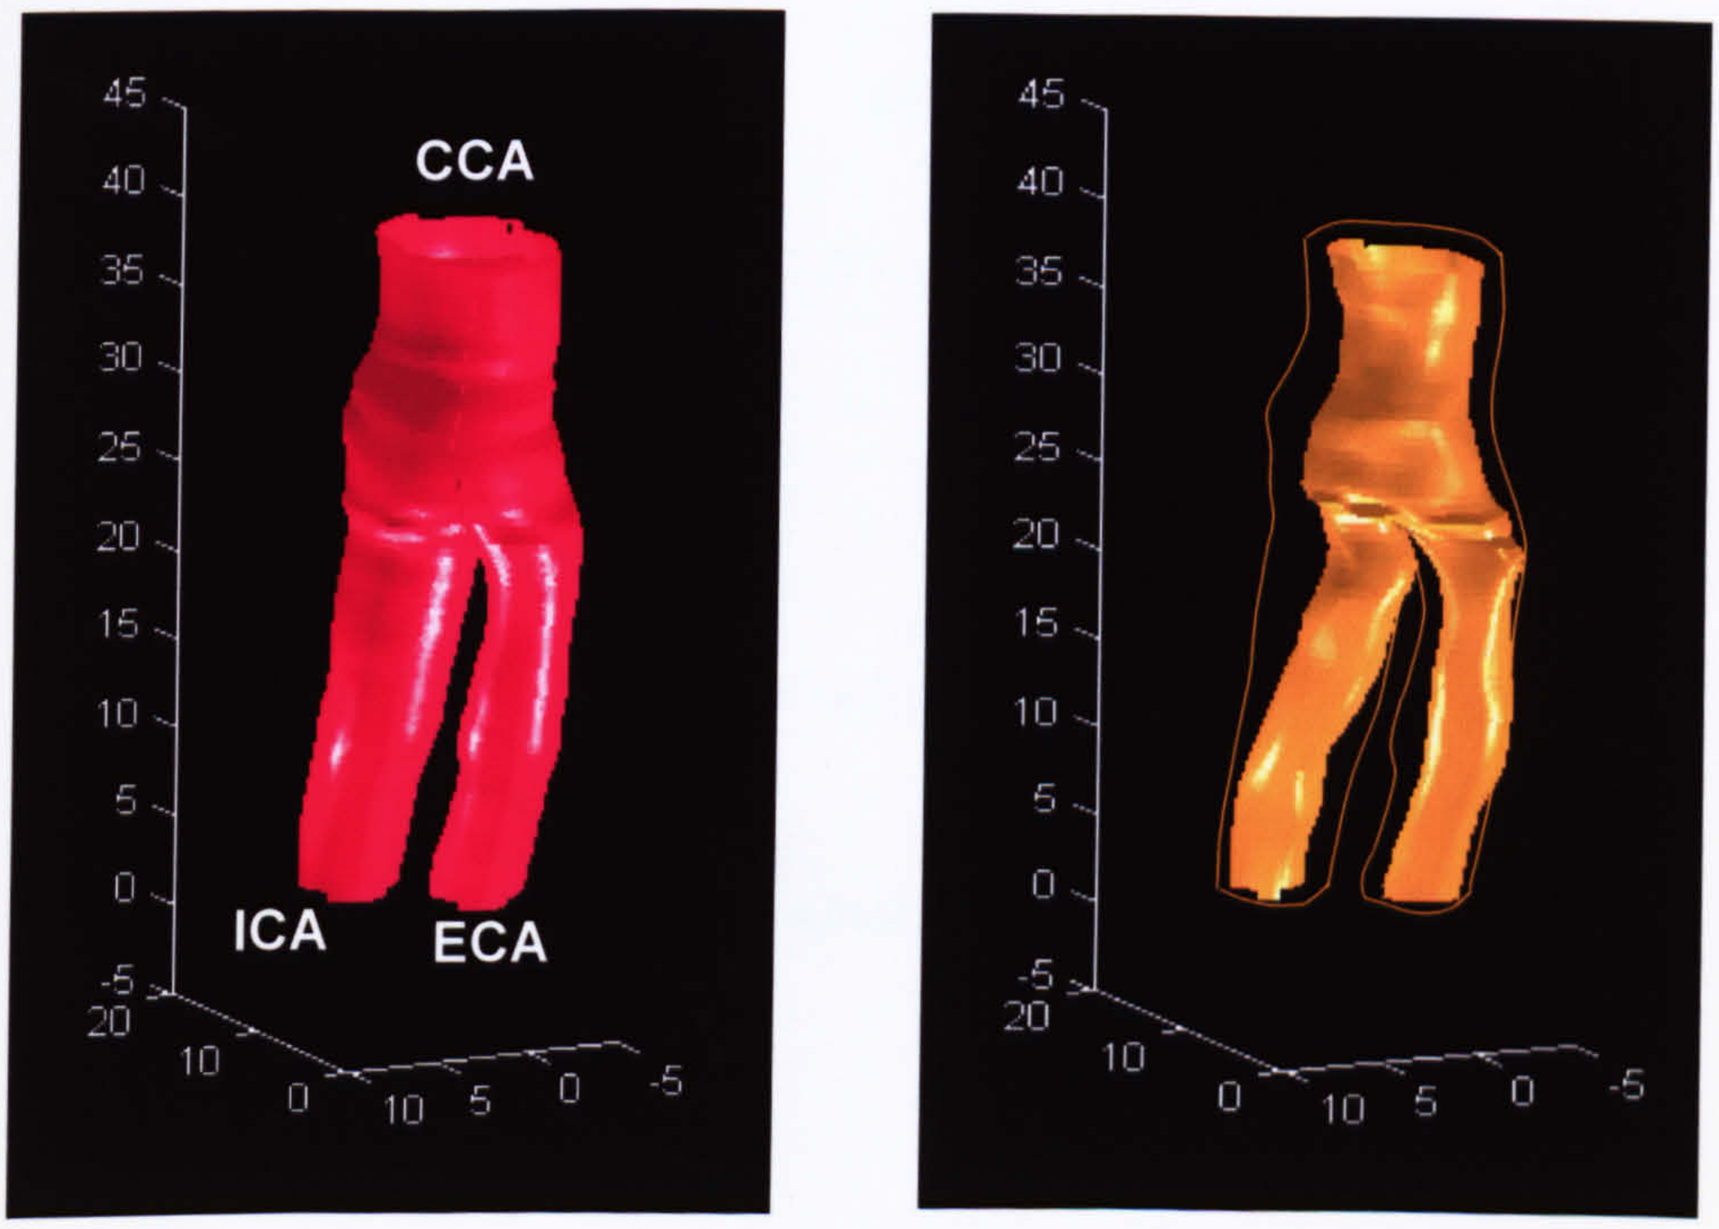
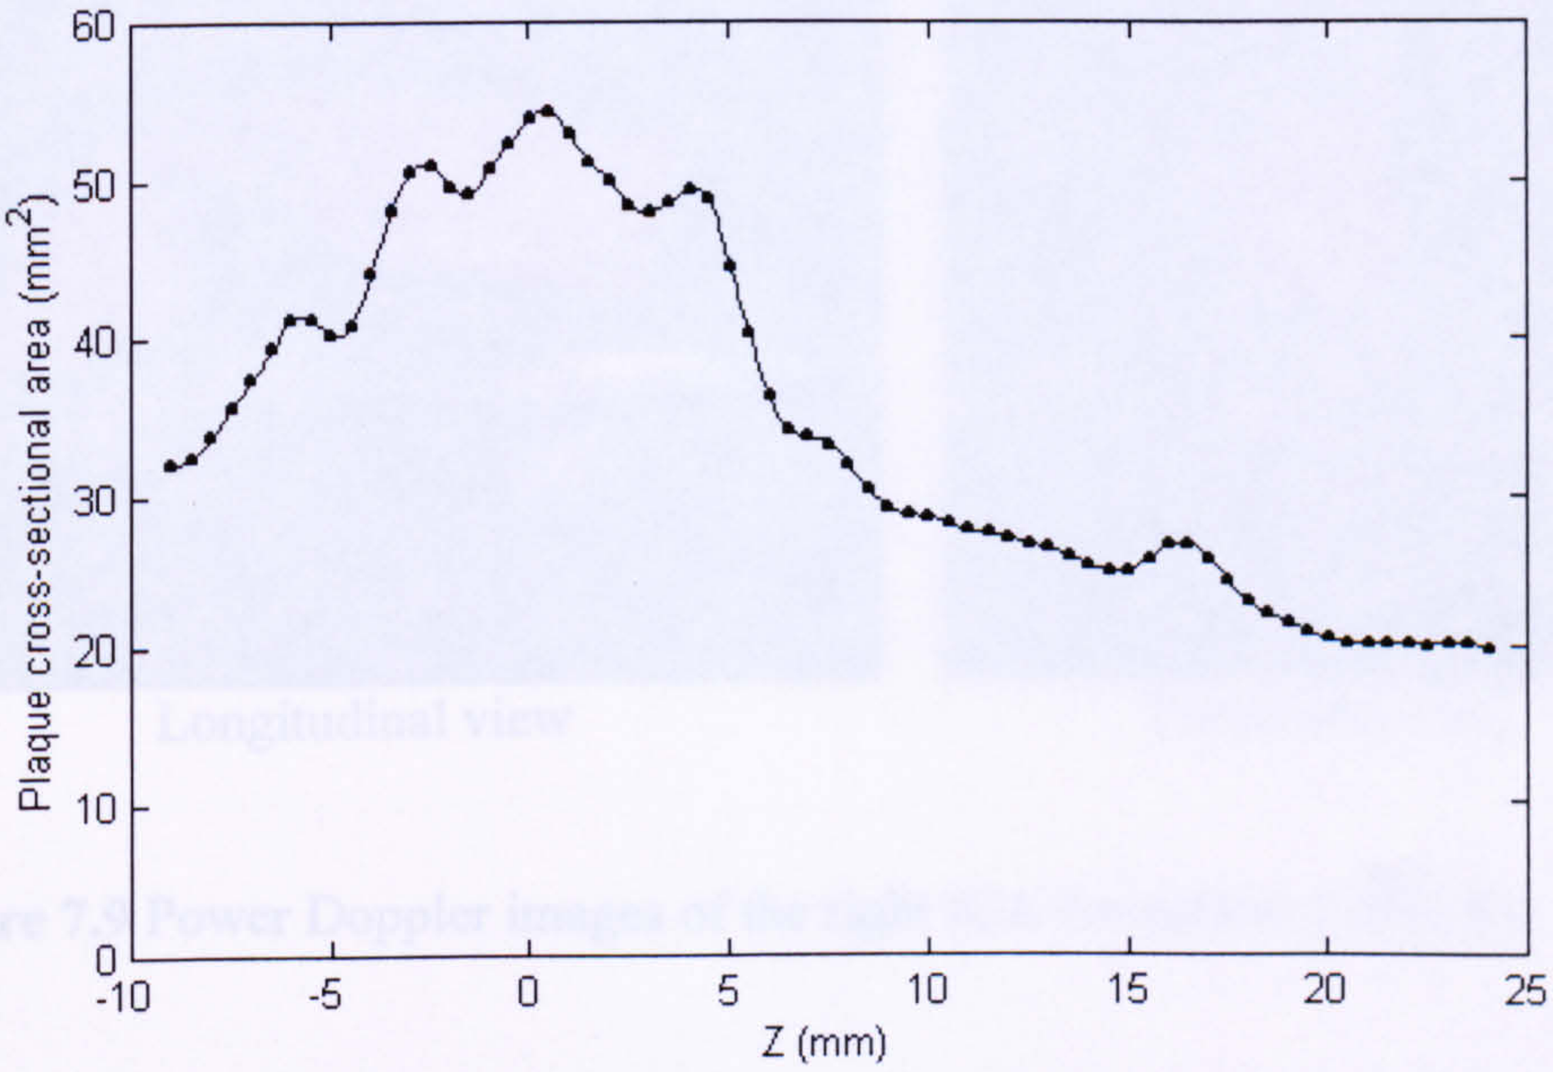
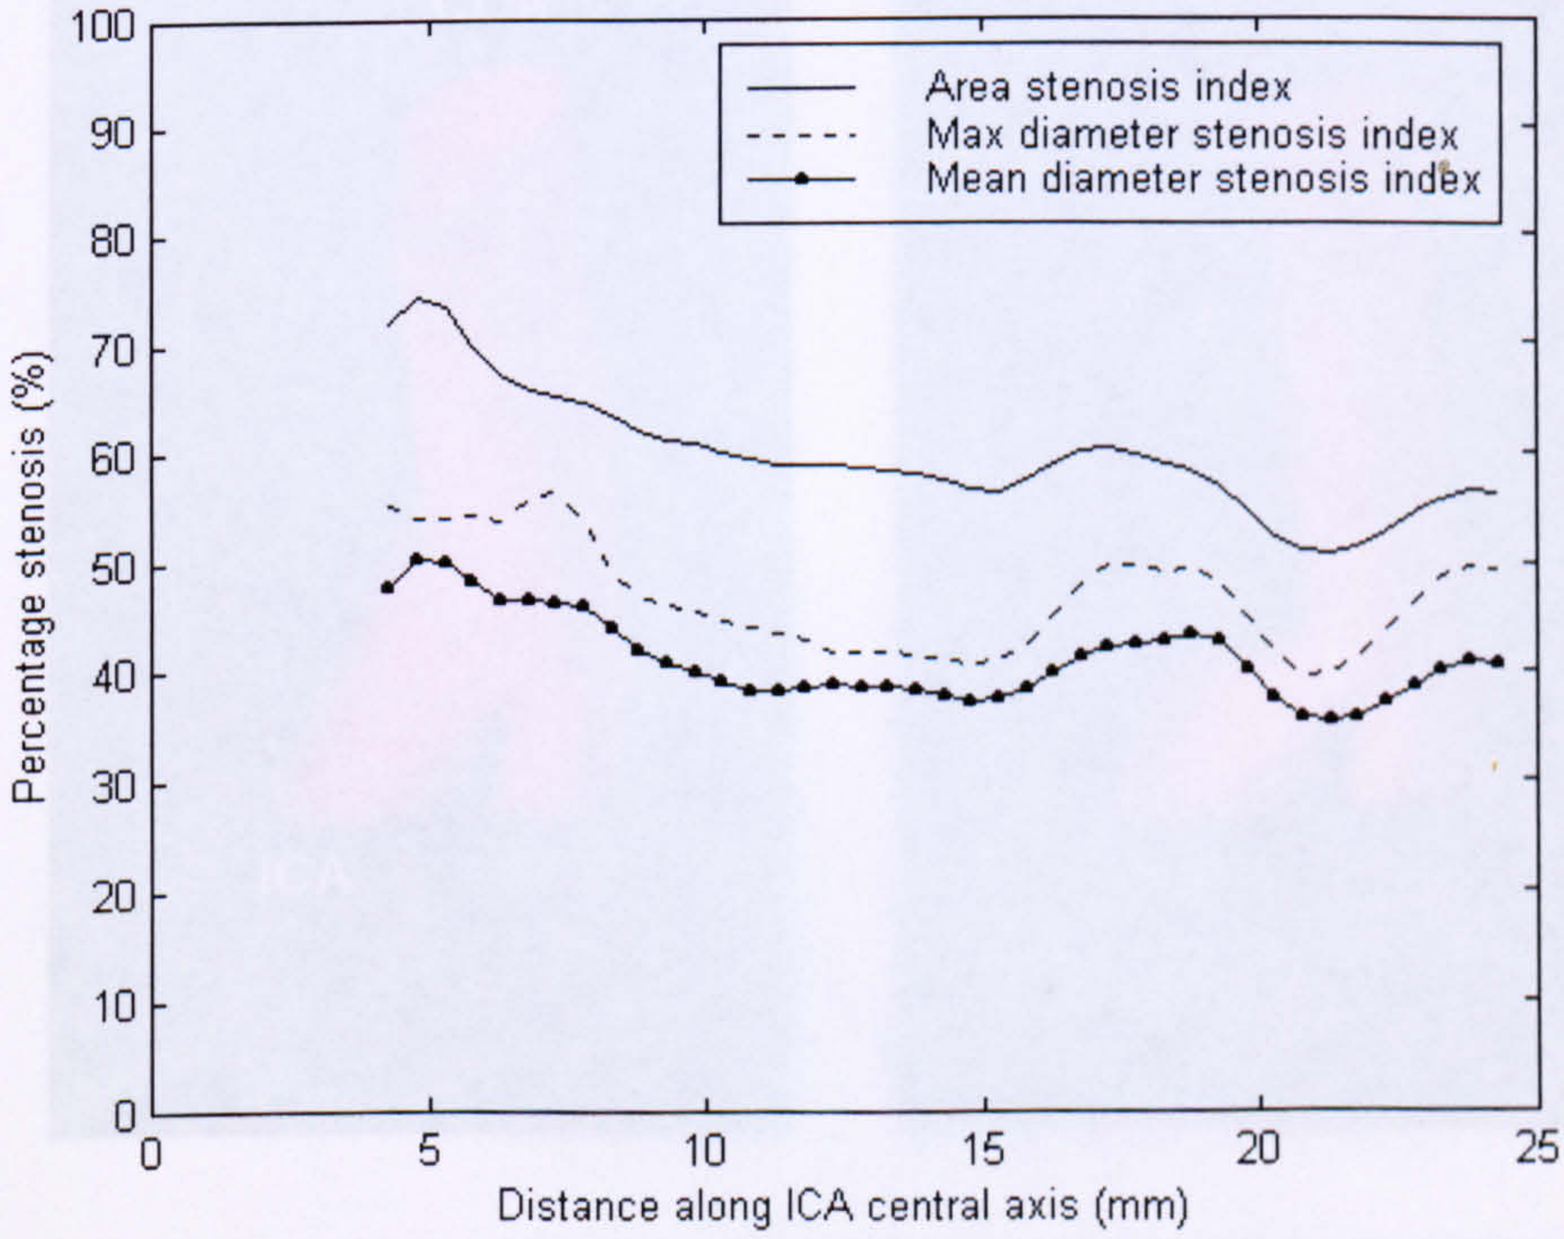


Figure 7.7 3D surface renderings of the reconstructed wall (left) and lumen (right) of the left carotid bifurcation for patient 1 (PA-L).



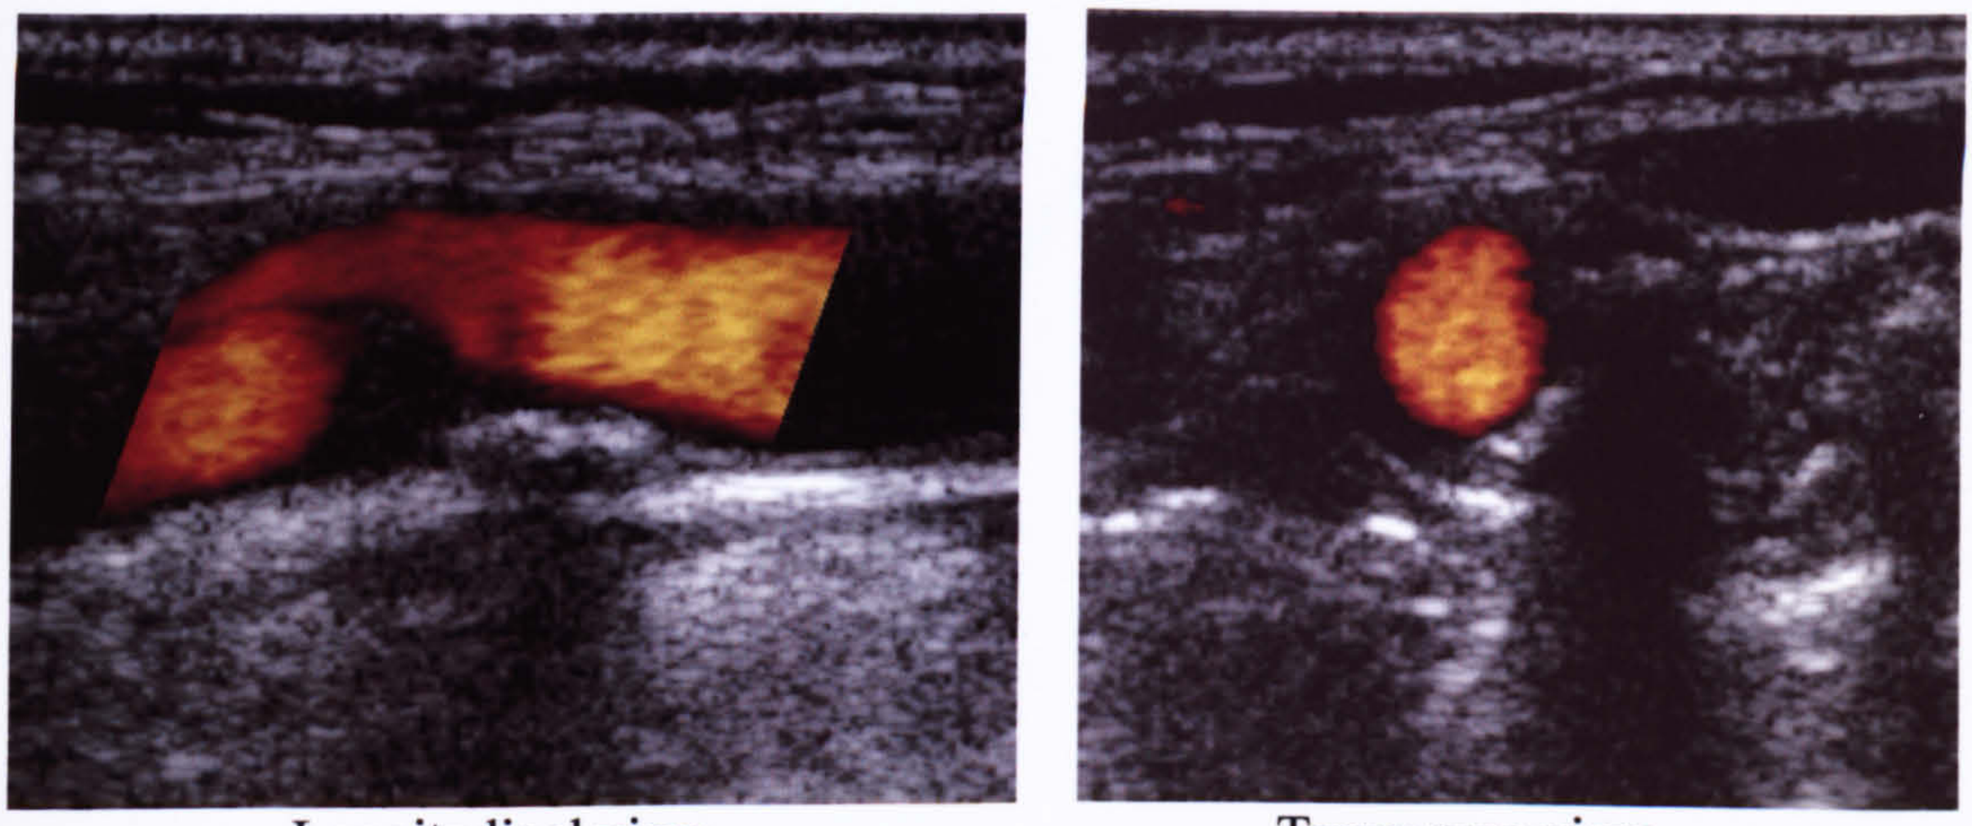
(a)



(b)

Figure 7.8 Results of 3D quantitative analysis for patient 1 (PA-L). (a) Plaque cross-sectional area as a function of length along the z-axis. (b) True transverse stenosis indices in the ICA.

7.2.2.2 Patient 2 (PG-R)



Longitudinal view | Transverse view

Figure 7.9 Power Doppler images of the right ICA for patient 2 (PG-R).

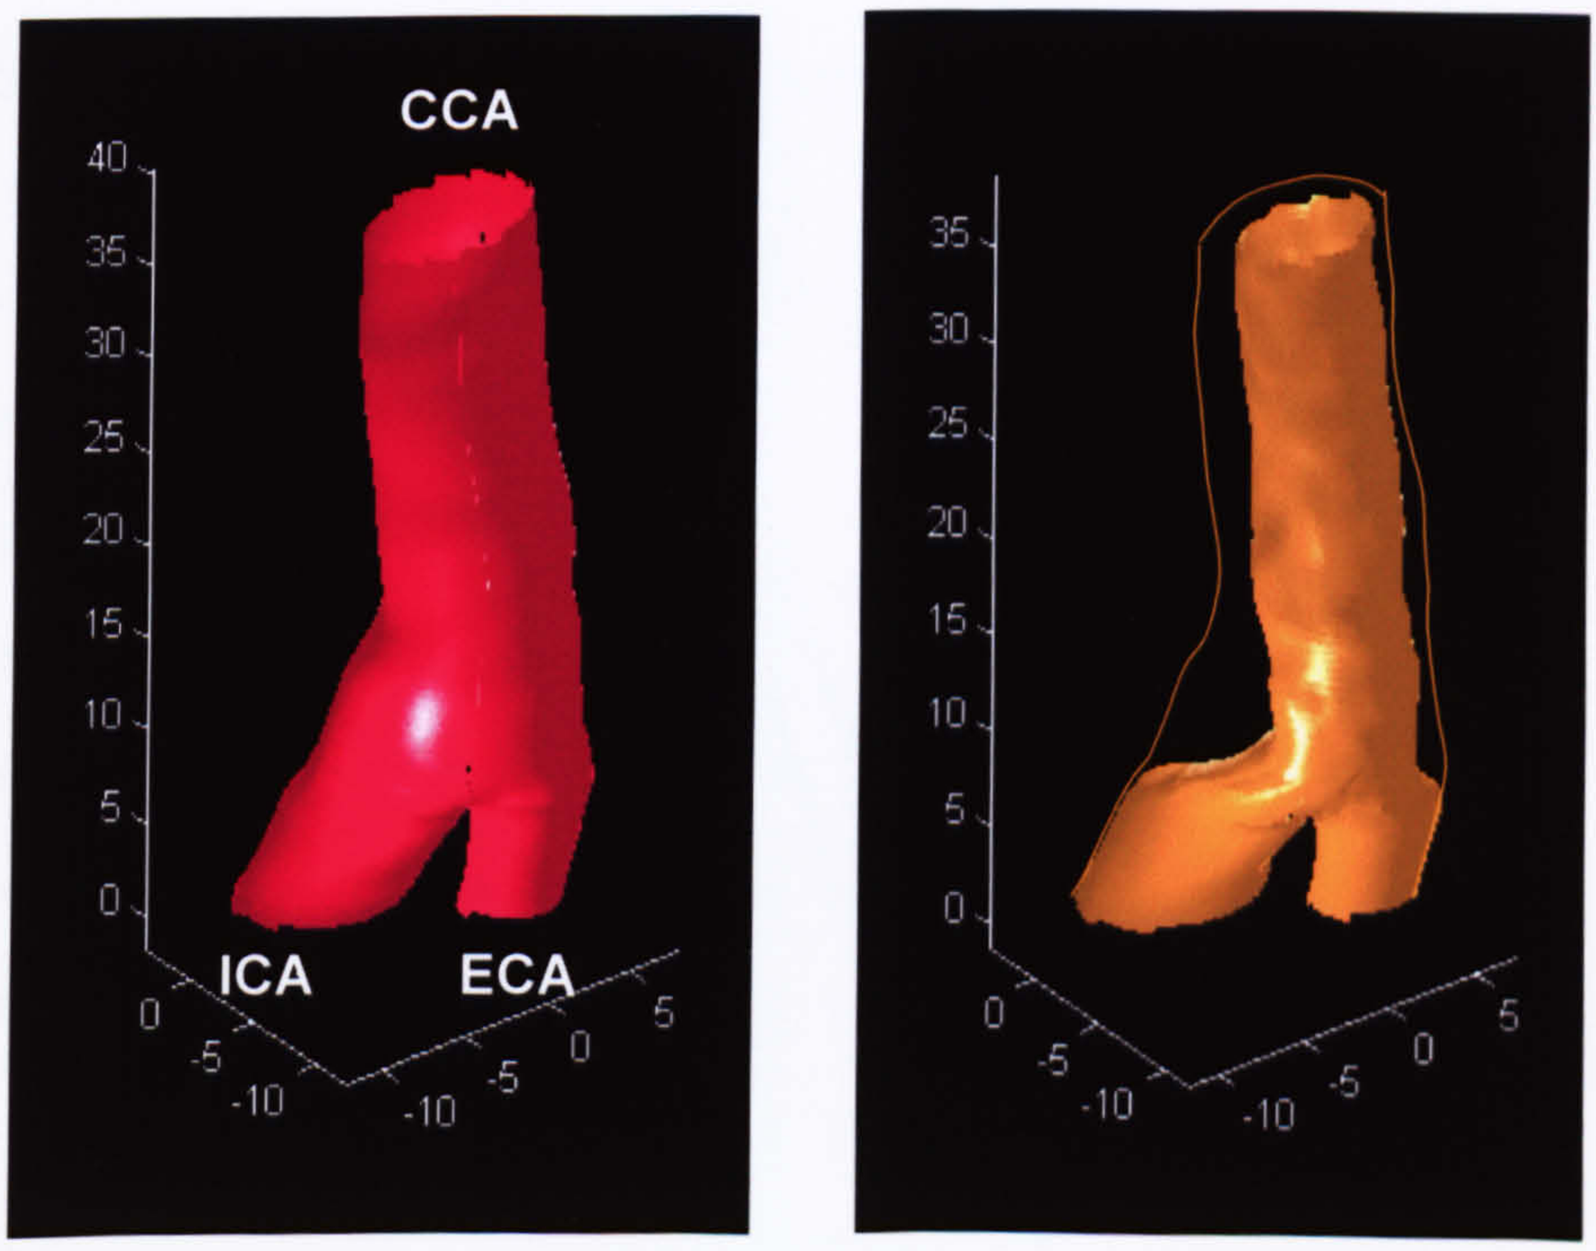
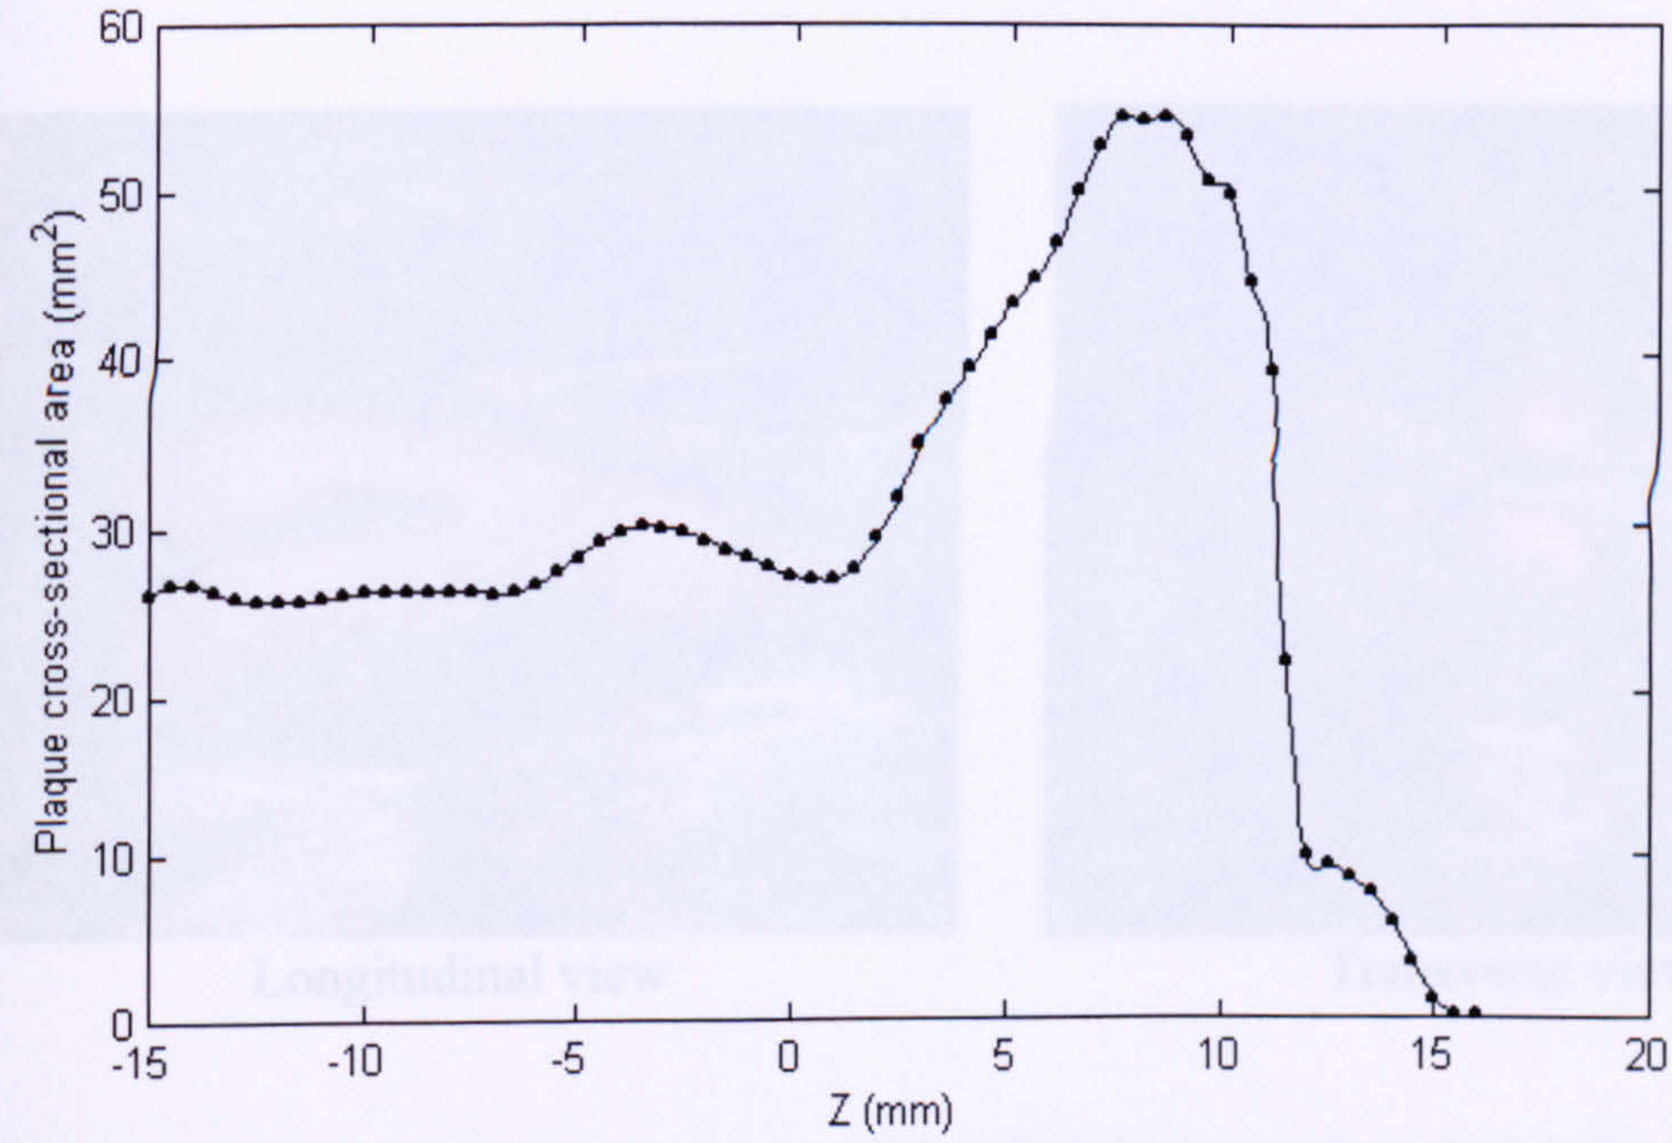
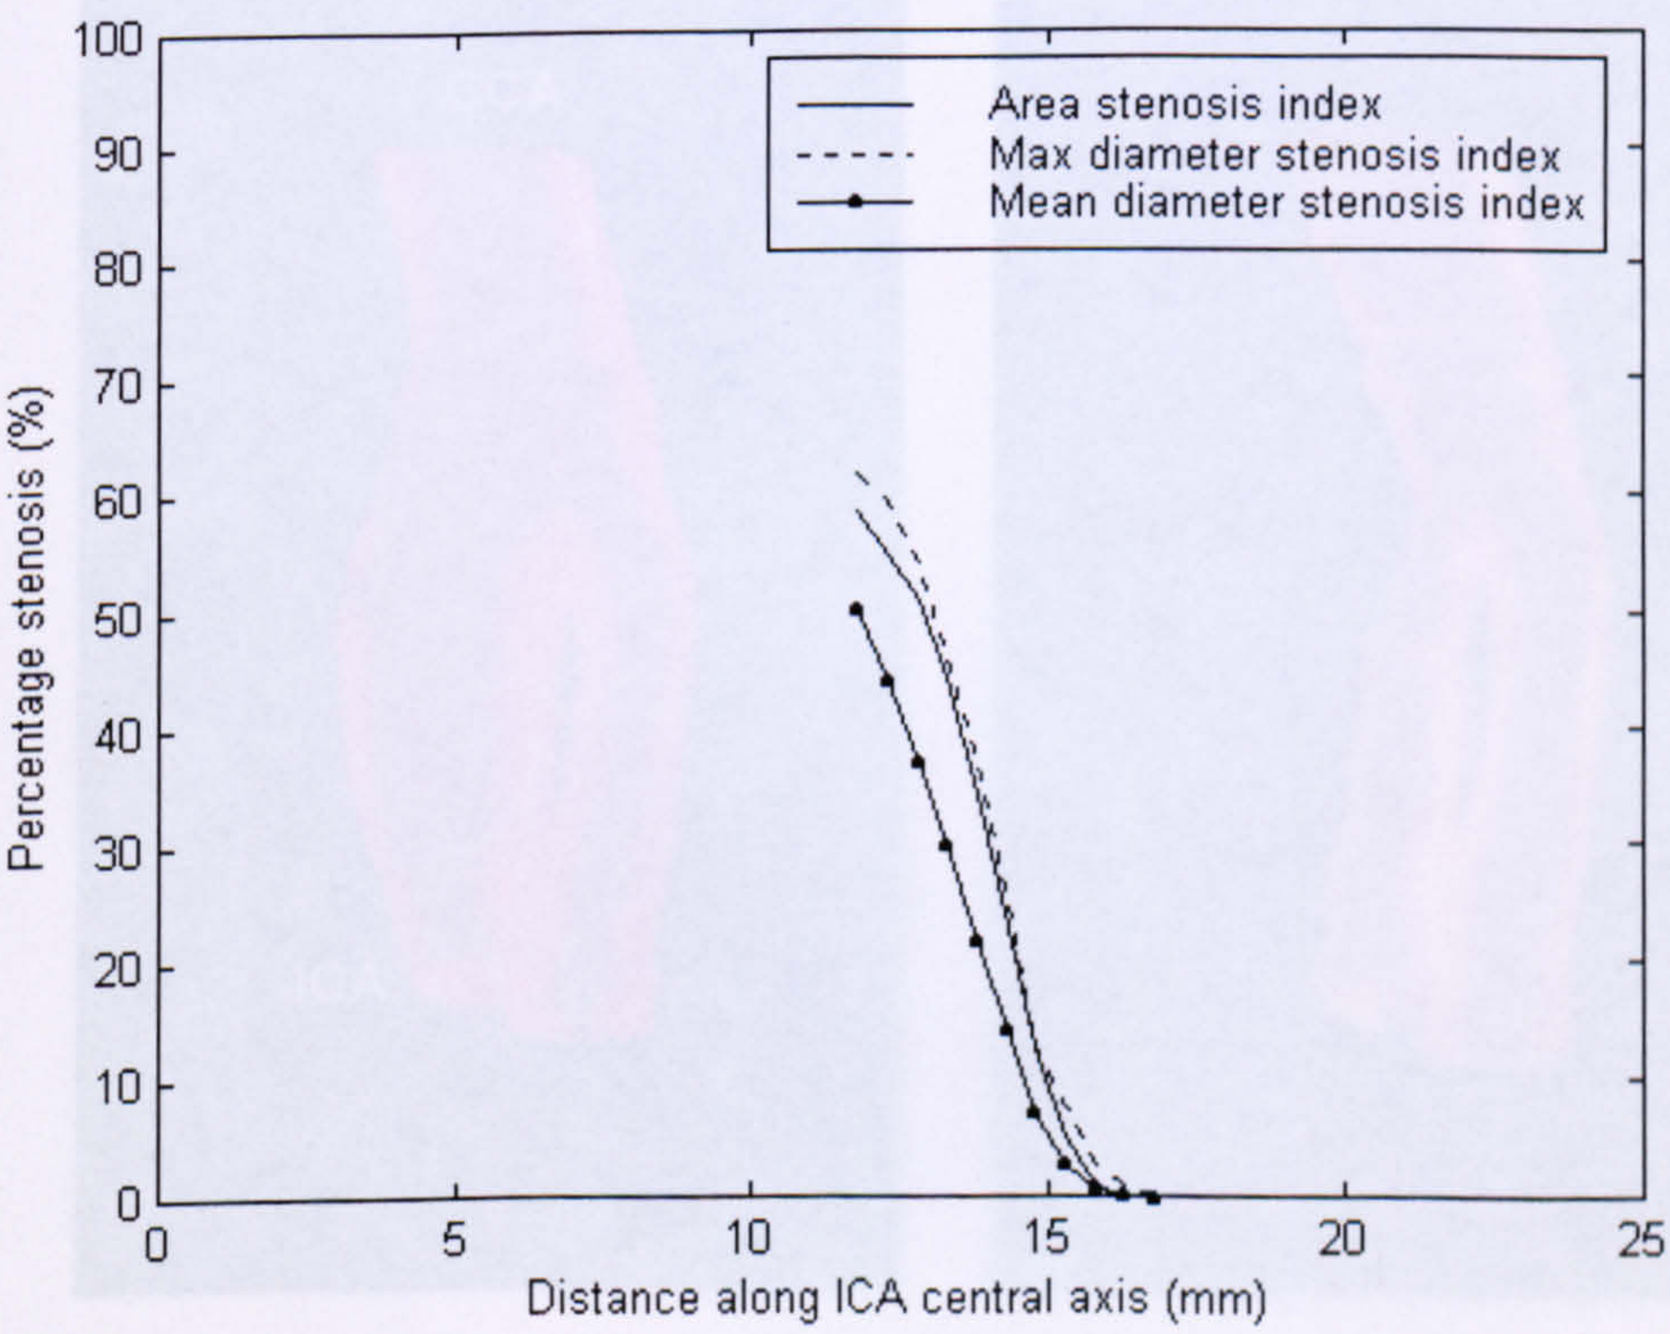


Figure 7.10 *3D surface* renderings of the reconstructed wall (left) and lumen (right) of the right carotid bifurcation for patient 2 (PG-R).



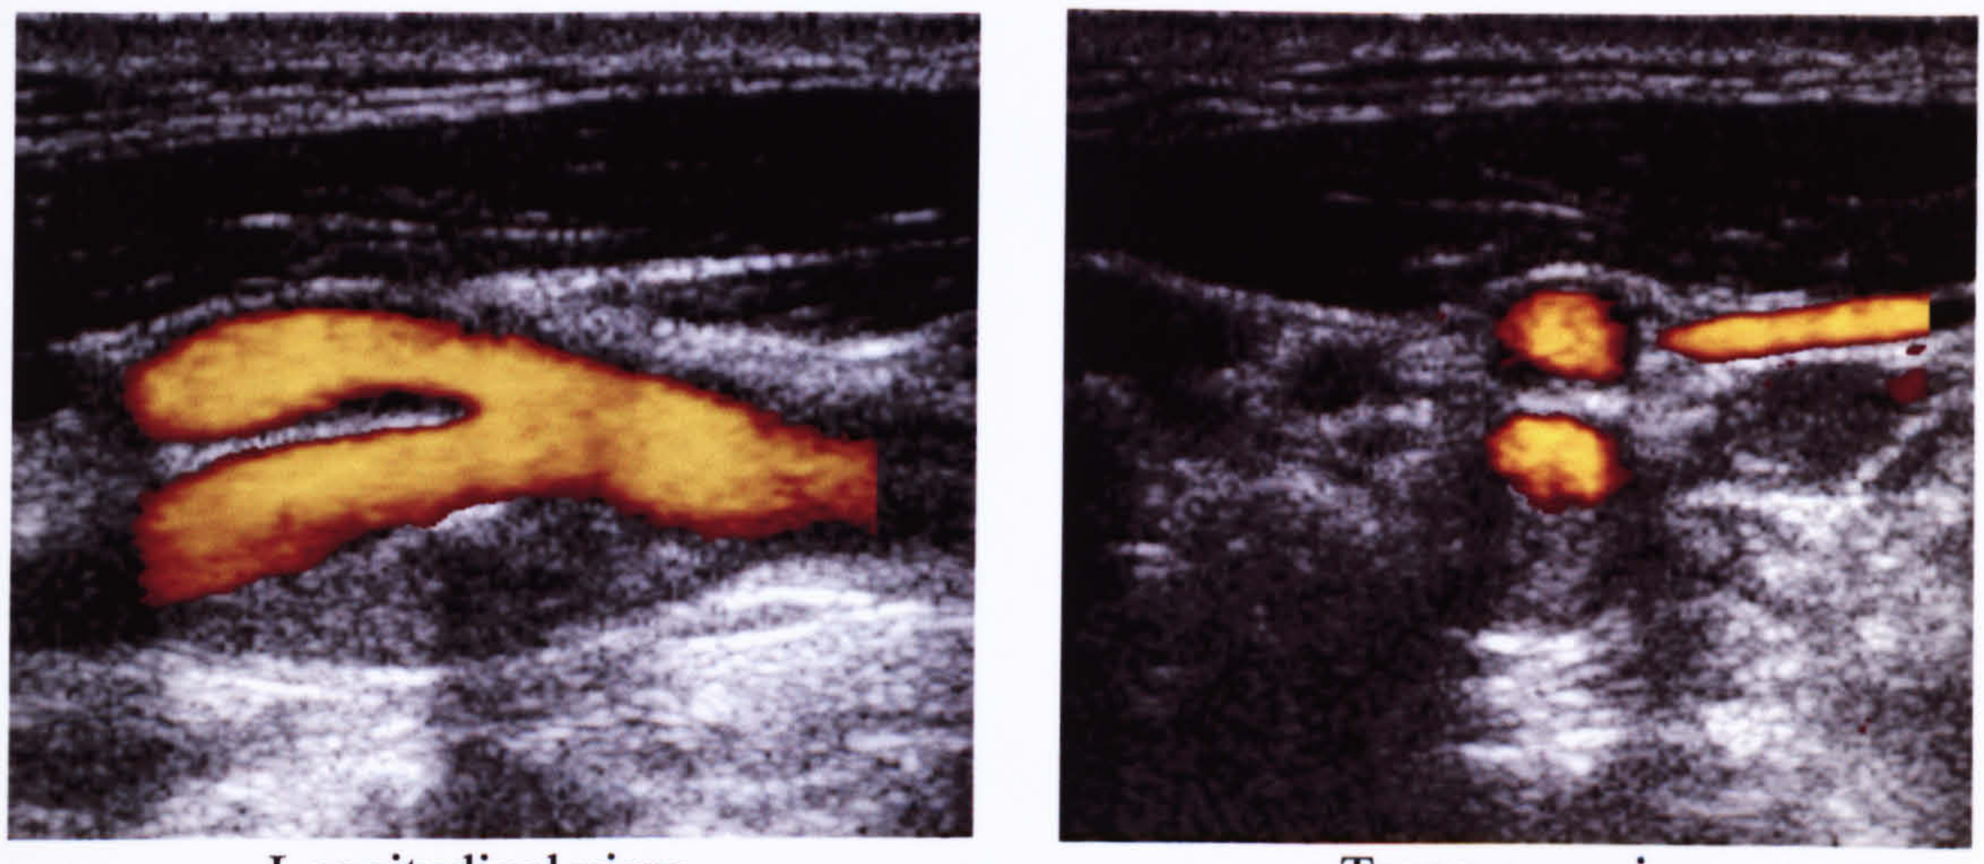
(a)



(b)

Figure 7.11 Results of 3D quantitative analysis for patient 2 (PG-R). (a) Plaque cross-sectional area as a function of length along the z-axis. (b) True transverse stenosis indices in the ICA.

7.2.2.3 Patient 3 (SA-R)



Longitudinal view

Transverse view

Figure 7.12 Power Doppler images of the right ICA for patient 3 (SA-R).

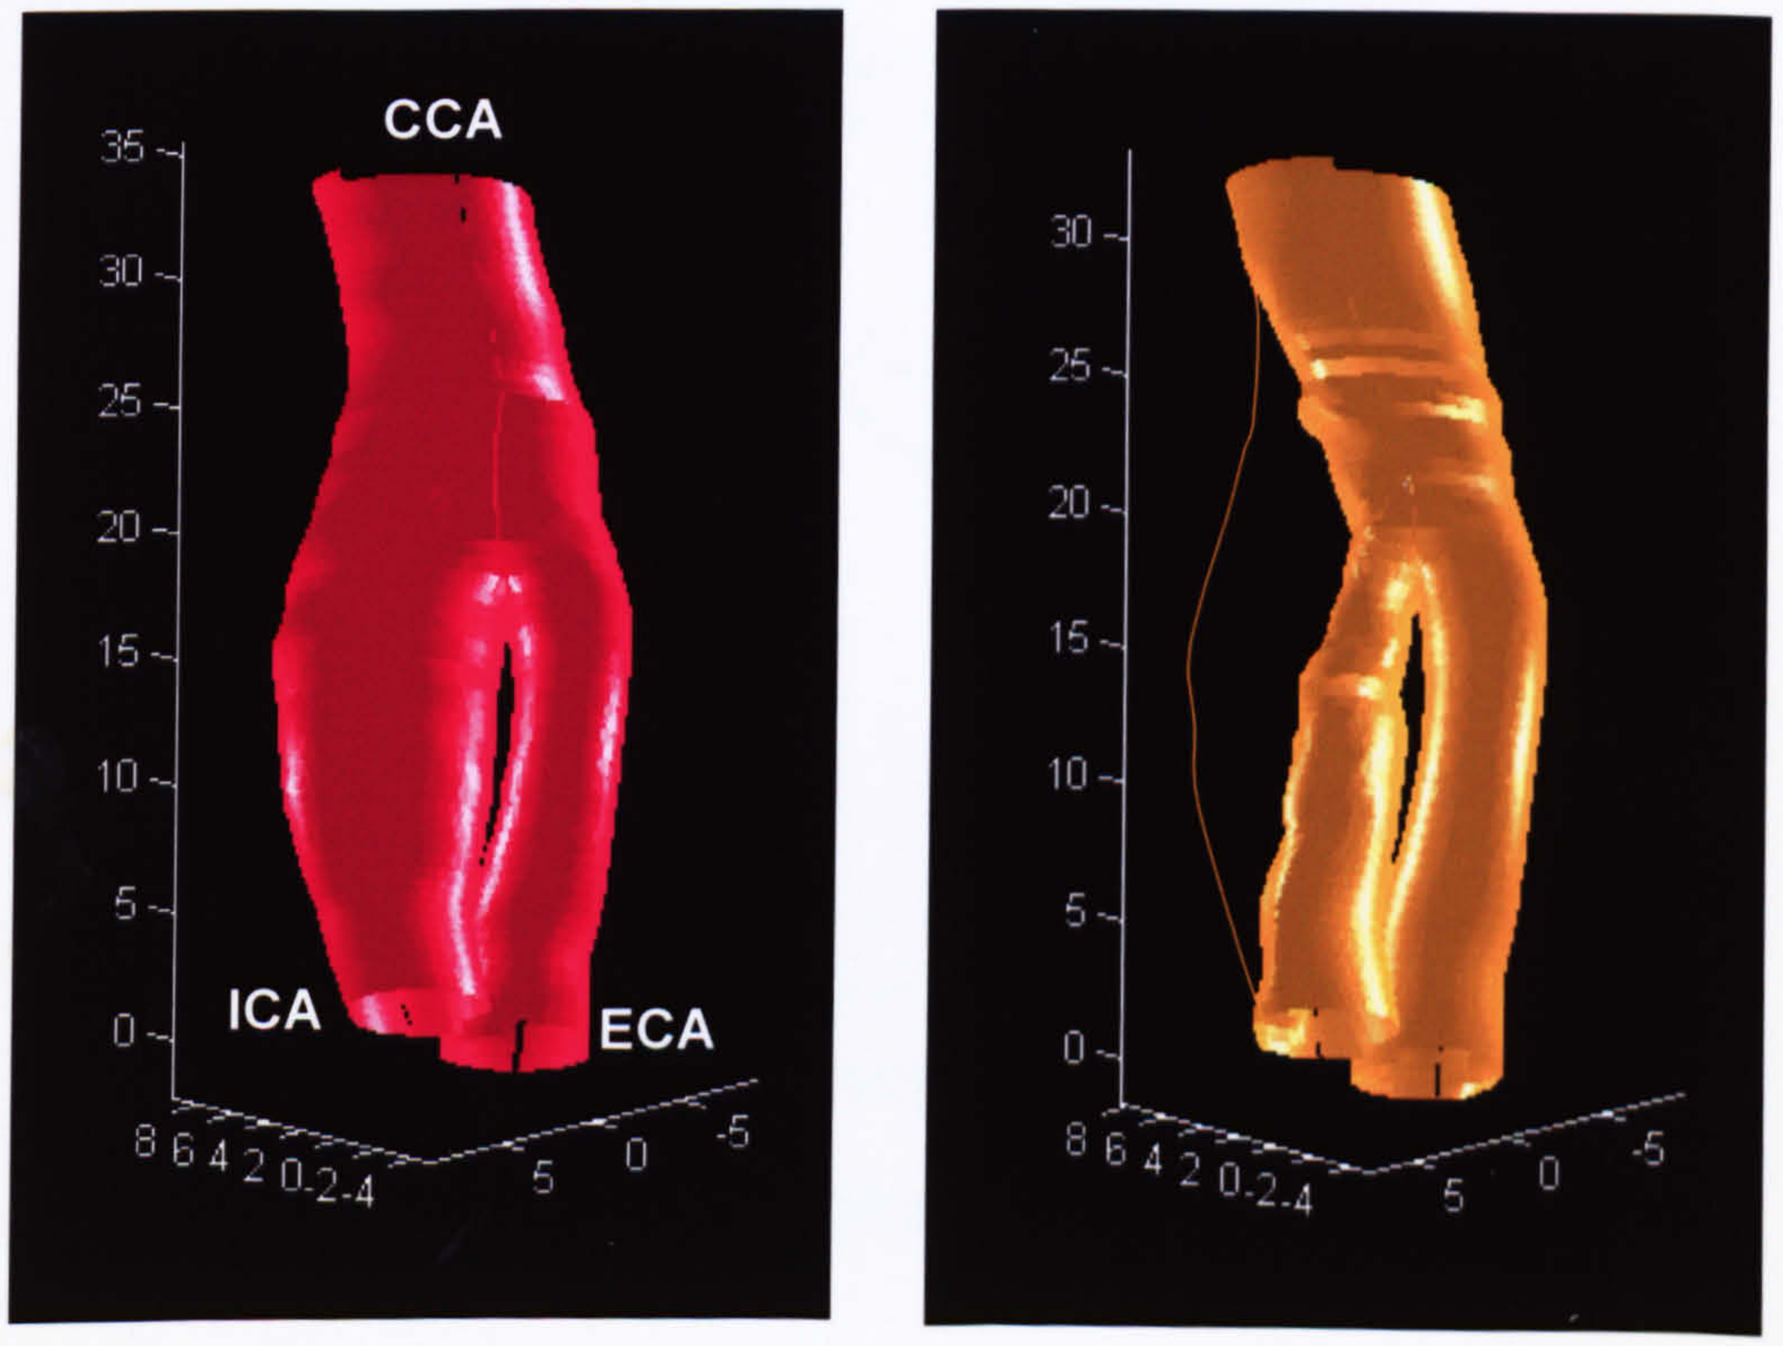


Figure 7.13 3D surface renderings of the reconstructed wall (left) and lumen (right) of the right carotid bifurcation for patient 3 (SA-R).

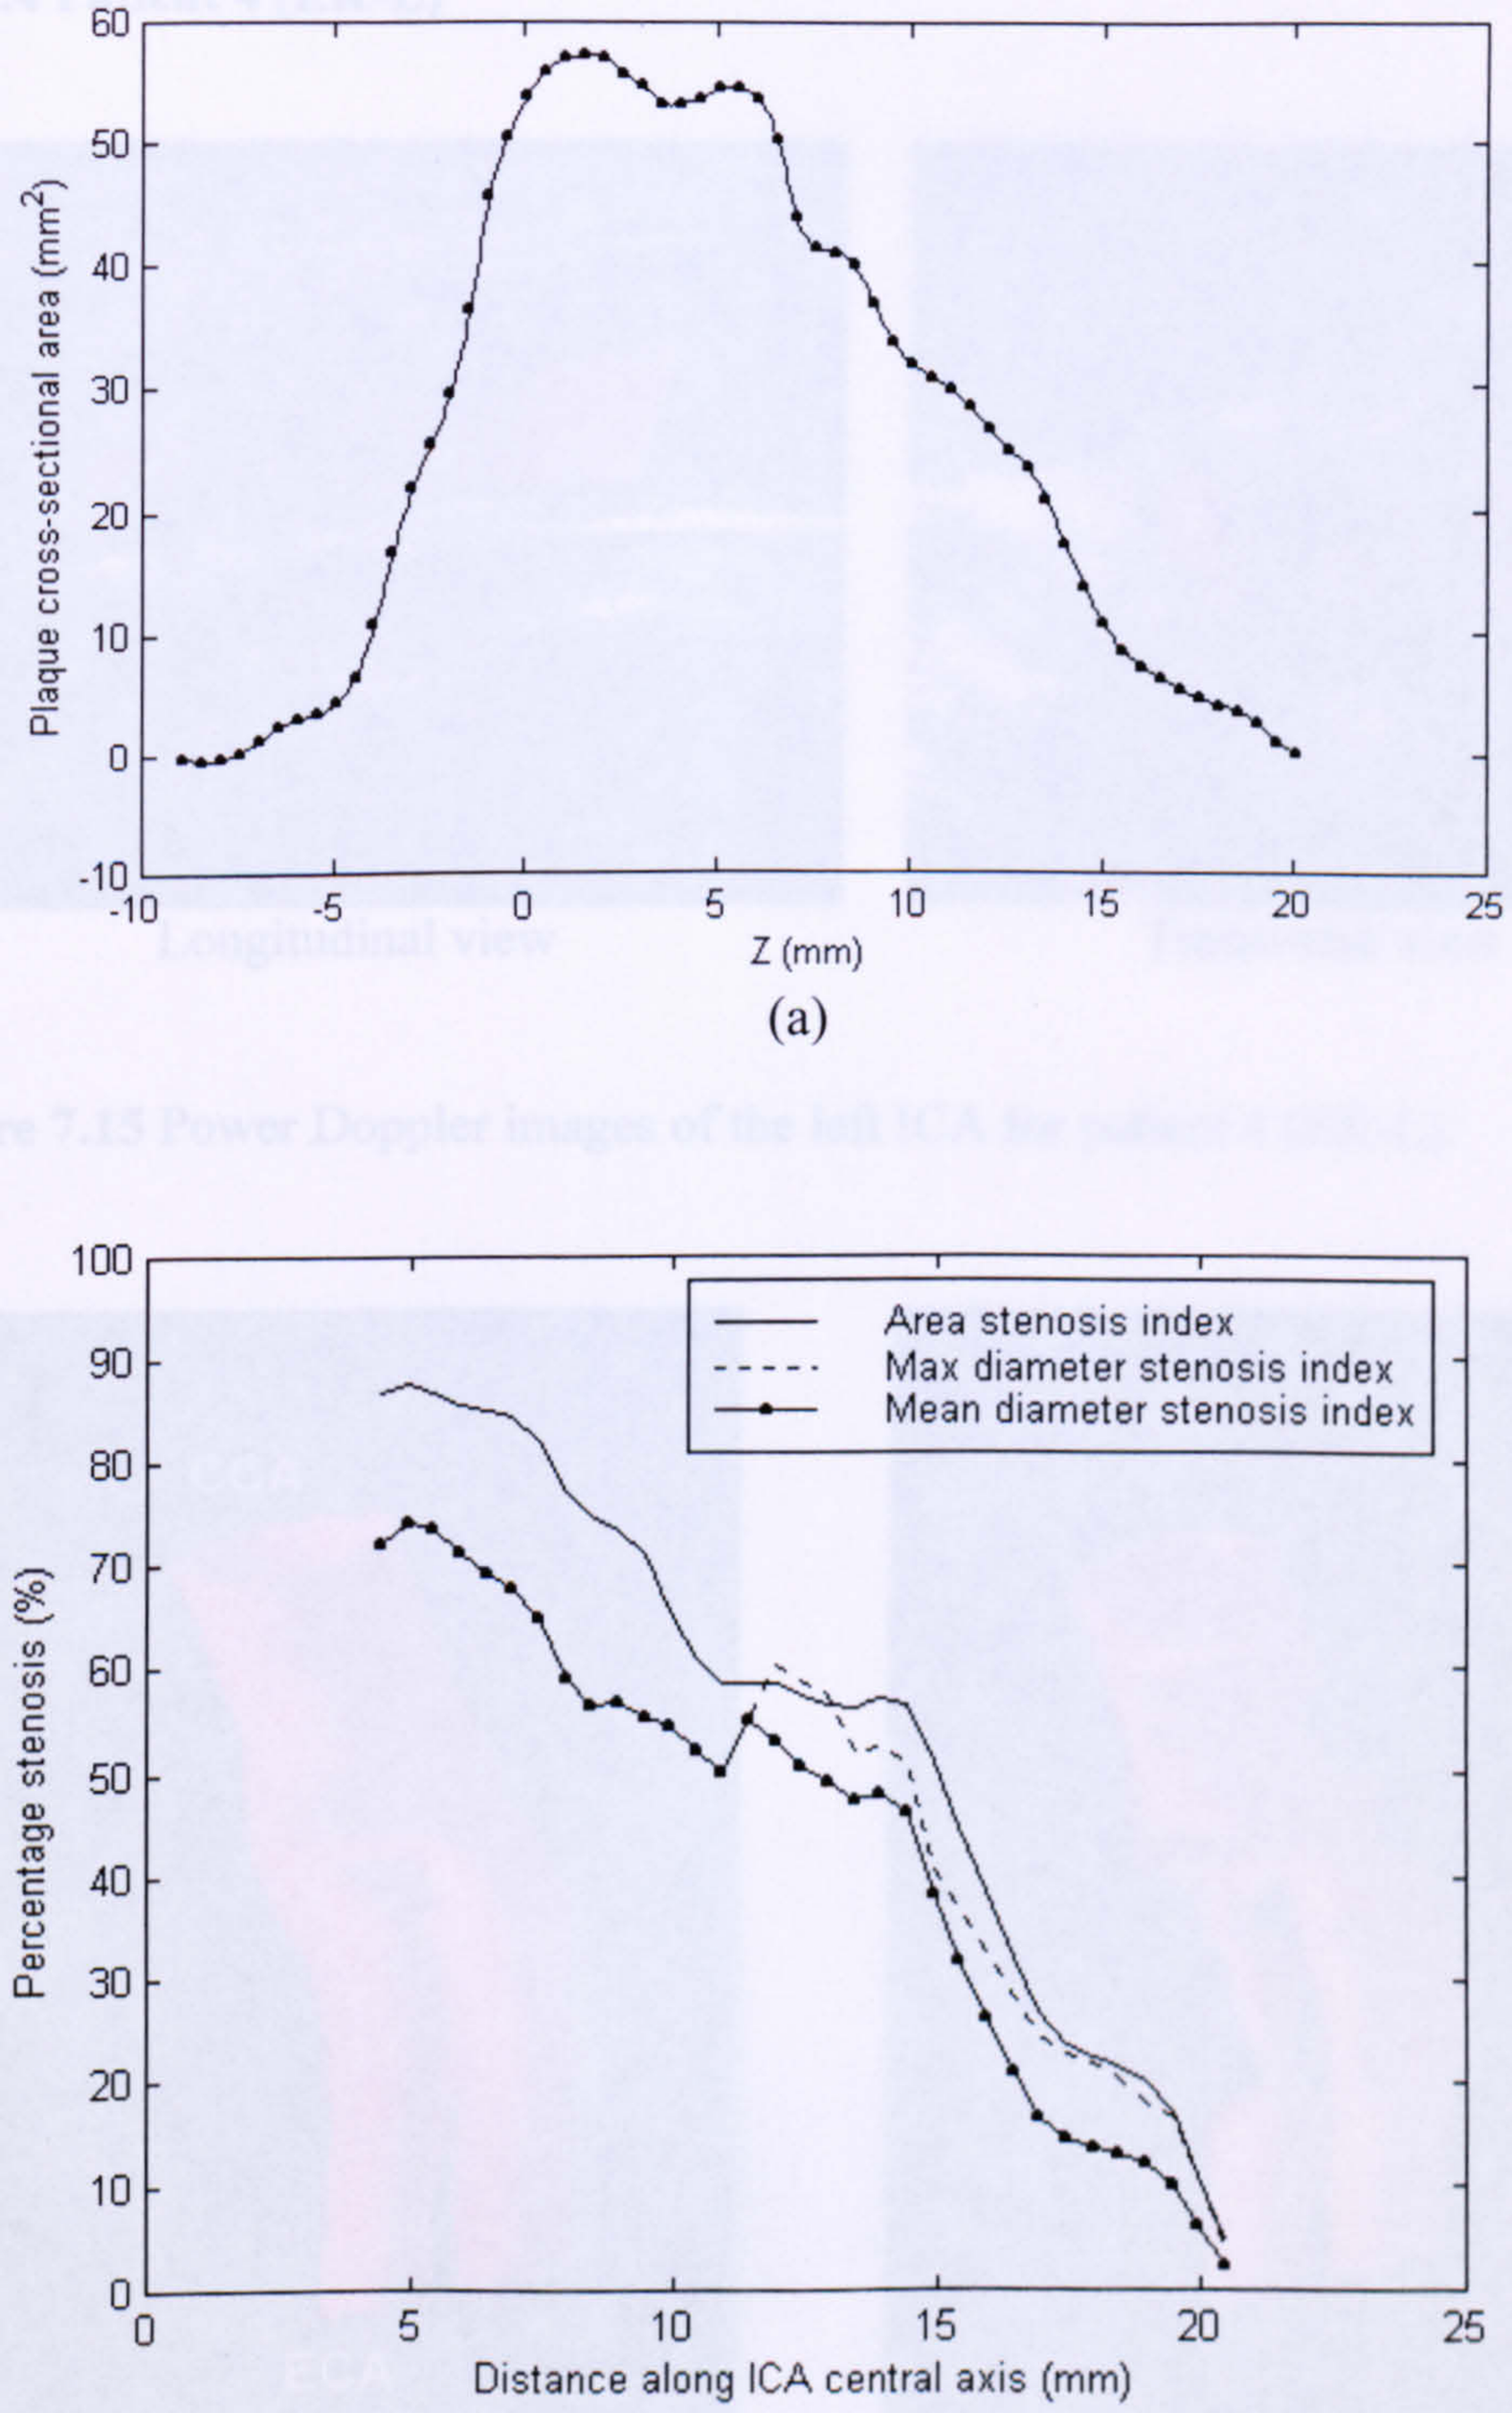
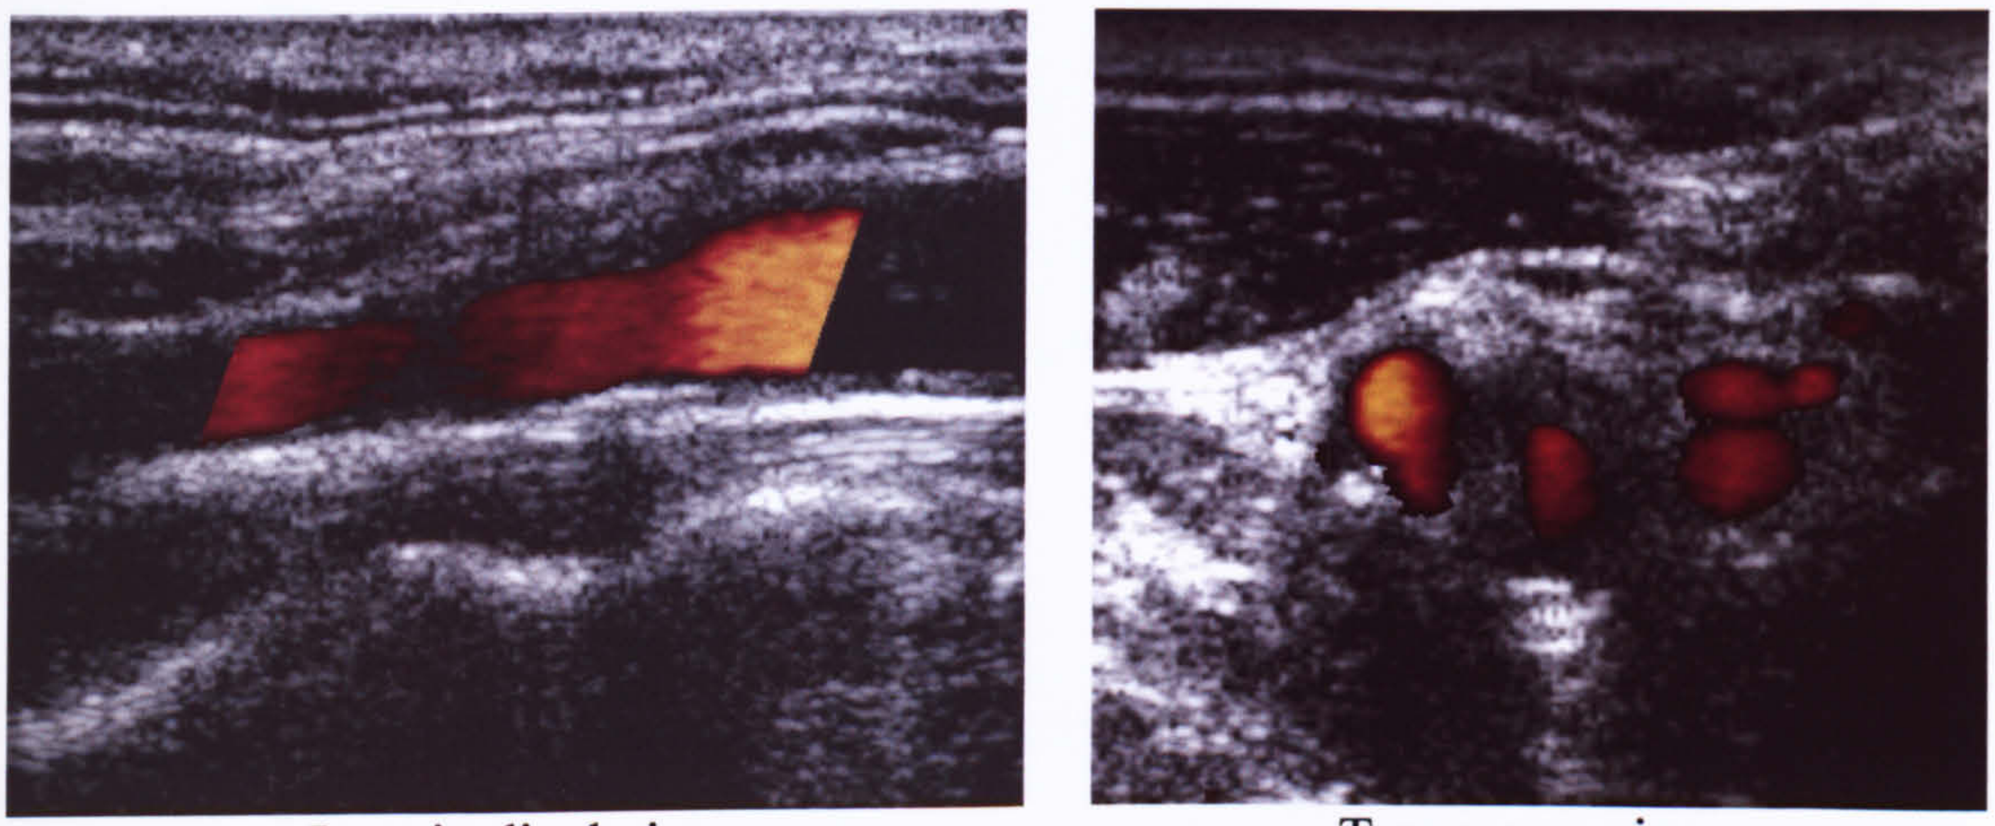


Figure 7.14 Results of 3D quantitative analysis for patient 3 (SA-R). (a) Plaque cross-sectional area as a function of length along the z-axis. (b) True transverse stenosis indices in the ICA).

7.2.2.4 Patient 4 (EK-L)



Longitudinal view | Transverse view

Figure 7.15 Power Doppler images of the left ICA for patient 4 (EK-L).

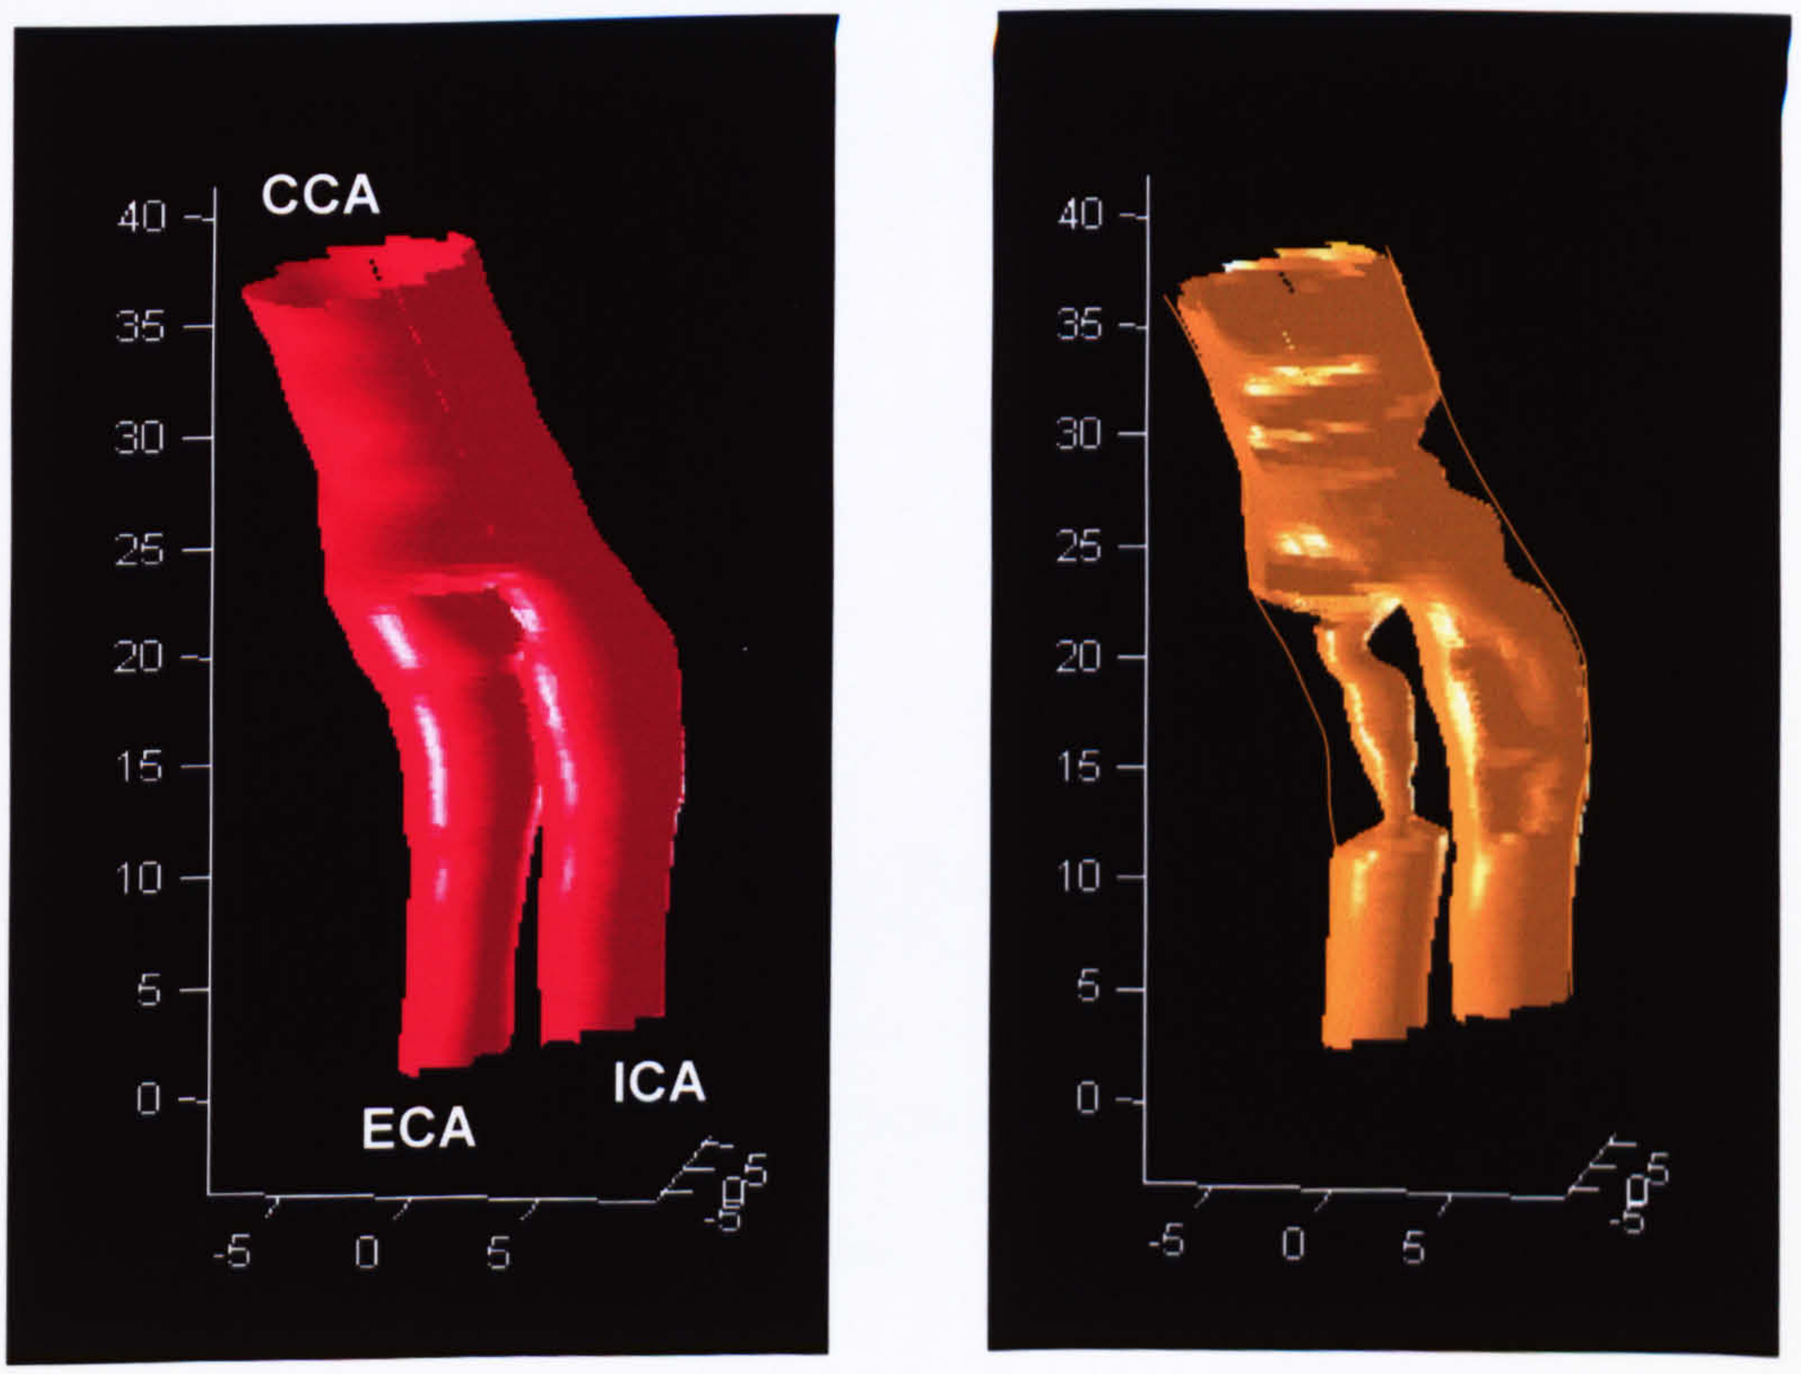
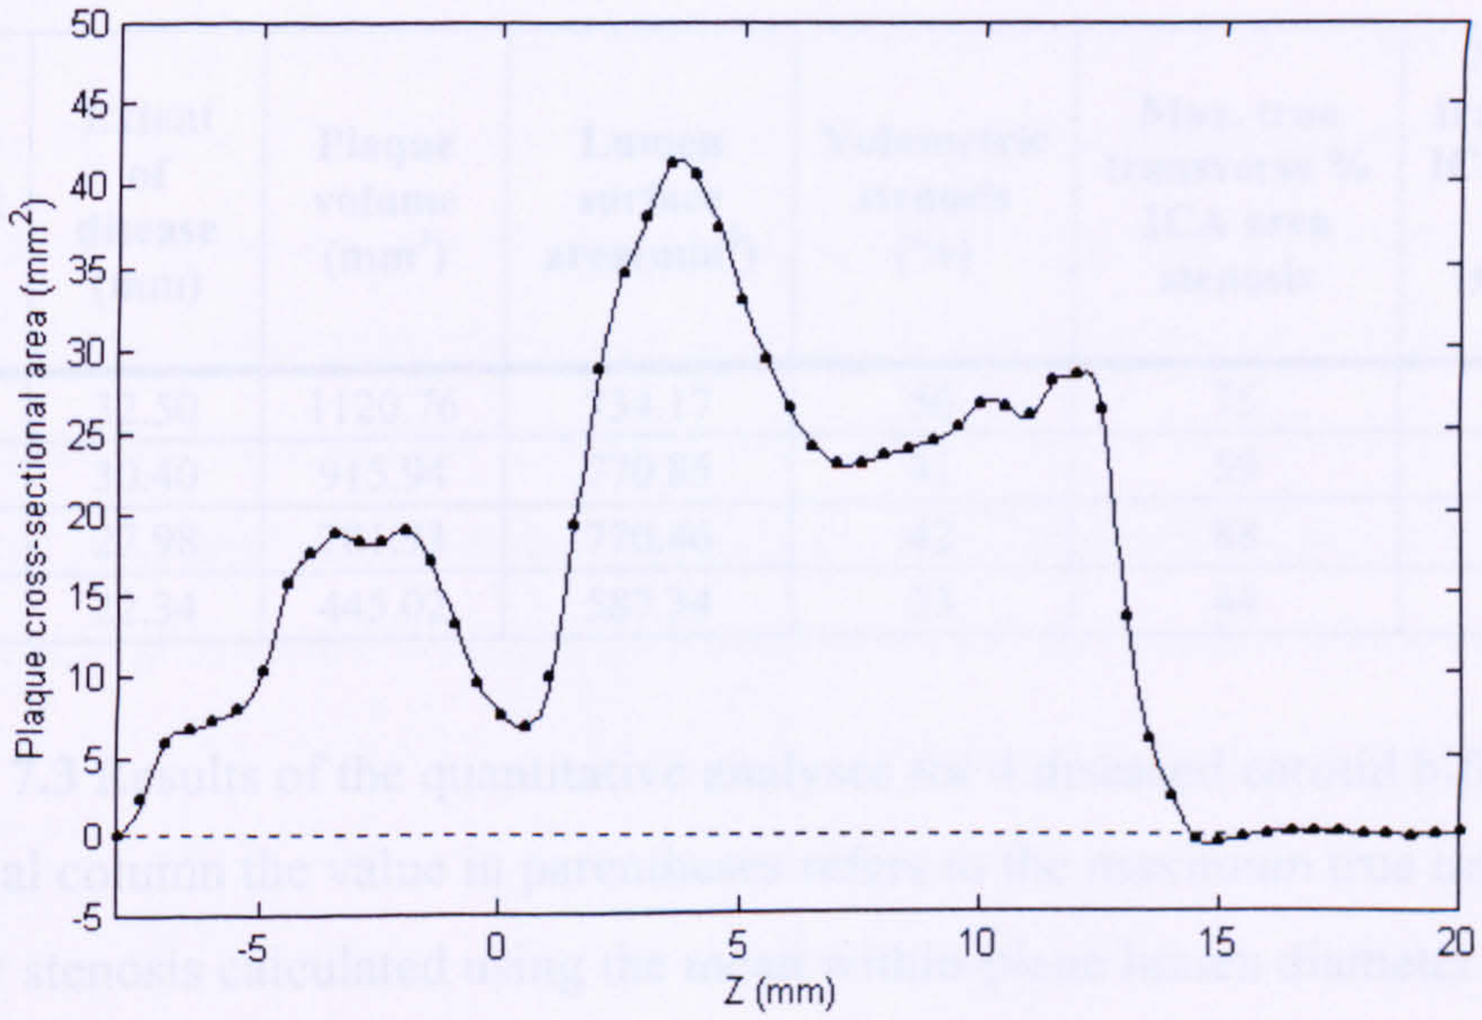
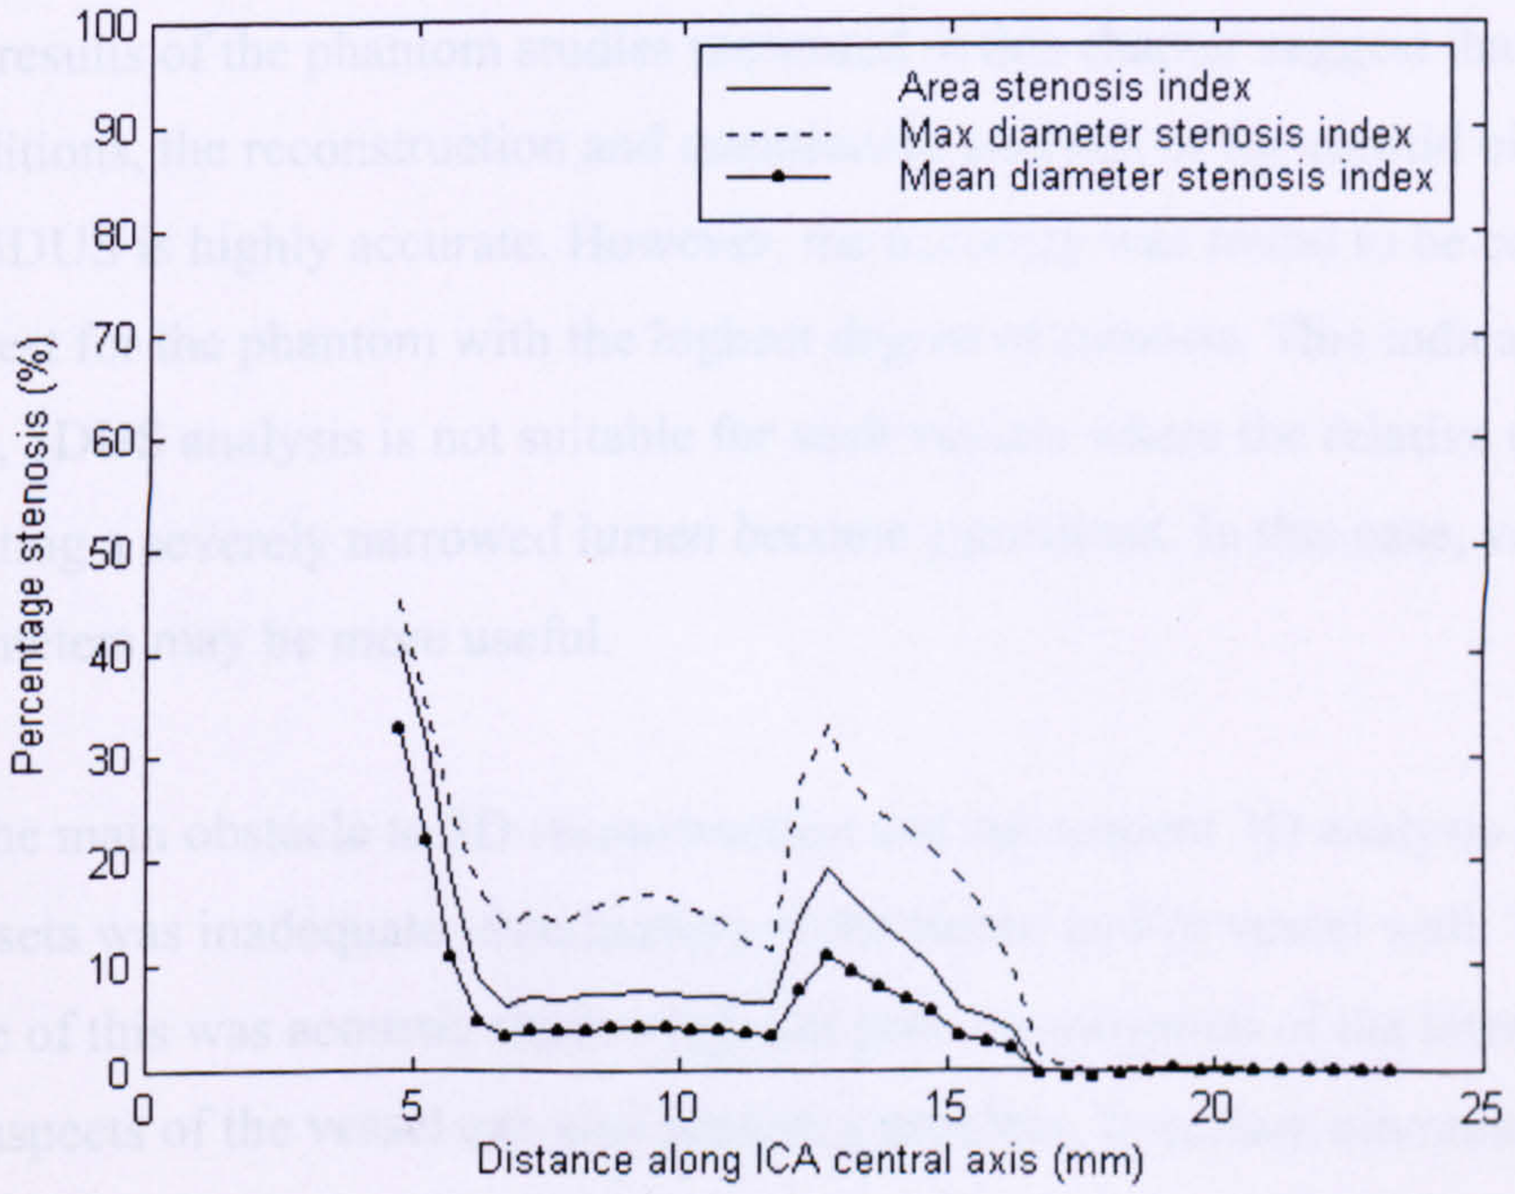


Figure 7.16 3D surface renderings of the reconstructed wall (left) and lumen (right) of the left carotid bifurcation for patient 4 (EK-L).



(a)



(b)

Figure 7.17 Results of 3D quantitative analysis for patient 4 (EK-L). (a) Plaque cross-sectional area as a function of length along the z-axis. (b) True transverse stenosis indices in the ICA.

Patient	Extent of disease (mm)	Plaque volume (mm^3)	Lumen surface area(mm^2)	Volumetric stenosis (%)	Max. true transverse % ICA area stenosis	Max. true transverse % ICA diameter stenosis (max. mean value)
PA-L	32.50	1120.76	734.17	56	75	57 (50)
PG-R	30.40	915.94	770.85	41	59	62 (51)
SA-R	27.98	781.93	770.46	42	88	74 (74)
EK-L	22.34	445.02	587.34	33	44	45 (33)

Table 7.3 Results of the quantitative analyses for 4 diseased carotid bifurcations. In the final column the value in parentheses refers to the maximum true transverse diameter stenosis calculated using the mean within-plane lumen diameter.

7.3 Discussion

The results of the phantom studies presented in this chapter suggest that, under ideal conditions, the reconstruction and quantitative analysis of the carotid bifurcation using 3DUS is highly accurate. However, the accuracy was found to be consistently poorest for the phantom with the highest degree of stenosis. This indicates that, perhaps, 3DUS analysis is not suitable for such vessels where the relative errors in delineating a severely narrowed lumen become significant. In this case, velocity-based parameters may be more useful.

The main obstacle to 3D reconstruction and subsequent 3D analysis of patient data sets was inadequate visualisation of the lumen and/or vessel wall. The principal cause of this was acoustic shadowing, but poor visualisation of the lateral and superior aspects of the vessel can also present a problem. In certain circumstances, such artefacts might be significantly reduced by using compound imaging (see Sections 1.3.7.4 and 3.5.3).

The requirement for manual image segmentation of cross-sectional contours for acquired transverse slices was a rate-limiting step, with sets taking at least 20 minutes to complete. However, in the patient examples presented here, more contours than necessary were defined. In practice, it is sufficient to reconstruct only the diseased portion of a vessel, and this may reduce the time required to perform a 3D analysis to less than 20 minutes in many cases.

A formal reproducibility study was not performed due to time constraints. Before this can be done, it is essential that successive analysis by the same or different observers are carried out in a consistent manner. This is easier when plaque is focal in nature, since quantitative measures, such as plaque volume, are confined to the diseased portion of the vessel. It is more difficult, however, for diffuse disease, and in this case the length of reconstructed and/or quantified vessel must be standardised. This is straightforward to implement using the central axis as a frame of reference.

Chapter 8

Conclusions and further work

The final chapter presents the conclusions for each of the main aspects of the work described in this thesis. The chapter ends with a brief discussion of future directions for the research.

8.1 Conclusions

8.1.1 3D ultrasound acquisition system

The 3DUS acquisition system developed for this thesis is a freehand system based on an electromagnetic sensor. A novel aspect of the system, compared to earlier designs is the utilisation digital image capture using a feature of the ATL HDI-5000 scanner. This feature allows pre-scan-converted image data to be stored in a cineloop and downloaded to a personal computer for off-line analysis. Previous systems have used earlier a VCR or frame grabber to capture frames from the video output port of the scanner. The digital system has the advantage that images are not subject to degradation. Images are also updated at the frame rate of the scanner, which on modern scanners can be much higher than the standard PAL or NTSC frame rates. This ensures that the image capture cycle is closely synchronised with scan-probe position and orientation measurements. The *HDILab* software, which controls downloading and post-processing images, allows several imaging parameters set, such as grey-scale and colour maps, and TGC, to be modified off-line. This permits a great deal of flexibility with regard to post-processing and optimisation of captured images before the segmentation and 3D reconstruction steps.

Because the 3DUS acquisition system does not require an external image capture device, and uses an inexpensive sensing device and personal computer, the component cost of the system (excluding the scanner) is low. A clinical ECG monitor, with a R-wave output trigger, was required to achieve ECG-gated acquisition. In future incarnations, however, it is intended that a modification to the HDI-5000 scanner will make available a scanner-generated R-wave trigger, thus obviating the need for a separate ECG monitor.

Although the susceptibility of EPOM devices to interference is well recognised, the experiments described in Chapter 4 revealed that this was not a concern for the EPOM device used in this project provided optimised device settings and scanning protocol were used. For the device tested, the RMS absolute error was estimated to be 0.2 mm (range: -0.7 to 0.5 mm) for position measurements, and 0.2° (range: -0.8 to 0.9°) for orientation measurements. These results are applicable to translations and rotations up to 90 mm and 30°, respectively, and, therefore, are appropriate for carotid imaging. These RMS errors are considerably lower than those quoted in the

manufacturers specifications, indicating that for limited translations and rotations compared with the operating range of the device, EPOM devices can be highly accurate in a clinical ultrasound environment. There are two provisos to this conclusion, however: first, a wooden examination couch was used as a precaution against distortion of pcBIRD measurements. Second, the device settings – in particular, the filters and measurement rate – as well as the scanning protocol were optimised on the basis of the results of noise tests within the environment. These tests revealed that changing the device settings arbitrarily with the intention of maximising the theoretical performance of the device is not advisable. For example, it was found that increasing the measurement rate to a high value introduced very large errors. The optimal parameter settings were found to include the default settings for the *pcBIRD*. Following optimisation, the accuracy of *pcBIRD* measurements was high using scan-probe movements appropriate to carotid artery scanning.

Very little modification to the clinical carotid scanning protocol was necessary to accommodate 3DUS investigations. The scanner could be placed close to the examination couch, within easy reach of the ultrasonographer, without causing errors in *pcBIRD* measurements. EPOM devices are generally considered to have an advantage over optical devices in that they do not require an unobstructed line-of-sight between the sensor and the reference. This is true, but it was noted from the experiments described in Chapter 5 that orientating the scan-probe, such the line-of-sight between the receiver and the transmitter was obstructed (by the scan-probe), resulted in poorer accuracy. For this reason, it was recommended that the scan-probe should be held so that the receiver faced the transmitter throughout a 3D acquisition sweep.

In spite of the excellent performance of the pcBIRD in the accuracy experiments carried out in the vascular laboratory of the Department of Clinical Pharmacology & Therapeutics at St. Mary's Hospital, replication of the 3D acquisition system in other laboratories without adequate validation does not guarantee the same robustness to measurement distortion. Interference might occur, for example, if a metallic examination couch is used in place of a wooden one. Different scan-probes may also give rise to unacceptable errors. From this point of view, EPOM devices are still inherently problematic. A solution to this problem might be to substitute the EPOM device with an alternative device that is not susceptible to environmental interference. The technology of position and orientation measurement devices is continually pro-

gressing, largely fuelled by motion capture applications within the computer graphics and virtual reality industries. Some commercial hybrid devices have recently appeared which offer an excellent compromise between cost and performance, and these may be ideally suited for 3DUS applications. Clearly, the accuracy of measurements produced by external registration devices determines the accuracy of 3D reconstruction and is, therefore, of prime importance. The use of high accuracy, low cost position tracking devices, which are insensitive to their environment, is likely to increase the popularity of freehand 3DUS systems as an inexpensive and flexible alternative to dedicated 3DUS scanners.

8.1.2 Image segmentation

Fast and reliable segmentation of acquired 2D images, or voxel-reconstructed 3D images, is currently the greatest challenge for 3DUS. To provide a degree of reliability, manual segmentation was adopted for the 3DUS data sets presented in this thesis. This was found to be limited in two respects: firstly, a significant investment of time was required to manually draw wall and plaque contours on acquired ultrasound slices – typically, 20-45 minutes per artery, depending on the extent of the disease. a few simple tools were incorporated into the segmentation software in order to simplify the segmentation task,. These include rapid elliptical contour fitting. A second limitation lies in the difficulty of visualising plaque and sometimes artery walls within acquired images. This is a direct consequence of common ultrasound artefacts, such as shadowing, but also the physical interaction between the ultrasound beam and tissue. For arteries, the latter has the effect that the anterior and lateral sides of the vessel are poorly visualised, with only the posterior wall producing a clear image. In many cases, combining several scan-views, each acquired using a different insonation angle, might improve visualisation, especially of plaque. A convenient method for achieving this, proposed in Chapter 2, would be to use within-plane B-mode compounding, although, at least using current commercial systems, colour or power Doppler information would be sacrificed.

The approach adopted for this thesis, where boundary contours are defined directly on a subset of the originally acquired ultrasound slices, avoids the voxel reconstruction step commonly carried out in other systems. Although voxel reconstruction followed by volume rendering is a useful automatic means of visualising 3D image

data, segmentation will inevitably still require manual contour delineation. This is typically done by first reslicing the voxel volume to obtain a series of 2D slices, and then segmenting each slice in turn. Therefore, from the point of view of quantification, voxel reconstruction prior to segmentation is likely to be of little advantage unless it makes segmentation easier. This might be the case if, for instance, 3D compounding is used, or if 3D segmentation methods are employed which perform better than a serial 2D segmentation approach.

Segmenting ultrasound images directly following acquisition has the advantage that artefacts peculiar to voxel reconstruction algorithms are avoided. This is important because such artefacts may affect subsequent quantification. Given the small size and irregularity of carotid plaque some post-reconstruction artefacts may be impossible to detect, and may significantly distort 3D measures.

Acquired ultrasound slices can be considered optimal because the ultrasonographer may to tilt the scan-probe during acquisition to obtain the 'best' transverse images. An interesting observation is that transverse images that are perceived as being optimal are actually most likely to be obtained from slightly 'off-transverse' views (in the sense that the image plane is not perpendicular to the vessel), as this view produces better definition of the lateral walls.

8.1.3 Surface reconstruction

One of the drawbacks of freehand 3DUS systems is that non-parallel image slices are captured. These require non-standard 3D reconstruction techniques, and a few algorithms for voxel or surface reconstruction from non-parallel slices can be found in the literature. The surface reconstruction method developed in this thesis adds to these and is applicable to both parallel and semi-parallel cross-sectional contours. In principle, the method can be adapted so that any 3D curve can be fitted to contour vertices. Smoothing splines were found to give good results, although it is important to bear in mind that the choice of the best smoothing parameters is rather subjective. This is an area that warrants investigation to further establish the effects of different smoothing parameters on the accuracy of reconstruction.

Correct 3D registration of cross-sectional contours is an important operation,

given that acquired slices can be easily misregistered by small patient movements during acquisition. Centroid registration proved an effective means of achieving this using the IWSS algorithm, although optimal choice of the parameters is again somewhat subjective; too much smoothing of the centroid curve will result in the vessel being artificially straightened, whilst too little will fail to adequately correct misregistered contours.

A merit of the longitudinal spline representation of the carotid bifurcation is efficiency, both in terms of computation time and storage. Using the present version of the reconstruction algorithm, the surface geometry of most diseased bifurcations can be approximated. As it stands, the algorithm is unable to reconstruct surfaces for which the surface splines do not have a monotonic z-component. So, large 'fish-bowl-shaped' ulcerations cannot be modelled, for example. Very divergent bifurcations are also problematic since it is generally not possible to obtain transverse vessel contours using a single scan sweep and oblique slices are more likely. It is likely, however, that these limitations may be resolved by further development of the algorithm.

The anthropomorphic phantom experiments presented in Chapter 7 demonstrate that, overall, the accuracy of carotid artery reconstructions is high. The absolute errors were found to be more or less uniform for all reconstructions. The magnitude of the relative reconstruction error increased, however, with the degree of stenosis; the 70% stenosis phantom produced the largest errors in cross-sectional area and circumference for the whole length of the vessel. This was attributed to the very narrow portion of the ICA lumen present in this phantom, which was both difficult to visualise and, by the same token, difficult to delineate accurately by hand. Further studies using the anthropomorphic phantoms are planned using a blood-mimicking fluid, pumped through the phantoms. This would enable the 3D reconstruction errors associated with power mode imaging to be estimated.

8.1.4 Quantification of carotid artery disease

The ease and speed with which carotid disease can be quantified using a 3D model of the wall and lumen surfaces means that a great deal of quantitative data can be de-

termined for a particular artery. The clinical utility of much of this data is unknown, but it is likely that new measures determined from 3D quantitative analysis might aid the identification of high-risk patients, individually or in combination,. The main advantages of the 3D quantification methods developed in this research are: (i) a larger search space – i.e. a very high number of measurements can be made at many different positions, thus increasing the chance of finding, for instance, the true minimum lumen diameter; (ii) well-established planar measurements can be found from true transverse planes which are precisely perpendicular to the central axis of a vessel; (iii) new quantitative measures, including plaque volume, volumetric stenosis, plaque surface area, and extent of carotid disease may be determined; and (iv) the *distribution* of disease can be characterised, for example, via a cross-sectional area stenosis profile.

The phantom experiments carried out suggest that, at least under ideal conditions, several indices of severity of carotid disease can be measured with high accuracy. An interesting observation was that, NASCET stenosis indices exhibited the largest errors. This is almost certainly due to the fact that the distal ICA diameter used in the expression is smaller than the diameter at the carotid bulb, where true-transverse diameter stenoses were calculated. Therefore, errors in minimal lumen diameter will result in larger percentage errors for the NASCET measures. Volumetric stenosis indices were found to be highly accurate in the phantom studies, but, in practice, it is important that the measure should be standardised with respect to the length of artery under consideration.

The work presented in this thesis demonstrates that accurate and comprehensive 3D quantification of carotid artery disease is feasible using 3DUS imaging. As commented above, however, the most challenging aspect of this remains the generation of a realistic 3D geometric model within a period which is acceptable to clinicians. How long is acceptable will, to a certain extent, depend on the effectiveness of a 3D analysis in identifying high risk patients. Currently, the extra time required to perform a 3DUS scan is minimal, since an ultrasonographer can easily switch to a 3D investigation. The 3D analysis can then be performed off-line at a later time, thus maximising the efficiency of the scanner usage. This step may take up to an hour, but if the information it provides is proven to be worthwhile in reducing the risk of stroke or TIA, then this time will be justifiable.

The methods developed in this thesis employ 3DUS as the chosen imaging modality because of its safety, low-cost, and ability to image soft-tissue. Although the segmentation and reconstruction techniques were designed specifically for 3DUS imaging, there is no reason why the same methods could not applied to images acquired using other techniques, such as MRI, since these processes are independent of the imaging technique.

8.2 Further work

8.2.1 Further validation of 3D ultrasound analysis

Before this can be done, however, it is essential to establish the errors associated with the technique. One approach to this would be to validate *in vivo* 3DUS measurements using another imaging 3D technique such as MRA, MRI, or CT. One of the problems with this is the choice of a suitable 'Gold Standard' modality as these techniques are subject to artefacts, finite resolution, and reconstruction errors. Also, assuming that adopting MRA or CT as a reference could be justified, these techniques only image the lumen. Nevertheless, comparison of lumen measurements alone is a useful exercise, as shown by the ultrasound phantom experiments described in Chapter 7. High resolution, 3D MRI might be the best choice, provided the lumen and vessel wall are distinguishable. Some researchers have compared *in vivo* plaque measurements determined from images with *ex vivo*, histological measurements following carotid endarterectomy. This method is, however, dependent on being able to remove the plaque in one piece and compensate for errors introduced by specimen shrinkage prior to fixing.

It is also very important to establish the reproducibility of 3DUS measurements both within and between observers. A formal reproducibility study was beyond the scope of this thesis, but work is already underway to find the inter- and intra-observer reproducibility of 3DUS measurements.

8.2.2 Clinical studies

Although the clinical feasibility of 3DUS quantification of carotid artery disease has been investigated in this thesis, the clinical utility of measures of the disease deter-

mined by 3D analysis has yet to be established. It is appropriate to answer the question of how 3D measures are related to patient risk using large clinical trials. In particular, the relative clinical utility of 2D versus 3D measures is of interest. To date, the ACST is the only study which includes stenosis measurements determined using ultrasound. Now that 3DUS techniques have been demonstrated, future trials are able to incorporate a range of 3D measures.

8.2.3 Image segmentation

The development of sophisticated segmentation tools is a critical area for further work. It is likely that the best approach to the problem of robust semi- or fully-automatic segmentation of blood vessels on ultrasound images will involve improved image quality as much as advances in image processing. Recent developments in ultrasound imaging technology have already delivered substantial improvements in image quality, and this trend appears to be set to continue. Possible solutions to the segmentation problem may employ a multi-mode approach, which harnesses information from a variety of ultrasound modes, including those which confer motion, such as TDI. With the advent of new tools for analysing 3D as well as 4D images, the clinical acceptance of 3D and 4D ultrasound imaging will undoubtedly grow.

8.2.4 Other applications

3DUS reconstructions of arteries may prove useful for detailed planning of vascular surgery. A pre-surgical inspection of a 3D surface model, like those presented in this thesis, may provide the surgeon with essential information needed to plan incision and/or clamping locations. This might have the effect of minimising risk to the patient by reducing the possibility of disturbing a plaque and producing unnecessary embolisation.

Another interesting application of 3DUS is in the field computational fluid dynamics (CFD) modelling of arterial blood flow. Realistic blood flow modelling requires 3D geometric data on the vessel, in addition to pressure and velocity data to specify boundary conditions. 3D MRA has been the most popular method for acquiring this data, but 3DUS is an inexpensive alternative which shows promise in this area [Augst 2001].

Appendix 1

3D matrix transformations

In the analysis that follows, a *left-handed* co-ordinate system is adopted as the convention (see Figure A1.1).

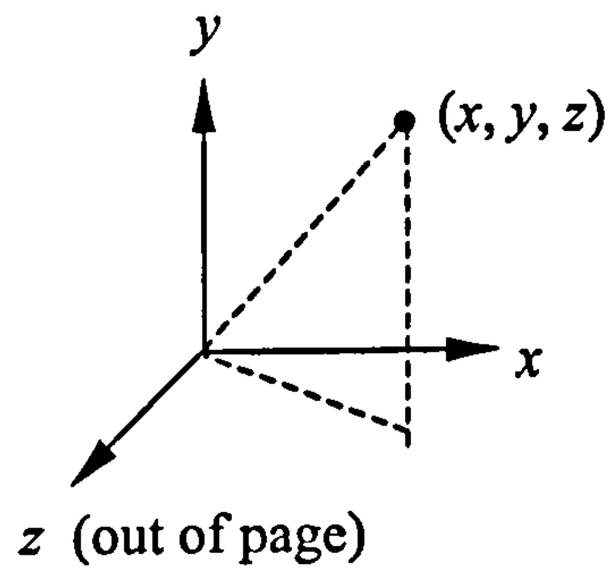


Figure A1.1 A left-handed 3D co-ordinate system with a point specified by the co-ordinates (x, y, z).

A point in 3D space with co-ordinates (x, y, z) within a co-ordinate system $\boldsymbol{C}$ is conveniently represented by a 3-element vector, $\mathbf{p}$:

$$\mathbf{p} = \begin{bmatrix} x \\ y \\ z \end{bmatrix}_C , \tag{A1.1}$$

where the subscript 'C' is used to denote the co-ordinate system relative to which the point is defined. Any 3D rigid-body transformation can be expressed as a matrix equation with the following form:

$$\mathbf{p}' = \mathbf{R}\mathbf{p} + \mathbf{s} , \tag{A1.2}$$

where $\mathbf{p}'$ is the vector specifying the new point following the transformation, $\mathbf{R}$ is a 3×3 rotation matrix, and $\mathbf{s}$ is a 3×1 translation vector. Any 3D rotation can be re-

duced to a sequence of basic rotations about each of the x-, y-, and z-axes. In particular, a *counterclockwise* rotation of α about the x-axis is given by the matrix:

$$\mathbf{R}_x = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\alpha & \sin\alpha \\ 0 & -\sin\alpha & \cos\alpha \end{bmatrix}; \tag{A1.3}$$

a counterclockwise rotation of β about the y-axis is given by:

$$\mathbf{R}_y = \begin{bmatrix} \cos\beta & 0 & -\sin\beta \\ 0 & 1 & 0 \\ \sin\beta & 0 & \cos\beta \end{bmatrix}; \tag{A1.4}$$

and a counterclockwise rotation of γ about the z-axis is given by:

$$\mathbf{R}_z = \begin{bmatrix} \cos\gamma & \sin\gamma & 0 \\ -\sin\gamma & \cos\gamma & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{A1.5}$$

Thus,

$$\mathbf{R} = \mathbf{R}_x\mathbf{R}_y\mathbf{R}_z. \tag{A1.6}$$

Substituting Equations (A1.3-5) into (A1.6) and multiplying leads to the general form:

.

$$\mathbf{R} = \begin{bmatrix} \cos\beta\cos\gamma & \sin\alpha\sin\beta\cos\gamma + \cos\alpha\sin\gamma & \sin\alpha\sin\gamma - \cos\alpha\sin\beta\cos\gamma \\ -\cos\beta\sin\gamma & \cos\alpha\cos\gamma - \sin\alpha\sin\beta\sin\gamma & \cos\alpha\sin\beta\sin\gamma + \sin\alpha\cos\gamma \\ \sin\beta & -\sin\alpha\cos\beta & \cos\alpha\cos\beta \end{bmatrix}$$

(A1.7)

Now **R** defines a counterclockwise rotation, α, about the x-axis; followed by a counterclockwise rotation, β, about the y-axis; followed by a counterclockwise rotation, γ,

about the z-axis.

The rotation and translation in Equation (A1.2) can be combined into a single matrix by the equation rewriting as

$$\mathbf{q}' = \mathbf{M}\mathbf{q}, \tag{A1.8}$$

where

$$\mathbf{q} = \begin{bmatrix} \begin{bmatrix} \mathbf{p} \end{bmatrix} \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}, \tag{A1.9}$$

and

$$\mathbf{M} = \begin{bmatrix} \begin{bmatrix} & \mathbf{R} & \end{bmatrix} & \begin{bmatrix} \mathbf{s} \end{bmatrix} \\ 0 \quad 0 \quad 0 & 1 \end{bmatrix}. \tag{A1.10}$$

The form of Equations (A1.9) and (A1.10) is called the homogeneous form, and it allows a general 3D rigid body transformation to be expressed by a single 4×4 transformation matrix comprising one rotation and one translation matrix.

The homogeneous matrix **M** provides a concise means of transforming points from one co-ordinate system to another. The notation adopted in this thesis is ${}^{C_2}\mathbf{M}_{C_1}$ to specify the transformation *from* co-ordinate system $\boldsymbol{C}_1$ *to* co-ordinate system $\boldsymbol{C}_2$. As an example, the *pcBIRD* electromagnetic tracking device used in this thesis consists of a receiver and a transmitter, each with its own 3D co-ordinate system. The orientation of the receiver co-ordinate system, $\boldsymbol{R}$, relative to the transmitter co-ordinate system, $\boldsymbol{T}$, is defined by:

1. a counterclockwise rotation, z, about the z-axis of $\boldsymbol{T}$, followed by
2. a counterclockwise rotation, y, about the y-axis of $\boldsymbol{T}$, followed by
3. a counterclockwise rotation, x, about the x-axis of $\boldsymbol{T}$.

Therefore, the rotation matrix is given by

$$
{}^{R}\mathbf{R}_{T} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos x & \sin x \\ 0 & -\sin x & \cos x \end{bmatrix} \begin{bmatrix} \cos y & 0 & -\sin y \\ 0 & 1 & 0 \\ \sin y & 0 & \cos y \end{bmatrix} \begin{bmatrix} \cos z & \sin z & 0 \\ -\sin z & \cos z & 0 \\ 0 & 0 & 1 \end{bmatrix}
$$

$$
= \begin{bmatrix} \cos y \cos z & \cos y \sin z & -\sin y \\ -\cos x \sin z + \sin x \sin y \cos z & \cos x \cos z + \sin x \sin y \sin z & \sin x \cos y \\ \sin x \sin z + \cos x \sin y \cos z & -\sin x \cos z + \cos x \sin y \sin z & \cos x \cos y \end{bmatrix}
$$

(A1.12)

and is the rotation matrix output by the pcBIRD. The matrix ${}^{T}\mathbf{R}_{R}$ is actually more useful, and is the matrix transpose of ${}^{R}\mathbf{R}_{T}$. The position of the receiver relative to the transmitter is output as 3 co-ordinates, (X, Y, Z). Hence, the homogeneous transformation matrix relating the co-ordinate systems *R* to *T* is:

$$
{}^{T}\mathbf{M}_{R} = \begin{bmatrix} \begin{bmatrix} & \left({}^{R}\mathbf{R}_{T}\right)^{\mathrm{T}} & \end{bmatrix} & \begin{bmatrix} \mathrm{X} \\ \mathrm{Y} \\ Z \end{bmatrix} \\ \begin{matrix} 0 & 0 & 0 \end{matrix} & 1 \end{bmatrix} =
$$

$$
\begin{bmatrix} \cos y \cos z & -\cos x \sin z + \sin x \sin y \cos z & \sin x \sin z + \cos x \sin y \cos z & \mathrm{X} \\ \cos y \sin z & \cos x \cos z + \sin x \sin y \sin z & -\sin x \cos z + \cos x \sin y \sin z & \mathrm{Y} \\ -\sin y & \sin x \cos y & \cos x \cos y & \mathrm{Z} \\ 0 & 0 & 0 & 1 \end{bmatrix}
$$

(A1.13)

Appendix 2

Miscellaneous algorithms

A2.1 The modified-Marquardt algorithm

The modified-Marquardt algorithm finds the least-squares solution to a system of non-linear equations [Nash 1990]. The essence of the method is to use the function of unknown parameters to generate a cost function which can be minimised subject to a set of real data. Then, the parameters which minimise the cost function are those for which the model function most closely fits the real data. For a non-linear problem, a suitable minimisation function, S, for a set of N non-linear functions, $f_i(\mathbf{q})$, $i = 1,2,\ldots N$, is the sum of the squares:

$$S(\mathbf{q}) = \sum_{i=1}^{N} f_i(\mathbf{q})^2 , \qquad \text{(A2.1)}$$

where $\mathbf{q}$ is the parameter vector. Collecting together the functions to form an N element vector function, $\mathbf{f}(\mathbf{q})$, gives:

$$\mathbf{f}(\mathbf{q}) = \begin{bmatrix} f_1(\mathbf{q}) \\ f_2(\mathbf{q}) \\ \vdots \\ f_N(\mathbf{q}) \end{bmatrix}. \qquad \text{(A2.2)}$$

Now, Equation (A2.1) can be rewritten as:

$$S(\mathbf{q}) = \mathbf{f}^T\mathbf{f} . \qquad \text{(A2.3)}$$

When the aim is to fit a model to a set of real data, the function $\mathbf{f}$ is called the *residual* and is defined as:

$$\mathbf{f}(\mathbf{q}) = \mathbf{g}_{\text{model}}(\mathbf{q}) - \mathbf{g}_{\text{real}}, \tag{A2.4}$$

where $\mathbf{g}_{\text{model}}(\mathbf{q})$ is the model function vector which is to be fitted to the real data contained in the vector $\mathbf{g}_{\text{real}}$. The modified-Marquardt algorithm is an iterative algorithm which produces an estimate for the parameter vector, $\mathbf{q}$, such that the $S(\mathbf{q})$ is minimised. On each iteration the algorithm updates the current (i^{th}) parameter vector, $\mathbf{q}_i$, according to the following scheme:

$$\mathbf{q}_{i+1} = \begin{cases} \mathbf{q}_i + \delta\mathbf{q}_i & \text{if } S(\mathbf{q}_{i+1}) < S(\mathbf{q}_i), \\ \mathbf{q}_i & \text{otherwise.} \end{cases} \tag{A2.5}$$

where $\delta\mathbf{q}_i$ is found by solving the following equation:

$$(\mathbf{J}^{\mathrm{T}}\mathbf{J} + e_i\mathbf{D}^2)\delta\mathbf{q}_i = -\mathbf{J}^{\mathrm{T}}\mathbf{f}. \tag{A2.6}$$

Here, $\mathbf{J} = \mathbf{J}(\mathbf{q}_i)$ is the Jacobian of $\mathbf{f}$, defined as:

$$\mathbf{J}_{\mathrm{jk}}(\mathbf{q}) = \frac{\partial f_{\mathrm{j}}(\mathbf{q})}{\partial q_{\mathrm{k}}}, \tag{A2.7}$$

where $\mathbf{J}_{\mathrm{jk}}$ is the matrix element in the j^{th} row and k^{th} column, and q_{k} is the k^{th} parameter of the parameter vector $\mathbf{q}$. $\mathbf{D}^2$ in Equation (A2.6) is a diagonal matrix defined by:

$$\mathbf{D}^2_{\mathrm{jk}} = \begin{cases} (\mathbf{J}^{\mathrm{T}}\mathbf{J})_{\mathrm{jk}} + \varepsilon & \text{if } \mathrm{j} = \mathrm{k}, \\ 0 & \text{if } \mathrm{j} \neq \mathrm{k}. \end{cases} \tag{A2.8}$$

where ε is a positive scalar which avoids zero diagonal elements in $\mathbf{D}^2$. A typical value of ε is 1.0. The scalar, e_i, in Equation (A2.6) is positive and updated on each iteration to stabilize convergence of the algorithm.

Since the matrix $(\mathbf{J}^{\mathrm{T}}\mathbf{J} + e_i\mathbf{D}^2)$ is positive definite, Cholesky decomposition and backsubstitution are used to solve Equation (A2.6) efficiently. The value of $S(\mathbf{q}_i)$ is used to monitor the convergence of the algorithm: if $S(\mathbf{q}_i+\delta\mathbf{q}_i) < S(\mathbf{q}_i)$ then the updated parameter is a better parameter estimate than $\mathbf{q}_i$. In this case, the updated

parameter, $\mathbf{q}_i+\delta\mathbf{q}_i$, is used as the estimate in the next iteration and the current value of e_i is decreased by some scale factor, say c^+. If $S(\mathbf{q}_i+\delta\mathbf{q}_i) \geq S(\mathbf{q}_i)$, however, then $\mathbf{q}_i$ is used in the next iteration and e_i is increased by some scale factor, c^-. Typical values for c^+ and c^- are 10.0 and 0.4, respectively, and typical starting value for e_0 is 0.001.

A2.2 Least-squares fitting two 3D point sets

A Matlab implementation of the algorithm described by Arun et al. [1987] to find the least-squares 3D transformation from one point set to another is given below. The Matlab function accepts two $3 \times N$ matrices, *P1* and *P2*, which contain the 3D position vectors for the two point sets. A 3×3 rotation matrix, *R*, and a 3×1 translation vector, *T*, are returned, such that $P2 \approx R*P1 + T$.

```
function [R,T] = LMS3DFit(P1,P2)

   c2 = mean(P1,2); c1 = mean(P2,2);
   Q1 = P1-c1; Q2 = P2-c2;

   H = Q2*Q1';
   [U,D,V] = svd(H);
   R = V'*U';
   T = c1 - R*c2;
```

A2.3 Direct least-squares ellipses fitting

Below is a Matlab implementation of the algorithm of Fitzgibbon et al. [1999] for direct, least-squares fitting of an ellipse to 6 or more points. The input points are contained in the $2 \times n$ matrix, *P*, and *N* is the number of points required in the fitted ellipse, returned in the $2 \times N$ matrix, *P_ellipse*.

```
function c = ncols(A)
[r,c] = size(A);


function P_ellipse = LMSEllipseFit(P, N)

if ncols(P) < 6
   warning('Not enough points to find an ellipse!');
   P_ellipse = [];
   return;
end

xx = P(1,:)'; yy = P(2,:)';
```

```
D = [xx.*xx xx.*yy yy.*yy xx yy ones(size(xx))];
S = D'*D;
C(6,6) = 0; C(1,3) = 2; C(2,2) = -1; C(3,1) = 2;

[gevec, geval] = eig(inv(S)*C);
[PosR, PosC] = find(geval > 0 & ~isinf(geval));
a = gevec(:,PosC);

A = a(1); B = a(2); C = a(3);
D = a(4); E = a(5); F = a(6);
Coeffs = [A B C D E F];

U = [-2*A -B; -B -2*C];
offset = inv(U)*[D;E];
cx = offset(1); cy = offset(2);
F = -(A*cx*cx + B*cx*cy + C*cy*cy + D*cx +E*cy +F);

Q = [A B/2; B/2 C];
[R,D] = eig(Q);

T = R'*Q*R;
alphaT = sqrt(F/T(1,1));
betaT = sqrt(F/T(2,2));

theta = linspace(0,2*pi,N+1);
pts = [alphaT*cos(theta); betaT*sin(theta)];

q = R*pts;
P_ellipse = q + [cx cy]'*ones(1,ncols(q));
```

A2.4 Iterative weighted-smoothing-spline

A Matlab function to generate an IWSS, described in Chapter 6, is reproduced below. The function requires the Matlab Spline Toolbox (version 3 or higher). The following input parameters are used: s is the arc-length parameter vector; C is a $3 \times N$ vector containing the sample points; p is the smoothing parameter in the interval [0, 1]; k is an exponent defined in Section 6.4.1.3; *W_first* is the weight for the first point; and *W_last* is the weight last point (*W_first* and *W_last* are generally set to high values, e.g. 100). The fitted spline is returned in PP-form in the structure *pp*.

```
function v = magc(M)

v = sqrt(sum(M.^2,1));


function pp = IWSSfit(s,C,p,k,W_first,W_last)

W = [W_first];

for n = 2:ncols(C)-1
   Cnew = C(:,[1:n-1, n+1:end]);
   snew = [0 cumsum(magc(diff(Cnew,1,2)))];
```

```
   pp_new = csaps(snew, Cnew, p, [], [W ones(1,ncols(C)-ncols(W)-
1)]);
   D = magc( C(:,n) - fnval(pp_new,s(n)) );
   W(1,n) = 1/(D^k);
end

W = [W W_last];

pp = csaps(s,C,p,[],W);
```

Appendix 3

A software tool for the manual segmentation of ultrasound images

---

The purpose of the software tool described is to enable cross-sectional contours to be manually defined on a set of ultrasound images. The main steps in the segmentation process are listed below. It is assumed throughout that the acquired images are obtained using B-mode with colour or power Doppler overlay, although it is perfectly possible to acquire B-mode images alone. This may be disadvantageous in certain situations, however; for instance, when echolucent plaque is present.

1. The cineloop containing the acquired ultrasound images is opened in the *HDILab* software provided by the manufacturers of the scanner.

2. The cineloop is 'trimmed' to remove frames captured before the first pcBIRD measurement is made. This is done by identifying the first image captured following the initial burst of trigger pulses shown in the ECG window (see Chapter 3 for details of the trigger sequence).

3. If necessary, the image rendering parameters may be modified so that images are optimised throughout the cineloop. It is possible, for example, to change the greyscale brightness, or decrease the sensitivity of the colour image to reduce obvious colour artifacts.

4. The cineloop is then saved in a movie file format (uncompressed AVI), with and without the colour pixels overlaid, thus producing two files. This format allows the captured images to be loaded *en bloc* into the segmentation software. Separating the B-mode and colour components of the images allows the colour overlay to be turned off during contour delineation

5. Both movie files are then opened using the segmentation software. By default, the first image is displayed with colour overlay, and the user can view any image within the sequence by using the forward and backward view buttons. In most cases, turning the colour overlay off is recommended before defining contours. The reasons for this are the relatively poor spatial resolution of, and the possibility of colour/power Doppler artifacts obscuring image features. In the case of echolucent plaque, however, the lumen of the artery may be invisible without the aid of colour information. Nevertheless caution must be exercised in the interpretation of such images.

6. The user then scrolls to an image and selects either to define a wall or a plaque contour. Once defined, a contour must be tagged with the name of the identifying artery (i.e. CCA, ICA, or ECA). This may be done at any time when the contour is active: i.e. during the definition of boundary points, or later, by selecting the contour.

7. Contours are defined for a number of images in the set, and it is not strictly necessary to draw contours on every acquired image. For example, in regions where contours change little between consecutive slices, or where the separation between slices is very small, images may be skipped.

8. When all contours have been defined, image contour points are saved in multiple data files for further processing. The co-ordinates (in mm) of each contour point are saved in a set of files, each having the '.ctr' file extension, and corresponding to an a single image slice. A separate file ('.ctl') contains the pixels co-ordinates specified by user for all slices. This file enables contours to be edited at a later date. The '.ctr' files may then be loaded directly into the reconstruction program.

The GUI for the segmentation software is shown in Figure A3.1.

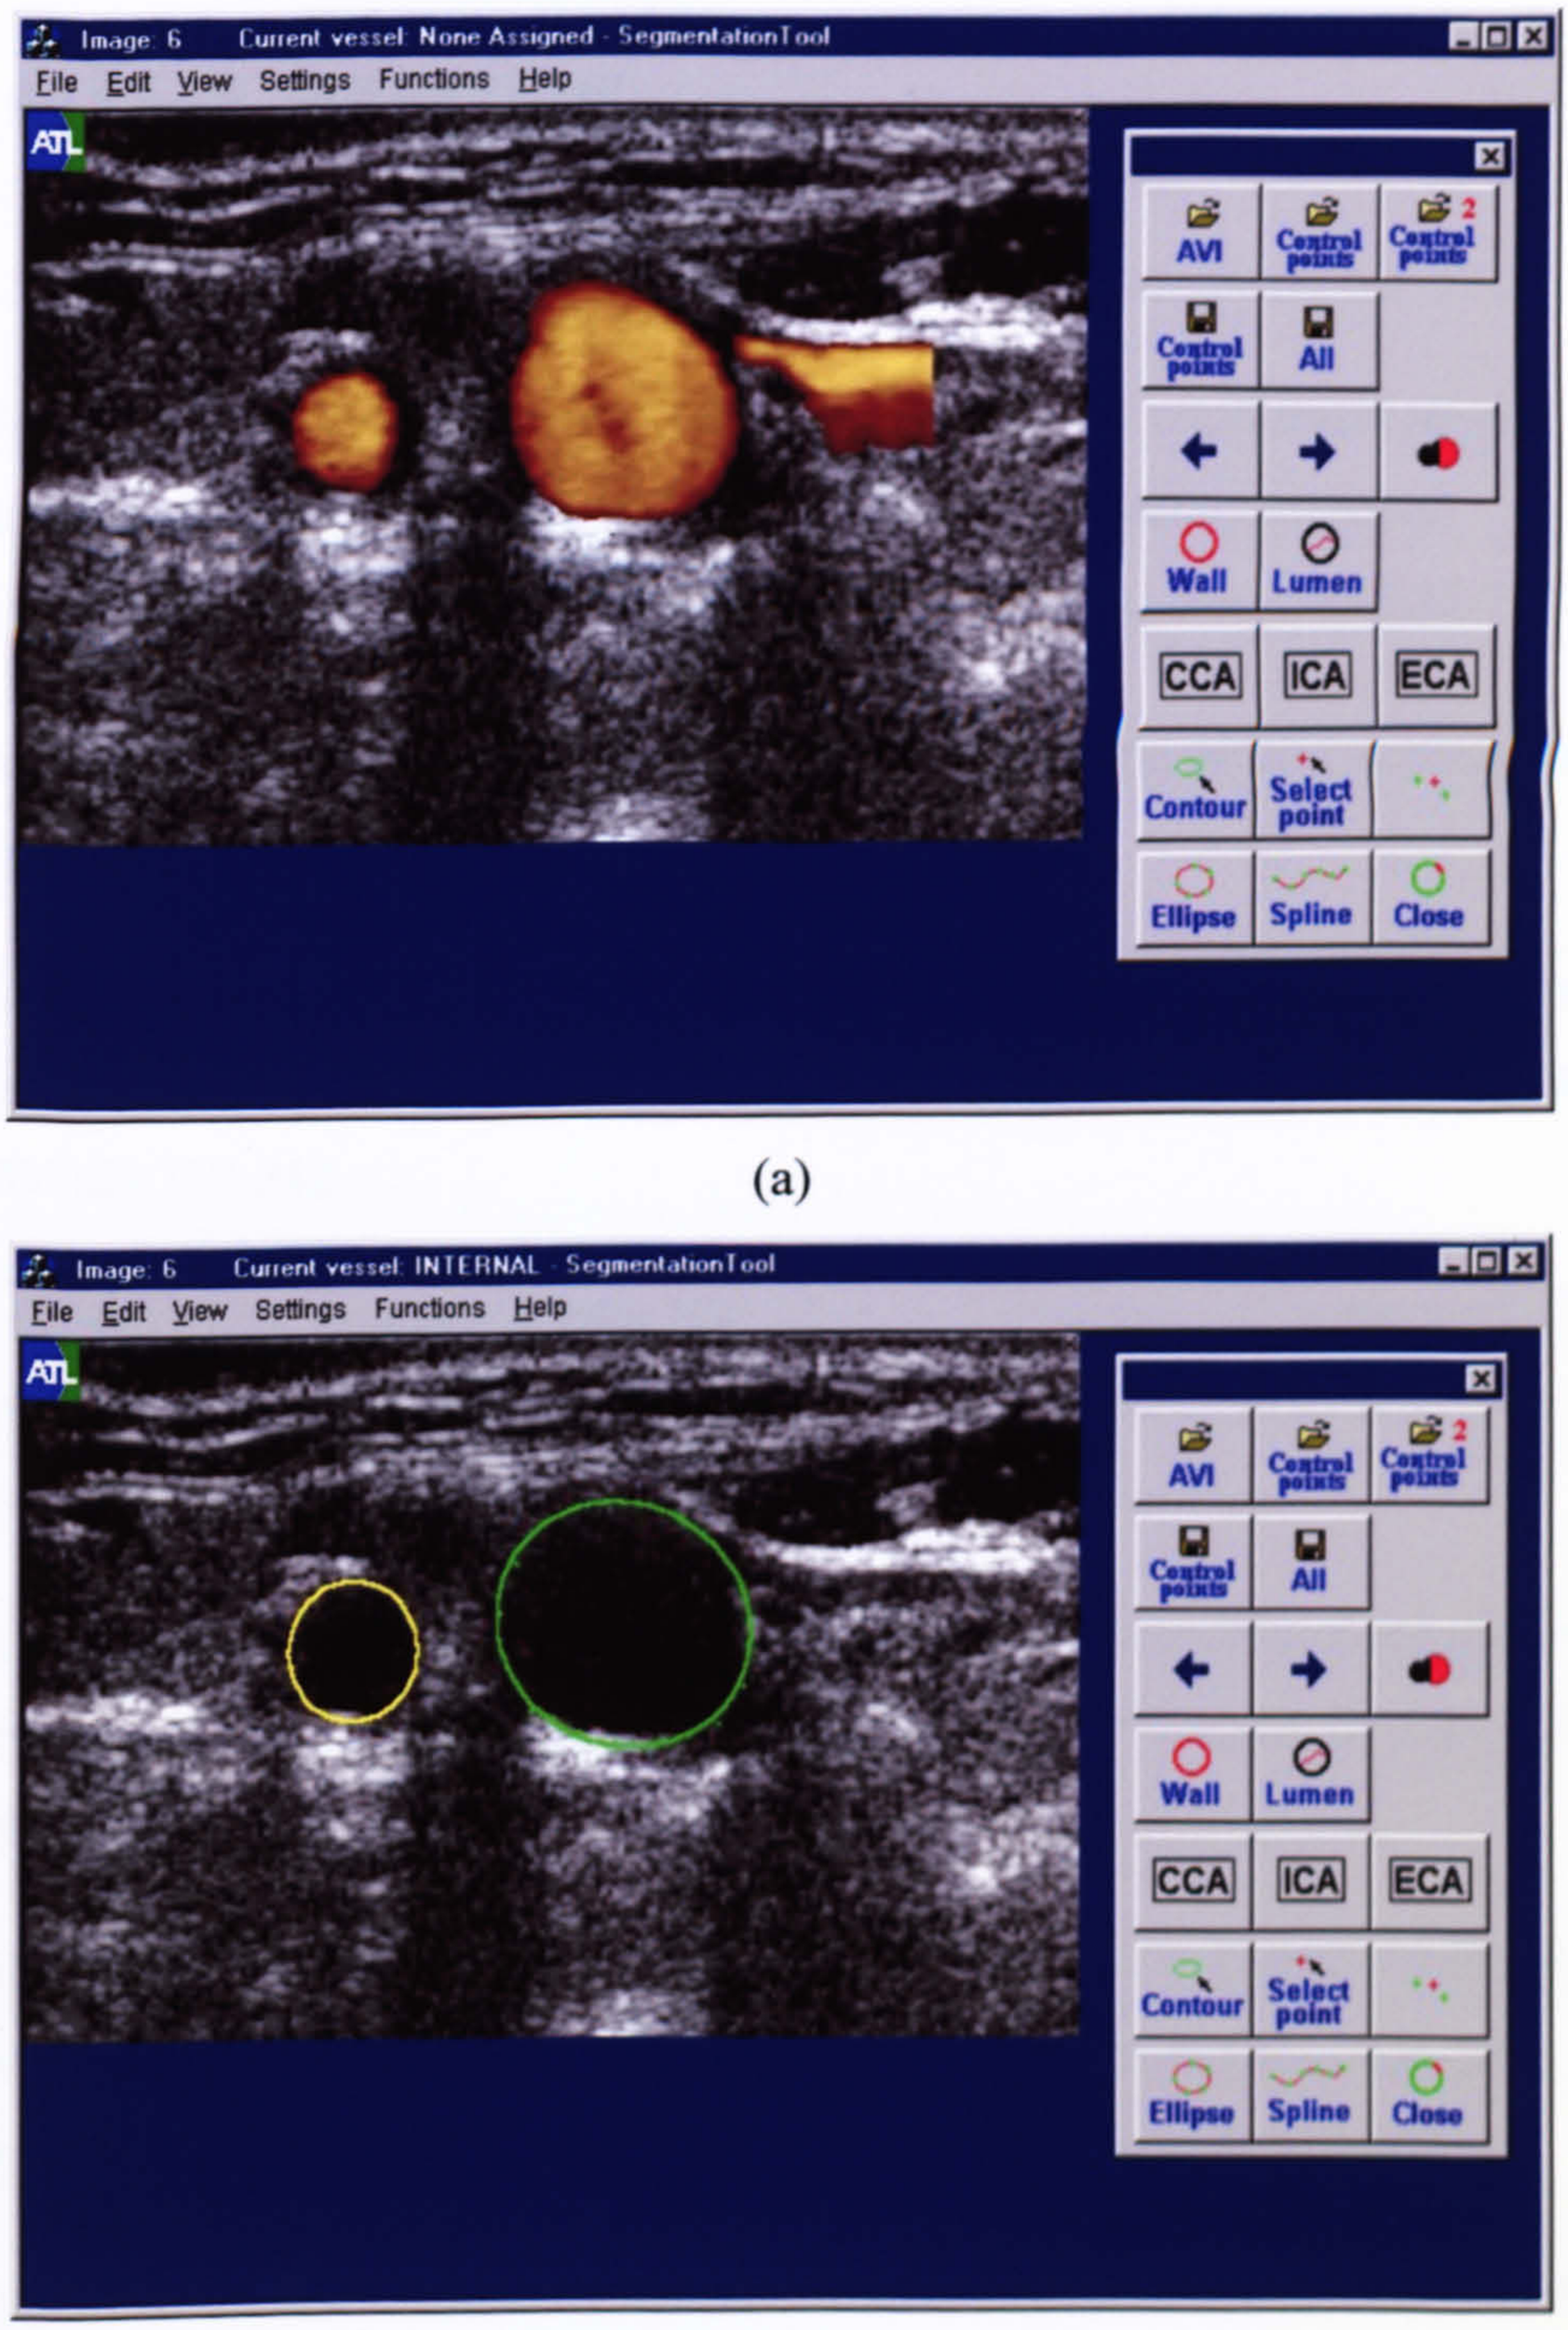


Figure A3.1 The GUI of the segmentation software used to define boundary contours on a series of transverse ultrasound images. (a) shows the user interface with colour overlay on (power Doppler in this case); (b) shows the interface with the colour overlay turned off, and wall contours drawn.

References

AbuRahma AF, Richmons B, Robinson P, et al. Effect of contralateral severe stenosis or carotid occlusion on duplex criteria of ipsilateral stenoses: comparative study of various duplex parameters. *J Vasc Surg* 1995;22:751-762.

Alexandrov A, Bladin C, Maggisano R, Norris J. Measuring carotid stenosis: time for reappraisal. *Stroke* 1993;24:1292-1296.

Alexandrov A, Brodie D, McLean A, Hamilton P, et al. Correlation of peak systolic velocity and angiographic measurement of carotid stenosis revisited. *Stroke* 1997;28:339-342.

Allot C, Barry CD, Bramley J, John NW, et al. Noninvasive 3-D ultrasound of atherosclerotic plaques in the Watanabe rabbit. *Ultrasound Med Biol* 1998;25:905-914.

Allott C, Barry C, Pickford R, Waterton J. Volumetric assessment of carotid artery bifurcation using freehand-acquired, compound 3D ultrasound. *Br J Radiol* 1999;72:289-292.

Arbeille P, Bouin-Pineau M, Herault S. Accuracy of the main Doppler methods for evaluating the degree of carotid (continuous wave, pulsed wave, and color Doppler). *Ultrasound Med Biol* 1999;25:65-73.

Arbeille P, Desombre C, Aesh B, Philippot M, et al. Quantification and assessment of carotid artery lesions: degree of stenosis and plaque volume. *J Clin Ultrasound* 1995;23:113-124.

Arbeille P, Eder V, Casset D, Quillet L, et al. Real-time 3-D ultrasound acquisition and display for cardiac volume and ejection fraction evaluation. *Ultrasound Med Biol* 2000;26:201-208.

Arroyo L and Lee R. Mechanisms of plaque rupture: mechanical and biologic interactions. *Cardiovascular Research* 1999;41:369-375.

Arun KS, Huang TS, Blostein SD. Least-squares fitting of two 3-D point sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1987;PAMI-9:698-700.

Athanasiou S, Khullar V, Cardazo L. Three-dimensional ultrasound in urogynecology. In: Baba K and Jukovic K, ed. *Three-dimensional ultrasound in obstetrics and gynecology*. New York: The Parthenon Publishing Group Inc, 1997:95-105.

Augst A, Barratt D, Hughes AD, Thom SA, et al. CFD model of a human carotid artery bifurcation reconstructed from 3D ultrasound data. *Proceedings of the 5th International Symposium on Computer Methods in Biomechanics and Biomedical Engineering, Rome* 2001;

Barequet G and Sharir M. Piecewise-linear interpolation between polygonal slices. *Computer Vision and Image Understanding* 1996;63:251-272.

Barratt DC. Volume measurement using 3D ultrasound. *Euroson School on 3D Ultrasound Imaging* 2001;

Barry CD, Allott CP, John NW, Mellor PM, et al. Three-dimensional freehand ultrasound: image reconstruction and volume analysis. *Ultrasound Med Biol* 1997;23:1209-1224.

Beckett W, Davis P, Hoffman J. Duplex Doppler sonography of the carotid artery: False-positive results in an artery contralateral to an artery with marked stenosis. *American Journal Of Roentgenology* 1990;155:1091-1095.

Beebe HG, SallesCunha SX, Scissons RP, Dosick SM, et al. Carotid arterial ultrasound scan imaging: A direct approach to stenosis measurement. *J Vasc Surg* 1999;29:838-844.

Belohlavek M, Foley D, Gerber T, Greenleaf J, et al. Three-dimensional ultrasound imaging of the atrial septum: normal and pathologic anatomy. *J Am Coll Cardiol* 1993;22:1673-1678.

Binder T, Globits S, Zangeneh M, Gabriel H, et al. Three-dimensional echocardiography using a transoesophageal imaging probe. Potential and technical considerations. *Eur Heart J* 1996;17:619-628.

Birkfellner W, Watzinger F, Wanschitz F, Enislidis G, et al. Systematic distortions in DC-pulsed electromagnetic position sensing devices. *Med Phys* 1998a;25:2242-2248.

Birkfellner W, Watzinger F, Wanschitz F, Ewers R, et al. Calibration of tracking systems in a surgical environment. *IEEE Trans Med Imaging* 1998b;17:737-742.

Bladin C, Alexandrov A, Murphy J, Maggisano R, et al. Carotid stenosis index: a new method of measuring internal carotid artery stenosis. *Stroke* 1995;26:230-234.

Bolle R and Vemuri B. On three-dimensional surface reconstruction methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1991;13:1-13.

Bonnefous O. Blood flow and tissue motion with ultrasound for vascular applications. *Comptes Rendus de l'Academie des Sciences (Series IV, Physics)* 2001;2:1161-1178.

Browman M, Cooperberg P, Harrison P, Marsh J, et al. Duplex ultrasonography criteria for internal carotid stenosis of more than 70% diameter reduction: angiographic correlation and receiver operating characteristic curve analysis. *Can Assoc Radiol J* 1995;46:291-295.

Bull A, Berkshire F, Amis A. Accuracy of an electromagnetic measurement device and application to the measurement and description of knee joint motion. *Proc Instn Mech Engnrs* 1998;212:347-355.

Burdin V, Roux C, Lefevre C, Stindel E. Modelling and analysis of 3-D elongated shapes with applications to long bone morphometry. *IEEE Trans Med Imaging* 1996;15:79-91.

Burns P, Powers J, Simpson D, Uhlendoft V, et al. Harmonic imaging - principles and preliminary results. *Angiology* 1997;47:563-573.

Byrd S, Robless P, Baxter A, Emson M, et al. Carotid duplex ultrasonography: importance of standardisation. *International Angiology* 1998;17:248-254.

Caplan L. *Stroke: A Clinical Approach*. Boston: Butterworth-Heinemann, 1993.

Carpenter J, Lexa F, Davis J. Determination of Duplex Doppler ultrasound criteria appropriate to the North American Symptomatic Carotid Endarterectomy Trial. *Stroke* 1996;27:695-699.

Chandran K, Vonesh M, Roy A, et al. Computation of vascular flow dynamics from intravascular ultrasound images. *Med Eng Phys* 1996;18:295-304.

Chang F, Hsu K, Ko H. Fetal heart volume assessment by three-dimensional ultrasound. *Ultrasound Obstet Gynecol* 1997;9:942-948.

Cheng G, Loree H, Kamm R, Fishbein M, et al. Distribution of circumferential stress in ruptured and stable atherosclerotic leasions: a structural analysis with histopathologic correlation. *Circulation* 1993;87:1179-1187.

Chikos P, Fisher L, Hirsch J, et al. Observer variability in evaluating extracranial carotid artery stenosis. *Stroke* 1983;14:885-892.

Coppini G and Poli R. Recovery of 3-D shape of the left ventricle from echocardiographic images. *IEEE Trans Med Imaging* 1995;14:301-317.

Correas J, Bridal L, Lesavre A, Méjean A, et al. Ultrasound contrast agents: properties, principles of action, tolerance, and artifacts. *European Radiology* 2001;11:1316-1328.

Curtis W, Janin A, Zikan K. A note on averaging rotation. *IEEE Virtual Reality Annual International Symposium* 1993;377-385.

Czerwinski R, Jones D, O'Brien W. An approach to boundary detection in ultrasound imaging. *1993 Ultrasonics Symposium* 1993;951-955.

Davies K and Humphrey P. Complications of cerebral angiography in patients with symptomatic carotid territory ischaemia screened by carotid ultrasound. *Journal of Neurology, Neurosurgery and Psychiatry* 1993;56:967-972.

de Boor C. *A Practical Guide to Splines*. New York: Springer-Verlag, 1978.

Deane C. Extended field-of-view and B-flow ultrasound: fashion or future? *Ultrasound Obstet Gynecol* 2000;15:96-97.

DeBakey M, Lawrie G, Glaeser D. Patterns of atherosclerosis and their surgical significance. *Ann Surg* 1985;201:115-131.

Delabays A, Sugeng L, Pandian N, et al. Dynamic three-dimensional echocardiographic assessment of intracardiac blood flow jets. *Am J Cardiol* 1995;76:1053-1058.

Delcker A and Tegeler C. Influence of ECG-triggered data acquisition on reliability for carotid plaque volume measurements with a magnetic sensor three-dimensional ultrasound system. *Ultrasound Med Biol* 1998;24:601-605.

Delcker A and Diener H. 3D-ultrasound of the carotid arteries. *Eur J Ultrasound* 1994a;1:337-344.

Delcker A and Diener H. Quantification of atherosclerotic plaques in carotid arteries by three-dimensional ultrasound. *Br J Radiol* 1994b;67:672-678.

Delcker A and Polz H. 3D ultrasound of the carotid artery using a sensor as a space orientation. *Cerebrovasc Dis* 1996;6(Suppl. 3):2-

Delcker A and Turowski B. Diagnostic value of three-dimensional transcranial contrast duplex sonography. *J Neuroimaging* 1997b;7:139-144.

Delcker A, Walter F, Caress J, Hunt C, et al. In-vitro measurement of muscle volume with 3-D ultrasound. *J Neuroimaging* 1997a;7:241-241.

Deng J. Novel technique for three-dimensional visualization and qualification of deformable, moving soft tissue body parts. *Lancet* 2000;356:127-131.

Deng J, Gardener JE, Rodeck CH, Lees WR. Fetal echocardiography in three and four dimensions. *Ultrasound Med Biol* 1996;22:979-986.

Derdeyn C, Powers W, Moran C, Cross D, et al. Role of Doppler US in screening for carotid atherosclerotic disease. *Radiology* 1995;197:635-643.

Detmer PR, Bashein G, Hodges T, Beach KW, et al. 3D ultrasonic image feature localization based on magnetic scanhead tracking - in-vitro calibration and validation. *Ultrasound Med Biol* 1994;20:923-936.

Devonald K, Ellwood D, Griffiths K. Volume imaging: Three-dimensional appreciation of the fetal head and face. *J Ultrasound Med* 1995;14:919-925.

Dhanjil S, Jameel M, Nicolaides A, et al. Ratio of peak systolic velocity of internal carotid to end diastolic velocity of common carotid: new duplex criteria for grading internal carotid stenosis. *J Vasc Tech* 1997;21:237-240.

Dierckx P. *Curve and Surface Fitting with Splines*. Oxford: Oxford University Press, 1993.

Dillon E, van Leeuwen M, Fernandez M, Eikelboom B, et al. Computed tomographic angiography for carotid imaging. *Lancet* 1992;340:1286-1286.

Downey D and Fenster A. Vascular imaging with a three-dimensional power Doppler system. *American Journal Of Roentgenology* 1995a;165:665-668.

Downey D and Fenster A. Three-dimensional power Doppler detection of prostatic cancer. *American Journal Of Roentgenology* 1995b;165:741-741.

Downey D, Nicolle D, Fenster A. Three-dimensional orbital ultrasonography. *Can J Ophthalmol* 1995c;30:395-398.

Downey D, Nicolle D, Levin M, Fenster A. Three-dimensional ultrasound imaging of the eye. *Eye* 1996;75-81.

Drebin RA, Carpenter L, Hanrahan P. Volume rendering. *Computer Graphics* 1988;18:65-71.

Dumoulin C. Phase contrast MR angiography techniques. *Magn Reson Imaging Clin N Am* 1995;3:399-411.

Dunmire B, Beach KW, Labs K-H, Plett M, et al. Cross-beam vector Doppler ultrasound for angle independent veloctiy measurements. *Ultrasound Med Biol* 2000;26:1213-1235.

Dzau V and Gibbons G. Vascular remodelling: mechanisms and implications. *Journal of Cardiovascular Pharmacology* 1993;21:

Edelman R, Mattle H, Wallner B, Bajakian R, et al. Extracranial carotid arteries: evaluation with "black blood" MR angiography. *Radiology* 1990;177:45-50.

Edwards WS, Deforge C, Kim Y. Interactive three-dimensional ultrasound using a programmable multimedia processor. *International Journal Of Imaging Systems And Technology* 1998;9:442-454.

Elatrozy T, Nicolaides A, Tegos T, Griffin M. The objective characterisation of ultrasonic carotid plaque features. *Eur J Vasc Endovasc Surg* 1998;223-230.

Elgersma O, van Leersum M, Buijs P, van Leeuwen M, et al. Changes over time in optimal duplex threshold for the identification of patients eligible for carotid endarterectomy. *Stroke* 1998;29:2352-2356.

Eliasziw M, Smith R, Singh N, Holdsworth D, et al. Further comments on the measurement of carotid stenosis from angiograms. *Stroke* 1994;25:2445-2449.

Elvins T. A survey of algorithms for volume visualization. *Computer Graphics* 1992;26:194-201.

European Carotid Surgery Trial Collaborators. MRC European Carotid Surgery Trial: interim results for symptomatic patients with severe (70-99%) or with mild (0-29%) carotid stenosis. *Lancet* 1991;337:1235-1243.

European Carotid Surgery Trialists' Collaborative Group. Randomised trial of endarterectomy for recently symptomatic carotid stenosis: final results of the MRC European Carotid Surgery Trial. *Lancet* 1998;351:1379-1386.

Evans D and McDicken W. *Doppler Ultrasound: Physics, Instrumentation and Signal Processing*. Chichester: John Wiley, 2000.

Eviatar H and Somorjai R. A fast, simple active contour algorithm for biomedical images. *Pattern Recognition Letters* 1996;17:969-974.

Executive Committee for the Asymptomatic Carotid Atherosclerosis Study. Endarterectomy for asymptomatic carotid artery stenosis. *JAMA* 1995;273:1421-1428.

Falcao A, Udupa J, Samarasekera S, Sharma S, et al. User-steered image segmentation pardigms: Live wire and live lane. *Graphical Models and Image Processing* 1998;60:233-269.

Faught W, Mattos M, vanBemmelen P, et al. Color-flow duplex scanning of carotid arteries: new velocity criteria based on receiver operator characteristic analysis for theshold stenoses used in the symptomatic and asymptomatic trials. *J Vasc Surg* 1994;19:818-828.

Feichtinger W. Transvaginal three-dimensional imaging. *Ultrasound Obstet Gynecol* 1993;3:375-378.

Fenster A, Lee D, Sherebrin S, Rankin R, et al. Three-dimensional ultrasound imaging of the vasculature. *Ultrasonics* 1998;36:629-633.

Ferguson G, Eliasziw M, Barr H, Clagett G, et al. The North American Symptomatic Carotid Endarterectomy Trial: Surgical results in 1415 patients. *Stroke* 1999;30:1751-1758.

Ferrara K and DeAngelis G. Color flow mapping. *Ultrasound Med Biol* 1997;23:321-345.

Fillinger MF, Baker R, Zwolak R, et al. Carotid duplex criteria for a 60% or greater angiographic stenosis: variation according to equipment. *J Vasc Surg* 1996;24:856-864.

Fine D, Perring S, Herbetko J, Hacking C, et al. Three-dimensional imaging of the gallbladder and dilated biliary tree: reconstruction from real-time B-scans. *Br J Radiol* 1991;64:1056-1057.

Fish P. Ultrasonic investigation of blood flow. *Proc Instn Mech Engnrs* 1999;213 Part H:

Fitzgibbon A, Pilu M, Fisher R. Direct least squares fitting of ellipses. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1999;21:476-480.

Foley J, van Dam A, Feiner S, Hughes JF. *Computer graphics: principles and practice.* Addison-Wesley, 1995.

Frangos S, Gahtan V, Sumpio B. Localization of atherosclerosis. Role of hemodynamics. *Arch Surg* 1999;134:1142-1149.

Frinking P, Bouakaz A, Kirkhorn J, ten Cate F, et al. Ultrasound contrast imaging: current and new potential methods. *Ultrasound Med Biol* 2000;26:965-975.

Fuchs H, State A, Pisano E, et al. Towards performing ultrasound-guided needle biopsies from within a head-mounted display. *Proceedings of Visualization in Biomedical Computing (VBC '96)* 1996;591-600.

Fürst G, Hofer M, Sitzer M, Kahn T, et al. Factors influencing flow-induced signal loss in MR angiography: an in vitro study. *Journal of Computer Assisted Tomography* 1995;19:692-699.

Gardener J, Lees W, Gillams A. Volume imaging with US. *Radiology* 1991;181(P):133-

Ghosh A, Nanda N, Maurer G. Three-dimensional reconstruction of echocardiographic images using the rotation method. *Ultrasound Med Biol* 1982;8:655-661.

Gibbons G and Dzau V. The emerging concept of vascular remodelling. *N Engl J Med* 1994;330:1431-1438.

Gil R, von Birgelen C, Prati F, et al. Usefulness of three-dimensional reconstruction for interpretation and quantitative analysis of intracoronary ultrasound during stent deployment. *Am J Cardiol* 1996;77:761-764.

Gilja O, Hausken T, Berstad A, Ødegaard S. Measurement of organ volume by ultrasonography. *Proc Instn Mech Engnrs* 1999;213:247-259.

Gilja O, Smievoll A, Thune N, et al. *In vivo* comparison of 3D ultrasonography and magnetic resonance imaging in volume estimation of human kidneys. *Ultrasound Med Biol* 1995;21:25-32.

Glagov S, Weisenberg E, Zarins C, Stanunavicius R, et al. Compensatory enlargement of human atherosclerotic coronary arteries. *N Engl J Med* 1986;316:1371-1375.

Gold D, Bartram C, Halligan S, Humphries K, et al. Three-dimensional endoanal sonography in assessing anal canal injury. *Br J Surg* 1999;86:365-370.

Gonzalez R and Woods R. *Digital image processing*. Reading, MA: Addison-Wesley, 1992.

Grant EG, Duerinckx AJ, ElSaden S, Melany ML, et al. Doppler sonographic parameters for detection of carotid stenosis: is there an optimum method for their selection? *American Journal Of Roentgenology* 1999;172:1123-1129.

Grant E, Duerinckx A, El Saden S, Melany M, et al. Ability to use Duplex ultrasound to quantify internal carotid arterial stenosis: fact or fiction? *Radiology* 2000;214:247-252.

Gray-Weale A, Graham J, Burnett J, et al. Carotid artery atheroma: comparison of preoperative B-mode ultrasound appearance with carotid endarterectomy specimen pathology. *J Cardiovasc Surg* 1988;29:676-681.

Griewing B, Schminke U, Morgenstern C, Walker ML, et al. Three-dimensional ultrasound angiography (Power mode) for the quantification of carotid artery atherosclerosis. *J Neuroimaging* 1997;7:40-45.

Griewing B, Morgenstern C, Driesner F, Kallwellis G, et al. Cerebrovascular disease assessed by color-flow and power Doppler ultrasonography: comparison with digital subtraction angiography in internal carotid artery stenosis. *Stroke* 1996;27:95-100.

Griffiths C, Murray A, Ramsden P. Accuracy and repeatability of bladder volume measurment using ultrasonic imaging. *J Urol* 1986;136:808-812.

Guo Z and Fenster A. Three-dimensional power Doppler imaging: A phantom study to quantify vessel stenosis. *Ultrasound Med Biol* 1996;22:1059-1069.

Guo Z, Moreau M, Rickey D, Picot P, et al. Quantitative investigation of *in vitro* flow using three-dimensional colour Doppler ultrasound. *Ultrasound Med Biol* 1995;21:807-816.

Haber I, Metaxas D, Axel L. Three-dimensional motion reconstruction and analysis of the right ventricle using tagged MRI. *Medical Image Analysis* 2000;4:335-355.

Halliday A. The asymptomatic carotid surgery trial (ACST). Rationale and design. *Eur J Vasc Surg* 1994;8:703-710.

Handschumacher M, Lethor J, Siu S. A new integrated system for three-dimensional echocardiographic reconstruction: development and validation for venticular volume with application in humans. *J Am Coll Cardiol* 1993;21:743-753.

Hankley G, Warlow C, Stellar R. Cerebral angiographic risk in patients with mild cerebrovascular disease. *Stroke* 1990;21:209-222.

Hastak S, Gammelgaard J, Holm H. Transrectal ultrasonic volume determination of the prostate — a preoperative and postoperative study. *J Urol* 1982;127:1115-1118.

Hayward J, Davies AH, Lamont P. Carotid plaque morphology: a review. *Eur J Vasc Endovasc Surg* 1995;9:368-374.

Hell B, Walter F, Schreiber S, Blase H, et al. Three-dimensional ultrasonography in maxillofacial surgery. A new diagnostic tool. *Int J Oral Maxillofac Surg* 1993;22:173-177.

Henderson R, Steinman D, Eliasziw M, Barnett H. Effect of contralateral carotid artery stenosis on carotid ultrasound velocity measurements. *Stroke* 2000;31:2636-2640.

Hennerici M, Kleophas W, Gries F. Regression of carotid plaques during low density lipoprotein cholesterol elimination. *Stroke* 1991;22:989-992.

Hennerici MG and Meairs SP. Cerebrovascular ultrasound. *Current Opinion In Neurology* 1999;12:57-63.

Hodges T, Detmer P, Burns D, Beach K, et al. Ultrasonic three-dimensional reconstruction: *in vitro* and *in vivo* volume and area measurement. *Ultrasound Med Biol* 1994;20:719-729.

Hood D, Mattos M, Monsour A, et al. Prospective evaluation of new duplex criteria to identify 70% internal carotid artery stenosis. *J Vasc Surg* 1996;23:254-262.

Hoskins PR. Accuracy of maximum velocity estimates made using Doppler ultrasound systems. *Br J Radiol* 1996;69:172-177.

Hoskins PR. A comparison of single- and dual-beam methods for maximum velocity estimation. *Ultrasound Med Biol* 1999a;25:583-592.

Hoskins PR, Fish PJ, Pye SD, Anderson T. Finite beam-width ray model for geometric spectral broadening. *Ultrasound Med Biol* 1999b;25:391-404.

Hoskins PR and McDicken W. Colour ultrasound imaging of blood flow and tissue motion. *Br J Radiol* 1997;70:878-890.

Howard G, Chambless L, Baker W. A multicenter validation study of Doppler ultrasound versus angiography. *J Stroke Cerebrovasc Dis* 1991;1:166-173.

Hughes S, D'Arcy T, Maxwell D, Chiu W, et al. Volume estimation from multiplanar 2D ultrasound images using a remote electromagnetic position and orientation sensor. *Ultrasound Med Biol* 1996;22:561-572.

Hughes S, D'Arcy T, Maxwell D, Saunders J, et al. The accuracy of a new system for estimating organ volume using ultrasound. *Physiol Meas* 1997;18:73-84.

Hunink M, Polak J, Barlan M, O'Leary D. Detection and quantification of carotid artery stenosis: efficacy of various Doppler velocity parameters. *American Journal Of Roentgenology* 1993;160:619-625.

Hünerbein M and Schlag P. Three-dimensional endosonography for the staging rectal cancer. *Ann Surg* 1997;225:432-438.

Kajiya JT and Von Herzen BP. Ray tracing volume densities. *Computer Graphics* 1984;10:165-173.

Kalman P, Johnston K, Auech P, et al. In vitro comparison of alternative methods for quantifying the severity of spectral broadening for the diagnosis of carotid arterial occlusive disease. *Ultrasound Med Biol* 1985;11:435-440.

Kass M, Witkin A, Terzopoulos D. Snakes: Active contour models. *International Journal of Computer Vision* 1988;321-331.

Kassam M, Johnston K, Cobbold R. Quantitative estimation of spectral broadening for the diagnosis of carotid arterial disease: method and in vitro results. *Ultrasound Med Biol* 1985;11:425-433.

Kavic M. Three-dimensional ultrasound. *Surg Endosc* 1996;10:74-76.

Keller P. Time of flight magnetic resonance angiography. *Neuroimaging Clin N Am* 1992;4:639-656.

Kelly I, Gardener J, Brett A, Richards R, et al. Three-dimensional ultrasound of the fetus. Work in Progress. *Radiology* 1994;192:253-259.

Kelly I, Gardener J, Lees W. Three-dimensional fetal ultrasound. *Lancet* 1992;339:1062-1064.

King DL, King DLJ, YiCi Shao M. Evaluation of in-vitro measurement accuracy of a three-dimensional ultrasound scanner. *J Ultrasound Med* 1991;10:77-82.

Klein H, Marquardt M, Polz H, Schuller G. 3D reconstruction of fetal images based on ultrasound. *Proceedings of the International Symposium on Computer Assisted Radiology* 1993;534-538.

Kovalski G, Beyar R, Shofti R, Azhari H. Three-dimensional automatic quantitative analysis of intravascular ultrasound images. *Ultrasound Med Biol* 2000;26:527-537.

Kroger K, Massalha K, Dobonici G, Rudofsky G. SieScape: a new sonographic dimension with fictive images. *Ultrasound Med Biol* 1998;24:1125-1129.

Ku D, Giddens D, Zarins C, Glagov S. Pulsatile flow and atherosclerosis in the human carotid bifurcation: positive correlation between plaque location and low and oscillating shear stress. *Arteriosclerosis* 1985;5:293-302.

Langlois Y, Greene F, Roederer G, et al. Computer-based pattern recognition of carotid artery Doppler signals for disease classification: prospective validation. *Ultrasound Med Biol* 1984;10:581-595.

Laudy J, Janssen M, Struyk P, Stijnen T, et al. Fetal liver volume measurement by three-dimensional ultrasonography: a preliminary study. *Ultrasound Obstet Gynecol* 1998;12:93-96.

Leclerc X, Godefroy O, Pruvo J, Leys D. Computed tomographic angiography for the evaluation of carotid artery stenosis. *Stroke* 1995;26:1577-1581.

Lee A. Visualisation of malformations of the fetal skeleton by volume rendering in three-dimensional ultrasound. In: Merz E, ed. *3-D Ultrasound in Obstetrics and Gynecology*. Philadelphia: Lippincott, Williams and Wilkins, 1999:121-125.

Lee A, Deutinger J, Bernaschek G. Three-dimensional ultrasound: abnormalities of the fetal face in surface and volume rendering mode. *Br J Obstet Gynaecol* 1995;102:302-306.

Lees W, Chong W, Lawrence R, Gardener J. Intravascular ultrasound with 3D reconstruction: a new way to image vascular disease. *Br J Radiol* 1991;64:655-655.

Lees W, Gardener J, Brett A, Davey M. Assessment of 3D US image acquisition. *Radiology* 1993;189(P):341-341.

Leotta DF, Primozich J, Beach KW, Berglein R, et al. Serial measurement of cross-sectional area in peripheral vein grafts using three-dimensional ultrasound. *Ultrasound Med Biol* 2001;27:61-68.

Leotta DF, Detmer PR, Martin RW. Performance of a miniature magnetic position sensor for three-dimensional ultrasound imaging. *Ultrasound Med Biol* 1997;23:597-609.

Leotta DF, Paun M, Beach KW, Kohler T, et al. Measurement of abdominal aortic aneurysms with three-dimensional ultrasound imaging: preliminary report. *J Vasc Surg* 2001;33:700-707.

Levine R, Weyman A, Handschumaker M. Three-dimensional echocardiography: techniques and applications. *Am J Cardiol* 1992;69:121-130.

Levoy M. Volume rendering, a hybrid ray tracer for rendering polygon and volume data. *IEEE Computer Graphics and Applications* 1990;10:33-40.

Li XN, Shiota T, Martin RW, Sahn DJ, et al. Digital three dimensional Doppler flow mapping: A new modality in echocardiography. *Circulation* 1997;96:1835-

Light E, Davidsen R, Hruschka T, Smith S. Progress in two-dimensional arrays for real time volumetric imaging. *Ultrason Imaging* 1998;20:1-16.

Lin W-C, Liang C-C, Chen C-T. Dynamic elastic interpolation for 3-D medical image reconstruction from serial cross sections. *IEEE Trans Med Imaging* 1988;7:225-232.

Link J, Brossman J, Grabener M, Mueller-Huelsbeck S, et al. Spiral CT angiography and selective digital subtraction angiography of internal carotid stenosis. *American Journal Of Neuroradiology* 1996;17:89-94.

Liu J-B and Goldberg B. 2-D and 3-D endoluminal ultrasound: vascular and nonvascular applications. *Ultrasound Med Biol* 1999;25:159-173.

Lubbers J and Graaff R. A simple and accurate formula for the sound velocity in water. *Ultrasound Med Biol* 1998;24:1065-1068.

Lyden P and Nelson T. Visualization of the cerebral circulation using three-dimensional transcranial power Doppler ultrasound imaging. *J Neuroimaging* 1997;7:35-39.

Mao F, Gill J, Downey D, Fenster A. Segmentation of the carotid artery in ultrasound images: Method, development and evaluation of the technique. *Med Phys* 2000;27:1961-1970.

Matre K, Stokke E, Martens D, Gilja O. In vitro volume estimation of kidneys using three-dimensional ultrasonography and a position sensor. *Eur J Ultrasound* 1999;10:65-73.

Meairs S, Beyer J, Hennerici M. Reconstruction and visualisation of irregularly sampled three- and four-dimensional ultrasound data for cerebrovascular applications. *Ultrasound Med Biol* 2000;26:263-272.

Merz E, Bahlmann F, Weber G. Volume scanning in the evaluation of fetal malformations: a new dimension in prenatal diagnosis. *Ultrasound Obstet Gynecol* 1995a;5:222-227.

Merz E, Bahlmann F, Weber G, Macchiella D. Three-dimensional ultrasonography in prenatal diagnosis. *J Perinat Med* 1995b;23:213-222.

Meyer C, Boes J, Kin B, Bland P, et al. Semiautomatic registration of volumetric ultrasound scans. *Ultrasound Med Biol* 1999;25:339-347.

Meyer J, Khalil R, Obuchowski N, Baus L. Common carotid artery: variability of Doppler US velocity measurements. *Radiology* 1997;204:339-341.

Meyer K, Applewhite H, Biocca F. A survey of position trackers. *Presence* 1992;1:173-200.

Mikkonen R, Kreula J, Virkkunen P. Reproducibility of Doppler ultrasound measurements. *Acta Radiologica* 1996;37:545-550.

Milne AD, Chess DG, Johnson JA, King GW. Accuracy of an electromagnetic tracking device: A study of the optimal operating range and metal interference. *J Biomech* 1996;29:791-793.

Moneta G, Edwards J, Chitwood R, et al. Correlation of North American Symptomatic Carotid Endarterectomy Trial (NASCET) angiographic definition of 70% to 99% internal carotid artery stenosis with duplex scanning. *J Vasc Surg* 1993;17:152-159.

Moneta G, Edwards J, Papanicolaou G, et al. Screening for asymptomatic internal carotid artery stenosis: Duplex criteria for discriminating 60% to 90% stenosis. *J Vasc Surg* 1995;21:989-994.

Moore J, Steinman D, Ethier C. Computational blood flow modelling: Errors associated with reconstructing finite element models from magnetic resonance images. *J Biomech* 1998;31:179-184.

Moore W and et al. Guidelines for carotid endarterectomy (a multidisciplinary consensus statement from the ad hoc committe, American Heart Association). *Stroke* 1995;26:188-201.

Moritz W, Pearlman A, McCabe D, Medema D, et al. An ultrasonic technique for imaging the ventricle in three dimensions and calculating its volume. *IEEE Trans Biomed Eng* 1983;30:482-492.

Murie J and McKay A. Radiographic assessment of the extracranial internal carotid artery. *J Cardiovasc Surg* 1986;27:154-157.

Nash J. *Compact numerical methods for computers - linear algebra and function minimisation*. Bristol: Adam Hilger, 1990.

Neales M, Chambers J, Kelly A, et al. Reappraisal of duplex criteria to assess significant carotid stenosis with special reference to report from the North American Carotid Endarterectomy Trial and the European Carotid Surgery Trial. *J Vasc Surg* 1994;20:642-649.

Nelson TR and Pretorius DH. Three-dimensional ultrasound of fetal surface features. *Ultrasound Obstet Gynecol* 1992;2:166-174.

Nelson TR, Pretorius D, Sklansky M, Hagen-Ansert S. Three-dimensional echocardiographic evaluation of fetal heart anatomy and function: acquisition, analysis, and display. *J Ultrasound Med* 1996;15:1-9.

Nelson TR and Pretorius DH. Three-dimensional ultrasound imaging. *Ultrasound Med Biol* 1998;24:1243-1270.

Nelson T and Elvins T. Visualisation of 3D ultrasound data. *IEEE Computer Graphics and Applications* 1993;13:50-57.

Nelson T and Pretorius D. Visualization of the fetal thoracic skeleton with three-dimensional sonography: A preliminary report. *American Journal Of Roentgenology* 1995;164:1485-1488.

Nerem R. Haemodynamics and the vascular endothelium. *J Biomech Eng* 1993;115:510-514.

Ng K, Evans J, Vonesh M, et al. Arterial imaging with a new forward-viewing intravascular ultrasound catheter. *Circulation* 1994;89:718-723.

Nikravesh P, Skorton D, Chandran K, Attarwala Y, et al. Computerized three-dimensional finite element reconstruction of the left ventricle from cross-sectional echocardiograms. *Ultrason Imaging* 1984;6:48-59.

North American Symptomatic Carotid Endarterectomy Trial (NASCET) Steering Committee. North American Symptomatic Carotid Endarterectomy Trial: methods, patient characteristics, and progress. *Stroke* 1991;22:711-720.

North American Symptomatic Carotid Endarterectomy Trial Collaborators. Beneficial effect of carotid endarterectomy in symptomatic patients with high-grade carotid stenosis. *N Engl J Med* 1991;325:445-453.

O'Donnell M. Coded excitation system for improving the penetration of real-time phased-arry imaging systems. *IEEE Trans Ultrason Ferroelc Freq Contr* 1992;39:341-351.

O'Leary D, Bryan F, Goodison M, et al. Measurement variability of carotid atherosclerosis: real-time [B-mode] ultrasonography and angiography. *Stroke* 1987;18:1011-1017.

Ohbuchi R and Fuchs H. Incremental 3D ultrasound imaging from a 2D scanner. *Proceedings of the First Conference on Visualisation in Biomedical Computing* 1990;360-367.

Paciaroni M, Eliasziw M, Kappelle L, Finan J, et al. Medical complications associated with carotid endarterectomy. *Stroke* 1999;30:1759-1763.

Pagoulatos N, Haynor D, Kim Y. A fast calibration method for 3-D tracking of ultrasound images using a spatial localizer. *Ultrasound Med Biol* 2001;27:1219-1229.

Palombo C, Kozakova M, Morizzo C, Andreuccetti F, et al. Ultrafast three-dimensional ultrasound - Application to carotid artery imaging. *Stroke* 1998;29:1631-1637.

Parker D, Wang K, Taylor C. Image-based geometric modelling of the human vasculature in computational haemodynamics. *ASME Proceedings* 1999;

Pauletzki J, Sackmann M, Holl J, Paumgertner G. Evaluation of gall-bladder volume and emptying with a novel three-dimensional ultrasound system: comparison with the sum-of-cylinders and the ellipsoid methods. *J Clin Ultrasound* 1996;24:277-285.

Picot P, Rickey D, Mitchell R, Rankin R, et al. Three-dimensional Doppler imaging. *Ultrasound Med Biol* 1993;19:95-104.

Prager RW, Rohling RN, Gee AH, Berman L. Rapid calibration for 3-D freehand ultrasound. *Ultrasound Med Biol* 1998;24:855-869.

Pretorius D and Nelson T. Prenatal visualisation of cranial sutures and fontanelles with three-dimensional ultrasonography. *J Ultrasound Med* 1994;13:871-876.

Pretorius D, Nelson T, Baergen R, Pae E, et al. Imaging placental vasculature using three-dimensional ultrasound and color power Doppler. *Ultrasound Obstet Gynecol* 1998;12:45-49.

Pretorius D, Nelson T, Jaffe J. 3-dimensional sonographic analysis based on color flow Doppler and gray scale image data: a preliminary report. *J Ultrasound Med* 1992;11:225-232.

Price D, Wallbridge D, Stewart M. Tissue Doppler imaging: current and potential clinical applications. *Heart* 2000;84 Suppl.2:II1-II8.

Prince M. Body MR angiography with gadolinium contrast agents. *Magn Reson Imaging Clin N Am* 1996;4:11-24.

Pryor J, Setton A, Nelson P, Bernstein A. Complications of diagnostic cerebral angiography and tips on avoidance. *Neuroimaging Clin N Am* 1996;6:751-757.

Raab F, Blood E, Steiner T, nes HR. Magnetic position and orientation tracking systems. *IEEE Trans Aero Elec Sys* 1979;709-718.

Raichlen J, Trivedi S, Herman G, St.John Sutton M, et al. Dynamic three-dimensional reconstruction of the left ventricle from two-dimensional echocardiograms. *J Am Coll Cardiol* 1986;364-370.

Ranke C, Creutzig A, Becker H, Trappe H-J. Standardization of carotid ultrasound. A hemodynamic method to normalize for interindividual and interequipment variability. *Stroke* 1999;30:402-406.

Ranke C, Hendrickx P, Roth U, Brassel F, et al. Color and conventional image-directed Doppler ultrasonography: accuracy and sources of error in quantitative blood flow measurements. *J Clin Ultrasound* 1992;20:187-193.

Riccabona M, Nelson T, Pretorius D. Three-dimensional ultrasound: accuracy of distance and volume measurements. *Ultrasound Obstet Gynecol* 1996b;7:429-434.

Riccabona M, Nelson T, Pretorius D, Davidson T. Distance and volume measurement using three-dimensional ultrasonography. *J Ultrasound Med* 1995;14:881-886.

Riccabona M, Nelson T, Pretorius D, Davidson T. Three-dimensional sonographic measurement of bladder volume. *J Ultrasound Med* 1996a;15:627-632.

Richard W, Grimmell C, Bedigian K, Frank K. A method for three-dimensional prostate imaging using transrectal ultrasound. *Comp Med Imaging Graph* 1993;17:73-79.

Richardson P, Davies M, Born G. Influence of plaque configuration and stress distribution on fissuring of coronary atherosclerotic plaques. *Lancet* 1989;2:941-944.

Rickey D, Picot D, Holdsworth D, Drangova M, et al. Quantitative three-dimensional true velocity colour doppler imaging. *IEEE Ultrasonics Symposium* 1991;1277-1280.

Ritchie C, Edwards W, Mack L, Cyr D, et al. Three-dimensional ultrasonic angiography using power-mode Doppler. *Ultrasound Med Biol* 1996;22:277-286.

Robert B, Richard B, Nicolas J-M. An interactive tool to visualize three-dimensional ultrasound data. *Ultrasound Med Biol* 2000;26:133-142.

Roelandt J, ten Cate F, Vletter W, Taams M. Ultrasonic dynamic three-dimensional visualisation of the heart with a multiplane transesophageal imaging transducer. *J Am Soc Echocardiogr* 1994;7:217-229.

Rohling R, Gee A, Berman L. Automatic registration of 3-D ultrasound images. *Ultrasound Med Biol* 1998;24:841-854.

Rosenfield K, Boffetti P, Kaufman, et al. Three-dimensional reconstruction of human carotid arteries from images obtained during noninvasive B-mode ultrasound examination. *Am J Cardio* 1992;30:379-384.

Ross R. The pathogenesis of atherosclerosis: a perspective for the 1990s. *Nature* 1993;362:801-809.

Rothwell P, Gibson R, Slattery J, Sellar R, et al. Equivalence of measurements of carotid stenosis. A comparison of three methods on 1001 angiograms. *Stroke* 1994;25:2435-2439.

Rothwell P, Gibson R, Villagra R, Sellar R, et al. The effect of angiographic technique and image quality on the reproducibility of measurement of carotid stenosis and assessment of plaque surface morphology. *Clinical Radiology* 1998;53:439-443.

Rothwell P, Pendlebury S, Wardlaw J, Warlow C. Critical appraisal of the design and reporting of studies of imaging and measurement of carotid stenosis. *Stroke* 2000;31:1444-1450.

Rotten D, Levaillant J, Zerat L. Analysis of normal breast tissue and of solid breast masses using three-dimensional ultrasound mammography. *Ultrasound Obstet Gynecol* 1999;14:114-124.

Sabella P. A rendering algorithm for visualizing 3D scalar fields. *Computer Graphics* 1988;22:51-57.

Sakas G, Schreyer L, Grimm M. Preprocessing and volume rendering of 3-D ultrasonic data. *IEEE Computer Graphics and Applications* 1995;15:47-54.

Sarti A, Lamberti C, Erbacci G Pini R. Volume rendering for 3-D echocardiography visualization. *Comput Cardiol* 1993;209-212.

Sawader H, Fujii J, Kato, Onoe M, et al. Three-dimensional reconstruction of the left ventricle from multiple cross-sectional echocardiograms. Value for measuring left ventricular volume. *Br Heart J* 1983;50:438-442.

Schmidt-Trucksäss A, Grathwohl D, Schmid A, Boragk R, et al. Assessment of carotid wall motion and stiffnes with tissue Doppler imaging. *Ultrasound Med Biol* 1998;24:639-646.

Schminke U, Motsch L, Griewing B, Gaull M, et al. Three-dimensional power-mode ultrasound for quantification of the progression of carotid artery atherosclerosis. *J Neurology* 2000;247:106-111.

Schwartz S, Chambless L, Baker W, Broderick J, et al. Consistency of Doppler parameters in predicting arteriographically confirmed carotid stenosis. *Stroke* 1997;28:343-347.

Scissons R, Salles-Cunha S, Altenburg L, et al. Doppler peak systolic velocity measurements: correction of angle-dependent errors. *J Vasc Tech* 1997;21:233-236.

Sehgal C, Broderick G, Gorniak R, Arger P. Three-dimensional imaging of prostates with transrectal sonography. *Radiology* 1993;189(P):342-342.

Serfaty J, Chirossel P, Chevallier J, Ecochard R, et al. Accuracy of three-dimensional gadolinium-enhanced MR angiography in the assessment of extracranial carotid artery disease. *American Journal Of Roentgenology* 2000;175:455-463.

Seward J, Belohlavek M, O'Leary P, Foley D, et al. Congenital heart disease: wide-field, three-dimensional, and four-dimensional ultrasound imaging. *Am J Cardiac Imag* 1995;9:38-43.

Sidhu P and Allan P. Ultrasound assessment of internal carotid artery stenosis. *Clinical Radiology* 1997;52:654-658.

Simeone A, Carriero A, Armillotta M, Scarabino T, et al. Spiral CT angiography in the study of carotid stenoses. *J Neuroradiol* 1997;24:18-22.

Sklansky M, Nelson TR, Strachan M, Pretorius D. Real-time three-dimensional fetal echocardiography: initial feasiblilty study. *J Ultrasound Med* 1999;18:745-752.

Skolnick M. Estimation of ultrasound beam width in the elevation (section thickness) plane. *Radiology* 1991;180:286-288.

Smith R, Rutt B, Fox A, Rankin R. Geometric characterization of stenosed human carotid arteries. *Acad Radiol* 1996;3:898-911.

Smith R, Rutt B, Holdsworth D. Anthropomorphic carotid bifurcation phantom for MRI applications. *J Magn Reson Imaging* 2001;10:533-544.

Smith S, Trahey G, von Ramm O. Two-dimensional arrays for medical ultrasound. *Ultrason Imaging* 1992;14:213-233.

Solberg L and Eggen D. Localization and sequence of development of atherosclerotic lesions in the carotid and vertebral arteries. *Circulation* 1971;43:711-724.

Spencer M and Reid J. Quantitation of carotid stenosis with continuous-wave (C-W) Doppler ultrasound. *Stroke* 1979;10:326-330.

Spoor C and Veldpaus F. Rigid body motion calculated from spatial co-ordinates of markers. *J Biomech* 1980;13:391-393.

State A, Livingston M, Garrett W, et al. Technologies for augmented reality systems: realizing ultrasound-guided needle biopsies. *Proceeedings of SIGGRAPH '96* 1996;439-446.

Steen E and Olstad B. Volume rendering of 3D medical ultrasound data using direct feature extraction. *IEEE Trans Med Imaging* 1994;13:517-525.

Steiner H, Spitzer D, Weiss-Wichert P, Graf A, et al. Three-dimensional ultrasound in prenatal diagnosis of skeletal dysplasia. *Prenat Diagn* 1995;15:373-377.

Steiner H, Staudach A, Spinzer D, Schaffer H. Three-dimensional ultrasound in obstetrics and gynaecology: technique, possibilities and limitations. *Human Reprod* 1994;9:1773-1778.

Steinke W, Meairs S, Ries S, Hennerici M. Sonographic assessment of carotid artery stenosis: comparison of power Doppler imaging and color Doppler flow imaging. *Stroke* 1996;27:91-94.

Steinke W, Ries S, Artemis N, Scwartz A, et al. Power Doppler imaging of carotid artery stenosis. *Stroke* 1997;28:1981-1987.

Stone M. A 3-dimensional model of tongue movement based on ultrasound and x-ray microbeam data. *J Acoust Soc Am* 1990;87:2207-2217.

Taylor R, Lavalle S, Burdea G, Mosges E. *Computer Integrated Surgery.* Massachusets: MIT Press, 1996.

Tegos TJ, Elatrozy TS, Sabetai M, Kalodiki E, et al. Identification of the unstable carotid plaque. *Br J Surg* 1998;85:705-

Teicholz L, Cohen M, Sonnonblick E, Gorlin R. Study of left ventricular geometry and function by B-scan ultrasonography in patients with and without asynergy. *N Engl J Med* 1974;291:1220-1226.

Tong S, Downey D, Cardinal H, Fenster A. A three-dimensional ultrasound prostate imaging system. *Ultrasound Med Biol* 1996;22:735-746.

Toole J and Castaldo J. Accurate measurement of carotid stenosis. Chaos in methodology. *J Neuroimaging* 1994;4:222-230.

Trobaugh JW, Trobaugh DJ, Richard WD. 3-Dimensional imaging with stereotaxic ultrasonography. *Comp Med Imaging Graph* 1994;18:315-323.

van der Wal A and Becker A. Atherosclerotic plaque rupture - pathologic basis of plaque stability and instability. *Cardiovascular Research* 1999;41:334-344.

Vanninen R, Manninen H, Koivisto K, Partanen K, et al. Carotid stenosis by digital subtraction angiography: reproducibility of the European Carotid Surgery Trial and the North American Symptomatic Carotid Endarterectomy Trial measurement methods and visual interpretation. *Journal of Neuroradiology* 1994;15:1635-1641.

von Birgelen C, de Vrey E, Mintz G, Nicosia A, et al. ECG-gated three-dimensional intravascular ultrasoúnd: feasibility and reproducibility of the automatied analysis of coronary lumen and atherosclerotic plaque dimensions in humans. *Circulation* 1997;96:2944-2952.

von Birgelen C, Di Mario C, Reimers B, et al. Three-dimensional intracoronary ultrasound imaging. Methodology and clinical relevance for the assessment of coronary arteries and bypass grafts. *J Cardiovasc Surg* 1996;37:129-139.

Walker M, Souza S, Dumoulin C. Quantitative flow measurement in phase contrast MR angiography. *J Comput Assist Tomogr* 1988;12:304-313.

Warnock N, Gandhi M, Bergvall U, Powell T. Complications of intraarterial digital subtraction angiography in patients investigated for cerebral vascular disease. *Br J Radiol* 1993;66:855-858.

Watanabe Y. A method for volume estimation by using vector areas and centroids of serial cross-sections. *IEEE Trans Biomed Eng* 1982;29:202-205.

Watt A and Watt M. Volume rendering techniques. In: Wegner P, ed. *Advanced Animation and Rendering Techniques: Theory and Practice*. New York: ACM Press, 1992:297-321.

Waugh J and Sacharias N. Arteriographic complications in the DSA era. *Radiology* 1992;182:243-246.

Weinraub Z, Maymon R, Shulman A. Three-dimensional saline contrast hysterosonography and surface rendering of uterine cavity pathology. *Ultrasound Obstet Gynecol* 1996;8:277-282.

Weismann C. Three-dimensional sonography of the breast. In: Kurjak A and Kupesic S, ed. *Clinical Application of 3D Sonography*. New York: Parthenon Publishing, 2000a:

Weismann C, Forstner R, Prokop E, et al. Three-dimensional targeting: a new three-dimensional ultrasound technique to evaluate needle position during breast biopsy. *Ultrasound Obstet Gynecol* 2000b;16:359-364.

Wells P. *Advances in ultrasound techniques and instrumentation*. New York: Churchill Livingstone Inc., 1993.

Weskott H. B-flow: a new method for detecting blood flow [German]. *Ultraschall in Der Medizin* 2001;21:59-65.

Wildes D, Chiao R, Daft C, Rigby K, et al. Elevation performance of 1.25D and 1.5D transducer arrays. *Ultras Ferro and Freq Control* 1997;44:1027-1037.

Williams M and Nicolaides A. Predicting the normal dimensions of the internal and external carotid arteries from the diameter of the common carotid. *Eur J Vasc Surg* 1987;1:91-96.

Wilterdink J, Feldmann E, Easton J, Ward R. Performance of carotid ultrasound in evaluating candidates for carotid endarterectomy is optimized by an approach based on clinical outcome rather than accuracy. *Stroke* 1996;27:1094-1098.

Withers C, Gosink B, Keightley A. Duplex carotid sonography peak systolic velocity in quantifying internal carotid artery stenosis. *J Ultrasound Med* 1990;9:345-349.

Wollschlager H, Zeiher A, Klein H, Kasper W, et al. Transesophageal echo computer tomography: A new method for dynamic 3-D imaging of the heart. *J Am Coll Cardiol* 1989;13:68A-

Yao J, van Sambeek M, Dall'Agata A, van Dijk L, et al. Three-dimensional ultrasound study of carotid arteries before and after endarterectomy: analysis of stenotic lesions and surgical impact on the vessel. *Stroke* 1998;29:2026-2031.

Yuan C, Beach KW, Smith LH, Hatsukami TS. Measurement of atherosclerotic carotid plaque size in vivo using high resolution magnetic resonance imaging. *Circulation* 1998;98:2666-2671.

Yucel E, Anderson C, Edelman R, Grist T, et al. Magnetic resonance angiography: update on applications for extracranial arteries. *Circulation* 1999;100:2284-2301.

Zarins C, Giddens D, Glagov S. Carotid bifurcation atherosclerosis: quantitative corrleation of plaque localization with flow velocity profiles and shear stress. *Circ Res* 1983;53:502-514.

Zhang J, Xue Q, Ogawa S. 1.5-dimensional arrays for effective dynamic focussing and receiving. *Proc. IEEE Ultrasonics Symposium* 1996;1531-1534.

Zhu S and Yuille A. Region competition: Unifying snakes, region growing, and Bayes/MDL for multiband image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1996;18:884-900.

Zwiebel W. Doppler evaluation of carotid stenosis. In: Zwiebel W, ed. *Introduction to vascular ultrasonography*. Philadelphia: W.B. Saunders Company, 1999:137-154.

LSU. LOUISIANA UNIV.