

Automated Visual Measurement of Body Shape in Scoliosis

by

Jeremy D. Pearson

A thesis submitted in partial fulfilment of the requirements of

Liverpool John Moores University for the Degree of Doctor of Philosophy

July 1996

School of Electrical and Electronic Engineering
Liverpool John Moores University
in collaboration with
Department of Human Anatomy and Cell Biology
The University of Liverpool

For Vivienne and Imogen.

The truth is too simple: one must always get there by a complicated route.

George Sand, letter to Armand Barbes, 1867.

Abstract

Automated Visual Measurement of Body Shape in Scoliosis

Jeremy D. Pearson

Doctor of Philosophy

This thesis describes the content and progression of research into automated non-contact methods for measuring the three-dimensional shape of the human back in scoliosis. Scoliosis is a condition in which the spine becomes distorted and a rib-hump appears on the surface of the back. The research was driven by the needs of the scoliosis clinician and was supported by the Royal Liverpool Children's Hospital, Merseyside.

A number of optical methods for measuring back surface shape are considered. Moiré contouring and Fourier transform profilometry are investigated through practical research in the laboratory. Stereophotogrammetry, phase stepping profilometry, optical scanning and raster pattern contouring are investigated through consideration of theory and literature review. However, none of these approaches is found to be free from limitations.

The main novel content of the work presented in this thesis lies in the research into a new method for reconstructing back shape. A new optical method is proposed in which a modified multi-stripe structured light pattern is projected onto the surface of the back. Image processing operations, specialised for this application, process the image of the pattern to reconstruct three-dimensional shape.

Further research demonstrates that the computer reconstruction can be interrogated to measure parameters of clinical significance such as Angle of Trunk Inclination and Standardised Trunk Asymmetry Score. A working clinical system was implemented and tested on scoliosis patients at the hospital. The method is evaluated in terms of technical qualities and as a usable clinical tool and was found to satisfy the criteria for a successful automated system.

Contents

| | |
|---|-----------|
| List of Figures | vi |
| List of Tables | ix |
| Abbreviations | x |
| Symbols | xi |
| Chapter 1. Introduction | 2 |
| 1.1 Background | 2 |
| 1.1.1 Objectives of the Investigation..... | 3 |
| 1.1.2 Scope of the Investigation..... | 4 |
| 1.1.3 Collaboration..... | 4 |
| 1.1.4 Structure of this Thesis | 4 |
| 1.2 A Background to Scoliosis..... | 5 |
| 1.2.1 Anatomical Planes of the Body..... | 6 |
| 1.2.2 The Basic Anatomy of the Spine | 6 |
| 1.2.3 Scoliosis: a curvature of the spine..... | 9 |
| 1.2.3.1 Rib Cage Distortion | 10 |
| 1.2.3.2 Vertebral Rotation..... | 10 |
| 1.2.3.3 The Three-dimensional Deformity..... | 10 |
| 1.2.4 The Incidence of Scoliosis | 10 |
| 1.2.5 The Cause of Scoliosis | 12 |
| 1.2.6 Treatment of Scoliosis | 12 |
| 1.3 Conventional Radiographic Assessment of Scoliosis..... | 13 |
| 1.3.1 Angle of Curvature..... | 14 |
| 1.3.2 Vertebral Rotation..... | 14 |
| 1.4 From Spinal Deformity to Surface Shape | 16 |
| 1.5 Assessment of Scoliosis from Surface Shape | 19 |
| 1.5.1 Angle of Trunk Inclination and the Scoliometer..... | 20 |
| 1.5.2 Trunk Asymmetry Score | 24 |
| 1.6 School Screening..... | 26 |
| 1.7 The Need for an Automated System to Measure Back Shape | 27 |
| 1.8 The Criteria for an Automated Measurement System..... | 31 |
| Chapter 2. Optical Methods for Measurement of Back Surface Shape | 35 |
| 2.1 Introduction..... | 35 |
| 2.1.1 Background | 35 |
| 2.1.2 Suitability of Optical Methods..... | 36 |

| | |
|---|----|
| 2.2 Stereophotogrammetry | 37 |
| 2.2.1 The Stereophotogrammetric Principle | 38 |
| 2.2.2 A Case Study with Stereophotogrammetry: Thompson | 38 |
| 2.2.3 Discussion of Stereophotogrammetry | 40 |
| 2.3 From Stereophotogrammetry to Structured Light..... | 40 |
| 2.4 Moiré | 41 |
| 2.4.1 The Principle of Moiré | 42 |
| 2.4.1.1 Shadow Moiré | 43 |
| 2.4.1.2 Projection Moiré..... | 43 |
| 2.4.1.3 Other Configurations..... | 44 |
| 2.4.2 Moiré Fringe Formation..... | 45 |
| 2.4.3 Application of Moiré To Back Shape Measurement..... | 48 |
| 2.4.4 Discussion of Moiré | 51 |
| 2.5 Phase-Measuring Methods | 53 |
| 2.5.1 The Relationship between Phase and Surface Shape..... | 54 |
| 2.5.2 Phase Extraction..... | 56 |
| 2.5.3 Phase Extraction using Phase Stepping Profilometry | 57 |
| 2.5.3.1 The Principle of PSP | 57 |
| 2.5.3.2 A Case Study with PSP: Halioua | 59 |
| 2.5.3.3 Discussion Of PSP | 61 |
| 2.5.4 Fourier Transform Profilometry..... | 62 |
| 2.5.4.1 The Principle of FTP | 62 |
| 2.5.4.2 Application to Back Shape Measurement | 65 |
| 2.5.4.3 Discussion on FTP | 68 |
| 2.5.5 Summary of Phase-Measuring Methods | 69 |
| 2.6 Optical Scanning Methods | 70 |
| 2.6.1 Principle of Scanning | 70 |
| 2.6.2 Clinical Application of Scanning | 72 |
| 2.6.2.1 Configurations for Scanning Systems..... | 73 |
| 2.6.2.2 A Case Study with Scanning: The ISIS System (Turner-Smith) | 73 |
| 2.6.2.3 Discussion of Scanning | 76 |
| 2.7 Raster Pattern Contouring | 77 |
| 2.7.1 Point Raster Projection..... | 78 |
| 2.7.1.1 Principle of Point Raster Projection..... | 78 |
| 2.7.1.2 Case Study with Point Raster Projection: Hill | 79 |
| 2.7.1.3 Discussion of Point Raster Scanning | 80 |
| 2.7.1.4 Bitwise Binary Encoding | 81 |
| 2.7.2 Line Raster Projection..... | 82 |
| 2.7.2.1 Principle of Line Raster Projection | 82 |
| 2.7.2.2 A Case Study with Line Raster Projection: the Frobin System | 83 |
| 2.7.2.3 Discussion of Line Raster Projection..... | 86 |

| | |
|---|-----------|
| 2.7.3 Discussion of Raster Contouring | 87 |
| 2.8 Comparison and Discussion of Optical Methods..... | 88 |
| 2.8.1 The Optical Methods Considered..... | 88 |
| 2.8.2 Comparison of Methods..... | 89 |
| 2.8.3 An Appropriate Method for Further Investigation..... | 91 |
| 2.9 Summary | 96 |
| Chapter 3. A Structured Light Method for Measuring Back Shape..... | 98 |
| 3.1 Introduction..... | 98 |
| 3.2 Operational Considerations..... | 99 |
| 3.3 The Optical Subsystem | 100 |
| 3.3.1 Basic Concept of the Optical Subsystem | 101 |
| 3.3.2 Choice of Structured Light Projection Pattern..... | 103 |
| 3.3.3 Measurement Procedure..... | 105 |
| 3.3.4 Coordinate System..... | 107 |
| 3.3.5 Optical Geometry..... | 108 |
| 3.3.6 Measurement of Pixel Size..... | 109 |
| 3.4 Image Processing | 111 |
| 3.4.1 Outline of the Image Processing Subsystem..... | 113 |
| 3.4.2 Source Image Data | 115 |
| 3.4.3 Contrast Stretching..... | 119 |
| 3.4.4 Fringe Location using Thresholding and Edge Detection..... | 120 |
| 3.4.5 Thresholding | 121 |
| 3.4.6 Edge Detection..... | 129 |
| 3.4.7 Point Filtering | 131 |
| 3.4.8 Fringe Ordering..... | 132 |
| 3.4.8.1 Searching for the Reference Fringes | 133 |
| 3.4.8.2 Indexing The Central Vertical Column of the Image..... | 136 |
| 3.4.8.3 Indexing the Bulk of the Image..... | 136 |
| 3.4.8.4 Indexed Image Editing | 139 |
| 3.4.9 The Fringe Table..... | 141 |
| 3.4.10 Representing Three-dimensional Coordinates | 144 |
| 3.4.11 Calculation of z Coordinates..... | 146 |
| 3.4.12 Defining the Boundary of the Back | 148 |
| 3.4.13 Interpolation | 150 |
| 3.4.14 Simplifying the Status Values..... | 152 |
| 3.4.15 Processing the Scaling Image..... | 153 |
| 3.5 System Implementation..... | 153 |

| | |
|---|------------|
| 3.6 Discussion of the Proposed Method..... | 154 |
| 3.6.1 The Optical Subsystem | 155 |
| 3.6.1.1 The Data Acquisition Problem..... | 155 |
| 3.6.1.2 The Framework and Mounting of Components | 156 |
| 3.6.1.3 The Projector and Gratings | 156 |
| 3.6.1.4 The Camera and Frame-store (Digitiser) | 157 |
| 3.6.2 The Image Processing Subsystem..... | 158 |
| 3.7 Summary | 158 |
| Chapter 4. System Evaluation and Clinical Application..... | 161 |
| 4.1 Introduction..... | 161 |
| 4.2 Measurement of the Patient in the Standing position..... | 161 |
| 4.3 A Non-Contact System | 164 |
| 4.4 Accuracy..... | 164 |
| 4.4.1 An Experiment to Measure Accuracy | 164 |
| 4.4.2 Results of the Accuracy Experiment..... | 166 |
| 4.4.3 Comments on Accuracy | 168 |
| 4.5 Measurement of Parameters of Clinical Significance | 168 |
| 4.5.1 Vertebral Positioning | 169 |
| 4.5.2 Measurement of Angle of Trunk Inclination | 170 |
| 4.5.3 Measurement of Standardised Trunk Asymmetry Score | 173 |
| 4.5.4 Provision of Other Clinically Useful Information | 174 |
| 4.6 Repeatability | 176 |
| 4.7 Speed of Data Acquisition and Processing | 177 |
| 4.7.1 Speed of Data Acquisition | 177 |
| 4.7.2 Processing Time..... | 178 |
| 4.8 Provision of Long Term Storage of Data | 179 |
| 4.9 Simplicity of Clinical Use..... | 181 |
| 4.10 Cost | 182 |
| 4.11 Portable and Collapsible | 183 |
| 4.12 Summary of the Evaluation Exercise..... | 183 |
| 4.13 Summary | 184 |
| Chapter 5. Conclusions and Recommendations for Further Work..... | 186 |
| 5.1 Introduction..... | 186 |
| 5.2 Conclusions..... | 187 |
| 5.3 Recommendations for Further Work..... | 190 |

| | |
|---|------------|
| 5.3.1 Fundamental Approaches..... | 190 |
| 5.3.2 A Composite Tool for Assessing Scoliosis..... | 191 |
| 5.3.3 Clinical Research | 193 |
| 5.3.4 Improving the Set-up and Acquisition Time..... | 194 |
| 5.3.5 Improving the Fringe Ordering | 195 |
| 5.3.6 Image Processing Performance | 195 |
| 5.4 Other Applications | 197 |
| 5.5 Dissemination of this Research..... | 197 |
| References | 198 |
| Acknowledgements..... | 205 |
| Appendix I Glossary of Medical Terms | 207 |
| Appendix II System Implementation Details..... | 212 |
| Appendix III Publications | 223 |

List of Figures

| | |
|---|----|
| Figure 1.1 The anatomical planes of the body. | 6 |
| Figure 1.2 The spine and the groups of vertebrae. | 7 |
| Figure 1.3 The seventh cervical vertebra. | 8 |
| Figure 1.4 A lateral view of a thoracic vertebra. | 8 |
| Figure 1.5 A single structural curve with compensatory curves. | 9 |
| Figure 1.6 A double structural curve with compensatory curves. | 9 |
| Figure 1.7 Radiograph of a thoracic structural curve (from James 1976). | 13 |
| Figure 1.8 Cobb's methods for measuring angle of curvature. | 14 |
| Figure 1.9 The Nash and Moe method for measuring vertebral rotation. | 15 |
| Figure 1.10 The surface deformity of scoliosis. | 17 |
| Figure 1.11 Using the Debrunner Kyphometer. | 19 |
| Figure 1.12 Angle of Trunk Inclination (ATI). | 20 |
| Figure 1.13 Bunnell's application of Adams' test. | 21 |
| Figure 1.14 The Scoliometer. | 21 |
| Figure 1.15 ATI map for a typical scoliosis patient. | 22 |
| Figure 1.16 Patient being measured in the prone position using the Scoliometer. | 23 |
| Figure 1.17 The FBCT in use to measure a patient in the forward bending position. ... | 24 |
| Figure 1.18 Measurement of TAS. | 25 |
| Figure 2.1 The principle of stereophotogrammetry. | 38 |
| Figure 2.2 The stereophotogrammetric arrangement used by Thompson. | 39 |
| Figure 2.3 The principle of structured light projection. | 41 |
| Figure 2.4 Moiré topogram of a scoliotic back. | 42 |
| Figure 2.5 Typical configuration for shadow moiré contouring. | 43 |
| Figure 2.6 Configuration for projection moiré contouring. | 44 |
| Figure 2.7 Moiré fringe formation. | 45 |
| Figure 2.8 Plan of shadow moiré geometry for contouring the back. | 49 |
| Figure 2.9 Elevation of shadow moiré geometry for contouring the back. | 49 |
| Figure 2.10 Selection of moiré topograms of scoliosis patients. | 50 |
| Figure 2.11 Reconstruction from moiré topogram. | 50 |
| Figure 2.12 Structured light illumination of surface. | 54 |
| Figure 2.13 Geometry of shift in surface position. | 56 |
| Figure 2.14 Flow chart for PSP. | 58 |
| Figure 2.15 The configuration used by Halioua. | 60 |
| Figure 2.16 Flowchart for the FTP method of phase extraction. | 64 |
| Figure 2.17 Filtering in the Fourier domain. | 65 |
| Figure 2.18 Test configuration for FTP. | 66 |

| | | |
|-------------|---|-----|
| Figure 2.19 | Image of plaster-cast of back with resulting phase image. | 67 |
| Figure 2.20 | Optical Scanning..... | 71 |
| Figure 2.21 | Optical scanning (a) Pre-objective (b) Post objective. | 72 |
| Figure 2.22 | A typical configuration for optical scanning. | 73 |
| Figure 2.23 | The ISIS Scanner (a) Clinical configuration (b) Mechanical lay-out..... | 74 |
| Figure 2.24 | Dot or point raster that might be used for projection..... | 79 |
| Figure 2.25 | Arrangement used by Hill..... | 80 |
| Figure 2.26 | Simple grating for Line Raster Projection..... | 83 |
| Figure 2.27 | Experimental set-up used by Frobin. | 84 |
| Figure 2.28 | Photograph of an image from the Frobin system..... | 85 |
| Figure 2.29 | Intensity profile observed by Frobin (1988)..... | 86 |
| Figure 3.1 | Area of the back to be measured..... | 100 |
| Figure 3.2 | Concept of the optical subsystem. | 101 |
| Figure 3.3 | The CCD camera with variable focal length lens. | 103 |
| Figure 3.4 | Back and reference images showing reference fringe. | 104 |
| Figure 3.5 | Patient positioned in frame for measurement. | 106 |
| Figure 3.6 | Patient standing in frame. | 107 |
| Figure 3.7 | The coordinate system for images. | 107 |
| Figure 3.8 | Optical geometry. | 108 |
| Figure 3.9 | Dark circle used to calculate pixel scaling factors..... | 110 |
| Figure 3.10 | The sequence of image processing operations..... | 114 |
| Figure 3.11 | Intensity profile through an image..... | 117 |
| Figure 3.12 | Depth-of-field for camera and projector. | 118 |
| Figure 3.13 | Typical failure of global thresholding. | 122 |
| Figure 3.14 | Thresholding based on a 40 x 40 rectangular area..... | 124 |
| Figure 3.15 | Comparison of local area threshold algorithms (m=40). | 126 |
| Figure 3.16 | Effects of altering the range threshold (for m=40, $n_g=64$)..... | 128 |
| Figure 3.17 | The effects of altering the size of the local area (for $T_r=15$, $n_g=64$). | 129 |
| Figure 3.18 | The binary edge image (inverted to improve visibility on printed page). | 130 |
| Figure 3.19 | Back and reference images after edge detection..... | 130 |
| Figure 3.20 | The back image before and after point filtering. | 131 |
| Figure 3.21 | Finding the reference fringes. | 133 |
| Figure 3.22 | Data structure for defining R_s | 135 |
| Figure 3.23 | The search matrix for the left half of the image. | 136 |
| Figure 3.24 | The search matrix for the right half of the image. | 138 |
| Figure 3.25 | Result of the indexing algorithm. | 138 |
| Figure 3.26 | Back image before and after trim editing. | 140 |
| Figure 3.27 | The back image before and after pixel editing. | 141 |
| Figure 3.28 | Surface representation..... | 145 |

| | | |
|-------------|--|-----|
| Figure 3.29 | The C data structure defining a z coordinate. | 145 |
| Figure 3.30 | The C declaration of the surface array for a 512 x 256 pixel image..... | 146 |
| Figure 3.31 | The status array after z calculation | 147 |
| Figure 3.32 | The status array after boundary location..... | 148 |
| Figure 3.33 | Assignment of pixels outside the back surface..... | 149 |
| Figure 3.34 | The status array after interpolation. | 151 |
| Figure 3.35 | A projection of the surface..... | 152 |
| Figure 3.36 | Status array after simplification..... | 153 |
| Figure 3.37 | Simplified layered system structure..... | 154 |
| Figure 4.1 | Foot-plate design. | 163 |
| Figure 4.2 | Configuration for accuracy measurement..... | 165 |
| Figure 4.3 | Graph of true versus measured values for the accuracy experiment..... | 167 |
| Figure 4.4 | Error in the mean measurement versus z position. | 167 |
| Figure 4.5 | Vertebral positions as displayed by the system. | 170 |
| Figure 4.6 | Measurement of ATI..... | 171 |
| Figure 4.7 | ATI map for the patient shown in Figure 4.5. | 171 |
| Figure 4.8 | ATI values obtained for the three sets of readings compared..... | 172 |
| Figure 4.9 | STAS readings produced by the system for the patient from Figure 4.5.... | 174 |
| Figure 4.10 | Projection of a back surface..... | 175 |
| Figure 4.11 | Contour maps of produced by the system..... | 175 |
| Figure 4.12 | ATI values for two independent measurements. | 176 |

List of Tables

| | | |
|------------|--|-----|
| Table 1.1 | Incidence of referrals by Burwell..... | 11 |
| Table 1.2 | Degree of rotation by category for the Nash and Moe method..... | 16 |
| Table 2.1 | Comparison of PSP and FTP phase measuring methods. | 69 |
| Table 2.2 | Comparison of methods. | 92 |
| Table 3.1 | Computational task for a square local area threshold. | 125 |
| Table 3.2 | Computational task for square versus linear local area thresholding..... | 127 |
| Table 3.3 | Approximate number of comparisons and execution time for thresholding. | 127 |
| Table 3.4 | Fringe colour coding for Figure 3.25. | 139 |
| Table 3.5. | Status values associated with surface coordinates. | 145 |
| Table 4.1 | Approximate breakdown of processing time. | 179 |
| Table 4.2 | File sizes used by the system for a typical dataset..... | 180 |
| Table 4.3 | Breakdown of approximate cost of the system. | 182 |
| Table 4.4 | Comparison of desired functionality with actual functionality..... | 184 |

Abbreviations

| | |
|--------|--|
| API | Application programming interface. |
| ARM | Acorn RISC machine. |
| ATI | Angle of trunk inclination. |
| BBE | Bitwise binary encoding. |
| BSS | British Scoliosis Society. |
| CCD | Charge Coupled Device. |
| CTAS | Crude trunk asymmetry score. |
| ECU | European currency unit. |
| FBCT | Formulator body contour tracer. |
| FFT | Fast Fourier transform. |
| FTP | Fourier transform profilometry. |
| GUI | Graphical user interface. |
| IFFT | Inverse fast Fourier transform. |
| ISIS | Integrated shape investigation system. |
| LCD | Liquid crystal display. |
| PSP | Phase stepping profilometry. |
| RISC | Reduced instruction set computer. |
| RISCOS | RISC Operating System. |
| SRS | Scoliosis Research Society (USA) |
| STAS | Standardised trunk asymmetry score. |
| SWI | Software interrupt. |
| TAS | Trunk asymmetry score. |
| VGA | Video graphics array. |

Symbols

| | |
|---------------------------|---|
| γ | Angle between the normal to a surface and the incident light. |
| Γ | A constant related to the angle between the normal to a surface and the incident light. |
| θ | Angle between incident and reflected light. |
| $\Delta\phi$ | Phase shift. |
| $\phi(x,y)$ | Phase of a fringe pattern as a function of x and y. |
| θ_{ATI} | Angle of trunk inclination (ATI). |
| θ_i | Angle of incident light. |
| θ_r | Angle of reflected light. |
| θ_1, θ_2 | Angles of light incident on imaging plane of camera. |
| Δx | Difference between the position of a fringe on the surface and the position of a fringe on the reference plane, in the x direction. |
| $\Delta x'$ | Equivalent to Δx , but measured at the imaging plane of the camera. |
| $\Delta x_1 - \Delta x_5$ | Fringe spacings for fringes surrounding a candidate for reference fringe. |
| Δx_{hi} | The mean separation of the two fringe pairs above the reference fringe. |
| Δx_{lo} | The mean separation of the two fringe pairs below the reference fringe. |
| Δx_m | The separation of the middle pair of the set of 6 fringes used in locating the reference fringe. |
| Δz | Fringe depth separation. |
| a | Incident light ray. |
| A | Incident light source. |
| A | Back-projection of the position of a fringe on the object onto the reference plane. |
| a,b,c,d,e,f,g,h | Pixel locations used for the fringe tracking search matrix. |
| A/D | Analogue to digital. |
| A' | Position of a fringe on the object in the imaging plane of the camera. |
| A(x,y) | Background intensity. |
| b | Reflected light ray. |
| B | Position of observer. |
| B | Position of a fringe on the reference plane. |
| B(x,y) | Amplitude of cosine carrier. |
| B' | Position of a fringe on the reference plane in the imaging plane of the camera. |
| C | Optical centre of the camera. |
| D | Separation of observer and light source. |

| | |
|---------------------------------|---|
| d | Displacement associated with a phase shift. |
| D | Diameter of the dark circle used for the measurement of pixel scaling factors. |
| D_C | Distance from projector/camera to reference plane. |
| D_P | Separation of camera and projector. |
| D_x | x dimension of an image in pixels. |
| D_y | y dimension of an image in pixels. |
| f | Frequency. |
| f_0 | Fundamental frequency of fringes. |
| h | The distance from the source/observer to the grating. |
| $I(x,y)$ | Intensity, usually in an image, as a function of x and y. |
| $I_0, I_{90}, I_{180}, I_{270}$ | Phase shifted intensity values. |
| I_A | Intensity of light striking a grating in moiré. |
| I_B | Intensity of light at the observer in moiré. |
| $I_d(x,y)$ | Intensity of a pixel at position (x,y) in a destination digital image, following an image processing operation. |
| $I_s(x,y)$ | Intensity of a pixel at position (x,y) in a source digital image, prior to an image processing operation. |
| I_{smax} | The maximum intensity in a source image, prior to an image processing operation. |
| I_{smin} | The minimum intensity in a source image, prior to an image processing operation. |
| L_x | Scaling factor relating the size of a pixel in the x direction to real measurements in mm. |
| L_y | Scaling factor relating the size of a pixel in the y direction to real measurements in mm. |
| m | Dimension of a local area of interest in an image. |
| n | Fringe order. |
| n_c | Numbers of comparisons required during an image processing operation. |
| n_g | Number of grey levels used to represent a digital image. |
| O | Origin of coordinate system. |
| p | Period of a grating. |
| P | Optical centre of the projection system. |
| P | The current pixel in an image processing operation. |
| $P(x,y,z)$ | A point on the surface with coordinates (x,y,z). |
| p_i | Period of incident fringe pattern. |
| p_R | Period of reflected fringe pattern. |
| p_x | Period of light on a reference plane, in the x direction. |
| R | Constant related to distance from object to light source. |

| | |
|-------------------------------------|--|
| R | Radius of the dark circle used for the measurement of scaling factors. |
| r | Range of intensity values in a local rectangular area of interest in an image. |
| $r(x,y,z(x,y))$ | Distance of object from light source. |
| R_s | Ratio of the reference fringe spacing to the mean spacings of the two fringe pair immediately above and below it. |
| S | A surface element. |
| T | Intensity threshold value. |
| T_g | Global intensity threshold value. |
| $T(x)$ | Transmission of a grating. |
| $T(x,y)$ | A local area intensity threshold. |
| T_r | Intensity range threshold. |
| x, y, z | Cartesian coordinates. |
| $X[I(x,y),y]$, $X[I,y]$ | The fringe table. |
| $x_1, x_2, x_3, x_4, x_5,$ x_6 | x positions of the six fringes used in locating the reference fringe. |
| x_1, y_1, x_2, y_2 | Positions of the edges of the dark circle used for the measurement of pixel scaling factors. |
| x_1, y_1, x_2, y_2 | Boundaries to rectangular area of interest in an image. |
| $X_{back}[]$ | Fringe table entry for back surface. |
| $X_{reference}[]$ | Fringe table entry for the reference plane. |
| $z(x,y)$ | The depth coordinate of a surface as a function of the x,y position on the surface. |
| z_n | Depth coordinate for nth fringe in moiré contouring. |
| z_s | z, or depth, coordinate of a surface element, measured from the reference plane. |
| z_s | The z, or depth, coordinate of a surface element measured from the reference plane. |
| $z_s(x,y)$ | The z, or depth, coordinate of a surface element, measured from the reference plane, as a function of x,y position. Includes both the actual depth value and a status value. |
| $z_s(x,y).s$ | Status field of $z_s(x,y)$. |
| $z_s(x,y).z$ | Depth coordinate field of $z_s(x,y)$. |
| $z_{SD}(x,y).z$ | The destination surface image in surface image processing operations. Has status and z fields. |
| $z_{SD}(x,y).s$ | |
| $z_{SS}(x,y).z$ | The source surface image in surface image processing operations. |
| $z_{SS}(x,y).s$ | Has status and z fields. |

Chapter 1

Introduction

1. Introduction

1.1 Background

Scoliosis is a condition in which the spine becomes distorted and exhibits curvatures of various types. Since the spine is one of the major supporting structures of the body the consequences for the individual may be serious. Scoliosis will impair quality of life and in severe cases will present a major risk to health. The condition affects only a small percentage of the population and little is known about its cause. It principally affects young children in infancy and girls in adolescence.

The curvature of the spine in scoliosis is usually accompanied by rotation of the vertebrae and a distortion of the rib cage. The result is that a rib hump appears in the thorax. Naturally, the individual concerned is often very aware of their rib hump and may be keen to undergo treatment to reduce both the rib hump and the underlying structural deformity.

In common with many conditions, if the scoliosis can be detected in its early stages then there is a likelihood that effective treatment may be administered. Special attention is given to children and especially adolescents because it is in this age range where the scoliosis is most likely to progress and become more severe. Some hospitals have specialist units for monitoring and treating children and adolescents. Treatment of scoliosis is often unpleasant for the child or adolescent because it typically involves either major surgery or wearing of an external brace for a considerable time. Mild scolioses are often left untreated unless there is evidence that the condition is progressing and is likely to result in a severe deformity. A scoliosis clinic will usually monitor a patient over a number of months or years and record any progression in the condition. Clearly there must be an objective means of assessing the extent of the scoliosis in a patient, preferably by the measurement of parameters of scoliotic deformity, so that changes in the condition can be monitored.

The conventional method for assessing scoliosis is to acquire a full length spinal radiograph [X-ray] in which the shape of the spine, the rotation of the vertebrae, and

sometimes the rib-cage distortion, can be seen. Objective parameters of scoliotic deformity can then be measured from the radiograph. The hazards of radiation are well known and it is undesirable, and in some cases unacceptable, to take full length spinal radiographs of a patient every few months to assess the progression of the condition.

An alternative approach to monitoring the deformity is to base the assessment on the surface shape of the back. This approach is taken by a number of scoliosis centres. The measurement of the back surface shape is performed using crude measurement devices that are placed on the surface and which, after the processing of measurements, produce quantified parameters of deformity. The measurement of back shape in this way is time consuming and laborious, and there is clearly a need for an automated system to replace the manual measurement process.

1.1.1 Objectives of the Investigation

In summary, the objectives of this investigation were:

- i. The primary purpose was to investigate automated, non-contact, optical and computer methods for measuring the shape of the human back to produce a complete three-dimension reconstruction of the back and, if appropriate, to propose a new method of measurement.
- ii. The secondary purpose was to demonstrate that the method selected is capable, with suitable further processing, of producing parameters of clinical significance that relate to scoliotic deformity.
- iii. In addition, to devise an automated system to implement the method selected in a working clinical system and to verify that the system is a usable clinical tool in the assessment of scoliosis.

The detailed criteria for an automated system are addressed in Section 1.8 at the end of this chapter. Clearly, these will complement the objectives and restrict methods of approach.

1.1.2 Scope of the Investigation

There were two main boundaries in the scope of the investigation. These were:

- i. The investigation was not to comment upon the nature, incidence, cause or treatment of scoliosis (except for background information) but was to focus on the measurement problem. It was likely that the system devised would be capable of producing clinical information of value in individual cases or statistically, but a detailed assessment of the clinical data lies in the medical, not the engineering, domain.
- ii. The investigation examines optical methods for acquiring data that is subsequently processed by computer to reproduce surface shape. Only those methods that can be implemented using a CCD camera/frame-store combination yielding a digital image that can be processed by a standard micro-computer are considered. The development of specialist sensors or computer hardware architectures was beyond the scope of the investigation.

1.1.3 Collaboration

The investigation has been a collaborative project between Liverpool John Moores University, School of Electrical and Electronic Engineering and Physics, and The University of Liverpool, Department of Human Anatomy and Cell Biology. The original identification of need for an automated system to measure body shape for scoliosis assessment came from Dr. P.H. Dangerfield (supervisor to this Ph.D. programme) of The University of Liverpool through his work in scoliosis at the Royal Liverpool Children's Hospital who have also collaborated with the investigation and provided facilities and patients for early clinical trials.

1.1.4 Structure of this Thesis

This chapter, Chapter 1, introduces scoliosis. The basic anatomy of the spine is discussed and the conventional methods for measuring the extent of the deformity, including the methods that rely on surface shape measurement, are reviewed. The need

for an automated system to measure body shape is justified. Finally, the criteria for a practical automated system are identified.

Chapter 2, *Optical Methods for Measurement of Back Shape*, focuses on the optical and image processing methods that can be used for the measurement of the shape of the human back including discussion of moiré, phase measuring, scanning and raster projection methods. The various methods are compared and the advantages and shortcomings of the methods are identified. The most suitable method is selected for further investigation.

Chapter 3, *A Structured Light Method for Measuring Back Shape* describes a new method that was devised during the investigation and which was developed to the level of a working clinical tool. The optical configuration and the image processing techniques that were used are described in detail.

Chapter 4, *System Evaluation and Clinical Application* evaluates the measurement system and describes its application to measuring the surface deformity of scoliosis. The chapter returns to the criteria that are identified at the end of Chapter 1 and evaluates the system by determining whether the criteria have been satisfied.

Chapter 5, *Conclusions and Recommendations for Further Work* reviews the investigation, draws conclusions from the work, and identifies areas for further research.

1.2 A Background to Scoliosis

Scoliosis is defined to be a lateral curvature of the spine (James 1976). Between 80% and 90% of cases are idiopathic [the cause is unknown], although the deformity may be associated with other spinal anomalies such as muscular dystrophy and polio (Moe 1978). Before examining the condition and methods for its assessment it is appropriate to familiarise the reader (especially the engineer or non-medical scientist) with the terminology and nature of the spine and the deformity.

1.2.1 Anatomical Planes of the Body

In discussion of deformities of the body, and in particular of scoliosis, it is necessary to define the anatomical planes in which the deformity occurs. The three planes that are commonly used are the transverse (or lateral) plane, the sagittal plane, and the coronal plane which are mutually perpendicular and are shown in Figure 1.1.




Figure 1.1 The anatomical planes of the body.

1.2.2 The Basic Anatomy of the Spine

The human spine is a flexible column formed by a series of bones called vertebrae. There are thirty three vertebrae which are divided into five general groups: the cervical, thoracic, lumbar, sacral and coccygeal vertebrae. Figure 1.2 shows the human spine and the five groups of vertebrae.

The vertebrae in the cervical, thoracic and lumbar regions of the spine are completely separate but in the lower two regions, the sacral and coccygeal vertebrae are fused to form two bony units: the sacrum and the coccyx. Inter-vertebral disks separate the individual vertebrae enabling spinal articulation and also act as dynamic shock absorbers. The seven cervical vertebrae are usually indexed from the top as C1 ... C7, the twelve thoracic vertebrae similarly T1 ... T12 and the lumbar vertebrae L1 ... L5.




Figure 1.2 The spine and the groups of vertebrae.

The spine is a major structural component of the body, supporting the trunk and skull. It also protects the spinal cord and blood vessels which pass within the spinal foramen, a hollow in each vertebra along the spine. Each vertebra has a number of protruding processes, such as the transverse processes, which serve, among other functions, for the attachment of muscles and ligaments. Figure 1.3 shows the seventh cervical vertebra which has an unusually large spinous process.

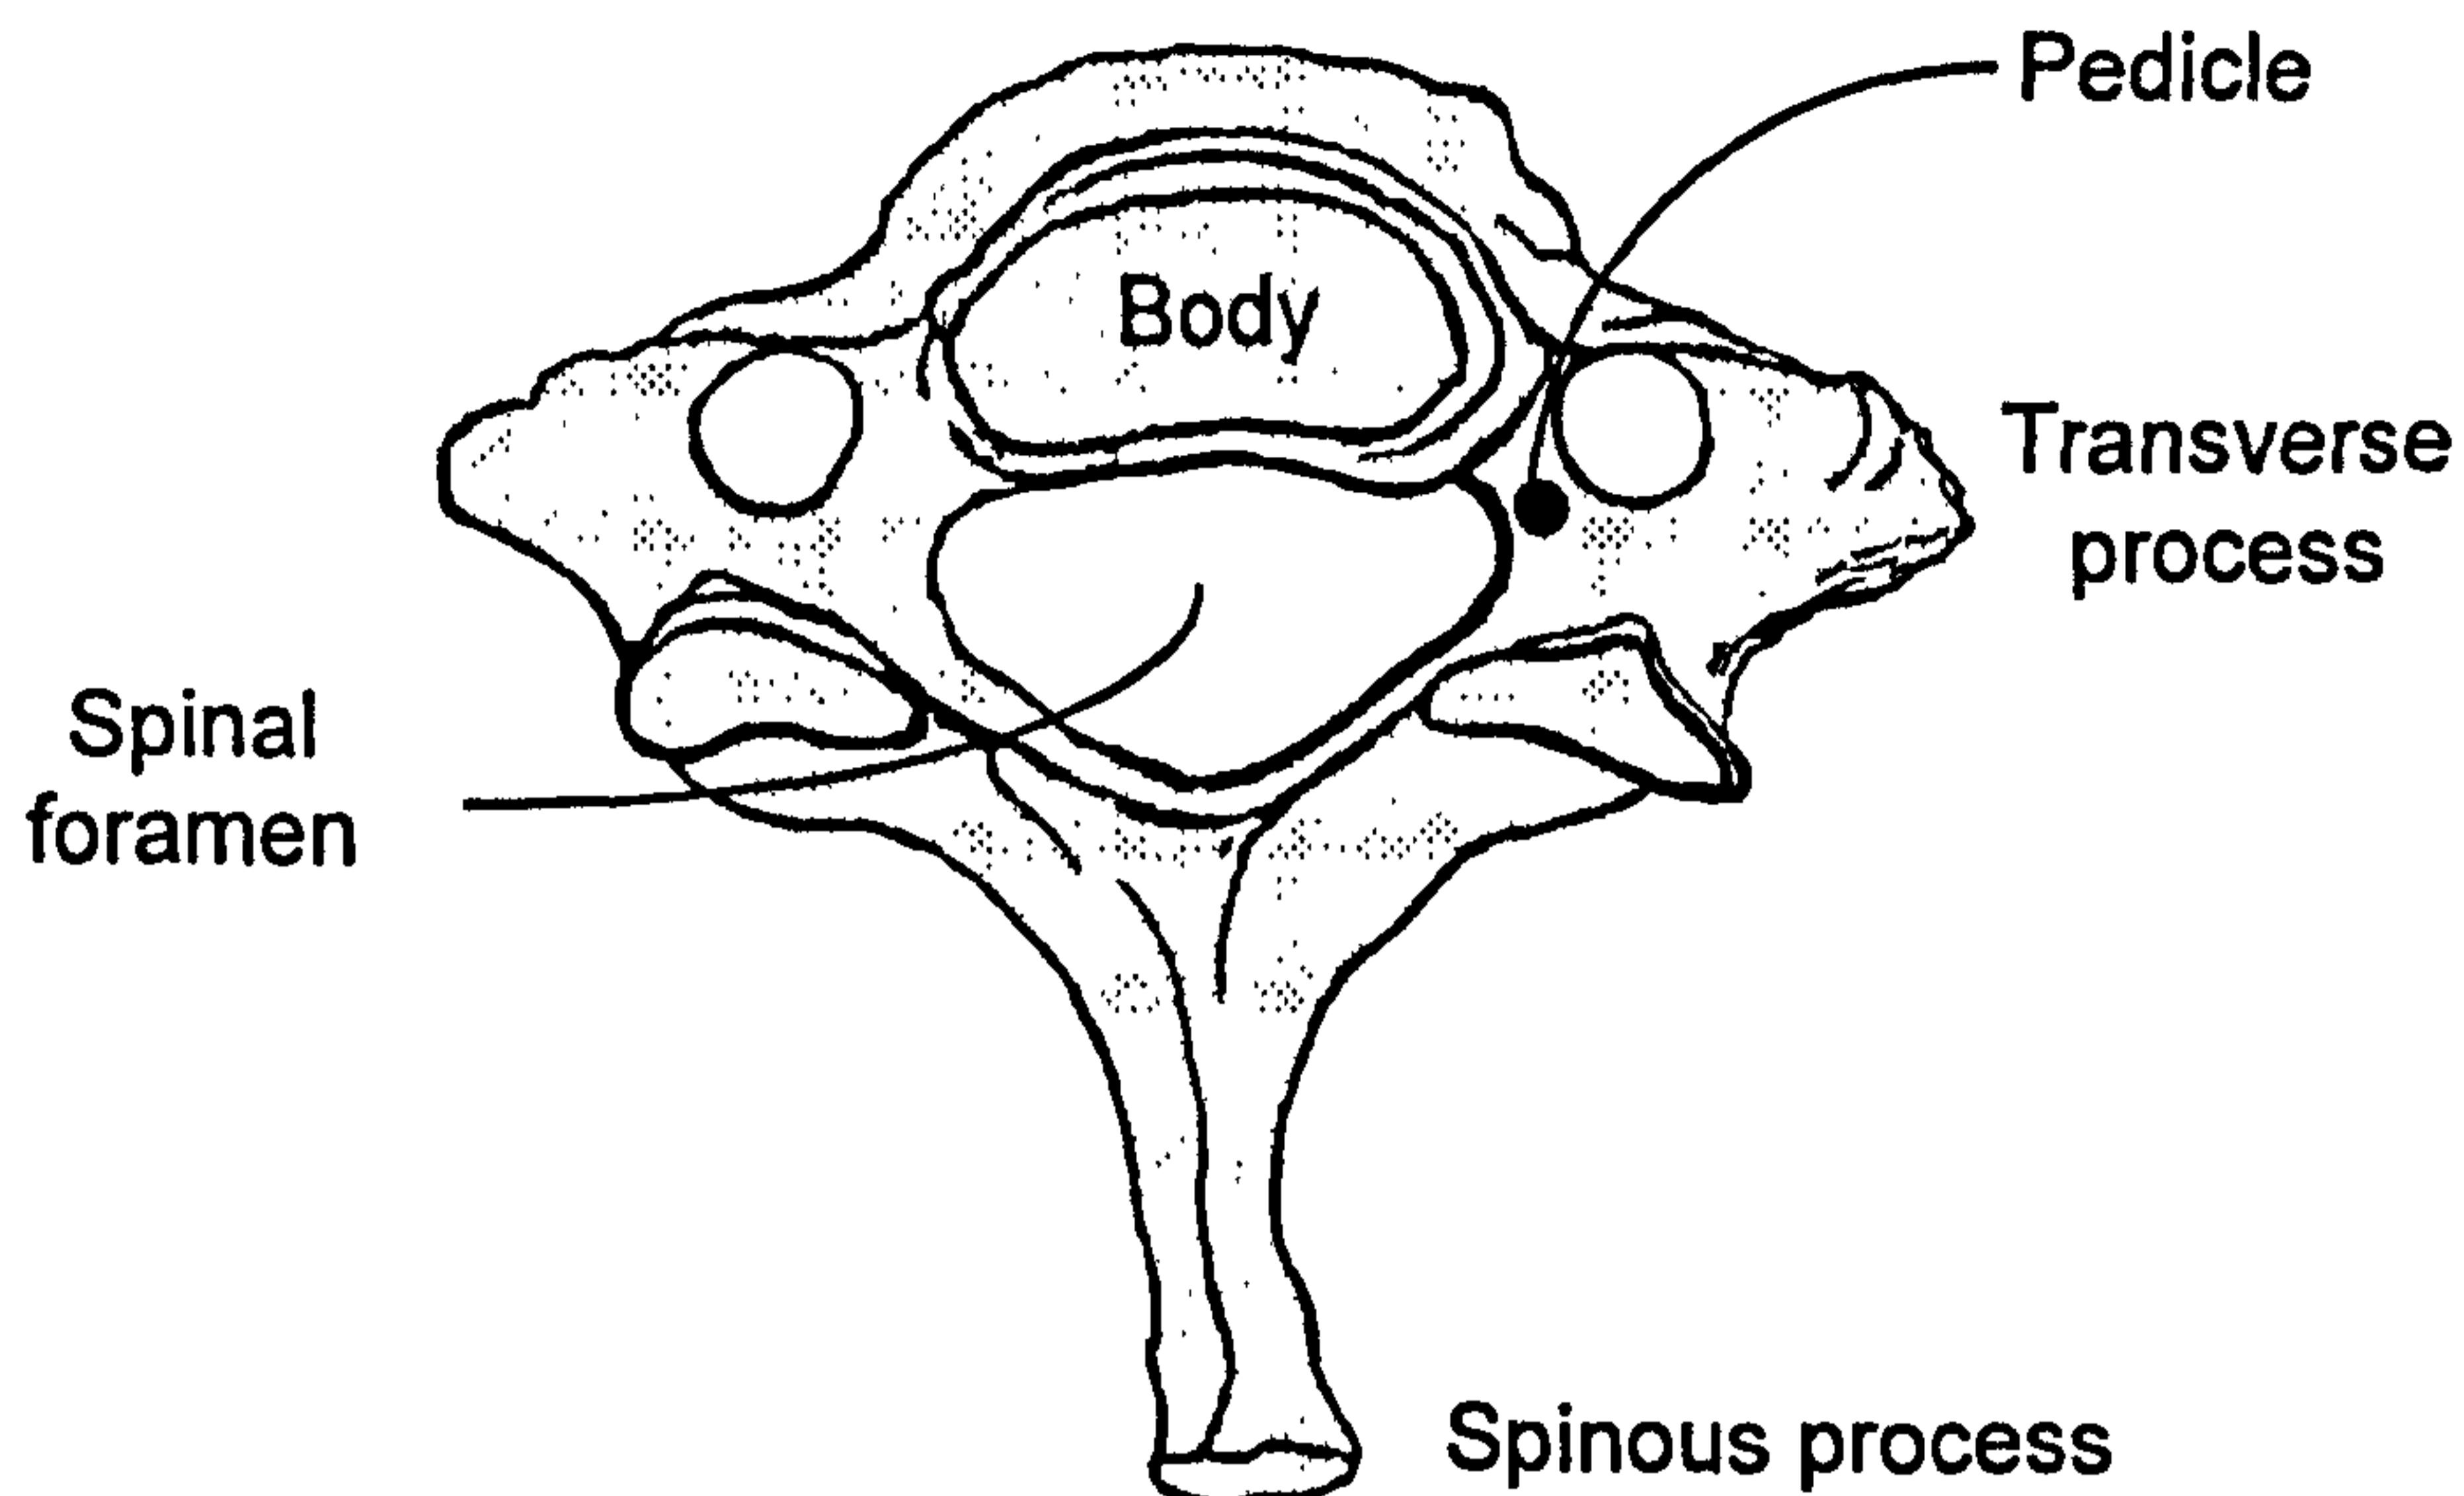


Figure 1.3 The seventh cervical vertebra.

The spinous processes of vertebrae are of particular interest in this investigation because they are the structures of the spine which may be most easily identified on the surface of the back. Often they will protrude slightly and the end of the process will form a small convex bump on the back which can be seen visibly or at the very least located by palpation. A lateral (or side) view of a vertebra, in this instance a thoracic vertebra, is shown in Figure 1.4 where the spatial relationship between the body of the vertebra and the spinous process can be more clearly seen.




Figure 1.4 A lateral view of a thoracic vertebra.

Simple consideration of the human form suggests that the spine will be subject to substantial and complex biomechanical loading and that deformities or abnormalities of the spine may have serious consequences for the individual concerned.

1.2.3 Scoliosis: a curvature of the spine

Strictly, scoliosis is defined to be a lateral curvature of the spine (James 1976).




Figure 1.5 A single structural curve with compensatory curves.

Typically, the scoliosis will exhibit either a single or double structural curve and will be accompanied by compensatory curves above and below the structural curves to maintain the general upright position of the individual. Figure 1.5 and Figure 1.6 show single and double structural curves respectively.

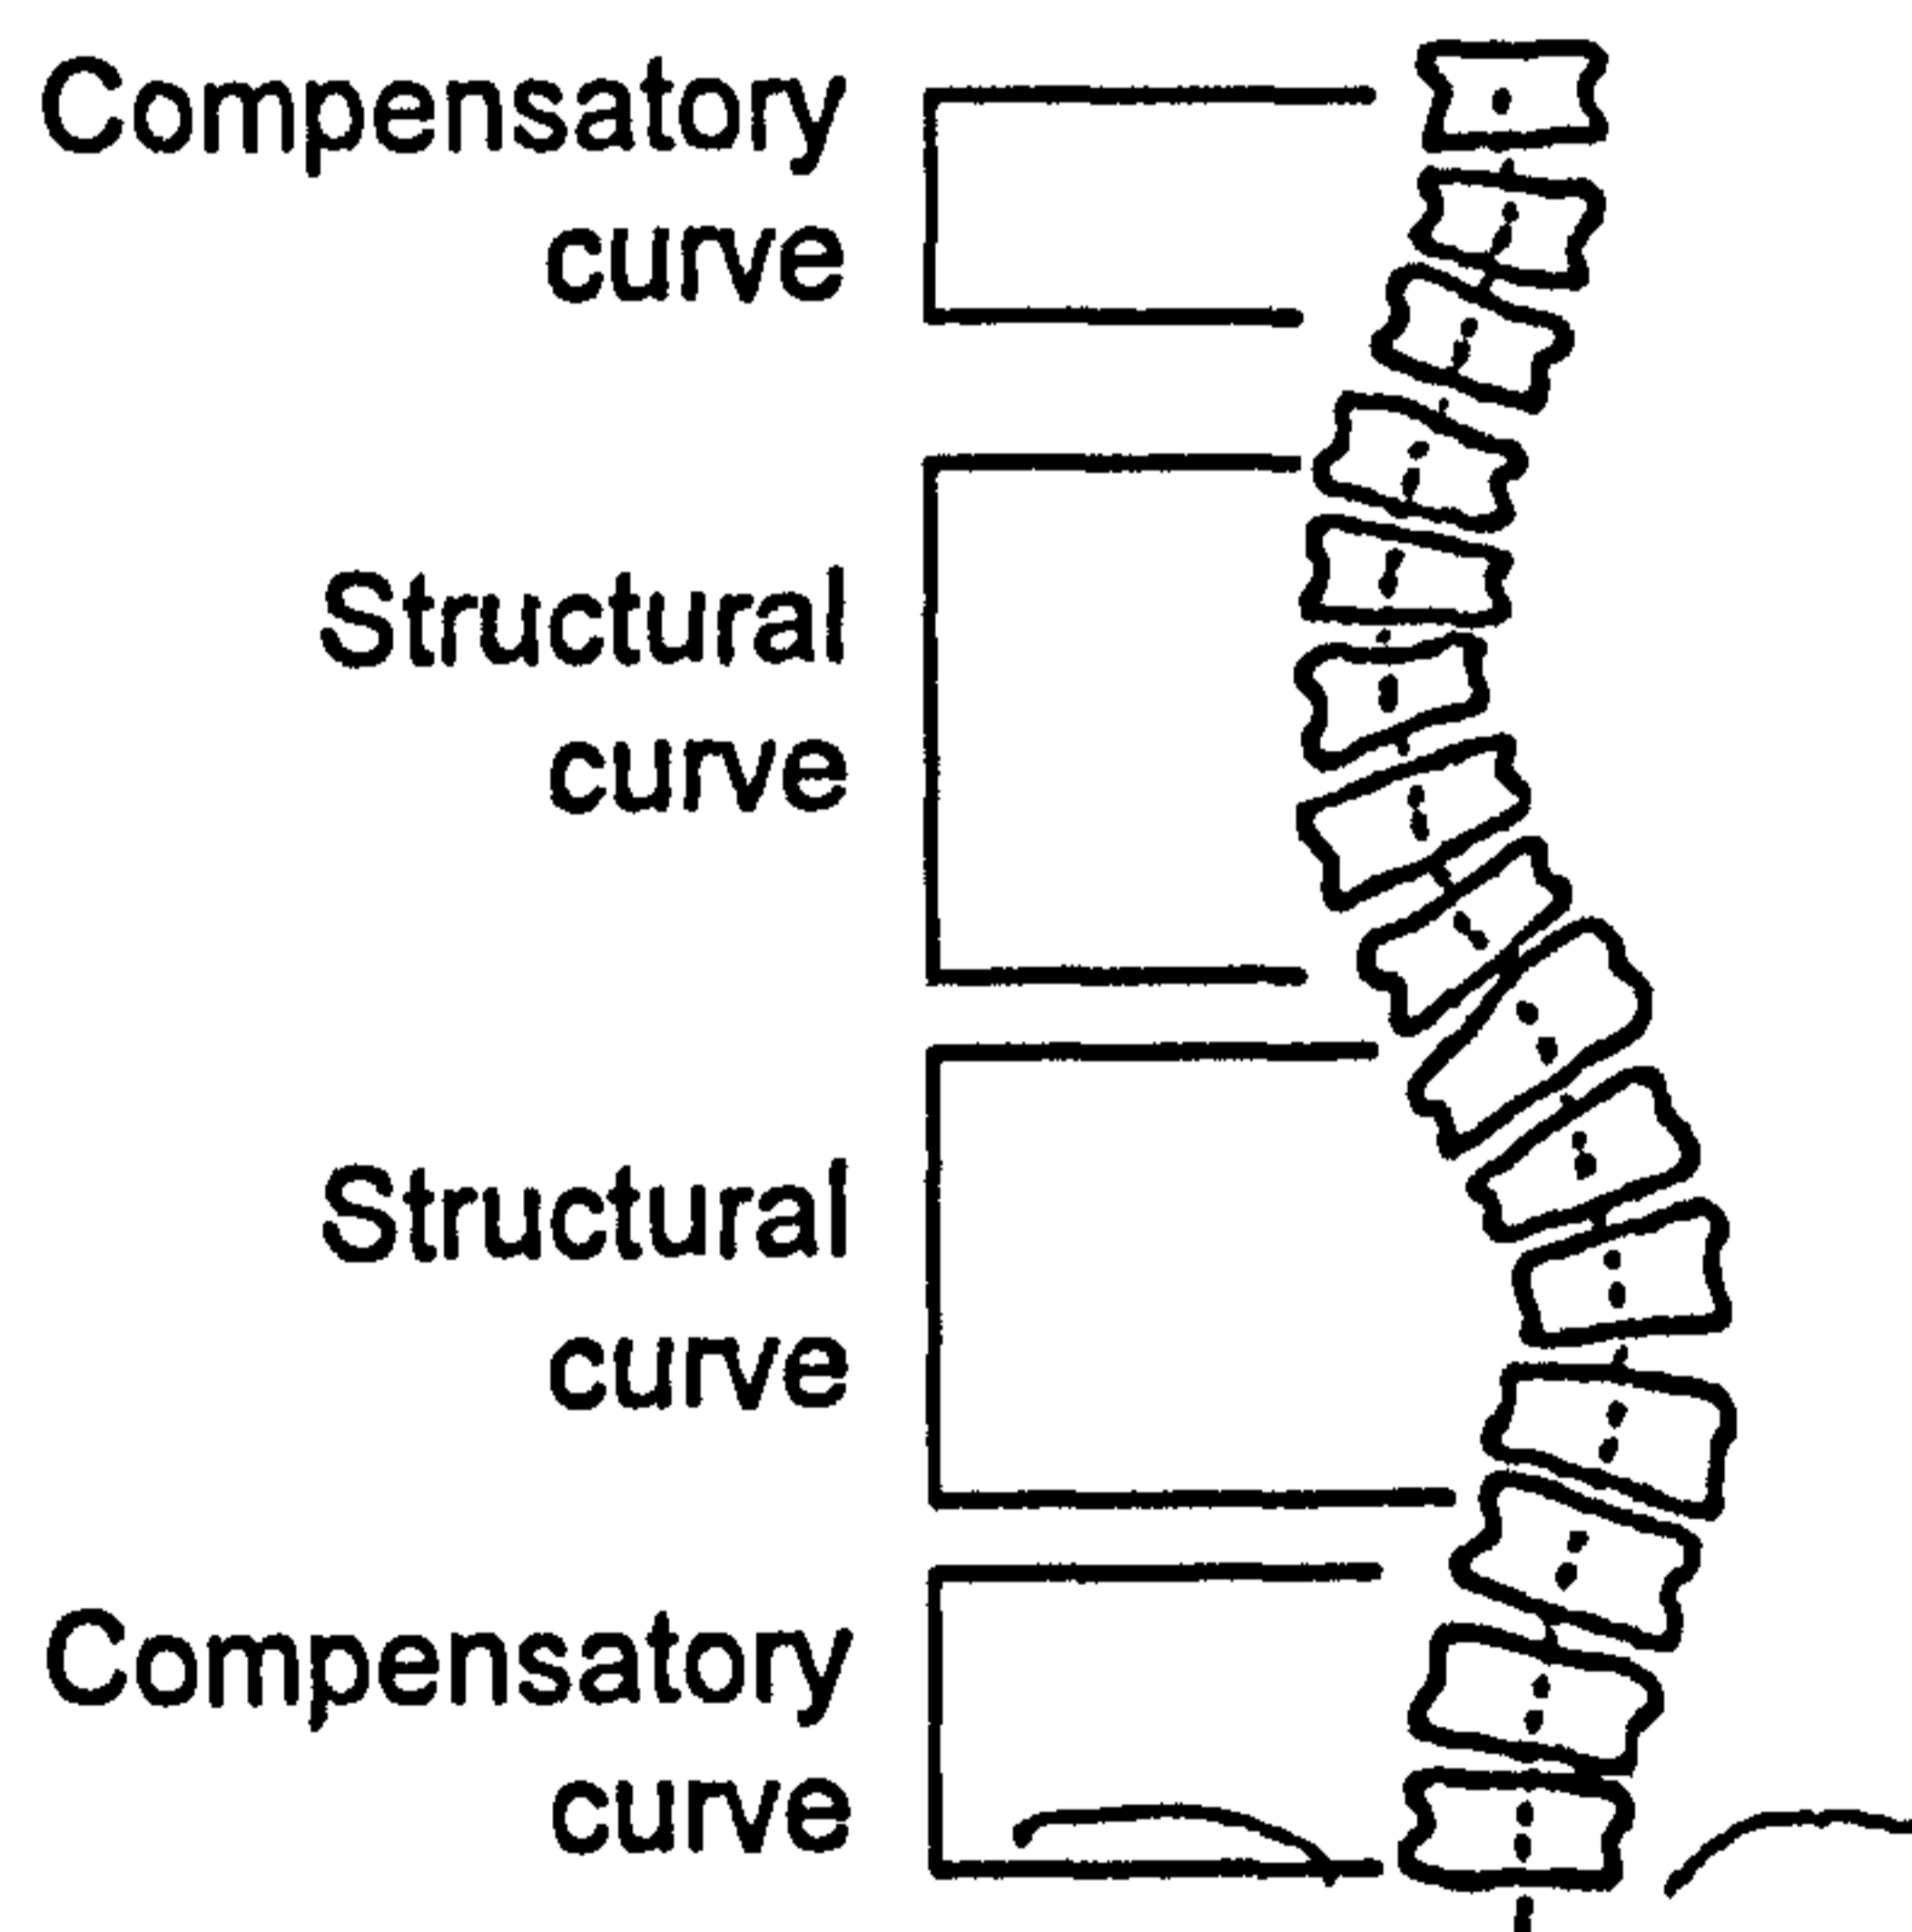


Figure 1.6 A double structural curve with compensatory curves.

This simple two-dimensional description of scoliotic deformity does not provide a complete representation of the condition. A better understanding may be gleaned from a

consideration of the complete three-dimensional distortion of the spine and associated structures. Two additional phenomena are important in the three-dimensional assessment of scoliosis: vertebral rotation and rib cage distortion.

1.2.3.1 Rib Cage Distortion

Because the ribs are attached to the spine the lateral distortion of the spine must cause a corresponding distortion to the rib cage. The distortion to the rib cage usually manifests most obviously as a hump on one side of the patient's back which is often more prominent in the forward bending position. The surface measuring methods for assessing scoliosis that are described later in this chapter rely on measuring the asymmetry of the back due to the hump.

1.2.3.2 Vertebral Rotation

Scoliosis is accompanied by a rotation of the vertebrae about the axis of the spine. The rotation of the vertebrae is linked to the distortion of the rib cage and will contribute to the asymmetry in the shape of the back. Some methods for assessing scoliosis depend upon measuring vertebral rotation (see Section 1.3.2).

1.2.3.3 The Three-dimensional Deformity

There is an increasing opinion in the scoliosis research community that the condition should be considered from a three-dimensional perspective both for the purposes of research into the causes of the condition and in routine clinical practice (Dansereau 1992). The simple definition of scoliosis as a lateral curvature cannot be considered in isolation as a basis for the assessment of scoliosis. Indeed, methods exist (see Sections 1.3 and 1.5) for assessing the extent of scoliosis from vertebral rotation and/or rib-cage distortion.

1.2.4 The Incidence of Scoliosis

An assessment of the incidence of scoliosis is difficult because the criteria for determining when a spine is classed as scoliotic, rather than within the bounds of normality, are vague. Bunnell (1984) notes that reports on the incidence of scoliosis in

schoolchildren have produced figures varying from 1 to 21 percent of the sampled population.

The incidence may vary with age group and locale. This investigation has been in collaboration with the Royal Liverpool Children's Hospital and consequently it is the incidence of scoliosis among children in the United Kingdom that is of most interest.

One recent study was conducted by Burwell (1990), who examined 4890 schoolchildren in an experimental screening programme in Nottingham, UK. The children were in their second year of comprehensive education and, therefore, in the age range 12-13 years. There was a total of 2392 boys and 2498 girls who comprised 96% of the school population in the age range. The study was particularly pertinent to this investigation since it was based almost entirely on surface shape measurements taken using devices such as the Scoliometer (see Section 1.5.1).

| | Tier 1 | Tier 2 | Tier 3 |
|--------|--------|------------|--------------|
| Boys | 2392 | 142 (5.9%) | 8* (0.3%) |
| Girls | 2498 | 180 (7.2%) | 32** (1.3%) |
| Totals | 4890 | 322 (6.6%) | 40*** (0.8%) |

* Not including 1 boy with juvenile idiopathic scoliosis, already being treated

** Not including 1 girl with congenital scoliosis, already being treated.

Table 1.1 Incidence of referrals by Burwell.

Three tiers of examination were adopted in which children identified as having possible scoliosis were referred to the next tier. Tier 1 involved 48 school nurses who were specially trained in the use of simple measurement instrumentation. The scoliosis screening was incorporated into the standard health appraisal of 12-year-old children. There were 322 children marked for further examination and referred to Tier 2.

Tier 2 involved six specialist school nurses with additional training in scoliosis and the instrumentation. Similar, but more probing, tests to those conducted at Tier 1 were used and 40 of the 322 children seen at this tier were referred on to Tier 3. Tier 3 was a detailed clinical examination performed by a specialist at a scoliosis clinic. Table 1.1 summarises Burwell's results and gives a breakdown of the proportion of boys and girls referred to the various stages.

The results at Tier 3 were not necessarily cases of scoliosis requiring treatment or further action, but they represent an indicator of the likely incidence of scoliotic children in the United Kingdom. In summary, it can be said that 0.8 % of school children in the 12-13 year old age group have some degree of scoliosis with a ratio girls: boys of 4:1. Burwell commented that the incidence of referrals to Tier 3 was low and that high ATI thresholds had been used to reduce the number of false-positive assessments.

1.2.5 The Cause of Scoliosis

The vast majority (about 90%) of scolioses are idiopathic [their cause is unknown]. Byrd (1988) summarises current knowledge on the cause of the condition as follows:

The cause of idiopathic scoliosis remains unknown, although research has possibly eliminated some hypothetical causes. Abnormalities of disc, bone, muscle, and collagen do not appear to be aetiological factors but, rather, reflect the effects of scoliosis on normal tissues. Although most patients with idiopathic scoliosis are thought of as tall and slender, it has not yet been proven that growth in this condition differs from normal. It is possible that idiopathic scoliosis is caused by multiple factors rather than a single factor. Perhaps there are different factors for curve initiation and curve progression. It appears that a brain stem or equilibrium abnormality does exist with patients with idiopathic scoliosis, although more research is needed to confirm and define the problem. There may also be a genetic basis for idiopathic scoliosis, as it does seem to occur within patients' families more frequently than in the general population. Whether these two factors form the basis for a multi-factorial aetiology of idiopathic scoliosis remains to be proven.

[Aetiology is the science of the causes of a disease or condition]. It appears, therefore, that very little is known with any certainty about the causes of the condition.

1.2.6 Treatment of Scoliosis

A number of treatment options are available. One option is to place the body in a rigid brace (such as the "Milwaukee brace") which places pressure on one side of the body in an attempt to correct the deformity. The alternative, but more severe option is to intervene with spinal surgery to correct the deformity more directly. A variety of operations are performed: in some cases vertebrae are fused and in others (such as the

Harrington-Luque or Cotrel-Dubousset operations) rigid metal rods are used to support the spine. The treatment of scoliosis was not of major concern to this investigation.

1.3 Conventional Radiographic Assessment of Scoliosis

Traditionally, clinical assessment of scoliosis often includes repeated radiographic examination of the spine using postero-anterior radiographs. The dangers of radiation are well known. Although the radiation dose can be minimised, patients can be subject to considerable dosage when the progression of the deformity is to be assessed (perhaps at three or six month intervals) over many months or years during monitoring, treatment and follow-up.




Figure 1.7 Radiograph of a thoracic structural curve (from James 1976).

Figure 1.7 shows an x-ray of a patient with a thoracic structural curve. The radiograph is a valuable tool that can be used by the clinician to see the underlying deformity of the spine. However, the radiograph alone does not provide a measurable parameter for scoliotic deformity. The radiograph must be processed manually by the clinician to measure appropriate parameters. Two types of parameter are typically measured: the angle of curvature and the vertebral rotation.

1.3.1 Angle of Curvature

The most popular method for measuring the angle of curvature for the spine is that proposed by Cobb (1948). Using Cobb's method the angle of curvature is measured by drawing lines parallel to the upper border of the upper vertebral body and to the lower border of the lowest vertebra of the structural curve, then erecting perpendiculars from these lines to cross each other, the angle between the perpendiculars being the Cobb angle of curvature. This procedure is shown in the schematic in Figure 1.8.

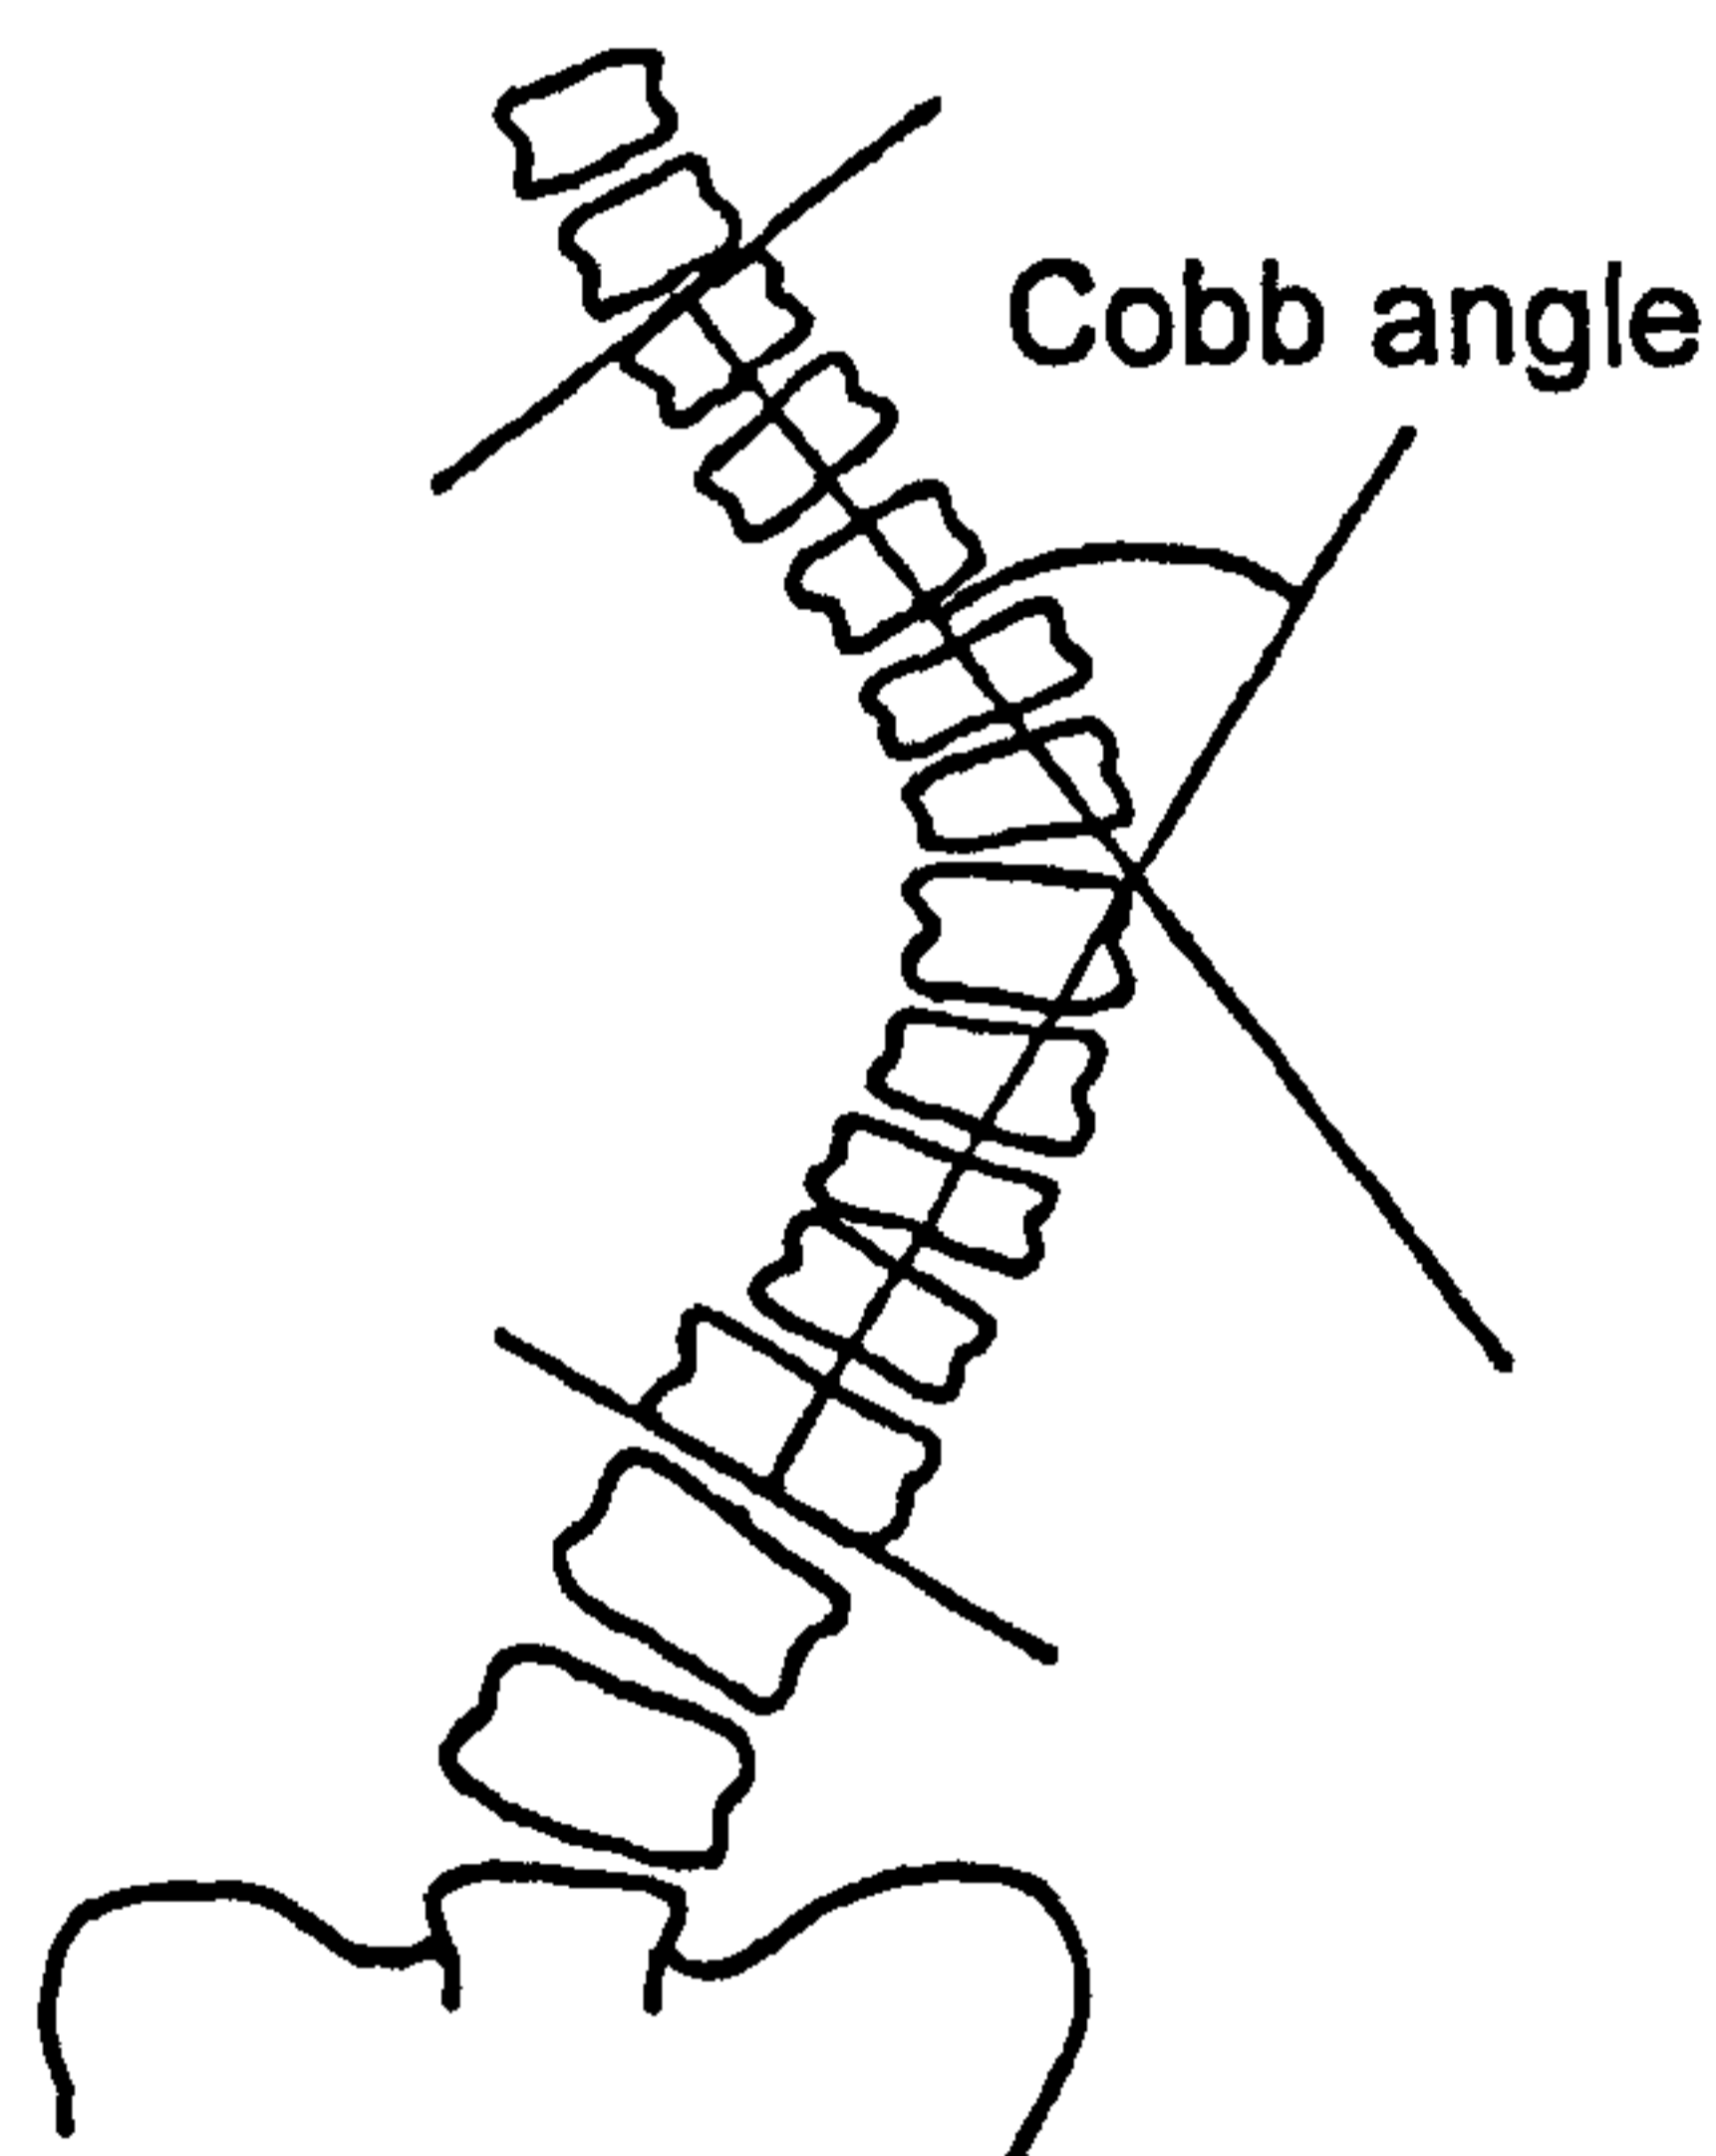


Figure 1.8 Cobb's methods for measuring angle of curvature.

James (1976) noted that the difference in measured angle when successive measurements are taken from the same radiograph may be as much as 5°. Some clinicians may observe an even greater difference.

1.3.2 Vertebral Rotation

Vertebral rotation is an alternative (and often supplementary) method for assessing the extent of scoliosis. Since the radiograph is a two dimensional projection of the bone structure and has limited visibility, the measurement of the rotation of the vertebrae is difficult. A number of methods have been proposed (Cobb 1948, Perdriolle 1979).

Nash and Moe (1969) based their method on the position of the pedicles relative to the vertebral body. [The pedicles are two short, thick pieces of bone which project backward, one on each side, to fuse in the mid-line and form the posterior components

of the vertebral arch surrounding the spinal cord. They can be clearly seen on a radiograph].

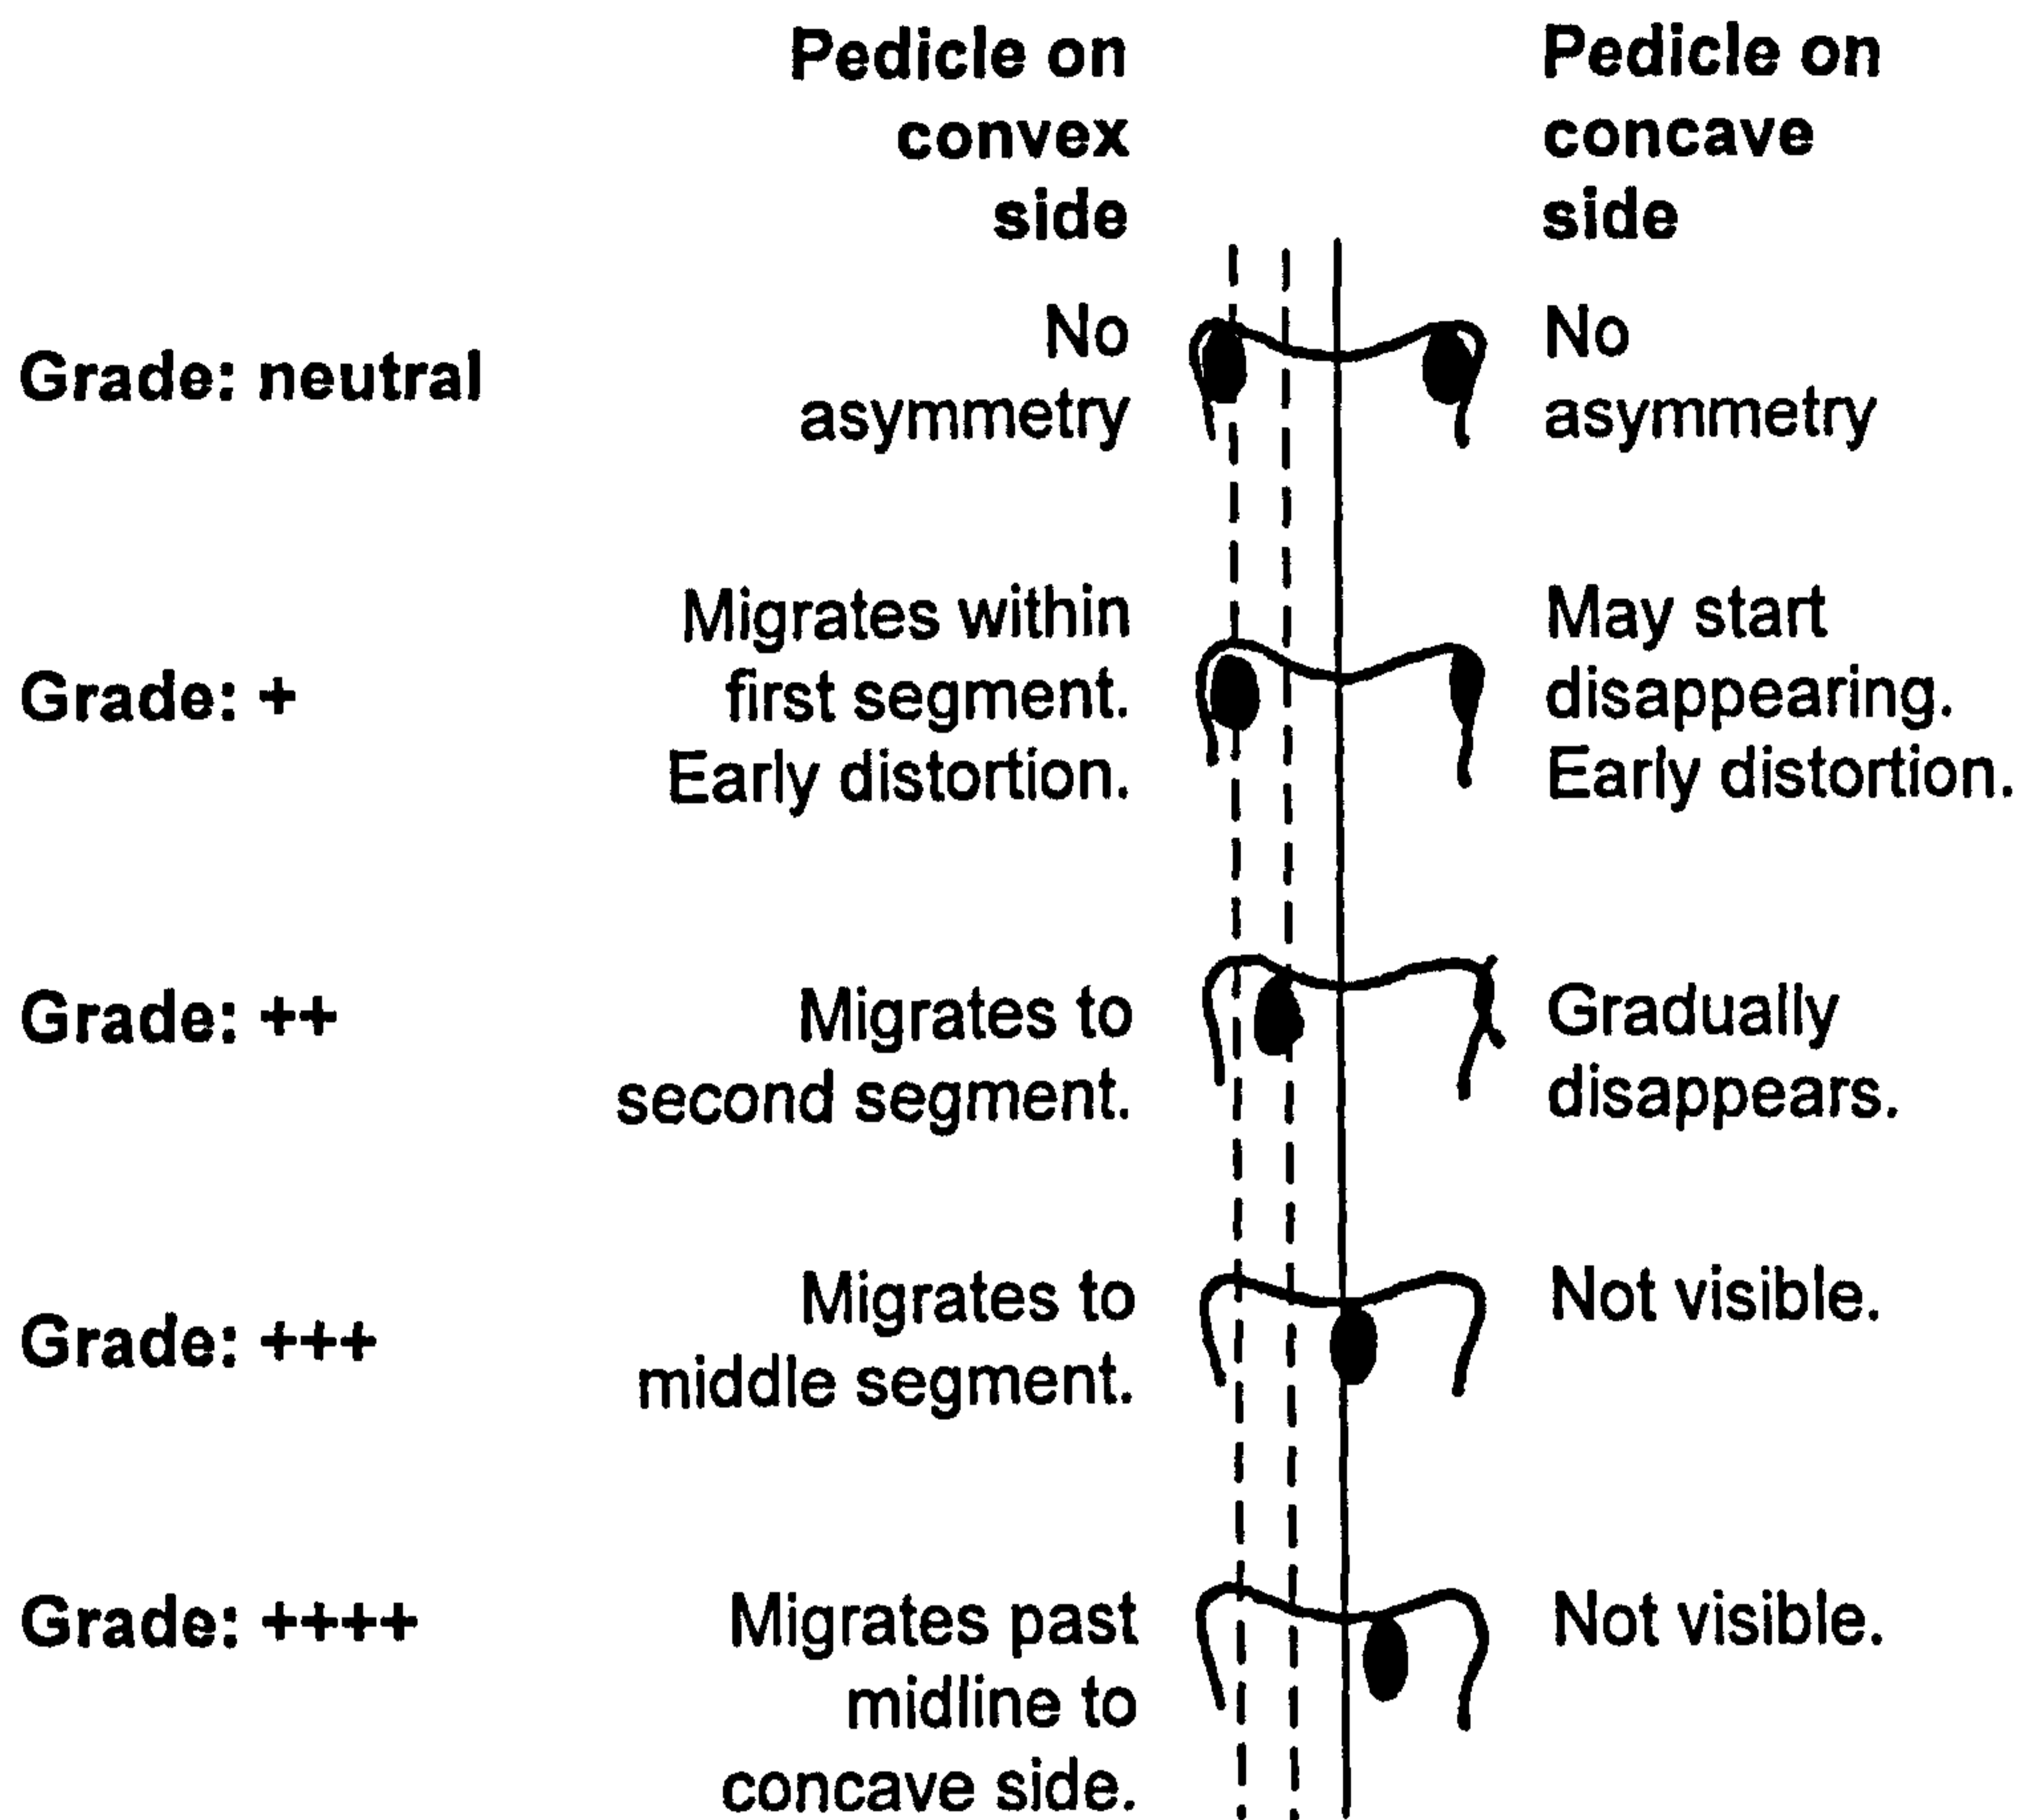


Figure 1.9 The Nash and Moe method for measuring vertebral rotation.

On a straight spine, the pedicles are projected symmetrically and ovals in the upper corners of the vertebral bodies. When the vertebrae rotate the pedicle on the convex side of the curve shifts towards the concave side of the curve and the pedicle on the concave side of the curve disappears. Figure 1.9 shows the migration of the pedicles. The method places the rotation into five categories: Neutral, +, ++, +++ and ++++. The degree of rotation assigned to each category is shown in Table 1.2.

Perdriolle (1979) used a variation on the technique for measuring pedicle position. He devised a simple protractor, called a Torsiometer, that consisted of a number of line constructions drawn on translucent plastic that could be placed on the radiograph and the rotation read directly from a scale.

| | | | | | |
|-----------|---------|--------|---------|---------|------|
| Grade* | Neutral | + | ++ | +++ | ++++ |
| Rotation* | 0° | 5°-10° | 15°-30° | 35°-40° | >50° |

* Some measurements are intermediate between grades

Table 1.2 Degree of rotation by category for the Nash and Moe method.

The methods used by Nash and Moe, and Perdriolle, depend crucially upon the ability to identify the pedicles in radiographs which may sometimes be difficult. However, they do permit a consistent and objective measurement of vertebral rotation.

1.4 From Spinal Deformity to Surface Shape

The best understanding of scoliosis in a particular individual will require the use of radiographs or other imaging modalities such as Computed Tomography (Brunie 1992), magnetic resonance imaging (Brunie 1992) or ultrasound (Suzuki 1989) to gain a complete picture of the skeletal deformity. Of these alternatives, radiography is the most widely used because most ordinary hospitals have a radiography department. The other imaging modalities are generally precluded from use in routine clinical sessions because of lack of availability, expense, or that the images produced are not of the required type. However, magnetic resonance imaging is now becoming more widely available as a diagnostic tool because it permits a more detailed examination of the vertebrae and internal structures such as the spinal cord.

Whilst radiographic imaging will often be performed prior to surgery or treatment, it is not advisable on dose grounds to take a complete spinal radiograph of a patient routinely, say every 3-6 months, over a period of several years to monitor the progression of the deformity.

It is, therefore, necessary to find other means of quantifying the extent of the deformity. The three-dimensional aspect of the deformity has already been discussed. In almost all cases, the underlying scoliosis produces abnormalities in the surface shape of the back.

Figure 1.10 is a sketch showing the typical back surface shape of a patient with scoliosis. The underlying deformity of the spine is visible as the line of the spine on the surface. In addition, there are other surface abnormalities. The scapula may protrude

unusually severely from the surface and the patient has a hump on his/her back. The hump is usually more visible as the patient bends forward.




Figure 1.10 The surface deformity of scoliosis.

Given that there exists a surface deformity in scoliosis, and that to some extent it will characterise the underlying structural deformity, there may be valid reasons to assess the extent of scoliosis from surface information. Indeed, this is now the practice in some scoliosis clinics and there have been many studies presenting data that has been acquired from surface shape (over the last 14 years, for example, there has been a series of international conferences entitled “Spinal Deformity and Surface Topography”, with proceedings published by Gustav Springer Verlag, which have been devoted almost exclusively to the assessment of scoliosis from back surface shape).

Specifically, the reasons for assessing scoliotic deformity from back shape are:

i. Perception of the Deformity by the Patient

Although scoliosis is a structural deformity of the spine, it is often not the underlying structural deformity that is the primary concern of the patient. The patient is most aware of the hump or other visible shape deformity of their back. This is particularly significant for adolescent girls (recall that the weighting of incidence is typically 4:1 in favour of girls) who are, perhaps, more conscious than others about the shape of

their body. It is sometimes suggested informally that the clinician is treating the hump, not the underlying structural scoliosis. It might be that the scoliosis in a young patient has stabilised and is unlikely to progress and cause future problems, but treatment is desirable for pseudo-cosmetic reasons.

ii. Surface shape versus radiological assessment

The main shortcoming of radiological assessment is the ionising radiation dose to the patient if he/she is examined on a regular basis over a number of years, even though modern digital systems have reduced this dose considerably. However, there are other problems associated with radiological assessment. First, the radiograph is a two-dimensional projection of the underlying structures of the body. Whilst it is possible to infer some three-dimensional information from the radiograph, three-dimensional information is only available if at least a second radiograph is acquired, at an angle typically perpendicular to the first, and this may double the radiation dose. Secondly, radiographs are often difficult to interpret due to the relatively poor quality of the image. Thirdly, radiography is an expensive procedure requiring high technology equipment and trained staff. Lastly, although radiography is available in a hospital environment, it is not possible off-site, for example in school screening programmes.

iii. Ease of measurement

The surface shape of the body is relatively straightforward to measure using simple, portable, non-invasive low-cost instruments such as those discussed in Section 1.5. In general, however, if these tools are used then there is an overhead in the time required to record and manually process the acquired data.

The scoliosis clinic at the Royal Liverpool Children's Hospital is a national centre of excellence and uses surface measurements as a key element in the routine monitoring of scoliosis patients. Specific devices are used to measure the shape of the back and these are outlined in Section 1.5.

1.5 Assessment of Scoliosis from Surface Shape

A variety of conventional methods and instruments exist for the measurement of back shape in scoliosis. In general, each method uses a simple tactile measuring device that is placed on the back and produces either the direct measurement of a specific parameter of scoliotic deformity or provides data from which parameters can be calculated.

1.5.1 Angle of Trunk Inclination and the Scolometer

The Scoliosis Research Society (SRS) defines the angle of thoracic inclination to be "the angle between the posterior rib cage at the greatest prominence




Figure 1.11 Using the Debrunner Kyphometer.

Instruments exist for measuring both scoliotic deformity and other related spinal deformities (Pun 1987, Goetz 1973). For example, the Debrunner Kyphometer (Debrunner 1972) measures the angle of kyphosis. [The angle of kyphosis, is not strictly a measure of scoliotic deformity because it relates to the curve in the sagittal plane: very roughly equivalent to the Cobb angle]. The Kyphometer is shown in Figure 1.11.

The shape measuring methods and instruments that were of most interest for this investigation were those used for scoliosis assessment at the Royal Liverpool Children's Hospital. The need for automated systems and the criteria for automated systems to measure back shape in scoliosis are addressed in Section 1.7 and 1.8 respectively. However, it is perhaps worth noting here that an automated system to measure back shape and parameters of clinical significance should be able to produce measurements

similar to the simple manual measurements that are conventionally taken by the clinician, if it is to be accepted by the clinical community.

Two parameters of scoliotic deformity are of particular importance in this investigation because they are used at the Royal Liverpool Children's Hospital: the Angle of Trunk Inclination (ATI) and Trunk Asymmetry Score (TAS).

1.5.1 Angle of Trunk Inclination and the Scoliometer

The Scoliosis Research Society (USA) originally defined the angle of thoracic inclination to be "the angle between the horizontal and a plane across the posterior rib cage at the greatest prominence of the rib hump" (SRS 1980).




Figure 1.12 Angle of Trunk Inclination (ATI)

Bunnell (1984) develops the concept further to define the Angle of Trunk Rotation, also known as the Angle of Trunk Inclination (ATI) as the angle between the horizontal plane and a plane across the posterior aspect of the trunk at the point or points of maximum deformity with the patient in the forward bending position. In other words the definition is expanded to encompass more than simply the thoracic area and specifies precisely the position of the patient. Figure 1.12 shows the defined angle. The diagram shows a patient with a severe hump bending forward. Bunnell suggested examining the ATI for the patient in three forward bending positions as shown in Figure 1.13: a development of Adams' Test (Adams 1865). In this case the ATI is no longer the angle subtended to the horizontal at the point of maximum deformity, but the angle subtended to the horizontal at one of three positions on the back. ATI is now generally specified with further qualification indicating how the measurement has been taken.

The spine must be viewed at eye level from behind, with the patient in at least three different degrees of flexion to evaluate all levels of the spine adequately.

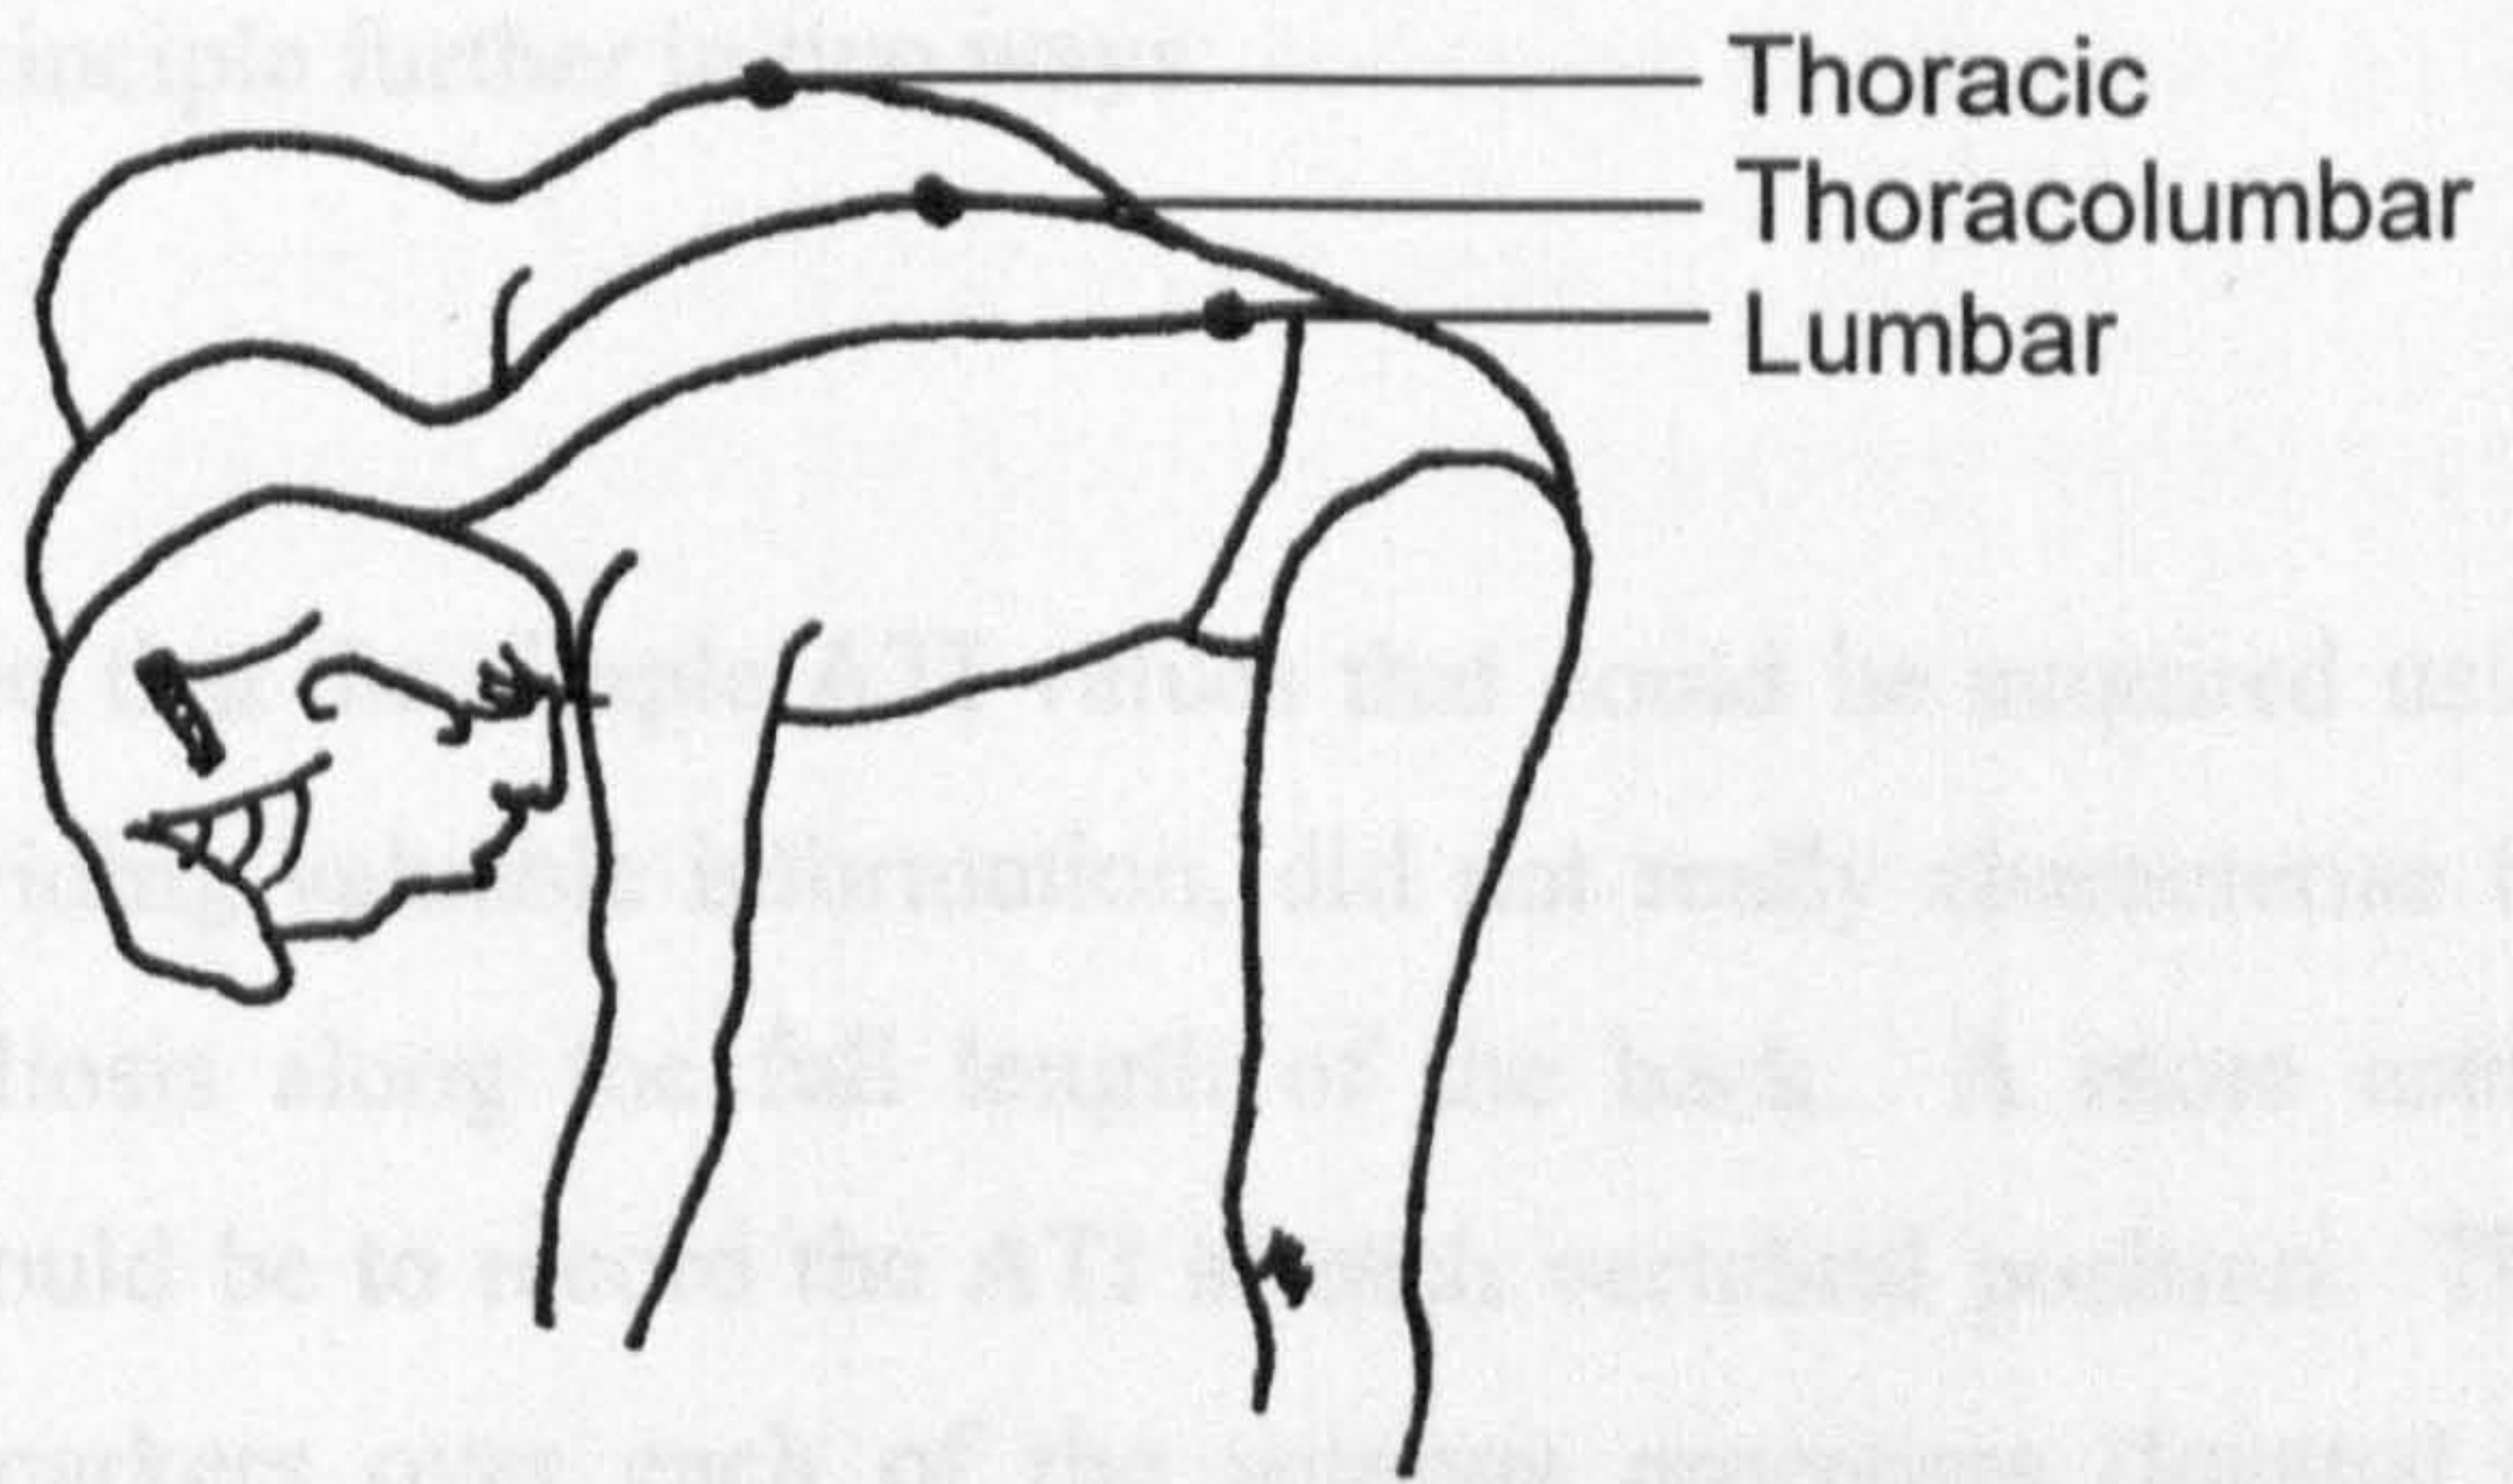


Figure 1.13 Bunnell's application of Adams' test.

A simple device, the Scoliometer, was developed to measure the ATI directly as shown in Figure 1.14. The Scoliometer is a single radius U-shaped tube that is filled with liquid to dampen the motion of a ball bearing contained by the tube. The ball quickly seeks the point that is lowest in the tube and the angle of rotation can be read directly.



Figure 1.14 The Scoliometer

Dangerfield developed the principle further in two ways:

i. Multiple spinal level ATI.

Dangerfield (1992a) observed that the simple ATI values that could be acquired using the Scoliometer, whilst providing valuable information, did not really characterise the extent or nature of the scoliosis along the full length of the back. A more useful measure of the deformity would be to record the ATI at each vertebral position. This was achieved by placing markers over each of the spinous processes (located by palpation) and then recording the ATI at each level. The resulting data set could then be plotted on a graph of ATI versus vertebral position and the graph is then a representation of the scoliosis along the length of the spine. Figure 1.15 shows a typical ATI map for a scoliosis patient.




Figure 1.15 ATI map for a typical scoliosis patient.

ii. Measuring ATI in the standing and prone positions.

Dangerfield also observed that ATI values taken with the patient in the forward bending positions characterise the deformity only in the that position. The scoliosis

patient is usually most concerned about the visible deformity of their back in the standing position (this has already been commented upon in Section 1.4). So perhaps, the ATI should be measured in a position that more closely resembles the standing position. Because the Scoliometer is an inclinometer that operates using the pull of gravity, it is not possible to measure ATI using the Scoliometer in the standing position. One alternative was to measure the ATI with the patient lying face down on a flat surface: the prone position. In this position, the patient is straight and measured ATI values will more closely relate to the shape of the back when the patient is standing. Figure 1.16 shows a patient being measured in this position using the Scoliometer (Dangerfield 1992b).




Figure 1.16 Patient being measured in the prone position using the Scoliometer.

The ATI readings acquired with the patient in the prone position were also recorded at multiple spinal levels. The multi-level ATI map is constructed and stored in patient records.

The measurement of ATI using the Scoliometer is now used extensively in the scoliosis research community (it was, for example, used by Burwell in the school screening study discussed in Section 1.2.4). ATI, or a similar reading measured in a true standing position, was selected to be one of the parameters to be included in the automated system developed as part of this investigation.

1.5.2 Trunk Asymmetry Score

The simple angular measurements obtained using the Scoliometer are indirectly providing information about the asymmetry of the back surface. An alternative, and possibly better, method was devised by Burwell (1983) that enabled the recording of a more accurate representation of trunk asymmetry.

Burwell proposed that a coronal cross-section of the back surface should be taken from which the parameter Trunk Asymmetry Score (TAS) can be calculated. A device called the Formulator Body Contour Tracer (FBCT) is used to record a coronal cross-section of the back. The FBCT is simple matrix of rods that are placed parallel to the coronal plane across the surface of the back. The rods assume the profile of the back. Figure 1.17 shows the FBCT in use on a patient in the forward bending position.




Figure 1.17 The FBCT in use to measure a patient in the forward bending position.

Markers are placed over the centre of the spinous process at the vertebral position to be measured. In positioning the apparatus, a centrally marked rod is placed exactly on the skin marker at the spinous process. To ensure that the transverse beam of the FBCT is horizontal, a spirit level is attached to it. The rods are then allowed to fall under gravity to assume the contour of the back and are then locked by tightening a winged nut on their housing and removed from the patient. The FBCT is then placed on a flat surface and the profile is transferred onto graph paper by drawing around the upper edges of the

rod matrix. The position of the central rod (the midline) is also recorded on the graph paper.

The recorded shape on paper is then processed manually to calculate TAS. First, each back contour tracing is marked into five divisions on both sides of the midline as shown in Figure 1.18. The distance between divisions is made equal to 5% of the corresponding lateral chest diameter. The figure of 5% is chosen because it was found that the fifth division at 5% included the most prominent part of the rib humps.

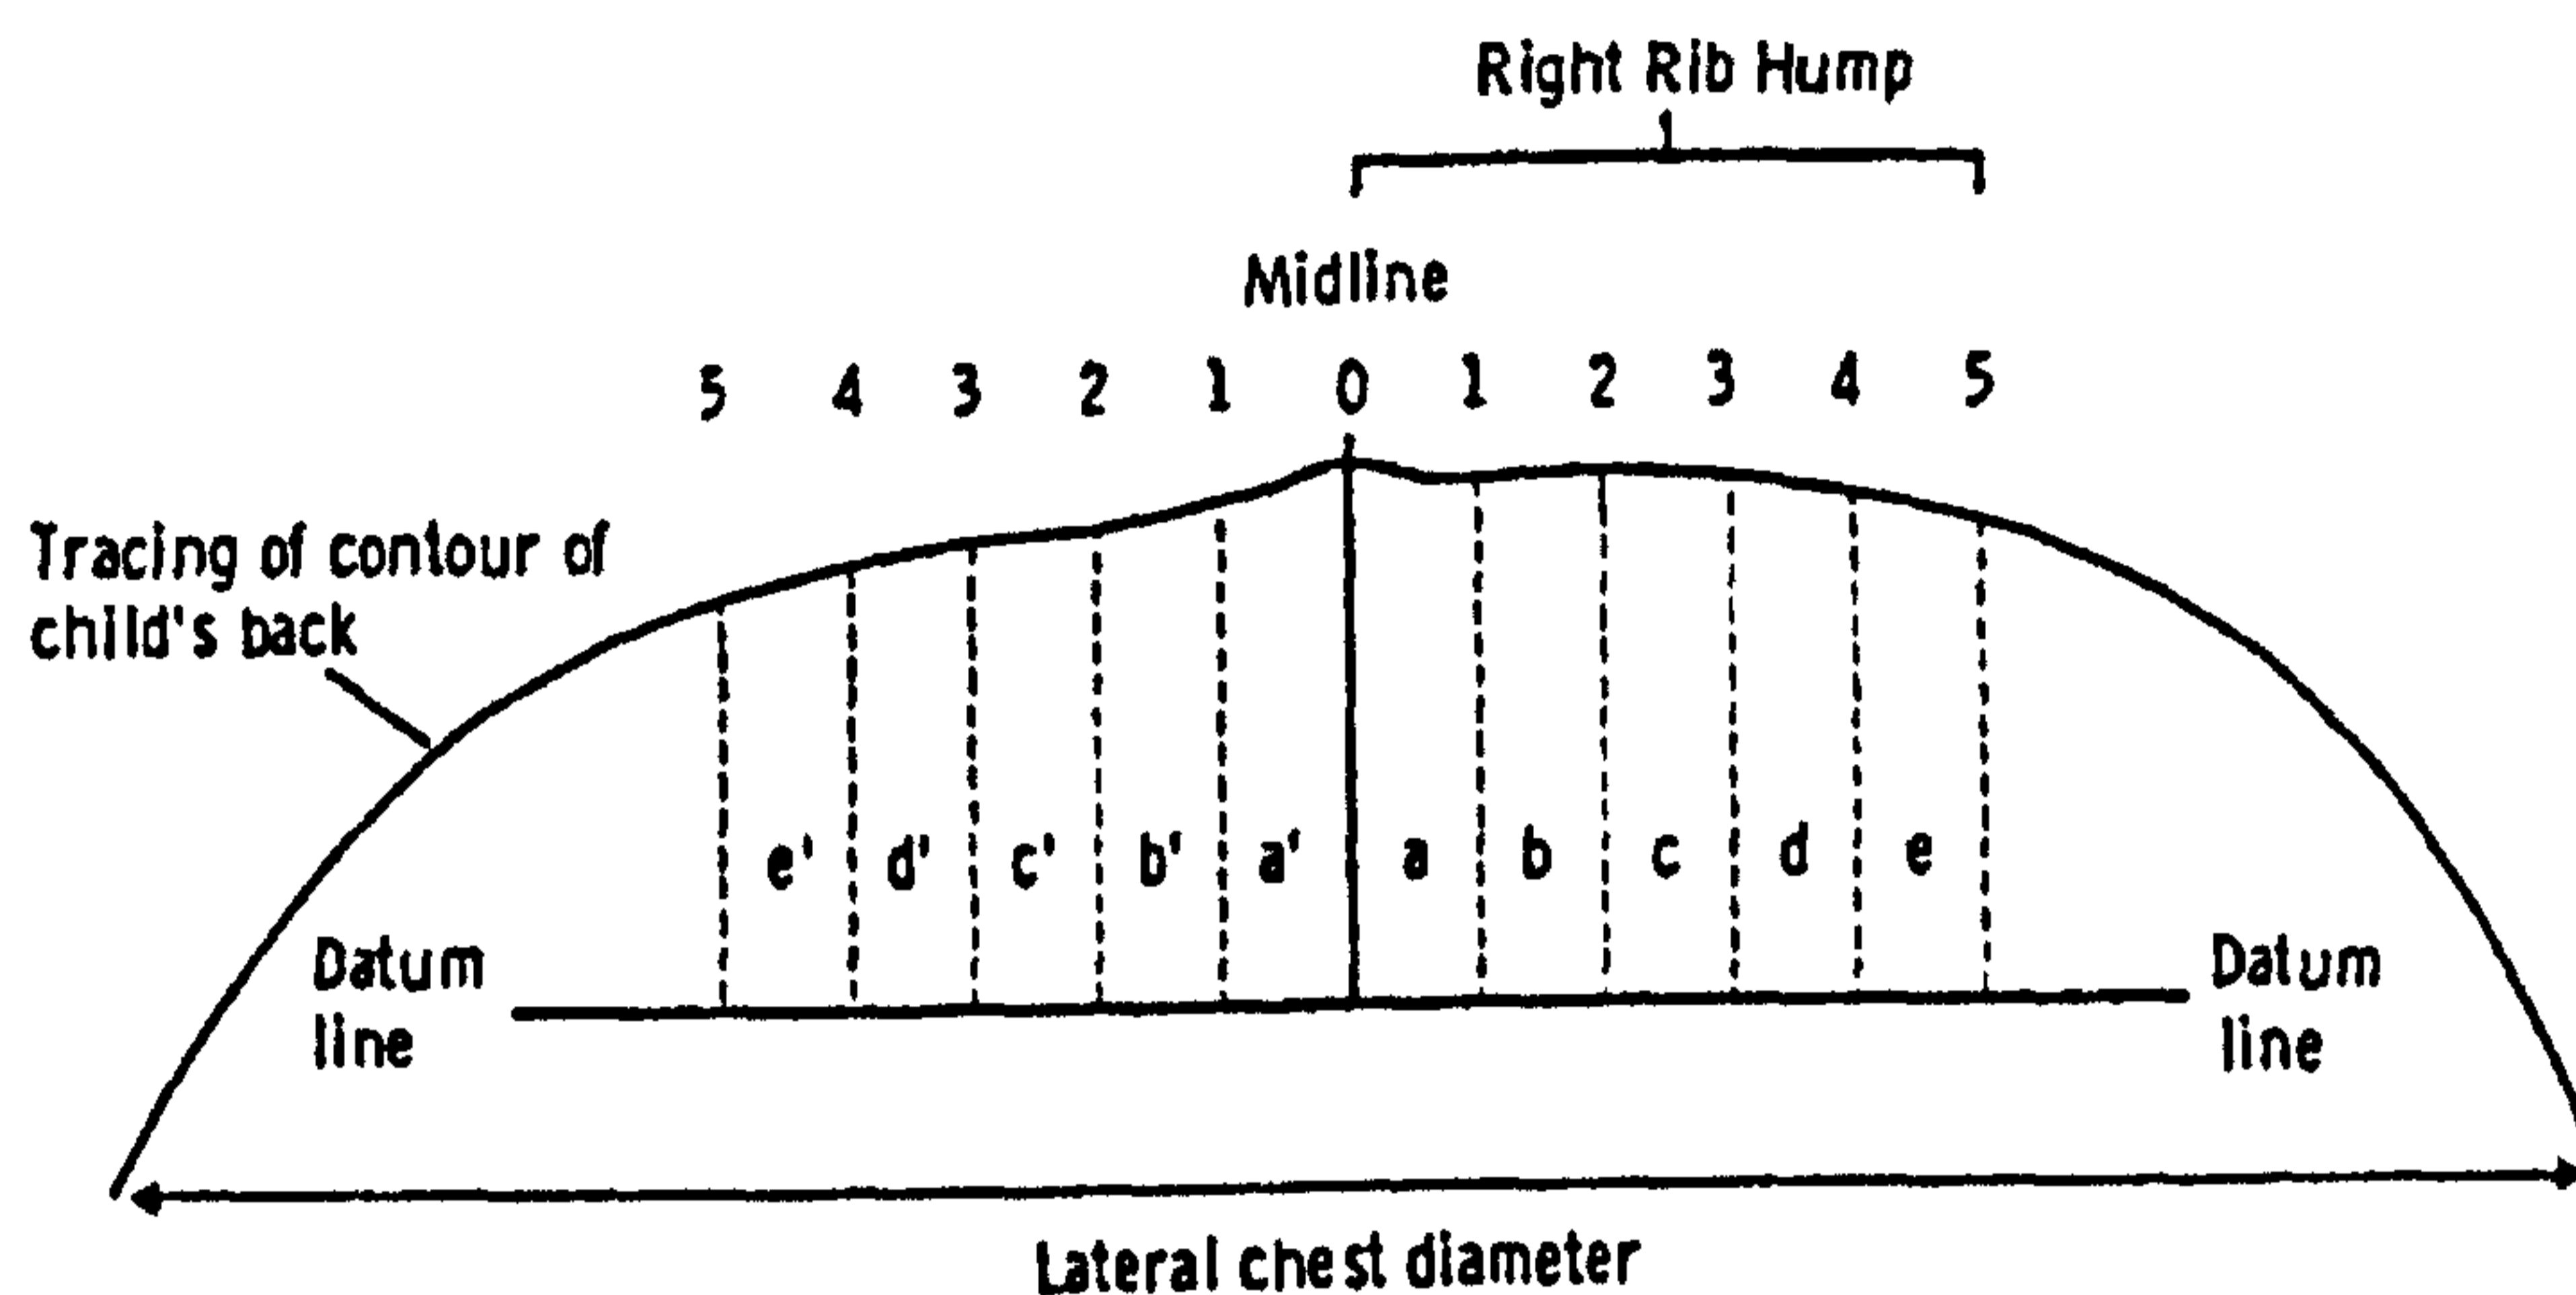


Figure 1.18 Measurement of TAS

At each of the five divisions the distance from the datum line to the back contour line is measured (e' , d' ... a' , a , b , ... e) and the differences from the two sides of the midline summed to produce a Crude Trunk Asymmetry Score (CTAS):

$$\text{CTAS} = (a - a') + (b - b') + (c - c') + (d - d') + (e - e') \quad [\text{Equation 1.1}]$$

This process was semi-automated. The tracings were digitised using a bitpad and processed by a program to produce the TAS values.

The crude TAS scores are then standardised to a trunk diameter of 21 centimetres. This figure was selected by Burwell because it was the overall mean lateral chest diameter of the group of schoolchildren with whom he was working at the time. The calculation of Standardised Trunk Asymmetry Score (STAS) from crude CTAS measurements is performed using:

$$\text{STAS} = \text{CTAS} \times \frac{21}{x} \times \frac{5}{s} \quad [\text{Equation 1.2}]$$

where x is the lateral chest diameter and s is the step size (e.g. 5%, 4% etc.). The STAS measurements enabled an objective, standardised, method for calculating the asymmetry of the back. Burwell originally measured STAS at three spinal levels (T8, T12 and L3). The method was used to examine a sample of schoolchildren in a statistical study.

Other workers (e.g. Dangerfield 1992) have also used the STAS method and developed it further by measuring STAS at all spinal levels (S1 to C7) and producing STAS maps: a graphical representation of the change in STAS at different spinal levels similar to the ATI maps discussed in the previous section. STAS readings acquired using the FBCT can also be acquired with the patient in the prone position and this type of STAS recording is currently performed at the Royal Liverpool Children's Hospital on an experimental basis, but is not yet part of standard clinical procedure.

In general there is less published work on using STAS-type measurements than ATI based studies. It seems likely that the reason for this is that the acquisition and processing, even when limited automated processing support is available, is more demanding and time consuming.

1.6 School Screening

Like many diseases and conditions, if scoliosis can be detected early, ideally shortly after the onset of the condition, then there is a greater probability that effective treatment can be administered. It is often proposed, therefore, that all schoolchildren should be screened for scoliosis. The argument in favour of school screening is obvious: scoliosis can be detected in its early stages and appropriate action taken. The argument against school screening is principally the cost and inconvenience of mass screening programs. Any conclusion must evaluate the benefit of screening against the cost and inconvenience of a screening program.

A number of screening studies have been undertaken in many countries (Burwell 1990, Bremberg 1986, Dickson 1980, Pin 1985, Willner 1984a,b). Most studies demonstrate

that there is some value to having a screening programme, but to date there is no widely agreed consensus as whether screening ought to be implemented.

In 1986 a meeting of the British Scoliosis Society produced the following statement of policy:

Statement of current policy of British Scoliosis Society on school screening

1. *The Natural History Committee of the British Scoliosis Society is unanimous in the opinion that screening for scoliosis should continue as:-
(a) an epidemiological study, also involving
(b) case finding.*
2. *It should still not be a national policy, at present, to screen children routinely throughout the UK.*
3. *More data are required.*
4. *National screening may be possible after further investigation*

(Source: Burwell 1989)

It seems likely, therefore, that there will continue to be interest in school screening for scoliosis. A screening programme might make use of automated systems such as those described and proposed in this investigation. However, due to changes in the financing of the National Health Service and the imposition of an internal market, it is now unlikely that a national screening programme will be implemented in the UK at present.

1.7 The Need for an Automated System to Measure Back Shape

In the majority of cases, automated systems for most applications are beneficial. Tedious work can often be largely eliminated and, in the case of data acquisition systems, data can be stored for future retrieval and processing and the time, date and activities of the system can be logged for auditing purposes.

To some extent the need for an automated system to measure human back shape and produce parameters of clinical significance will vary according to the environment that is to be considered. In some cases, for example school screening, the need may be to increase throughput, whilst in others, for example, the *research* scoliosis clinic, the need may be to investigate parameters of scoliotic deformity that could not be realistically measured using conventional methods.

Specifically, the current and potential advantages of the automated system identified in this investigation are:

i. Saving of time and effort in data acquisition and processing.

Current methods for measuring back shape in scoliosis using the techniques that have been outlined in this chapter are both laborious and time consuming. Suppose that the clinician is to measure both ATI (and produce the ATI map) and STAS (including the manual or semi-automated processing) for a patient. Measurement of ATI will include careful positioning of the patient, a series of Scoliometer readings, the drawing of the ATI map itself and finally the storage and filing of the results. This activity might typically take half an hour. The measurement of STAS is more lengthy still. For STAS, the clinician must place markers over the spinous processes, record the back shape with a FBCT (and transfer to graph paper), then do relatively complex measurements and calculations based on the recorded information. This measurement might typically take an hour if multiple level measurements are required. In an automated system it is possible that data acquisition can be performed in a fraction of a second and automated or semi-automated processing performed in a few seconds or minutes. The clinician's time is limited in almost all real clinical environments and the benefits of an automated system are clear.

ii. Long term storage and retrieval of data.

Currently, the results of clinical sessions are usually stored on paper and kept as part of patient records in a paper filing system. Retrieval of results for future sessions is complicated: the clinician must first request a file, which is fetched by administrative staff, then used and perhaps updated and finally dispatched for filing. An automated system under the control of a microcomputer will typically be able to store data locally on a hard disk for almost instantaneous recall and updating. The scoliosis clinician might select a patient by hospital number and be presented with a complete set of previously acquired data to which he can add. This type of change is already being implemented in other areas of the UK National Health Service.

iii. Storage of surface shape versus storage of specific parameters.

The measurement of specific parameters such as ATI and STAS severely limits the long term development of clinical programmes. Suppose a patient has been monitored with, say, ATI measurement and the clinic makes a strategic decision to change to, say, STAS as the primary measure in scoliotic monitoring. The new STAS measurements will be incompatible with earlier ATI measurements. If an automated system is capable of measuring and recording the complete three-dimensional shape of the back surface then earlier measurements can, with suitable re-programming or re-processing of data, be interpreted for the required parameter.

iv. Backwards compatible *and* future proof.

Methods for assessing scoliosis from back surface shape will undoubtedly continue to change and advance. Two issues arise. First, there is already an “installed base” of manually acquired and stored data (certainly at the Royal Liverpool Children’s Hospital). An automated system that measures back shape will be able to calculate measured parameters that are, to some extent at least, compatible with perhaps years of prior data acquisition (in the system proposed in this thesis, this is generally true except that data is always acquired with the patient in the standing position). It may be that patients have been processed and have left the hospital system but old data that is compatible with new methods is required for research purposes. Second, new parameters may be investigated in the clinical system. An automated system that records surface shape should be able to calculate new parameters provided they are functions of back shape that do not require the patient to be positioned any differently from the position in which the original data was acquired.

v. New parameters of scoliotic deformity.

The parameters of scoliotic deformity that are conventionally measured are limited by both the simplicity of the devices used to acquire data, the time taken to acquire and process that data and the low spatial density of the acquired data. Normally, an automated system will not be limited by these constraints and it is possible that new

parameters of scoliotic deformity that could not previously be realistically measured will become possible.

vi. Accuracy

Although conventional methods appear to provide an accuracy in measurement that enables quantitative measurement of the deformity, simple consideration of the techniques used dictate that accuracy must be limited in scientific or engineering terms. Although the degree of increased accuracy will depend upon the nature of the automated system used, it is likely that an automated system will produce more accurate results.

vii. Integration with other clinical systems.

The value of the long term storage of raw and processed data relating to scoliotic deformity has already been identified. However, an additional level of value is added if this data can, subject to security requirements, be made widely available through electronic media. Developments in general computer networking, wide area networking such as the UK National Health Service network and super wide area networking such as the Internet, have revolutionised thinking on data distribution. It is now easily possible for information and data stored on a simple microcomputer to be mapped into the databases of other hospital-wide, national or even international computer systems. This type of networking is currently being implemented in the UK, but there are issues of data confidentiality and security that remain to be solved. The Bangemann report to European Council of Ministers, "Europe and the Global Information Society" (Bangemann 1993), points to an increasing use of such communications in Europe. In addition, the European community has recently invested 800 million ECU in research into telematic applications of which a substantial proportion is allocated for medical projects.

viii. Statistical data acquisition

Another consequence of the long term electronic/magnetic storage of data is that statistical data relating to the incidence, type and trends in scoliosis can be easily produced. For example, in an automated system it should be relatively easy to, say,

calculate the mean ATI for all the patients currently being treated (obviously this is a simple example).

ix. School screening

The possibilities of school screening for scoliosis have already been discussed. Almost all of the aforementioned justification of need applies equally to school screening programmes. However, the importance of the automated system is elevated. For example, whilst the saving of time and effort is important in the scoliosis clinic it is even more important when a huge number of individuals are to be processed in as short a time as possible.

There is a clear need for an automated system that will acquire and store the back shape of scoliosis patients. The only system to have gained widespread use, at least in research orientated clinical programmes, is the ISIS system (Turner-Smith, 1988) which is now technologically outdated, does not completely record back surface shape and uses structured light scanning which may not be the best approach to back surface shape measurement (see Chapter 2).

1.8 The Criteria for an Automated Measurement System

Before research can proceed, or new methods be proposed, it is necessary to identify the properties that an automated system should exhibit. Based on the presentation of fact, justified opinion, and surrounding discussion in this chapter it may be concluded that the following criteria define the necessary requirements and boundaries for this research: a practical automated system for measurement of body shape should have the following properties:

i. Standing Position

The scoliosis patient is most aware of the deformity of his or her back in the normal standing position. Already the measurement of back surface has embraced this fact and some research workers are attempting to measure back shape accordingly. *The automated system should measure the shape of the back in the standing position.*

ii. Non-contact

A non-contact system will minimise the discomfort to the patient and will usually minimise the number and type of precision mechanical components required. *The automated system should be non-contact.*

iii. Accuracy

The automated system should produce measurements that are at least as accurate as conventional methods. *A notional target accuracy of 5mm in the measurement of surface shape was used for this research.*

iv. Clinical Parameters

If an automated system is ever to be accepted by the clinical community it should be capable of producing clinical parameters that are identical or similar to those currently measured using conventional methods as well as any new parameters. Clearly there will be restrictions to this condition if the patient is to be measured in the standing position and current methods require another position for measurement. *The system should, as far as is possible, measure parameters that relate closely to conventionally measured parameters for scoliotic deformity.*

v. Repeatability

The system should be capable of producing approximately the same measured surface shape and clinical parameters when applied to the same patient more than once (assuming that back shape does not change between measurements).

vi. Speed of Acquisition and Processing

The automated system should be capable of acquiring and processing data significantly faster than conventional methods. *A notional target time for the acquisition and processing of data in the method proposed in this investigation was set to be one minute.*

vii. Long term storage

The system should be capable of long term storage of back surface shape, or at least the data from which the surface shape can be reconstructed by processing.

viii. Ease of Use

The system must be usable by the clinician and should not require lengthy or complicated setup, acquisition or control activities.

ix. Cost

If an automated system, or the techniques upon which it is based, is ever to be a practical tool for clinical use or school screening then it should be affordable in those environments. Since these types of environments are usually restricted in their budget, the cost must be low. Although a successful system should, in the long term be cost effective because it will reduce expenditure on clinical effort, appropriate funds based on this rationale may not be available. *A notional cost for a target automated system for consideration in this investigation was set to be £2000.*

x. Portable/Collapsible

Few Scoliosis clinics have the resources to dedicate permanent space to an automated system for the measurement of back shape. Indeed, the clinic may only be used for scoliosis activities on certain days of the week, with other days allocated for other purposes. If the system is ever to be used for school screening then it will need to be completely portable. *The system should be portable, or if this is not possible it should at least be collapsible.*

These criteria form the broad functional specification for the automated system that is to be investigated. Chapter 2 describes the range of methods that may be suitable and Chapter 3 describes the research into a new method which aims to satisfy all the criteria. Chapter 4 returns to these criteria and uses them to evaluate whether the new method, or rather its implementation in a clinical system, is a suitable tool in the assessment of scoliosis.

Chapter 2

Optical Methods for Measurement of Back Surface Shape

2. Optical Methods for Measurement of Back Surface Shape

2.1 Introduction

This chapter outlines the fundamental methods of approach to back shape reconstruction that were investigated in the research programme. The methods considered were stereophotogrammetry, moiré topography, phase-stepping profilometry, Fourier transform profilometry, point raster projection, binary encoding and dot and line raster projection. It was possible to investigate many of these through the literature. However, two of the methods, moiré topography and Fourier transform profilometry, have been investigated through practical research in the laboratory.

Whilst the optical technique defines the fundamental method the acquired optical image must be processed to yield a surface reconstruction. Where appropriate, image processing methods are also discussed. Only those methods suitable for processing using digital image processing are considered.

2.1.1 Background

Optical methods are often the choice for data acquisition in automated three-dimensional shape measurement systems. The methods have been applied successfully to a wide range of industrial and biomedical applications including the measurement of objects as diverse as printed circuit board components (Pearson 1994), aircraft fuselages (Pirodda 1982), human dentures (Atkinson 1982) and the cornea of the eye (Kawara 1979). The optical methods in themselves do not usually directly provide three dimensional measured coordinates. Usually, light intensity or spatial information which relates to three-dimensional shape is recorded using an appropriate electro-optic transducer, readings from which can be processed to extract three-dimensional coordinates.

Early optical measurement methods relied upon photographic film for the recording of two-dimensional images which encoded three-dimensional information. Measurements of the two-dimensional images using travelling microscopes, microdensitometers, stereocomparators and other instrumentation produced numeric values from which the

depth, or z , coordinate for an (x,y) point in an image could be calculated. The development of digital electronic computers in the 1950's and 1960's provided fast semi-automated processing under the control of a program. Reconstruction of coordinates representing three dimensional shape was performed by the computer. Acquisition of data was initially performed either manually or using crude electro-optic devices.

The 1970's saw the advent of framestores: fast analogue-to-digital converters with associated memory capable of capturing, digitising and storing signals from conventional T.V./video cameras at frame rate. Both intensity and spatial information could be simultaneously recorded electronically. The digital image from a frame-store could be transferred to the memory of a computer for processing so that the computer and frame-store combination provided a complete image acquisition and processing system capable of fully automating many optical measurement procedures. This configuration has subsequently proved to be a powerful tool for data acquisition in modern optical metrology. The increase in power and sophistication of computers and framestores coupled with a decrease in cost has brought digital image processing into widespread use today.

2.1.2 Suitability of Optical Methods

Whilst a number of methods are possible for the measurement of back shape, there are several reasons why optical methods are an obvious choice:

i. Non-contact

Optical methods are non-contact: there is no physical contact of tactile probes or other devices with the back. This is desirable because the physical contact of a tactile device is often uncomfortable to the subject. In the case where the subject is a hospital patient under measurement, say following surgery, tactile measurement might even be painful. The contact of a probe with the back may cause the subject to become tense or flinch, in which case the shape of the back may change and the patient adopt a posture other than that in which he has been placed for measurement. In some hospital environments and for some patients there may also be issues of hygiene and safety to be addressed if a contact system is used.

ii. Limited mechanical requirements

Optical methods do not generally require moving mechanical or electro-mechanical elements such as motors to accurately position probes. Moving mechanical components are prone to wear and breakage and very often need to be handled with great care. The weight, bulk and sensitivity of these components may render them unsuitable for inclusion in a portable or collapsible system. Calibration of precision moving components may have to be performed in a specialist laboratory and the system may have to be regularly serviced, leading to a system with low availability. Some optical methods do require mechanical moving components, notably phase stepping profilometry (see Section 2.5.3) and optical scanning (see Section 2.6), in which case the mechanical element of the system is confined to a single, well protected and, where possible, maintenance-free element.

iii. Speed

Optical methods often use digital images as the recording medium for the measurement. Since digital images can typically be acquired in $1/25^{\text{th}}$ of a second (a typical frame acquisition time) the acquisition time for the measurement is short in comparison to manual methods or methods involving probes or other instrumentation. Even when a number of images must be acquired for the method, it is still likely that an optical approach will be faster than other approaches.

The combination of these advantages strongly suggest that optical/image processing methods are likely to be the most suitable approach for measurement of the shape of the back.

2.2 Stereophotogrammetry

Many optical three-dimensional measurement methods have some commonality with the well established principles of photogrammetry and stereo vision. Stereophotogrammetry uses two cameras that are offset from each other and would normally acquire images simultaneously. There is a similarity with the human vision system which relies heavily on stereo vision to judge distance: it is well known that the loss of vision in one eye severely impairs the ability to judge range.

2.2.1 The Stereophotogrammetric Principle

In stereophotogrammetry, the position of a point on the object is observed in the image plane of at least two cameras which are angularly offset from the object. The two images form a stereo image pair.




Figure 2.1 The principle of stereophotogrammetry.

Figure 2.1 illustrates the principle of stereophotogrammetry. The (x,y) position of a set of points on the object, in this case the human back, is recorded in each image and consideration of optical geometry yields the depth, or z coordinates, for each point. Note that the principle can be used either with the cameras having their image planes parallel (geometrically the most simple case) or otherwise. Stereophotogrammetry thus enables the reconstruction by triangulation of the three-dimensional shape of the object.

2.2.2 A Case Study with Stereophotogrammetry: Thompson

Thompson (1985) investigated the application of stereophotogrammetry to back shape measurement as a tool to evaluate scoliosis. The system was based upon a combination of photographic recording and semi-automated processing.

Thompson recorded the stereo pair using two metric photographic cameras at a separation of 1.5m with the scoliosis patient at an orthogonal distance 3.25m from the cameras as shown in Figure 2.2. Both cameras were mounted on theodolites to enable

accurate levelling and alignment. The films were exposed simultaneously using a double cable release with a typical exposure time of $1/60^{\text{th}}$ second.

The developed photographs were then transferred to a stereocomparator for measurements. The stereocomparator was capable of measuring coordinates to an accuracy of $\pm 10\mu\text{m}$ translating to an accuracy of $\pm 2\text{mm}$ at the object (the patient). Few details of the data processing system were published, but it is known that the extraction of three dimensional coordinates was partially automated using an early, 8-bit, microcomputer.




Figure 2.2 The stereophotogrammetric arrangement used by Thompson.

Thompson encountered difficulty in locating corresponding points in the stereo pair. The solution proposed was that “an artificial pattern was projected on to the patient’s back using a slide with geometrical shapes on it and an ordinary slide projector”. In fact, what Thompson was doing was introducing structured light (see Section 2.3) into the system. The coordinates selected for measurement were the structured light markers projected onto the back surface.

2.2.3 Discussion of Stereophotogrammetry

Although Thompson used photographic cameras and transferred the film to a stereocomparator, there is no reason why his technique cannot be implemented directly using CCD/video cameras. Stereophotogrammetry has a major shortcoming when measuring the shape of irregular objects such as the human back. The object is often very uniform and without structure, and finding corresponding points in a stereo image pair may be difficult, especially if such points are to be detected using automated digital image processing. This problem might not present a major drawback if the number of points to be measured is small. However, a large number of data points must usually be measured in order to obtain satisfactory shape replication.

It might be that, in Thompson's system, given that the system will then only be as accurate as the precision with which the structured light markers can be located, it was not necessary to use stereo vision at all. Depth information could have been extracted using simply the optical geometry of the projector and one of the cameras. Used in this configuration, the second camera could be used to provide redundancy, to increase the area of the surface that could be recorded, and to minimise the shadowing problem.

2.3 From Stereophotogrammetry to Structured Light

The uniform nature of the back which presents a problem in stereophotogrammetry can be overcome, as Thompson found, if a structured light pattern is formed on the object. One approach is simply to project the pattern onto the object, in which case the principle of structured light can be thought of as replacing one of the cameras in a stereo vision system with a projector which projects a known pattern of light, often a grating, onto the surface of the object as shown in Figure 2.3.

The distortion, or modulation, of the projected pattern by the object is a function of the three-dimensional shape of the object, and suitable processing of the image acquired by the camera enables extraction of the depth, or z , coordinates of the object. In these cases the uniform and structureless nature of the back can be, and often is, an advantage and one problem is to reduce the effects of severe changes in reflectivity and surface gradient in acquired images.

A number of variations of the technique can be applied. Instead of projecting a structured light pattern onto the surface, the pattern may be formed interferometrically using coherent sources (usually lasers). Another variation is to project a very simple pattern such as a single dot or a line and to scan this across the image, acquiring images at any points for which sampled depth or range values are required (see Section 2.6).




Figure 2.3 The principle of structured light projection.

The techniques that follow in this chapter are all variants of the structured light concept. Different patterns and acquisition modes are used and there are several methods for interpreting the acquired data.

2.4 Moiré

The optical technique which has gained most widespread clinical use in the measurement of body shape has been moiré (F. silk) topography (Adler 1984, Daruwalla 1992, Willner 1984, Neugerbauer 1995). A moiré topogram provides a contour map which can be easily understood by a human observer, and clinical application has often been confined to direct observation of the topogram or a photograph of it as shown in Figure 2.4.




Figure 2.4 Moiré topogram of a scoliotic back.

The deformity of scoliosis in the form of a hump will show clearly as an asymmetry in the topogram. Such a contour map, typically recorded on photographic film, is a clinically useful recording of body shape. Manual processing of the contour map may also be used to extract clinically significant parameters which relate to body shape deformity.

2.4.1 The Principle of Moiré

The principle of moiré topography (Takasaki 1970, Meadows 1970) is that a coarse grating is projected onto an object which is then viewed through the same or a similar grating. The period of the grating must be sufficiently coarse that the diffraction effects can be ignored. A shadow of the grating is formed on the object and is termed the primary fringe pattern. Geometric interference between the primary fringe pattern on the object and the grating through which the object is observed generates a moiré fringe pattern. The transmission of the grating need not necessarily be square: it may be sinusoidal or have other form. Very often it is convenient to select a Ronchi grating since these are easy to construct for most applications.

There are two popular methods for forming the fringe pattern that is produced when the primary fringes projected onto the surface of an object interfere with a grating of similar pitch: shadow moiré and projection moiré.

2.4.1.1 Shadow Moiré

In shadow moiré, the object is viewed through the same grating used to generate the primary fringes on the object as shown in Figure 2.5. The grating must be close to the object and approximately the same size as the object.




Figure 2.5 Typical configuration for shadow moiré contouring.

Geometric interference takes place between the primary fringes caused by the shadow of the grating on the object and the grating itself as light from the object passes through the grating to be imaged at the camera. The precise orientation of the camera and projector are not crucial for achieving the moiré effect.

2.4.1.2 Projection Moiré

In projection moiré, a grating mounted in a projector is projected onto the object to generate primary fringes and the object is then viewed through a second, similar, grating placed close to the imaging plane of the camera as shown in Figure 2.6. Geometric

interference takes place between the primary fringes on the object and the grating at the camera.

The grating at the camera should be equivalent to an image taken by the camera of the projection grating projected onto a flat reference plane. In practice, the camera's grating is often produced in precisely this way: by recording an image of the grating projected onto a flat plane. The analysis of moiré fringe formation is often derived using shadow moiré, but in all cases an equivalent explanation using projection moiré is possible.




Figure 2.6 Configuration for projection moiré contouring.

2.4.1.3 Other Configurations

A number of variations on the technique have been applied. The moiré pattern generated is independent of the position of the grating within its own plane so that if the grating is translated the moiré pattern remains static (Halioua et al 1983). Translation of the grating in this fashion, combined with a suitable long photographic exposure, or another device with long integration time, removes much of the noise and second order effects in the moiré fringe pattern. Scanning moiré (Idesawa 1977) uses a projection arrangement in which the scan lines of a T.V. camera used to acquire the image form the grating through which the object is observed.

2.4.2 Moiré Fringe Formation

Moiré contour patterns can be generated in a variety of different configurations. The fundamental principles of the fringe pattern generation can be applied in a similar fashion to most configurations. In addition, a moiré pattern will be produced for almost all types of grating. The theory described below applies specifically to shadow moiré contouring with a sinusoidal grating but has relevance to all moiré contouring configurations.




Figure 2.7 Moiré fringe formation.

Consider the optical geometry shown in Figure 2.7. A sine wave transmission grating of transmission $T(x)$, period p , is illuminated with light from a point source at A (ray a) onto a continuous surface, an element of which exists at S . The surface is then observed at B (ray b) through the grating. A and B are at perpendicular distance h from the grating. The plane of the grating defines the xy plane with the z axis pointing towards

the surface. The surface is described by the function $z(x,y)$ to be the perpendicular distance of the surface for each x,y coordinate from the plane of the grating. The coordinates of S are then $(x,y,z(x,y))$. The angle between the normal to the surface, n , and the incident light is a function of the surface shape $\gamma(x,y,z(x,y))$.

The transmittance of the grating can be expressed as

$$T(x) = \frac{1}{2} + \frac{1}{2} \sin\left(\frac{2\pi x}{p}\right) \quad [\text{Equation 2.1}]$$

If we assume that the gradient of the surface varies slowly enough that γ is constant then

$$\cos(\gamma(x,y,z(x,y))) = \Gamma, \quad (\Gamma \text{ constant}) \quad [\text{Equation 2.2}]$$

Further, if the depth of field of the object under illumination is small such that the distance $r(x,y,z(x,y))$ of the object from A does not vary significantly then we define a constant R to be

$$\frac{1}{r(x,y,z(x,y))^2} = R, \quad (R \text{ constant}) \quad [\text{Equation 2.3}]$$

If the intensity of light striking the grating is I_A then the intensity of light observed at B is given by

$$I_B = \frac{I_i R \Gamma}{4} \left\{ 1 + \sin\left(\frac{2\pi}{p} \left[\frac{hx}{h+z(x,y)} \right]\right) + \sin\left(\frac{2\pi}{p} \left[\frac{Dz(x,y) + hx}{h+z(x,y)} \right]\right) - \frac{1}{2} \cos\left(\frac{2\pi}{p} \left[\frac{Dz(x,y) + 2hx}{h+z(x,y)} \right]\right) + \frac{1}{2} \cos\left(\frac{2\pi}{p} \left[\frac{Dz(x,y)}{h+z(x,y)} \right]\right) \right\} \quad [\text{Equation 2.4}]$$

This equation has a number of terms including that which gives rise to the moiré pattern. The first and second significant terms in Equation 2.4, namely

$$\sin\left(\frac{2\pi}{p} \left[\frac{hx}{h+z(x,y)} \right]\right) \text{ and } \sin\left(\frac{2\pi}{p} \left[\frac{Dz(x,y) + hx}{h+z(x,y)} \right]\right) \quad [\text{Equation 2.5}]$$

represent the primary fringe pattern formed on the object and the grating itself through which the object is observed.

The penultimate term of the moiré equation,

$$\frac{1}{2} \cos \left(\frac{2\pi}{p} \left[\frac{Dz(x,y) + 2hx}{h + z(x,y)} \right] \right) \quad \text{[Equation 2.6]}$$

gives a secondary pattern. The last term of Equation 2.4,

$$\frac{1}{2} \cos \left(\frac{2\pi}{p} \left[\frac{Dz(x,y)}{h + z(x,y)} \right] \right) \quad \text{[Equation 2.7]}$$

depends *solely upon z* and gives rise to the moiré effect.

The primary fringe interval should be an order of magnitude less than the fringe interval for the fringes to be of suitable quality for most applications. Peaks (centres of bright fringes) occur at surface depths $z_n(x,y)$ when

$$\cos \left(\frac{2\pi}{p} \left[\frac{Dz(x,y)}{h + z(x,y)} \right] \right) = 1 \quad \text{[Equation 2.8]}$$

$$\Rightarrow \frac{2\pi}{p} \left[\frac{Dz(x,y)}{h + z(x,y)} \right] = 2n\pi, \quad n = 0, 1, 2, \dots \quad \text{[Equation 2.9]}$$

where n is the fringe order. The n^{th} fringe has a depth z_n given by

$$z_n(x,y) = \frac{nph}{D - np}, \quad n = 0, 1, 2, \dots \quad \text{[Equation 2.10]}$$

The separation between the n^{th} and $(n-1)^{\text{th}}$ fringes is then

$$\begin{aligned} \Delta z &= z_n(x,y) - z_{n-1}(x,y) \\ &= \frac{nph}{D - np} - \frac{(n-1)ph}{D - (n-1)p}, \quad n = 0, 1, 2, \dots \end{aligned} \quad \text{[Equation 2.11]}$$

Equation 2.11 shows that the fringe interval Δz is not constant. This implies a severe limitation in most practical applications since the exact fringe number would have to be available in order to calculate fringe depth interval: the surface would have to be in physical contact with the grating at some point. For this reason a further restriction is placed into the special case:

$$h \gg z(x, y) \quad \text{[Equation 2.12]}$$

So the distance from the source/observer to the grating is large in comparison to the distance of the grating from the surface. The approximation:

$$\frac{z(x, y)}{h + z(x, y)} \approx \frac{z(x, y)}{h} \quad \text{[Equation 2.13]}$$

can then be made. Placing this approximation into Equation 2.8, peaks occur at surface depths $z_n(x, y)$ when

$$\cos\left(\frac{2\pi}{p} \left[\frac{Dz_n(x, y)}{h} \right]\right) = 1 \quad \text{[Equation 2.14]}$$

$$\Rightarrow \frac{2\pi}{p} \left[\frac{Dz_n(x, y)}{h} \right] = 2n\pi, \quad n = 0, 1, 2, \dots \quad \text{[Equation 2.15]}$$

where n is the fringe order. The n th fringe has a depth z_n given by

$$z_n(x, y) = \frac{nph}{D}, \quad n = 0, 1, 2, \dots \quad \text{[Equation 2.16]}$$

The separation between the n th and $(n-1)$ th fringes is then

$$\Delta z = \frac{(n+1)ph}{D} - \frac{nph}{D} = \frac{ph}{D} \quad \text{[Equation 2.17]}$$

Equation 2.17 shows that the fringe interval Δz is constant for this configuration, provided that the distance from the source/observer and grating is large in comparison with the distance from the grating to the surface.

2.4.3 Application of Moiré To Back Shape Measurement

Since the moiré topogram is already in relatively widespread clinical use it was reasonable to investigate the use of moiré topography as a basis for a fully automated system. The moiré technique was the initial approach investigated during this research. The shadow moiré geometry shown in Figure 2.8 and Figure 2.9 was used. Moiré topograms of scoliosis patients, such as those shown in Figure 2.10, were acquired at the Royal Liverpool Children's Hospital.




Figure 2.8 Plan of shadow moiré geometry for contouring the back.

The shadow moiré grating used for acquiring the topograms was a commercially available grating of pitch approximately 0.3 lines/mm. The grating consisted of dark lines ruled onto a glass pane and mounted in a frame behind which the patient stands. A grating of this type suffers from unwanted reflections from the glass, which must be kept well polished in order to achieve acceptable topograms.




Figure 2.9 Elevation of shadow moiré geometry for contouring the back.

The problem of unwanted reflections was solved by Swain (1986), who formed the dark lines of the grating by stretching wires or lacing cord across a frame (Swain 1986).

However, a grating constructed in this way would be delicate and it may be difficult to ensure exact equi-spacing of the dark lines (wires).





Figure 2.10 Selection of moiré topograms of scoliosis patients.

As part of a feasibility study prior to automated image processing, the contour information in the topogram was recorded semi-manually. The moiré topograms, recorded on photographic prints were attached to a bitpad and the centres of fringes traced using a cross-hair cursor. Example images are shown in Figure 2.10.



(a) Photograph of moiré pattern



(b) Reconstruction

Figure 2.11 Reconstruction from moiré topogram.

The data stored then represented the centres of bright fringes. The fringes were ordered manually and the data processed to form a continuous reconstruction of the back shape. The three-dimensional coordinates of data points lying between fringes were

interpolated from local known values. Figure 2.11 shows a moiré topogram together with a projection of the resulting surface reconstruction.

Data acquisition was therefore manual, but the principle allowed crude acquisition of data which could be processed to extract clinically meaningful data. A publication arose from this work (See Appendix III, "Moiré Contouring").

2.4.4 Discussion of Moiré

The moiré method has very obvious advantages in that the moiré topogram encodes meaningful information in a form directly recognisable to the clinician. Asymmetries in back shape, for example, are often obvious from the fringe pattern and direct observation of the moiré topogram is sometimes used, for example, in the measurement of low back pain (Swain 1992). In addition the pattern can be processed to extract three-dimensional information.

A number of shortcomings of moiré topography which limit its useful application in automated surface shape recording systems were identified through experience of use, theoretical considerations and literature:

i. Directional ambiguity

Whilst the moiré topogram produces a contour map which encodes three dimensional information, it is difficult without *a priori* knowledge of the surface to determine the sign of the contour interval. This directional ambiguity is sometimes called the "Hills or valleys" problem: there are many surface shapes which give the same fringe pattern.

ii. Poor contrast

The presence of primary fringes on the object, and the grating through which the object is viewed cause the fringe pattern in a moiré topogram to have poor contrast.

iii. Systematic noise

Moiré topograms have an inherent systematic noise feature due to the presence of the grating and primary fringes in the pattern. Automated processing of these topograms

using a C.C.D. camera/computer configuration usually involves the processing of the two-dimensional image to extract edges or intensity maxima and minima. The grating itself presents local intensity peaks and sharp intensity edges and any automated signal processing must filter out these features.

A partial solution to the problem of noise is to form the moiré pattern using moving gratings (Halioua 1983). If a projection moiré arrangement is used, the projected grating is translated in its plane at constant velocity and the primary fringes on the object move correspondingly. The reference grating through which the object is viewed is also translated at constant velocity. Provided the velocities of the moving gratings are exactly synchronised, the desired moiré fringe pattern remains stationary and only the grating lines move. A long exposure photographic recording is acquired at the image plane of a camera that views the fringe pattern. The individual moving grating lines are "washed out" by the photographic analogue integration process leaving clear contour fringes. In addition any secondary moiré fringes are removed and pitch errors and imperfections averaged out. The fringes only remain stationary if the gratings are translated at precisely the same effective speed. If the speeds are not the same then the moiré fringes drift and satisfactory results are not obtained. A long exposure time is required so that the moving grating lines are removed. In an automated system, the photographic camera must be replaced with a video camera and an appropriate frame grab/store device. The frame cycle of the standard video camera (typically 1/25s) will often be too fast for the integration process to occur satisfactorily. It is necessary, therefore, to include specialised components. A specialised CCD video camera with a long integration time may be suitable (provided that the patient does not move significantly during the integration). Powerful and accurate mechanical devices may be able to move the gratings at high velocity to achieve integration within a single frame cycle. Multiple frame acquisition together with averaging at the A/D conversion stage may also achieve a satisfactory effect. The need for these specialist components exclude the use of this technique in a low cost automated system.

Moiré topograms may also be interpreted using phase measuring techniques such as those described in the next section. The problem of noise in the topogram remains and

moiré contouring is therefore not the best fringe formation method for most phase measuring applications.

The full automation of this semi-automated method for processing moiré topograms would rely heavily upon accurate and reliable detection of fringe centres, fringe edges, or phase (see next section) in a digital image of the topogram and resolution of the directional ambiguity problem.

2.5 Phase-Measuring Methods

A fringe pattern generated on the surface of an object, such as the human back, encodes information about the three-dimensional shape of the object. The task of reconstructing the shape in the memory of a computer system is then one of analysing the fringe pattern to extract three-dimensional coordinates.

A fringe pattern can be considered, as it often is by the human observer, to be simply a set of “stripes” for each of which the peak or edge is determined to be the location of the fringe. This concept is addressed in Section 2.7. Fringe location by this method, however, discards almost all of the phase information in the image. An alternative approach to fringe pattern analysis is to regard the fringe pattern as a periodic sinusoidal signal which is modulated by the presence of the object.

The modulation of a signal in engineering and science is often considered from the standpoint of frequency modulation because it has many applications including the frequency modulation of electromagnetic signals for radio and television. In this case the signal is frequency modulated to encode a second, different signal.

In phase measuring applications the signal is *phase* modulated by the surface. The objective is to demodulate the signal to extract its phase at all coordinates. The relationship between the phase of the signal and the surface shape of the object can then be determined from the optical geometry of the system.

There are several ways in which the sinusoidal intensity profile can be generated. Interferometric methods such as the use of a Michelson interferometer are often used where the depth of focus of fringes must be large or where very fine pitch fringes are

required for the measurement of small objects. In many cases, however, interferometric fringe pattern generation is not necessary and the pattern can be produced by direct projection of a suitable grating. In addition, if interferometric fringes were to be formed on a large object such as the back, then relatively high power coherent sources would be required and there may be complex ethical and safety considerations that preclude their use.

2.5.1 The Relationship between Phase and Surface Shape

For simplicity, consider an object illuminated with parallel light. Suppose that a plane surface is illuminated at an angle θ_i with a light source with sinusoidal intensity variation and then viewed at an angle θ_r , as shown in Figure 2.12.




Figure 2.12 Structured light illumination of surface.

If the incident sinusoidal intensity profile has a period p_i , then the period of the light on the plane surface, p_x , is given by

$$p_x = \frac{p_i}{\cos(\theta_i)} \quad [\text{Equation 2.18}]$$

and the period viewed by the camera will be

$$p_r = p_x \cos(\theta_r) = \frac{p_i \cos(\theta_r)}{\cos(\theta_i)} \quad [\text{Equation 2.19}]$$

So the relationship between the period of the incident and reflected light has been established for light striking a flat plane.

Consider now a surface element S which is a distance z from the flat reference plane as shown in Figure 2.13. The reflected light will shift by a distance d . The shift d can be expressed in phase terms to be

$$\frac{\Delta\phi}{2\pi} = \frac{d}{p_r} \Rightarrow \Delta\phi = 2\pi \frac{d}{p_r} \quad [\text{Equation 2.20}]$$

Now,

$$a = z \tan(\theta_i) \quad \text{and} \quad b = z \tan(\theta_r) \quad [\text{Equation 2.21}]$$

Adding Equations 2.21 gives

$$a + b = z [\tan(\theta_i) + \tan(\theta_r)] \quad [\text{Equation 2.22}]$$

Also,

$$d = (a + b) \cos(\theta_i) = z [\tan(\theta_i) + \tan(\theta_r)] \cos(\theta_i) \quad [\text{Equation 2.23}]$$

Substituting Equations 2.21 and 2.22 into Equation 2.20,

$$\Delta\phi = \frac{2\pi z [\tan(\theta_i) + \tan(\theta_r)] \cos(\theta_i)}{p_i} \quad [\text{Equation 2.24}]$$

Rearranging this equation for z ,

$$z = \frac{p_i \Delta\phi}{2\pi [\tan(\theta_i) + \tan(\theta_r)] \cos(\theta_i)} \quad [\text{Equation 2.25}]$$

If, as is often the case, the imaging lens of the camera is parallel to the plane of the object under illumination so that $\theta_r=0$ then Equation 2.25 simplifies to

$$z = \frac{p_i \Delta\phi}{2\pi \sin(\theta_i)} \quad [\text{Equation 2.26}]$$

Equations 2.25 and 2.26 demonstrate that the depth, z , at a point on a surface may be obtained by an evaluation of the change, or modulation, of phase, $\Delta\phi$, of a pattern of structured light that has been projected onto the surface.

A continuous surface is simply a collection of points like S which together form the coordinate set that defines the surface. Thus the surface can be reconstructed by evaluation of the phase of the fringe pattern for each sampled value. The sampled values will be pixel coordinates in the digital image.




Figure 2.13 Geometry of shift in surface position.

The problem of reconstructing three dimensional shape reduces to one of extracting the phase of the fringe pattern, and it is this stage which characterises the method selected. A number of methods exist and each has its own specific set of advantages and disadvantages and these are discussed later in this section.

2.5.2 Phase Extraction

For an image illuminated with a sinusoidal fringe pattern, the intensity, $I(x,y)$, at any given (x,y) pixel coordinate in the image can be expressed as:

$$I(x,y) = A(x,y) + B(x,y)\cos(\phi(x,y)) \quad [\text{Equation 2.27}]$$

where $A(x,y)$ is the additive background intensity, B is the amplitude of the cosine carrier ($B(x,y)/A(x,y)$ is the fringe contrast) and $\phi(x,y)$ is the phase of the signal. This equation illustrates concisely the challenge of phase extraction. The equation contains three unknowns, $A(x,y)$, $B(x,y)$ and $\phi(x,y)$: the single equation cannot easily be solved for $\phi(x,y)$. The phase extraction problem, then, can be further refined to be the problem of extracting $\phi(x,y)$ in this equation.

Most methods for extracting phase yield phase values which are the result of the arctan of some expression and as such they are wrapped phase values: phase values which are modulo 2π and lie between the values $\pm\pi$. Conceptually the resulting phase image has a saw-tooth profile. The investigation reported here has not sought to address the phase unwrapping problem because phase measuring methods are later shown to be inapplicable for use in the kind of automated system envisaged and this was obvious before the need to unwrap phase images arose.

It is perhaps worth noting at this point that the edge or peak detection methods for reconstructing surface shape described in Section 2.7 and in Chapter 3 in fact sometimes amount to attempting to sample a surface at a number of fixed phase values. If the surface does not vary significantly between the sampled phase values, and high accuracy is not required then it may not be necessary to track the continuous phase of the fringe pattern.

There are several methods available for phase demodulation. Possibly the two most common methods are Phase-Stepping Profilometry and Fourier Transform Profilometry and these are given special attention below. There are other methods possible for the phase demodulation of images, for example those of Womack (1984) and Ichioka (1972) that rely on direct filtering operations on the signal.

2.5.3 Phase Extraction using Phase Stepping Profilometry

By capturing a minimum of three images it becomes possible to solve a set of equations simultaneously and calculate the image phase. Although the method will work satisfactorily with three acquired images, four are commonly used since they produce a simple mathematical expression for phase and combine to eliminate systematic noise.

2.5.3.1 The Principle of PSP

A fringe pattern is generated on the object and an image I_0 is captured. The fringe pattern is then translated through exactly the distance equal to a quarter of a period, introducing a 90° shift in the phase of the pattern, and a second image I_{90} acquired. This process is repeated twice more to give images I_{180} and I_{270} . The phase $\phi(x,y)$ of each

pixel in the image can be calculated using pixel values from all four images. Figure 2.14 is a flow chart for the method.

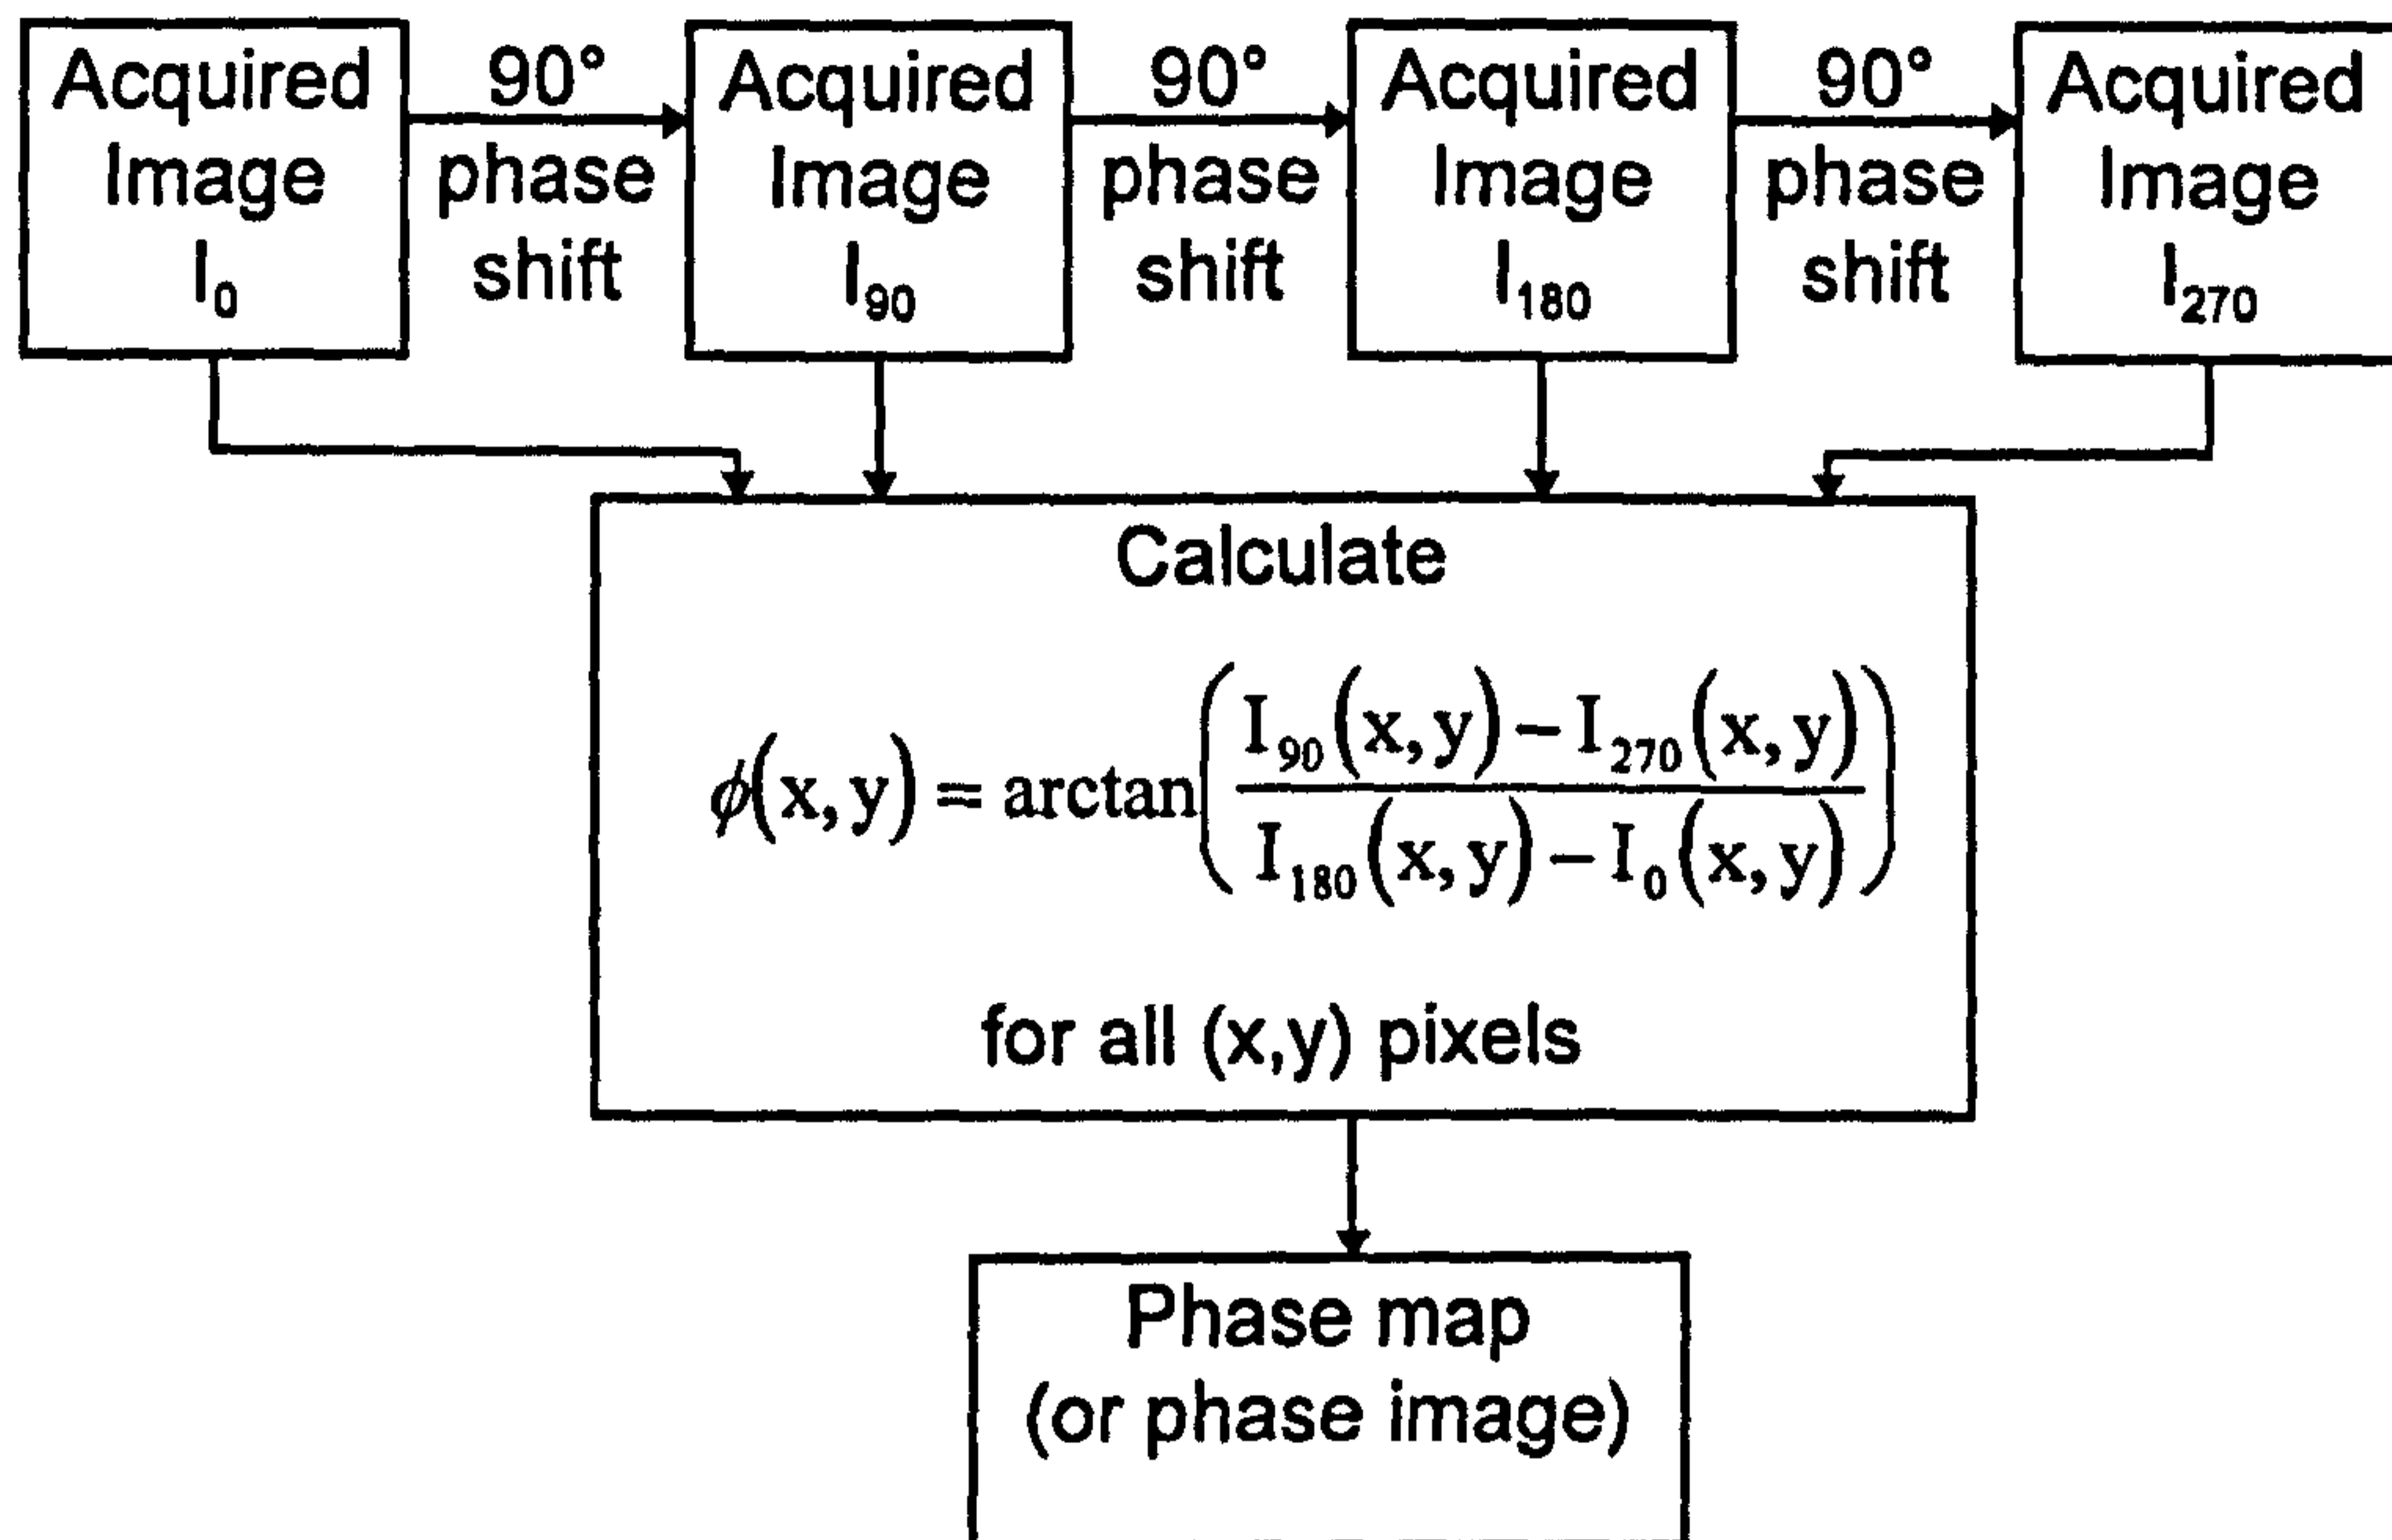


Figure 2.14 Flow chart for PSP

Each image may be represented by the generic Equation 2.27 having three unknowns, one of which is the required phase in the image. Suppose that four images have been stored, each of which have been illuminated by fringe patterns offset by known values of phase shift 0° , 90° , 180° and 270° respectively: the four images are represented by the equations:

$$I_0(x, y) = A(x, y) + B(x, y)\cos(\phi(x, y) + 0^\circ) \quad [\text{Equation 2.28}]$$

$$I_{90}(x, y) = A(x, y) + B(x, y)\cos(\phi(x, y) + 90^\circ) \quad [\text{Equation 2.29}]$$

$$I_{180}(x, y) = A(x, y) + B(x, y)\cos(\phi(x, y) + 180^\circ) \quad [\text{Equation 2.30}]$$

$$I_{270}(x, y) = A(x, y) + B(x, y)\cos(\phi(x, y) + 270^\circ) \quad [\text{Equation 2.31}]$$

Of the three unknowns $A(x, y)$, $B(x, y)$ and $\phi(x, y)$, we are required to find $\phi(x, y)$ and must eliminate $A(x, y)$ and $B(x, y)$. The $A(x, y)$ term may be removed by taking the equations in pairs and subtracting:

$$I_{180}(x, y) - I_0(x, y) = B(x, y)[\cos(\phi(x, y) + 180) - \cos(\phi(x, y))]$$

$$\Rightarrow I_{180}(x, y) - I_0(x, y) = -2B(x, y)\cos(\phi(x, y)) \quad [\text{Equation 2.32}]$$

and

$$I_{90}(x, y) - I_{270}(x, y) = B(x, y)[\cos(\phi(x, y) + 90) - \cos(\phi(x, y) + 270)]$$

$$\Rightarrow I_{90}(x, y) - I_{270}(x, y) = -2B(x, y)\sin(\phi(x, y)) \quad [\text{Equation 2.33}]$$

Dividing Equation 2.33 by Equation 2.32 the $B(x, y)$ term is eliminated:

$$\frac{I_{90}(x, y) - I_{270}(x, y)}{I_{180}(x, y) - I_0(x, y)} = \frac{-2B(x, y)\sin(\phi(x, y))}{-2B(x, y)\cos(\phi(x, y))} = \tan(\phi(x, y)) \quad [\text{Equation 2.34}]$$

giving an expression for phase

$$\phi(x, y) = \tan^{-1} \left[\frac{I_{90}(x, y) - I_{270}(x, y)}{I_{180}(x, y) - I_0(x, y)} \right] \quad [\text{Equation 2.35}]$$

It should be noted that Equation 2.35 gives phase values which are wrapped and further processing is necessary to remove 2π phase wraps and obtain a continuous phase map. The continuous phase map can then be transformed to a depth map of z coordinates through consideration of the optical geometry and three dimensional reconstruction is achieved.

2.5.3.2 A Case Study with PSP: Halioua

Halioua (1989) implemented a generic phase-stepping profilometry system and one of the applications for which the system was used was the measurement of body shape, in particular in application to scoliosis. The generic system was able to accept many different kinds of fringe pattern from a wide range of sizes and surface type. Other examples of the applications for the generic system included phase measuring polarisation interferometry for measurement of curved mirror surfaces and phase measuring holographic interferometry for measurement of dentures and restorative tooth inlays.

The system developed by Halioua for measurement of body shape used a white light projector which was in fact a standard 35mm slide projector modified to accept suitable sinusoidal gratings mounted on a computer controlled micro-stepping stage. The quarter period shifting of the fringe pattern was achieved by translating the grating by an appropriate amount. The configuration used (Halioua, 1990) is shown in Figure 2.15.

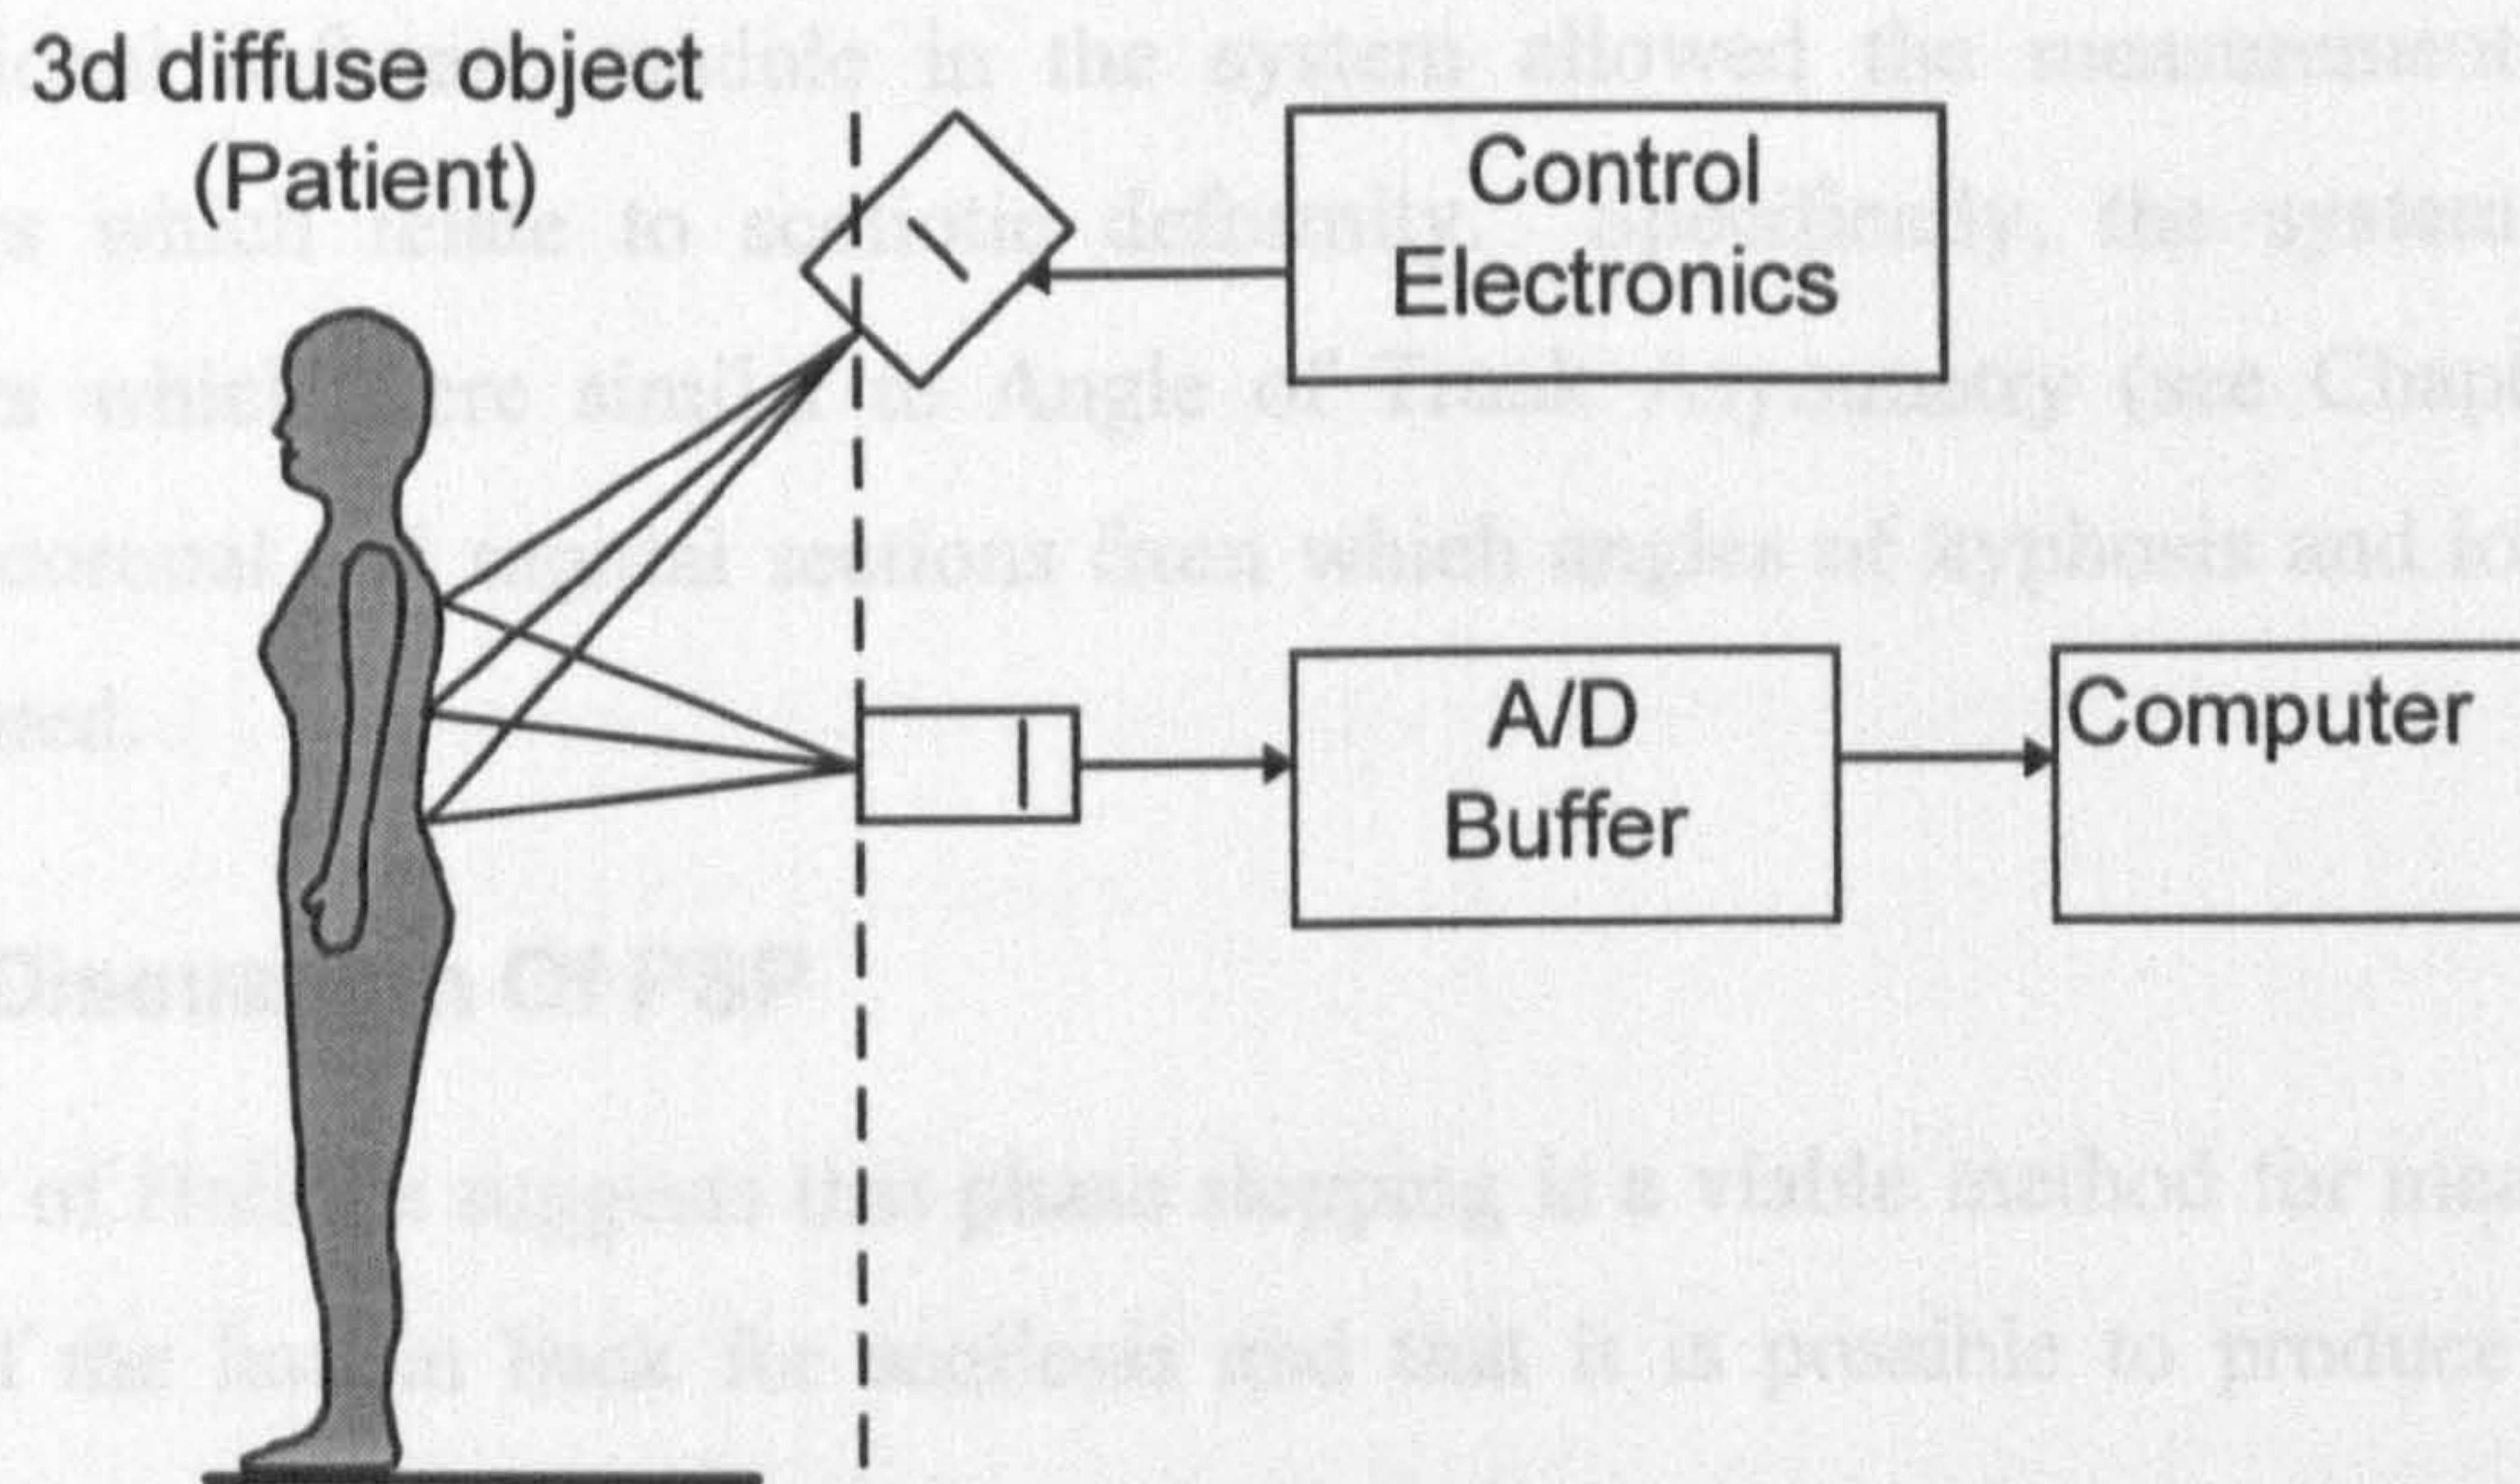


Figure 2.15 The configuration used by Halioua.

A mechanical frame permitted adjustment of depth as well as linear and angular separation of the projection and detection optics. The data acquisition system comprised a CCD camera, followed by A/D and storage electronics mounted onto a single frame buffer board in an AT compatible computer slot. Both the control and data processing operations were carried out by an IBM AT compatible microcomputer.

A plane surface was used for calibration. Depth calibration was achieved by measuring axially shifted positions of the reference plane. A phase mapping algorithm compared the phase values calculated on the body with their corresponding phase values on the reference plane and three-dimensional coordinates were extracted by triangulation.

The data acquisition time for Halioua's system is quoted to be less than one second. This is the time taken to acquire four images including the three shifts in grating position required between image acquisitions. The acquired images had a resolution of 500x500 pixels and the field of view of the camera was 60x80 cm; this translates to a

pixel size of approximately 1.4mm. The camera-object distance was 2m and the angle of illumination was approximately 15°. The processing time for shape measurement and surface reconstruction was approximately 15s using a 10MHz IBM AT microcomputer. The quoted accuracy of the system is “0.6mm RMS. ($<1/100$ th of wavelength $p/\sin\theta$)”.

An additional software module in the system allowed the measurement of some parameters which relate to scoliotic deformity. Specifically, the system provided parameters which were similar to Angle of Trunk Asymmetry (see Chapter 1) and provided coronal and sagittal sections from which angles of kyphosis and lordosis can be calculated.

2.5.3.3 Discussion Of PSP

The work of Halioua suggests that phase stepping is a viable method for measuring the surface of the human back for scoliosis and that it is possible to produce a reliable system which might be suitable for routine clinical use.

One important advantage of phase stepping is that it is usually possible to calculate three-dimensional coordinates for all points on the surface which are within the field of view of the camera. Each pixel is processed independently of other local values to extract phase (although subsequent phase unwrapping will require consideration of neighbouring phase values). There is no requirement that fringes should be tracked as may be the case in, say, line raster contouring, or that whole areas of continuous data values are processed as is required for Fourier Transform Profilometry (see next section).

Perhaps the most important shortcoming of PSP is that it requires the acquisition of at least three, and usually four, images. Between the acquisition of images the grating or other fringe generator must be shifted by a fraction of the fringe period. This usually takes a long time in comparison with the frame cycle of the camera and adds significantly to the acquisition time. It is axiomatic to the success of the technique that the object, in this case a scoliosis patient, should remain stationary during the acquisition of the frames. In the Halioua case study described above, the time taken for

acquisition was 1s, during which the patient must remain *perfectly* still. Whilst it has not been addressed in the literature as a specific problem, it must be difficult, if not impossible to maintain a patient in a perfectly still position for this length of time. The problem is compounded where the patient is a child (the group which are most carefully monitored in scoliosis clinics).

The presence of moving mechanical components in the optical system is undesirable. The moving components may require careful calibration, sometimes by the manufacturer and are often prone to damage and wear, especially if the system is to be portable, for example, for school screening. The computer demands and precision components add significantly to the cost of any ultimately commercial system and may preclude its use in clinical environments.

It is possible that the future objective of scoliosis measurement systems and other body shape measuring systems may be to move from current static measurement to the domain of dynamic measurement with real-time rendering of three-dimensional information. It is extremely unlikely that PSP-based measuring sensors would ever be able to be adapted to cope with this eventuality and that such systems will rely on a single frame for each measurement.

2.5.4 Fourier Transform Profilometry

The principal shortcoming identified in the phase stepping method described in the previous section was the number of images required and the consequent data acquisition time. The phase information which is extracted using PSP is available in a single modulated fringe pattern. The PSP method provides a simple method for extracting phase. There are a number of alternative methods available for extracting the phase of the fringe signal in a single image and one such method is Fourier Transform Profilometry (FTP), first proposed by Takeda (1982). This research into the suitability of FTP was conducted by investigating the method in the laboratory.

2.5.4.1 The Principle of FTP

The FTP technique relies on the calculation of the Fourier transform of the modulated fringe pattern. The Fourier transform is usually calculated using the Fast Fourier

Transform (FFT) algorithm (or one of many closely related algorithms) which dramatically reduces the number of calculations required for the discrete transform (Cooley 1965). A number of variations on the FTP technique exist: the fundamental concept is addressed here.

Unlike PSP, the method does not require any special optical resources such as multiple images or precision stepping and will work with many types of fringe pattern. Figure 2.16 shows the method in flow chart form. A single image of the modulated fringe pattern is acquired in the conventional way and stored in a computer. The Fourier transform of the image is then calculated. Since the acquired image holds only real values, optimised variations on the FFT for real data only can be used.

A series of filters are applied to the FT image and it is the filtering operations which determine the quality of eventual phase signal. The objective is to filter all but the principal maxima in the Fourier domain. Often this is conceptually separated into three steps: removing the DC component, removing negative frequency components and removing noise and harmonics in the positive spectrum. Figure 2.17 shows this concept graphically.

The first filtering operation is the removal (setting to zero) of the negative frequency values in the Fourier spectrum, sometimes called the *half-plane* filtering operation. This operation is axiomatic to the method. The second operation is to remove the zero frequency component. In some cases it has been found to be preferable, for testing and implementation reasons, to remove the direct component (analogous in electrical theory to the direct current level) in the real domain by subtracting the mean intensity value from all values in the image. The third operation is to remove high and low frequency noise and sometimes harmonics. An interesting feature of the FTP method is that that it does not necessarily require the original fringe pattern to have a sinusoidal intensity profile. A square-wave intensity profile, such as that produced by a Ronchi grating, will, for example, work adequately but harmonics will be produced in the Fourier domain which must also be removed.

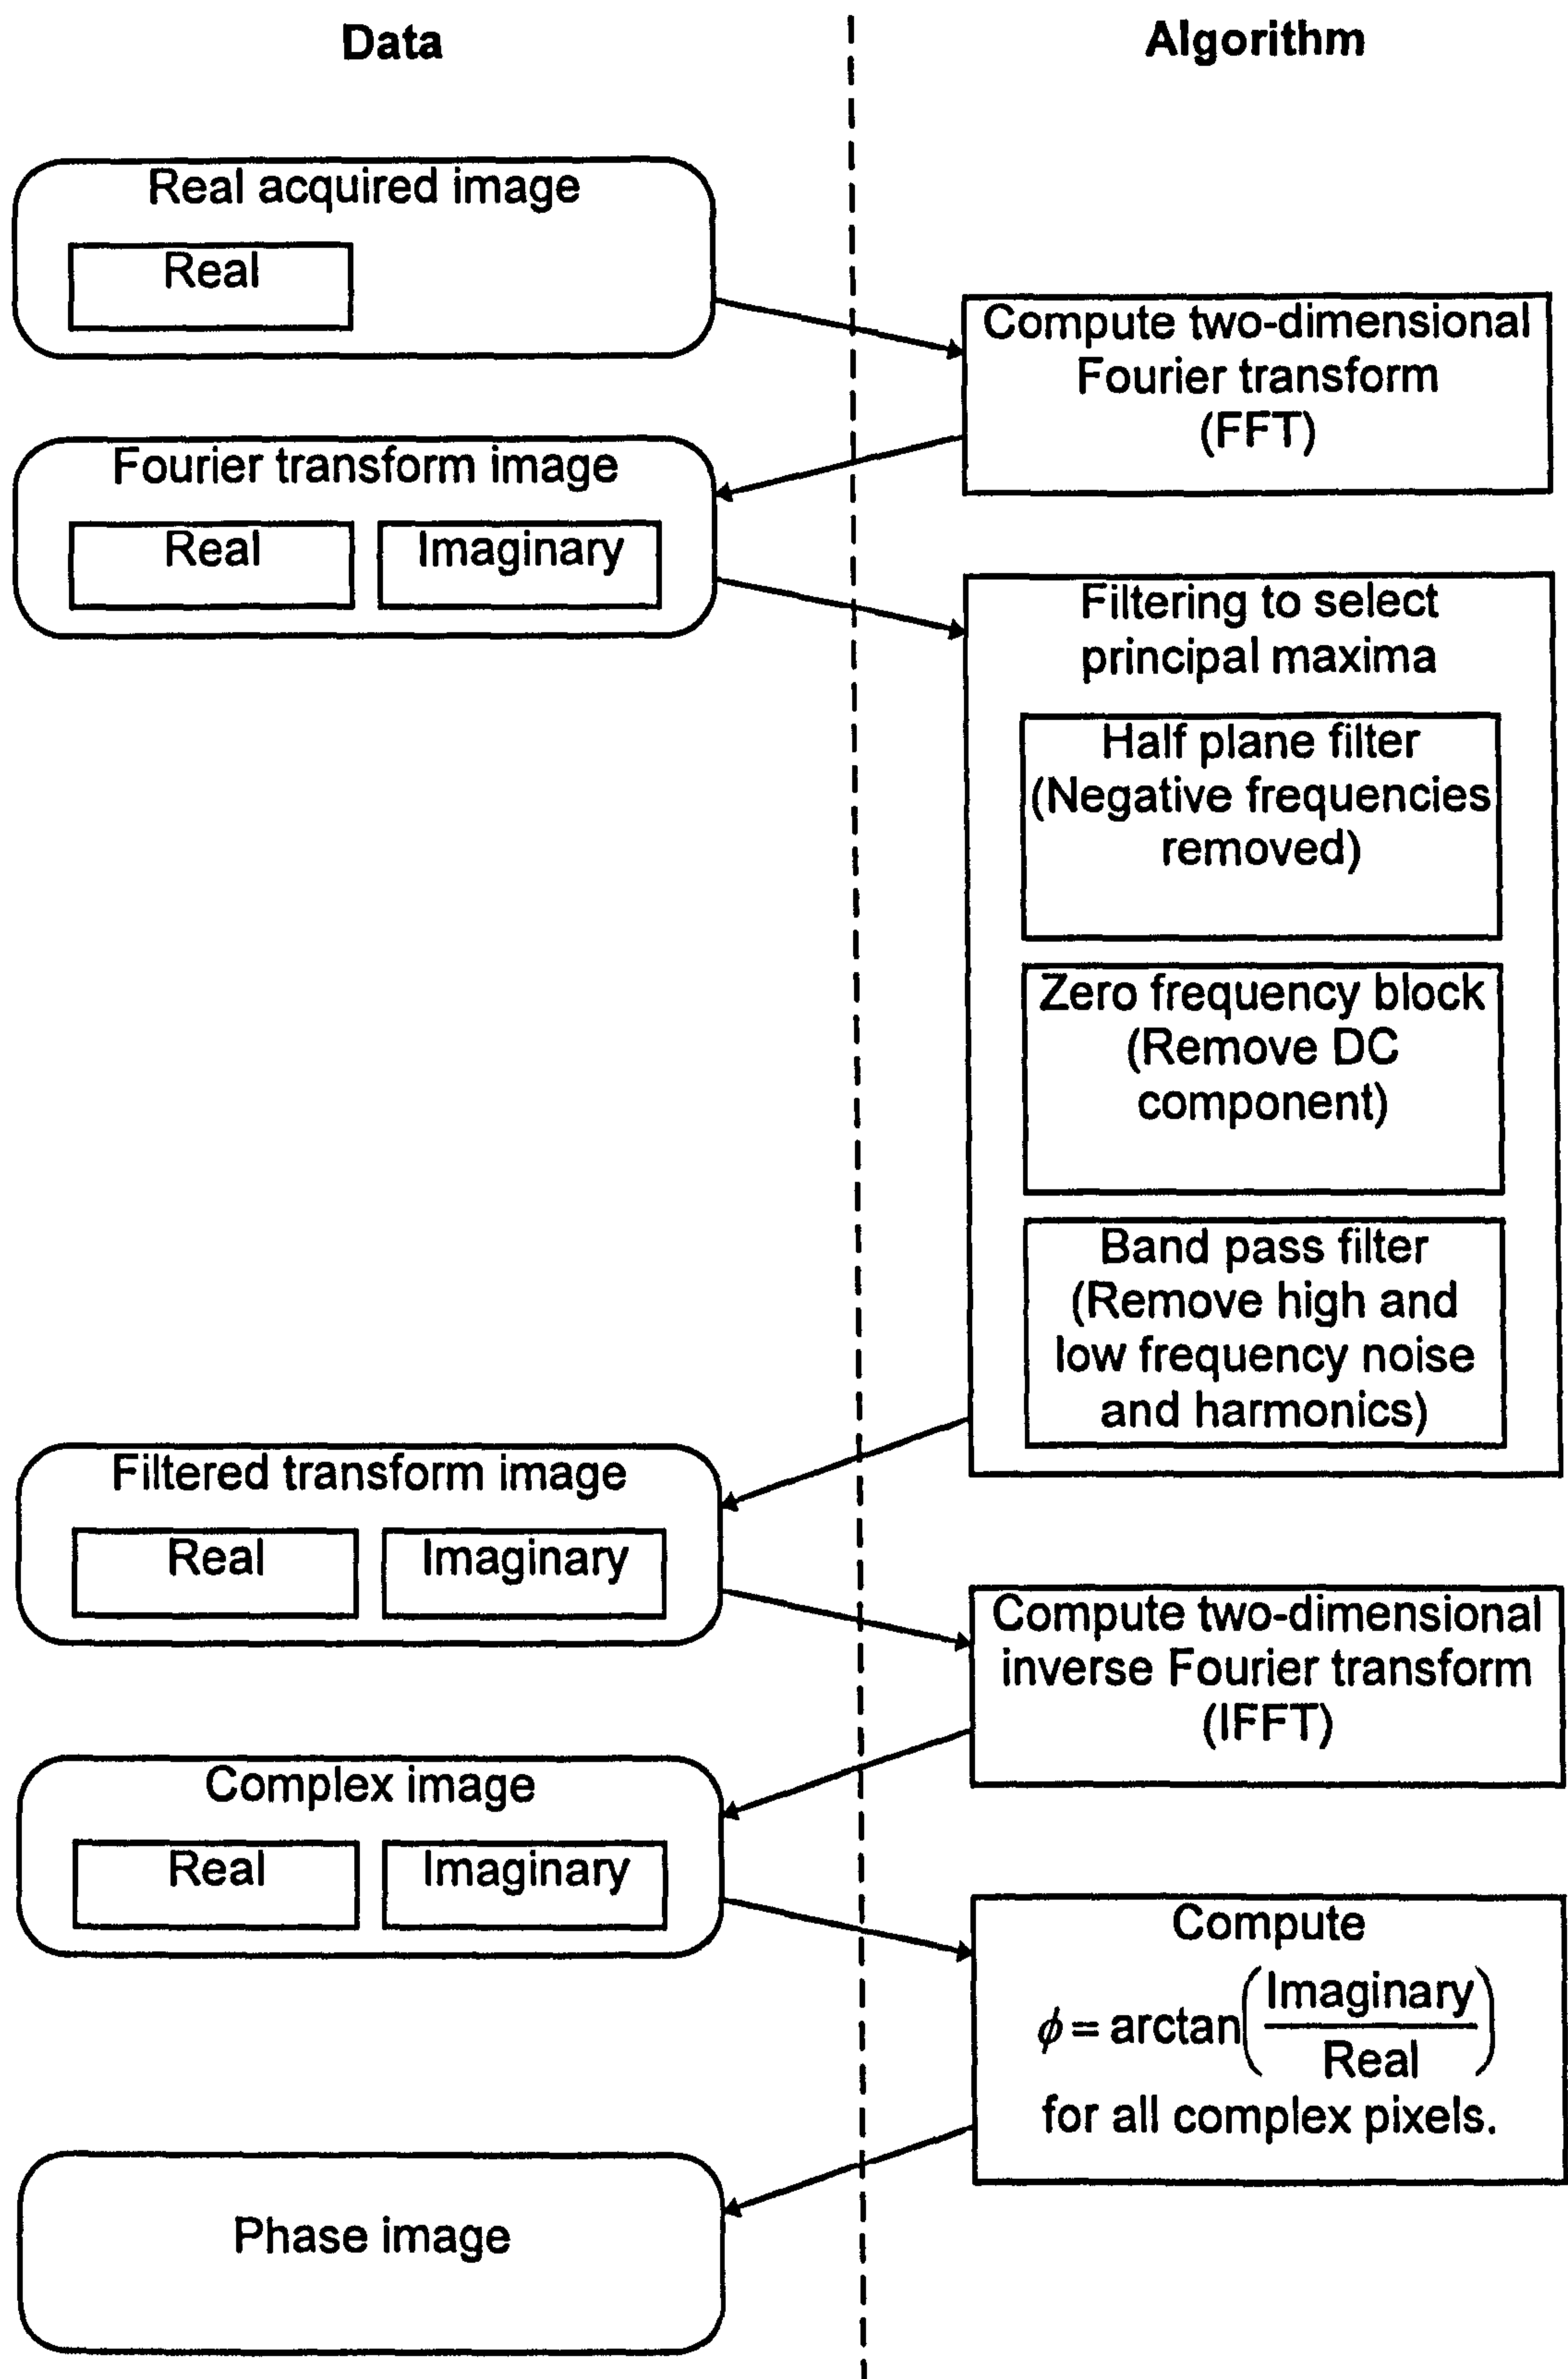


Figure 2.16 Flowchart for the FTP method of phase extraction.




Figure 2.17 Filtering in the Fourier domain.

After the image has been filtered, the inverse Fourier transform is applied to the image. The IFFT algorithm yields an image which has both real and imaginary parts. The phase of the fringe pattern for each pixel in the image is then calculated to be the angle between real and imaginary parts. The resulting image is then a phase map, or phase image of the original fringe pattern which should be similar to that which would be produced if the phase-stepping method were used. Again, phase values will be wrapped, and an unwrapping algorithm followed by conversion to depth coordinates must be applied for surface reconstruction.

2.5.4.2 Application to Back Shape Measurement

Perhaps surprisingly, little, if anything, is mentioned in the literature about the application of FTP to measurement of back shape. This may be due to a number of physical and computational limitations of the method. The computation of the FFT is notoriously demanding of floating point computing power and application of the FFT to areas of interest which are not full field in the image, or are not at least rectangular, is difficult.

An experimental arrangement to investigate the application of the FTP method was constructed as part of this investigation. A Ronchi grating of period approximately 1 line/mm was mounted as a 35mm slide and placed in a conventional slide projector. A test object was used for the investigation and this was a plaster model taken from a mould of a real human back.

Figure 2.18 shows the optical arrangement. For convenience, the object was chosen to be full field in the image. The camera-object separation, D_C was 2000 mm and the camera-projector separation D_P was 1000 mm giving an angular separation of 26.6° at the object. The camera had a resolution of 512x256 pixels and was mounted with a variable focal length lens.




Figure 2.18 Test configuration for FTP.

The FTP phase extraction method described above was applied to the acquired image shown in Figure 2.19. The DC component was removed in the image by subtracting the average image intensity. In addition to the half plane filter, a band pass filter with lower and upper limits set to be $f_0 - f_0/2$ and $f_0 + f_0/2$ respectively, where f_0 is the fundamental frequency of the fringes, was applied. This filter is often satisfactory as a 'rule of thumb' filter for FTP since it is generally broad enough to pass the modulation of the fundamental whilst blocking dc and low frequency elements due to background intensity changes and blocking high frequency harmonics which are present because the fringe pattern was not sinusoidal.

Figure 2.19 shows the original source image together with the resulting phase image. The phase image shown was produced simply to demonstrate the FTP method as part of an investigation into the feasibility of the technique. It can be seen that in some areas of the image, phase recovery has failed. More careful design of filters would probably have increased the effectiveness of the phase extraction.




Figure 2.19 Image of plaster-cast of back with resulting phase image.

The FTP phase extraction requires the computation of both forward and inverse Fourier transforms as well as filtering operations. The microcomputer used for this investigation (see Appendix II for details about the machine) took 21 minutes to extract the phase of the fringe pattern using an FFT algorithm with standard optimisations (look-up tables for sin, cos and bit-reversal operations etc.). The processing time necessary using the FTP method precludes it from use in a practical automated measurement system based on a standard microcomputer.

The sample images shown here have a rectangular area of interest, full field in the image. In many circumstances, especially for scoliosis assessment, the patient's back may not be full field in the image. If the object is not full field in the image then implementation of the technique becomes far more complex because we are no longer

calculating the FFT of a simple two-dimensional rectangular image. Either the FFT implementation or the optical configuration must be carefully modified.

2.5.4.3 Discussion on FTP

The principal advantage of the FTP method over PSP is that it requires only a single image for phase demodulation. The total data acquisition time is confined to a single frame cycle (1/25s). During this time a patient will not move significantly and the fringe pattern will be a good record of surface shape. FTP does not require a projection system in which the fringes must be shifted by a fraction of the period of the grating. This leads to a simple optical system in which there are few, if any, mechanical moving components which are subject to wear and damage.

FTP does however have shortcomings. The method requires the computation of both a forward and inverse Fourier transform. The high computational task of computing the FFT limits the time taken for a fringe pattern to be processed if high speed or real-time results are required. Specialist architectures have been developed for phase extraction using the FTP method such as the parallel array of digital signal processors developed by Kshirsagar (1995), but these are not in widespread use for the application and cost would almost certainly prohibit their use in the type of imaging system proposed. The general increase in computing power and decrease in cost of microcomputers does, however, increase the ease with which FTP can be realistically implemented.

Perhaps the most significant shortcoming of FTP in its application to measurement of back surface shape is that if the fringe pattern is not full field in the image then a number of problems can arise. Unlike PSP, FTP cannot perform phase extraction on single pixels: it requires a continuous, rectangular sub-image area. Human beings, sadly, do not come in this shape and it would be difficult, if not impossible, to reconstruct the shape of the entire back using the method.

Lastly, FTP requires considerable expertise in filtering in the Fourier domain and it is very difficult to use filters that will work across a broad range of surfaces.

2.5.5 Summary of Phase-Measuring Methods

Table 2.1 gives a comparison of the two principal methods for phase demodulation. The practical limitations of phase measuring methods render them unsuitable for an automated system of the kind defined by the criteria in Chapter 1.

| | Phase-stepping profilometry (PSP) | Fourier transform profilometry (FTP) |
|------------------------------|--|---|
| Fringe Pattern | Generally requires fringes with sinusoidal intensity profile. | Can be used with most types of fringe intensity profile |
| Optical Configuration | Requires precision moving mechanical elements to translate fringe pattern | Simple static optical system |
| Acquisition Time | Long acquisition time: four images are captured and the pattern must be translated three times | Short acquisition time: only a single image need be captured |
| Data Size | Large data size: four digital images | Smaller data size: single digital image, but floating point complex images required in processing. |
| Image Processing | Simple algorithm. | Complex algorithm: FFT plus careful filtering |
| Computational Task | Phase values can be obtained quickly | Heavy computational task mainly due to FFT and inverse FFT |
| Noise Resilience | Relatively resilient to effects of noise | High and low frequency noise must be filtered from the fringe pattern during phase evaluation |
| Area of Interest | Individual pixels or small areas of interest can be selected for processing | Entire image or sub-image area must be processed: can lead to problems if object is not full field in the image |

Table 2.1 Comparison of PSP and FTP phase measuring methods.

FTP requires powerful processing capabilities which cannot be found on standard microcomputers and is difficult to implement if the object is not full field in the image, as would be the case for back shape measurement. PSP requires impractically long data acquisition times as well as precision moving mechanical components which are both expensive and possibly prone to failure.

As well as these practical limitations there are absolute limitations of the techniques which are sometimes overlooked in the literature. For example, it is sometimes suggested that phase-measuring methods are resilient to the reflective properties of the object. Suppose, for example, in the case of PSP, that an object had a small specular area on its surface. This would not affect reflected intensity when illuminated with a dark phase area of a fringe. However, after phase-stepping through half the fringe period, the bright fringe would be unexpectedly intense and resulting phase information would be corrupted. As a second example, consider the case for FTP. Suppose a surface itself had a sinusoidal reflectivity distribution. If sinusoidal fringes were formed on the surface in anti-phase with the surface sinusoidal reflectivity, then no fringes would be seen. Both these examples are unlikely scenarios in the application of phase measuring methods to the measurement of back shape. However, they illustrate the fallibility of the methods and many real surfaces will occasionally exhibit unusual reflectivity properties to some extent.

2.6 Optical Scanning Methods

An optical scanner is a device which allows controlled angular deflection of an incoming beam, plane or other pattern of structured light. Typically, a mirror mounted on a galvanometer controls the deflection of the beam. Two mirrors are required: one mirror to deflect the beam in each direction. For two-axis deflection, galvanometers are mounted perpendicular to each other and the motion of, say, a beam can be controlled in horizontal and vertical directions.

2.6.1 Principle of Scanning

Figure 2.20 shows possible configurations for single and two-axis deflection systems. In the case of a straight forward beam, which may be a laser, the beam is focused at the target plane and a small spot is formed where the beam strikes the object. For measurement of human back shape the back would be placed at the target plane. In a measurement system, a camera, which is angularly offset from the incoming light, views the object and triangulation is used to extract three-dimensional coordinates.

An important decision in the design of a scanning system is the type of structured light to be used. The most popular structured light patterns are a simple spot and a line of light. The light source for the system is often a laser which produces a narrow beam which is almost directly suitable for spot scanning. The insertion of a cylindrical lens in the path of the laser will expand it in one direction only so that a "plane" of light is formed which produces a bright line on the surface of the object.




Figure 2.20 Optical Scanning.

There are two common lens configurations suitable for this type of optical scanning which are identified by the position of the lens in relation to the scanner (Pelsue 1990). The structured light may either be focused prior to striking the deflecting mirrors, in which case we refer to post-objective scanning; or it may be focused following deflection in which case we refer to pre-objective scanning. Figure 2.21 illustrates these configurations.

The target is typically observed using a camera mounted orthogonally or, at least, angularly offset from the incoming beam. The image is acquired and transferred to the memory of a computer where the image is processed. The coordinates of the spot or line of light are extracted and consideration of the geometry of the system yields a z coordinate for the pixel (x,y) and three-dimensional reconstruction is achieved.

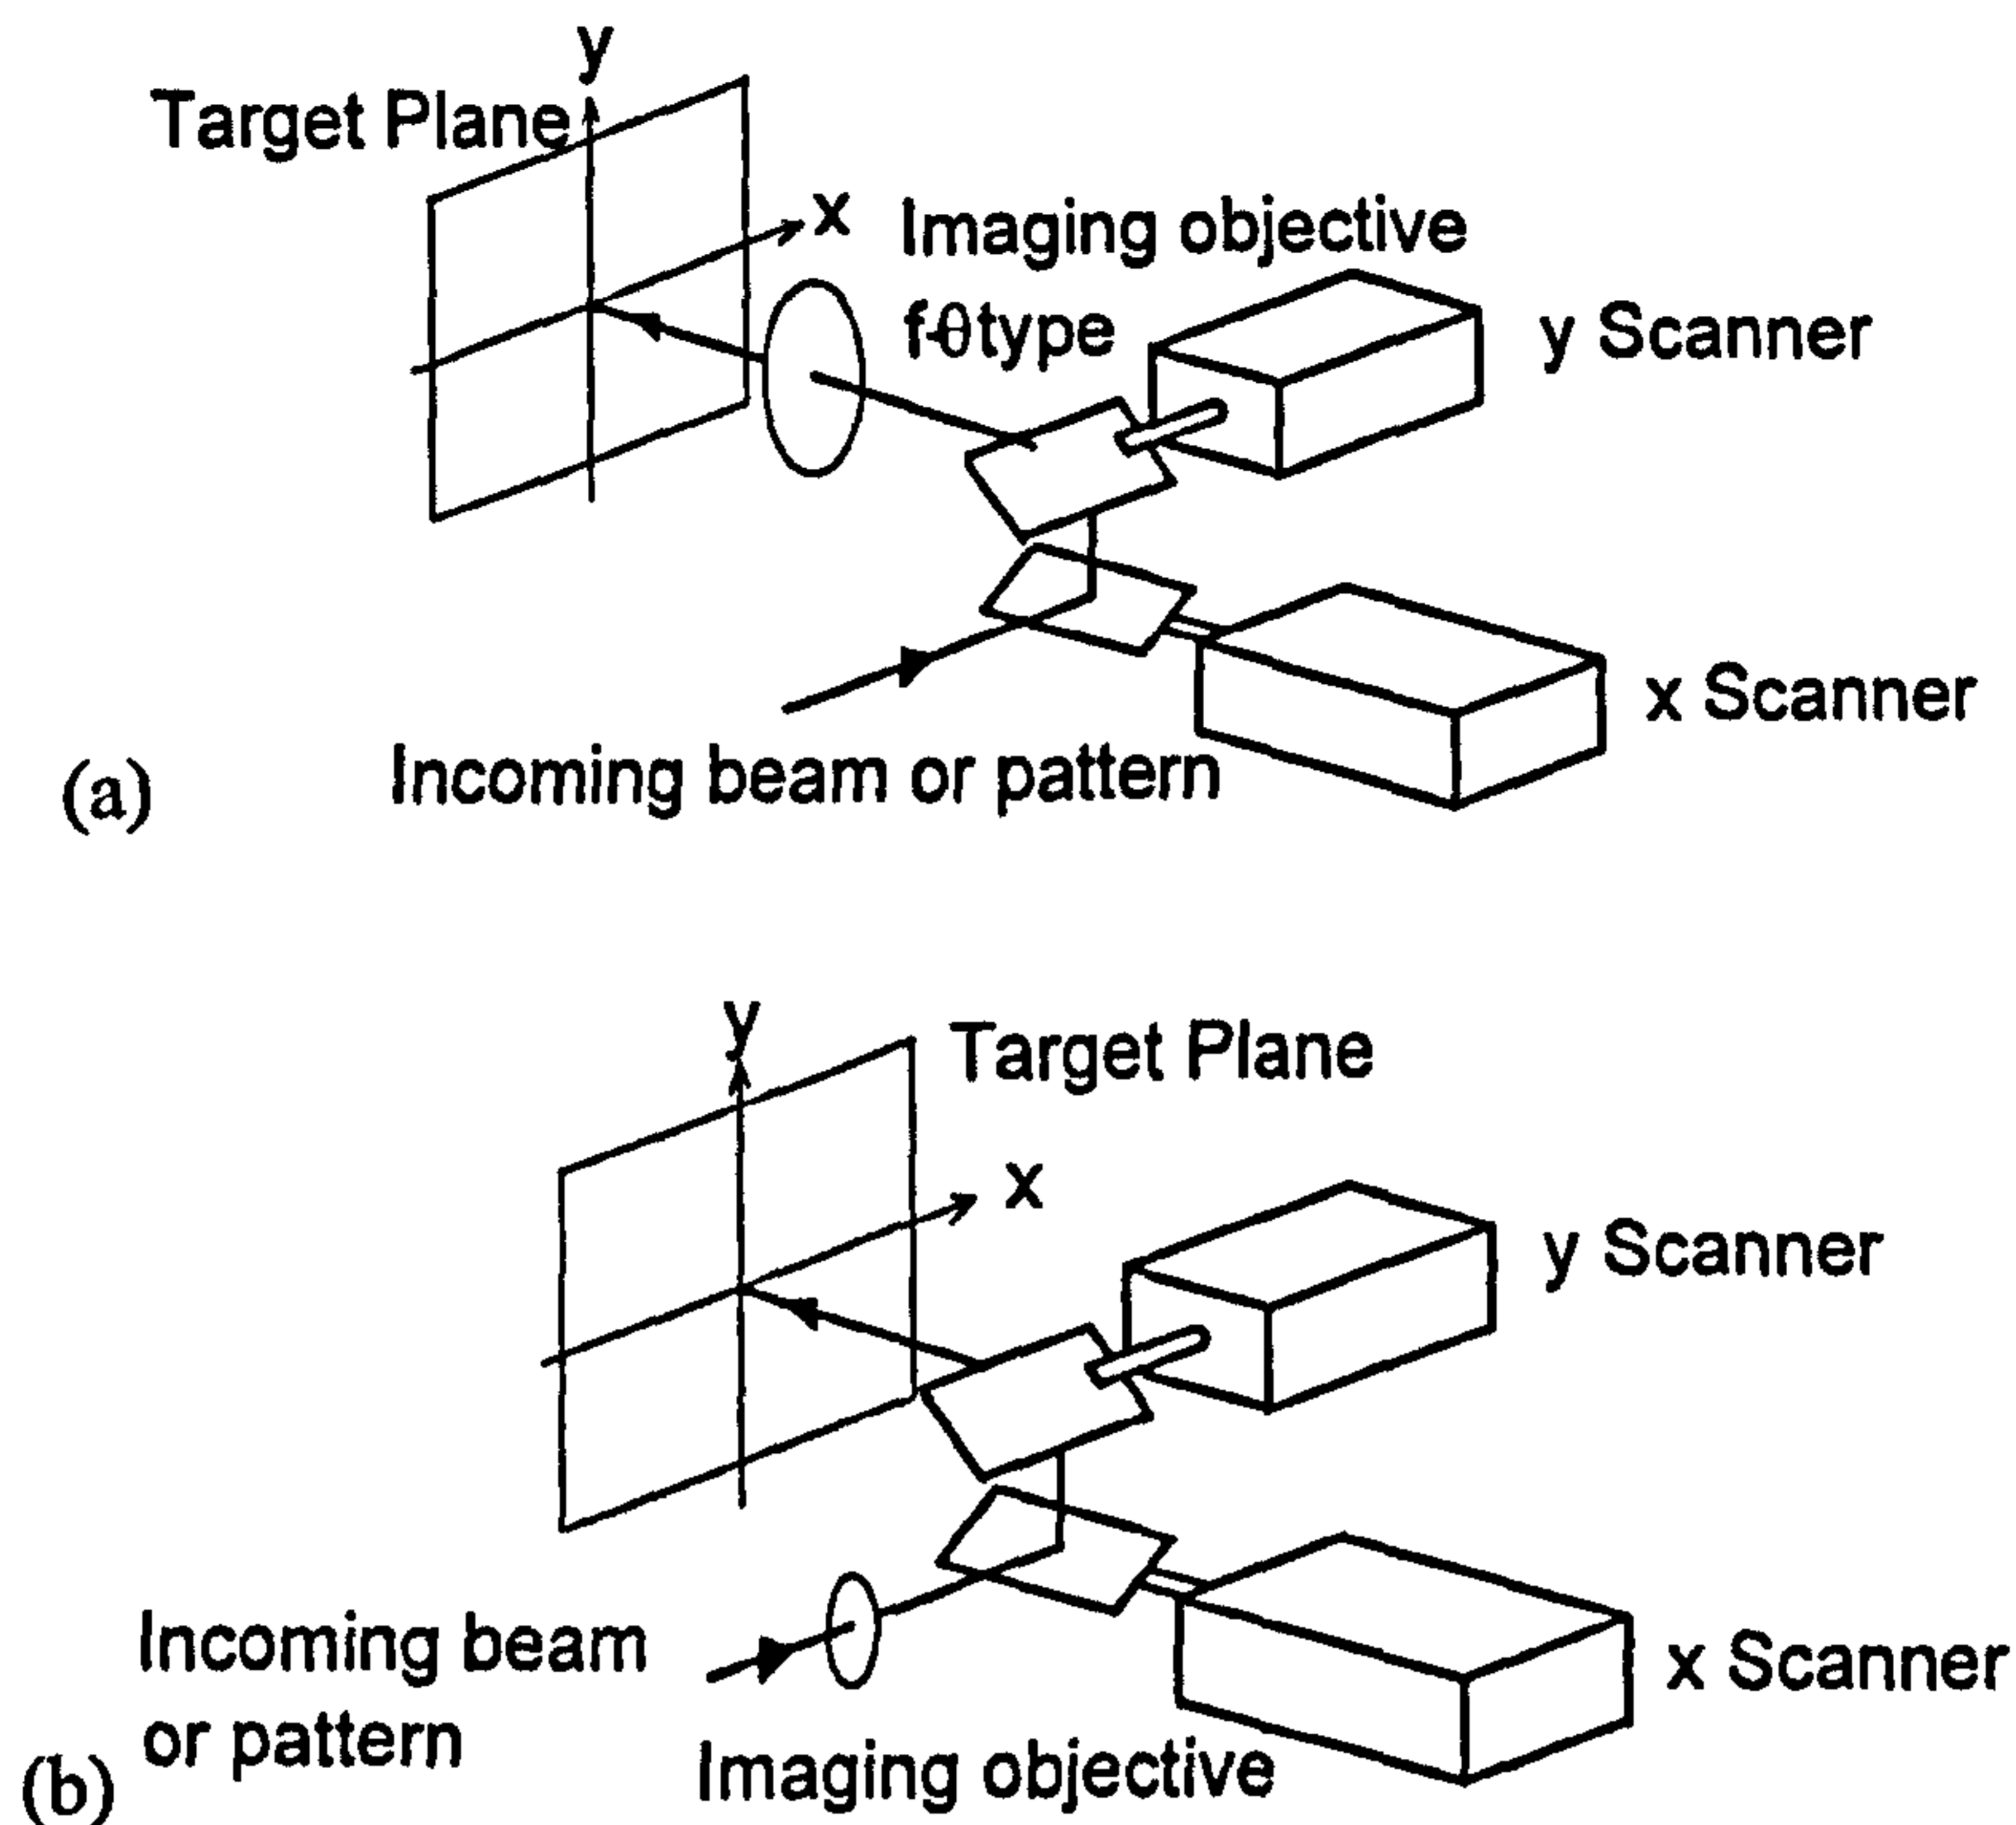


Figure 2.21 Optical scanning (a) Pre-objective (b) Post objective.

The image processing task is generally simple: all that is required is to extract the coordinates of the bright spot or other pattern from a dark background. The simplicity of processing of the image is, however, offset by the large number of images which must be acquired if a realistic number of sampled data values is to be collected. In addition the object must remain stationary during scanning if accurate results are to be achieved.

2.6.2 Clinical Application of Scanning

Despite the known problems of scanning, the technique has been successfully applied to the measurement of human body shape. Indeed the only body shape measurement system to achieve widespread clinical use uses this method for data acquisition. Lasers are often chosen for light sources in scanning systems. However, if the object is a human back, the ethical problems of using a laser source, especially since it is often focused to a spot or line of light are difficult and wherever possible the light source should be a simple high power white light projector. Moiré grating projectors often provide the necessary intensity.

2.6.2.1 Configurations for Scanning Systems

Since the human back is a gently varying surface with few discontinuities it is often convenient to use a simple stripe or plane of light as a structured light pattern. Spot scanning is not impossible but since a spot represents a single sampled data value in an image, the number of images required is large and this presents a large data set and a heavy computational task. The typical configuration for scanning of the human back is shown in Figure 2.22.

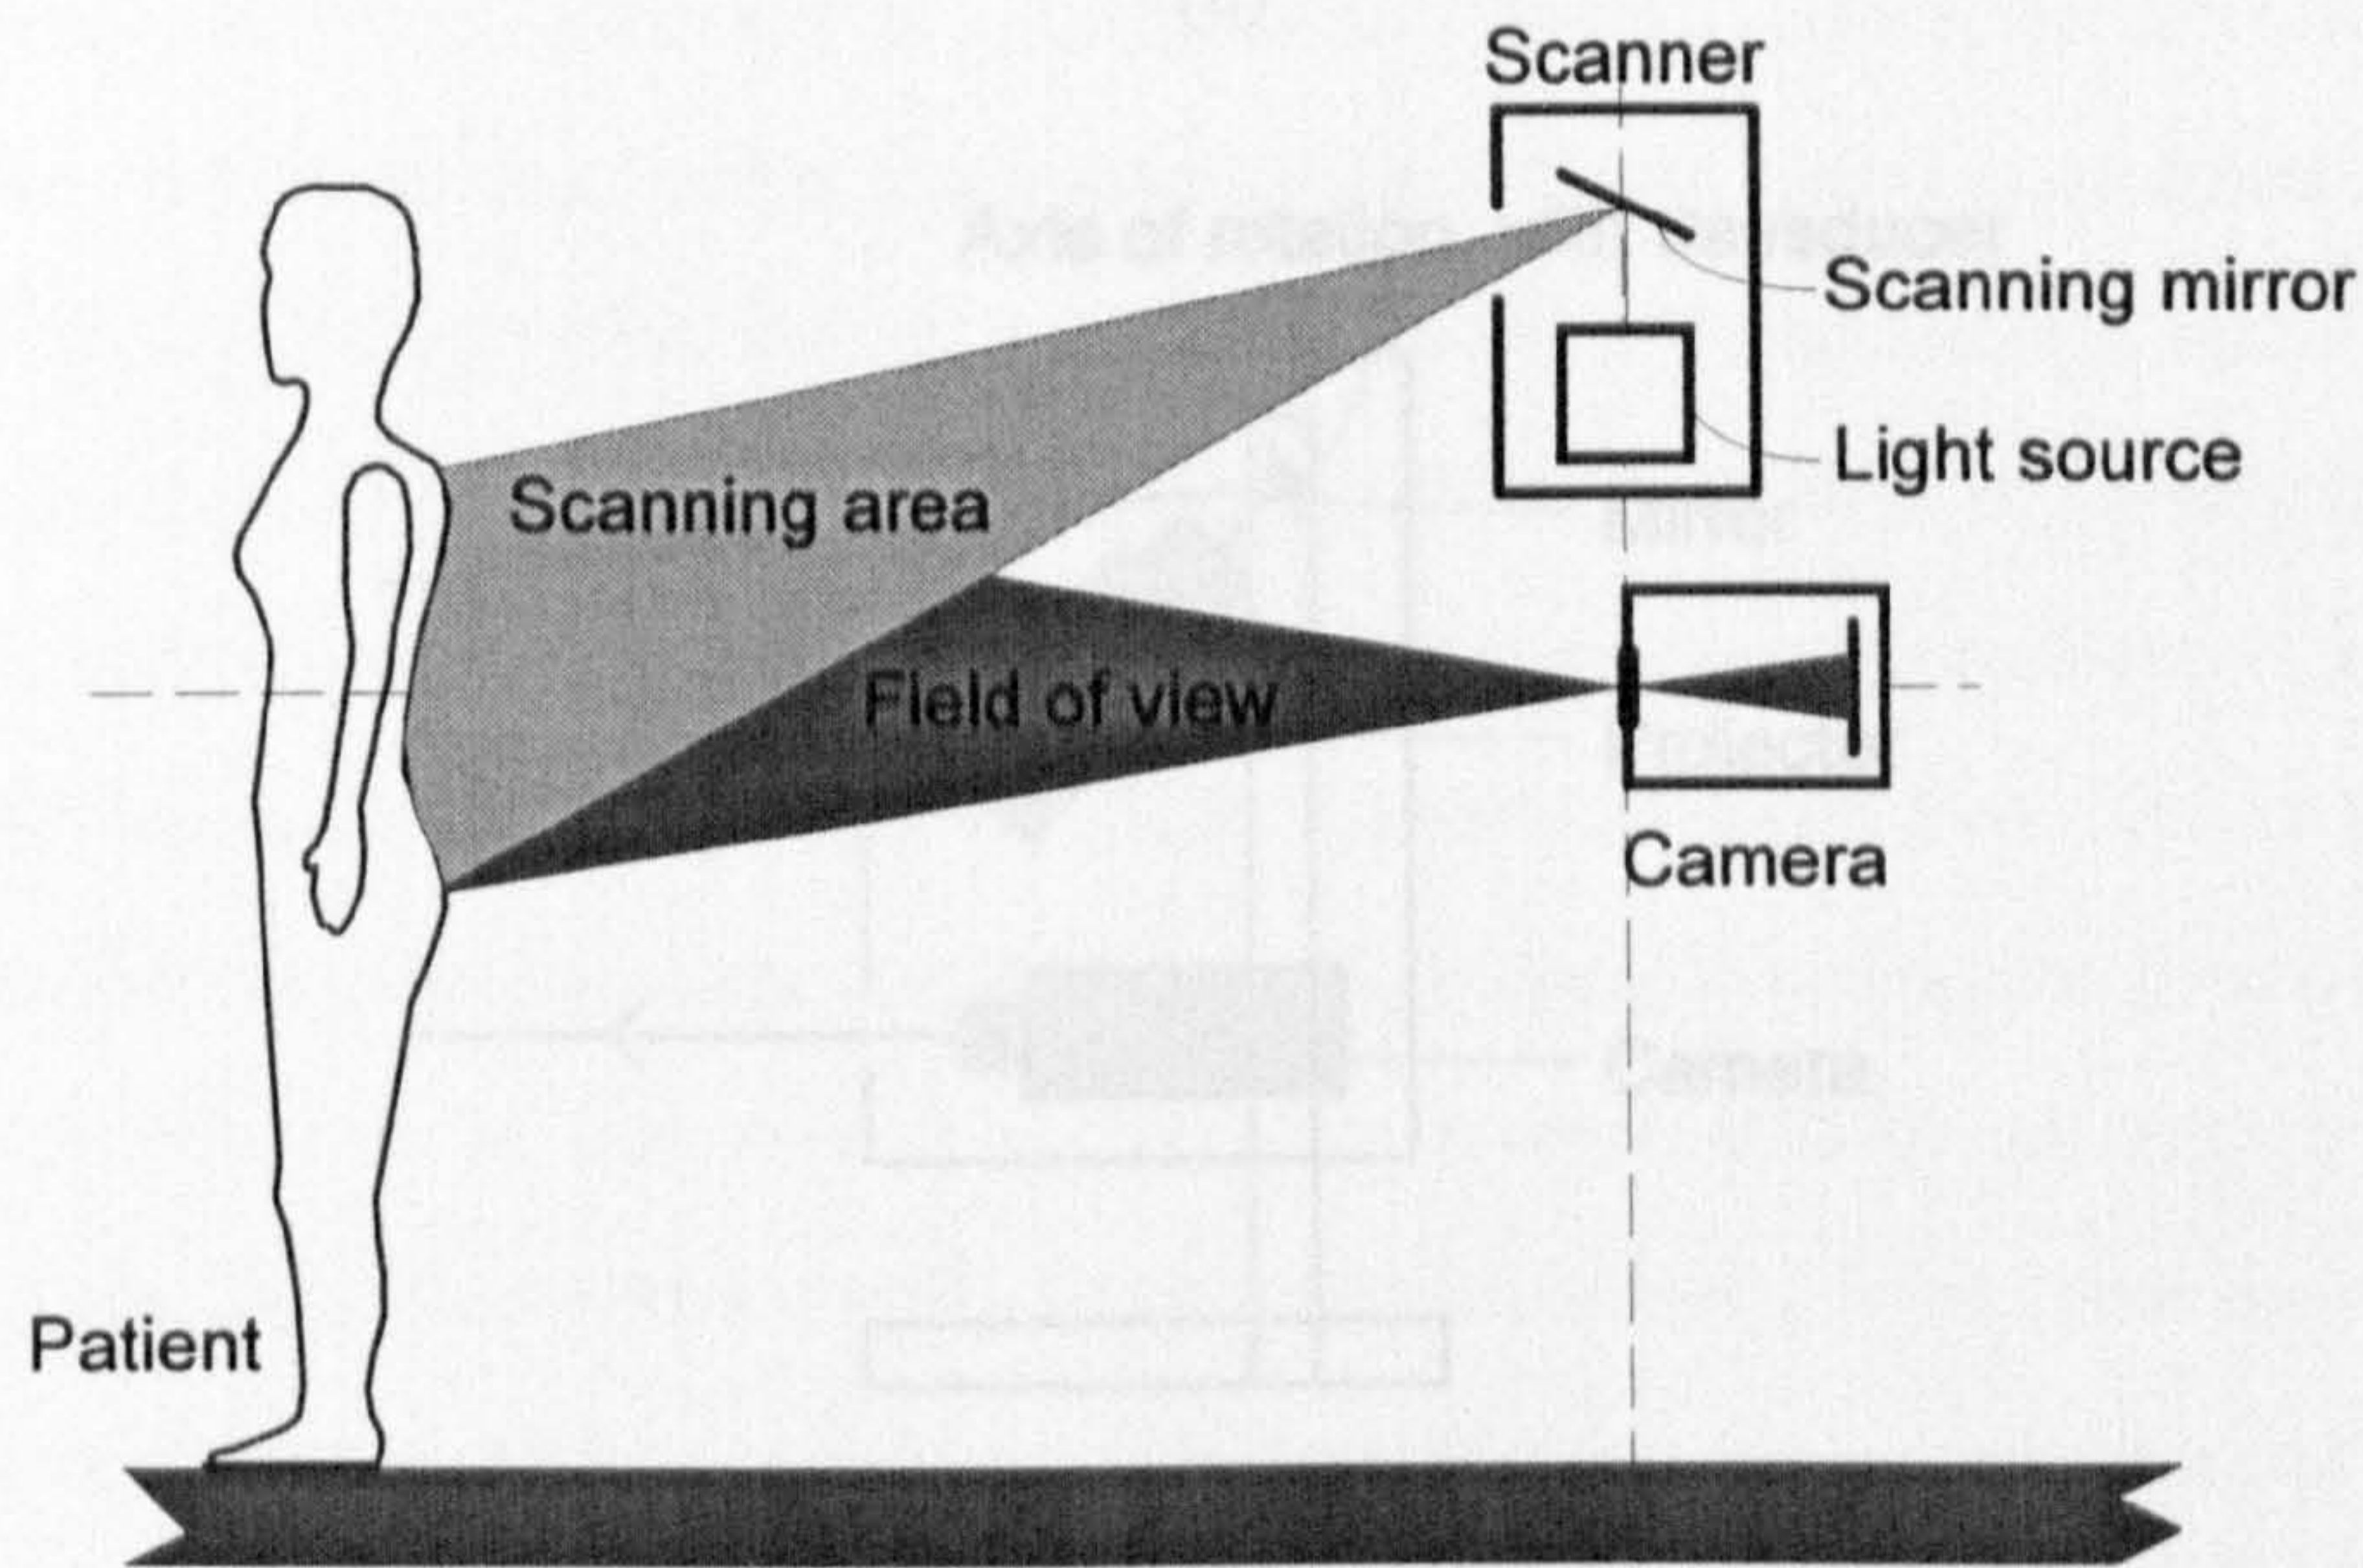



Figure 2.22 A typical configuration for optical scanning.


2.6.2.2 A Case Study with Scanning: The ISIS System (Turner-Smith)

The most extensively used automated system for measuring human back shape in a clinical environment is almost certainly the ISIS system which is commercially available (Turner-Smith, 1988). The ISIS system uses white light structured as a single line on the object. The design of the optical scanner, which incorporates a camera, is of particular interest. A schematic of the ISIS scanner is shown in Figure 2.23.

The principle of operation of the ISIS scanner is very different from that used in conventional scanning applications. There is no galvanometer based mirror scanning: instead the entire unit rotates about an axis.



(a)



(b)

Figure 2.23 The ISIS Scanner (a) Clinical configuration (b) Mechanical lay-out

A standard 625-line television camera and a standard 35mm slide projector are mounted in a fixed relationship within a unit which could rotate about a horizontal axis. The projector produces a single line, or horizontal plane, of light which the camera then views from below. The camera produces a TV signal and is mounted so that the scan lines of the TV signal are vertical. The angle between the angle of incidence and the viewing angle is about 26° and was chosen as a compromise between the necessary distance sensitivity and the required depth of field and the overall bulk of the scanner.

The line, or plane, of light was produced by a 35mm slide mounted in the projector. An interesting method was used to adapt the system for the specialist application of measurement of human back shape. It is desirable to measure as far around the body as

possible, but less light is reflected from the edges of the patients back, where the skin surface curves away from the projector, than from the centre of the back. To compensate, the 35mm projector contained two slides which were sandwiched together: one of a single white horizontal line, as expected, and the other of a central faint grey cloud to diminish the intensity in the centre of the object. This gave an approximately even distribution of intensity when viewed by the camera. The line was approximately 3mm wide on the patient.

All the components are mounted in a fixed configuration inside the scanner enclosure. To produce scanning, the entire enclosure is rotated by speed-controlled motors (not shown in Figure 2.23). The unit is counterbalanced by gas springs which reduce the power required from the motor and allow a smooth motion with rapid acceleration to an almost constant angular velocity. The light from the projector is reflected from a mirror close to the axis of rotation so that the light appears to come from the axis of rotation although calibration procedures ensures that this does not require precise alignment. The angular rotation of the unit about its axis is recorded by an angular transducer mounted at the axis (this was a simple linear potentiometer). The value measured by the transducer is electronically encoded and inserted into the TV signal as a bar in the top of the television picture using custom electronics.

The TV signal then holds all the information required for three-dimensional shape reconstruction. The two dimensional coordinates of the line are extracted using specialist circuitry which detects the location of black-white transitions for each TV line. The same circuitry is able to extract a value for the angular position of the scanner. The set of coordinates which is recorded and passed to a computer are the angle of the transducer signal, the television line and the position of the black-white transition along the television line. The coordinates are passed to a microcomputer which then reconstructs the three-dimensional shape based on the system optical geometry. The entire system is carefully calibrated and non-linearities of the camera and optics are corrected.

The scanning time for data acquisition for a single patient is approximately 3/4 of a second. The accuracy of the system is approximately $\pm 1.2\text{mm}$ in the x direction, $\pm 3.0\text{mm}$ in the y direction and $\pm 3.5\text{mm}$ in the z, or depth, direction.

It is generally agreed that the ISIS system is now technologically outdated. The TV camera, supporting electronics and microcomputer are the most obvious examples of obsolescence in the system. The method used remains valid and it is possible that a more accurate and faster system could easily be constructed using modern components. However, some of the shortcomings are unlikely to change with advancing technology. A principal shortcoming of the ISIS scanner is that it takes some 3/4 seconds, or thereabouts, to acquire a single data set. Perhaps surprisingly, the manufacturers and some of the users claim that this long acquisition time, during which the patient must remain perfectly still, does not effect the usefulness of the machine and satisfactory results are obtained. The single projected scan line means that each image can be easily and unambiguously interpreted to extract three-dimensional coordinates.

The microcomputer which performs the surface reconstruction activity also performs other tasks such as the measurement of specific parameters which relate to scoliotic deformity. Until recently, the ISIS system has been the only widely used and clinically proven automated system for measurement of body shape. Much work has been done by the international research community in the interpretation of its results and on assessment of its use as a clinical tool.

2.6.2.3 Discussion of Scanning

Optical scanning is undoubtedly a reliable method for data acquisition in three dimensional measurement systems. The simple light structures used, usually a spot or a line, enable coordinates to be extracted unambiguously: the indexing problem associated with fringe pattern or line raster contouring (see Section 2.7) does not exist because the projected line is the sole data carrier in each image. Furthermore if, for example, data is unavailable in a part of the surface due to shadowing then the absence of data is easily recognised and the need to use complicated algorithms to correct indexing of other local stripes or fringes does not arise.

The principal shortcoming of scanning for measurement of human body shape is the time taken to acquire data. Typical times are of the order of a second. Although this has been reported to be satisfactory for the application, the problem of the patient moving during acquisition must still exist: it is not really possible for the patient to remain motionless for this period of time. Also, as has been already mentioned, it is desirable that acquisition time should be minimised so that dynamic measurement of body shape can be possible in the future. Scanning systems cannot be adapted for dynamic measurement.

If scanning were to be implemented using the now widely used imaging principle of using a frame-store hosted in a microcomputer, the processing power and memory capabilities of the computer must be extensive because the number of images to be processed is large (50-200 frames are typically required). This problem can be overcome if either (a) the algorithms for location of coordinates can be performed in real time or (b) specialist hardware is used to extract coordinates as is the case with the ISIS system. (a) requires expensive and often non-standard hardware and (b) requires that the algorithm be fixed by hardware so that modification and development is difficult.

Scanning systems necessarily require moving mechanical components. For classic scanning this requires the use of delicate and precise mechanical components which must be carefully maintained. For scanning using configurations such as that used for the ISIS scanner there is still a motorised component. These moving mechanical components are undesirable and a system which is free from such components will be less susceptible to failure.

2.7 Raster Pattern Contouring

A raster pattern is simply a regular array of structured light objects that are usually projected onto the object. Raster pattern contouring involves the direct location of the structured light objects from which three dimensional shape can be inferred. The new method proposed in Chapter 3 borrows from raster pattern contouring where typical

configurations. This section will address the area from the standpoint of the various types of structured light pattern that might be selected.

It has already been observed that the main challenge of structured light projection is to select a structured light pattern in which individual structured light markers can be easily located for subsequent processing to calculate three-dimensional coordinates. Possibly the ideal solution would be to encode each individual structured light marker uniquely so that it can be unambiguously identified in the acquired image. Solutions of this type might include assigning each marker to a unique colour and to project an array of coloured markers onto the object, which would then be inspected using a colour camera and colour image processing hardware and software. This option has not been considered in this investigation because of the high cost of colour imaging cameras and processing hardware. Another solution that would provide unambiguous measurements would be to encode each marker using a unique shape for the marker, but this would require complex and possibly high resolution imaging and image processing if satisfactory results were to be achieved.

If the approach is *not* to uniquely encode each structured light marker, then similar or identical markers must be used together with image processing algorithms that will resolve ambiguities in marker position, usually by identifying the position of markers relative to each other. The range of possibilities for marker structure is almost infinite and many arrangements have been constructed. However, two patterns appear to have been successful in the literature: point raster projection and line raster projection; and these are discussed in the following subsections.

2.7.1 Point Raster Projection

2.7.1.1 Principle of Point Raster Projection

Point raster projection is possibly the simplest and most intuitive way to produce a structured light pattern on an object. A rectangular array of bright dots, on a dark background (for example the dot pattern shown in Figure 2.24) is usually projected onto the object.




Figure 2.24 Dot or point raster that might be used for projection.

The optical geometry will be similar to that described in Chapter 3. Typically the image will be acquired using a CCD camera and transferred to a frame-store where it may be manipulated by a computer. The simplest image processing approach is then to apply a threshold to the image to identify the dots and to order the dots so that matching pairs on the object and the projected pattern can be identified. Trigonometric consideration of the optical geometry will produce depth, z , coordinates for each measured (x,y) point.

2.7.1.2 Case Study with Point Raster Projection: Hill

Hill (1992) used point raster projection to measure trunk shape. He used 2mm diameter dots with 5mm separation between dot centres on the body. In addition, a number of control points were projected onto the body to aid dot location. The optical configuration, shown in Figure 2.25, used three cameras.

For any individual point, two cameras only were required to calculate coordinates, with the third camera providing redundancy in some areas of the image. The cameras had a resolution of 649x492 pixels and the frame grabber a resolution of 640x480 pixels. The field of view of the cameras was about 600x500mm using 16mm f/8 lenses and this translated to a spatial resolution of about 1mm/pixel at the object.

Initially, Hill experimented with a number of dot sizes but found that dot sizes less than 2.5mm in diameter at the object (the back) could not be reliably resolved. Control points in the image were added to compute the positions and attitudes of the three cameras. Eventually a dot size of 2.5mm with a separation of 6mm was satisfactory and

this gave an average error of 6mm in the z (depth) direction and a maximum error of 10mm.




Figure 2.25 Arrangement used by Hill

Hill found that lighting levels were critical, with the effective depth of field of the system increasing with higher lighting levels but excessive lighting causing blooming; conversely lower lighting levels reduce the effective depth of field of the system and caused pattern noise due to defects in the camera sensing unit. The algorithms for identifying dots in the two images were not published. The system did eventually progress into a local clinical system for measuring scoliotic deformity (Moreau, 1992).

2.7.1.3 Discussion of Point Raster Scanning

One problem that must inevitably be encountered using a point raster is that of identifying the same dot in two or more images. If a particular dot in one image is paired with a different dot in another image then the three-dimensional coordinate calculation will give an erroneous result for that sampled point. Moreover, since subsequent dots are ordered relative to each other, pair mismatching errors may propagate and subsequent coordinate calculations may be erroneous. One partial solution is to increase the separation of the spots. However this will decrease density of sampled data (density of sampled data \propto dot spacing²)

Point raster projection gives an approximately uniform density of sampled data in all directions. The ultimate objective of most body shape measuring systems is to calculate specific parameters which characterise the shape or quantify a deformity. Such parameters often require local measurements of shape to be taken in a particular direction on the body surface. It is then desirable to have a higher density of sampled data in the direction in which the measurement is to be made. Parameters describing scoliotic deformity such as ATI and TAS (see Chapter 1) require a high density of data in coronal sections. It is sometimes desirable, therefore, to have a high density of sampled data in the coronal plane.

Raster point projection allows the calculation of three-dimensional shape to be obtained from both the horizontal and vertical position of the dot in the acquired images. For the majority of sampled data points the three-dimensional coordinates can be measured independently from both horizontal and vertical data in the image thus allowing either redundancy of data or increased accuracy through averaging of coordinate results.

The practical problems of ordering point raster projections coupled with the low density of data in the direction of interest render this projection method inappropriate for use in a fully automated system to measure the parameters of interest.

2.7.1.4 Bitwise Binary Encoding

A development of raster point projection that has become practical for recent measurement systems (Goodall 1994) is, perhaps worthy of note here. Bitwise Binary Encoding (BBE) relies on an ability to alter the projected pattern dynamically. This can be achieved by replacing the dot pattern grating (conventionally produced on glass or a 35mm slide and mounted in a projector) with a Liquid Crystal Display (LCD) panel that can be easily controlled by a computer (often it is a VGA device).

Suppose that the image acquisition hardware uses 8 bits for intensity measurement at a pixel (256 grey levels). A rectangular grid of 256 points is constructed. Each point is assigned a unique index in the range 0 to 255. A set of images is acquired in which a regular rectangular grid (or raster) of points is projected onto the surface. The first image is a projection of bright pixels for those pixels that have bit 0 of their index set;

the second image is a projection of bright pixels for those pixels that have bit 1 of their index set ... and so on so that 8 images are acquired. After acquisition, each image is binary thresholded in the image processor and the appropriate bit is set in the composite image produced. After processing all the images, the composite image will have 256 measured coordinates with each coordinate indexed to correspond to the appropriate pixel in the projection system. Consideration of the optical geometry for the measured pixels will yield three-dimensional coordinates and a surface reconstruction is achieved. In practice it may be difficult to use single pixels for the operation, so often a local collection of pixels is used.

The advantage of BBE is that the problem of locating unique points is eliminated, however, the disadvantages include that the technique is limited to, typically, 256 measured data points and requires multiple images of the object. The acquisition of more than one image has already been identified as undesirable in a practical system for measuring back shape in the discussion of PSP earlier in this chapter. BBE was not considered seriously for use in the system that is the result of this investigation. To date no back surface measurement research based on the BBE technique has been published.

2.7.2 Line Raster Projection

Point raster projection has two major drawbacks. First, the density of sampled data points is low because the points must be separated by a significant distance in order to be reliably resolved and second, matching corresponding points in the projection system and the acquisition system is difficult. These problems can, to some extent, be overcome by projecting structured light lines. In this case the structured light pattern may be similar to the fringe patterns produced in interferometric optics and for this reason the structured light pattern produced is sometimes referred to as a *fringe pattern*.

2.7.2.1 Principle of Line Raster Projection

The principle of line raster projection is described in detail for the research in Chapter 3, but the key aspects of the basic method are reviewed here. Figure 2.26 shows a grating that might typically be projected to produce a line pattern projection. The raster pattern is projected onto the surface to be measured and an image is acquired by a camera

angularly offset from the projector. The shift position of the line on the object will obviously be curved for a back and the extent of the curvature encodes three-dimensional information. The shift in the position of the line in a direction perpendicular to the raster lines is measured by comparing the position that the line would have on either a physical or conceptual reference plane.




Figure 2.26 Simple grating for Line Raster Projection

The problem of finding matching locations in the projection system and the acquisition system is very much reduced in comparison to point raster projection: each line represents a single structured light object and there are far fewer lines than there might be points in a point raster configuration. Each line, however encodes a series of sampled data values. The density of sampled data values in a direction perpendicular to the lines may not be significantly different than for point rasters. However, in the direction parallel to the line the sampled data density is high and is only limited by the resolution of the imaging device.

2.7.2.2 A Case Study with Line Raster Projection: the Frobin System

Frobin et al. at the Munich Institute for Experimental Biomechanics proposed a line raster system that was first based on photographic film that was subsequently scanned as an image (Frobin 1981), then adapted for automated imaging using a CCD camera (Frobin et al. 1988) and finally produced as a clinical system by Jena (JenaOptik GmbH, 1993).

The system devised by Frobin (1981) modifies the basic principle of line raster projection by using a complex pattern of slightly broader and narrower lines to aid the matching of lines in the acquired image and the projection system.




Figure 2.27 Experimental set-up used by Frobin.

The experimental configuration is shown in Figure 2.27. The projector was mounted approximately orthogonally to the object. The projector held a grating (called a “raster diapositive” by the workers) with a bright line spacing of 0.5mm on a dark background. Two line widths, 40 μ m and 60 μ m, were used in the pattern. This separation translates to about 1cm at the object. The camera and projector are held in a fixed geometry at a distance approximately 2m from the object and with the camera 0.8m above the projector. Both the projector and camera have fixed focal length lenses and the area of measurement is approximately 60cm x 80cm. Figure 2.28 shows a photograph of an image captured using the Frobin system.

The camera had a resolution of 576 x 604pixels and the signal was digitised to 256 grey levels and held in a frame-store. The image was subsequently transferred across a VMEbus system for processing using a 16.7MHz 68020 processor with a 68881 coprocessor.

[The early system, based on scanning of photographic film, relied upon detecting the leading and trailing edges of a line and selecting a point half-way between to be the location of a peak, which was defined to be the location of the line at that point (Frobin 1983). The separation of the leading and trailing edges, the width of the peak, was used to determine the type of fringe. The broad lines were typically four times the width of the narrower lines.]



Figure 2.28 Photograph of an image from the Frobin system.

In the automated system, Frobin observed a different response to the lines than might at first be expected. The CCD sensor identified “broad” lines in the grating to have the same width as narrow lines but presenting a higher (brighter) peak in the intensity profile as shown in Figure 2.29. This was attributed to the limited light sensitivity and resolution of the sensor and to blooming. The differentiation of broad from narrow peaks was then based on the height of the peak in the intensity profile. The method for detecting the fringe peak was also changed for the automated system. Working on the assumption that the peak was approximately Gaussian shaped, Frobin fitted the peak to a Gaussian curve. The location of the line could then be easily determined by identifying the maxima of the Gaussian curve. This method also enabled sub-pixel estimation of the peak position. The data set (the set of sampled data values) was reduced to about one-third for processing.




Figure 2.29 Intensity profile observed by Frobin (1988)

Both the location of lines and the determination of their type (heavy or light) using these techniques depended heavily upon the quality of the image signal. Careful control of gain and offset at the camera and in the frame-store were required to maximise contrast and minimise flooding. The optimisation of the system was ensured using an “exposure meter” program which automatically adjusts the system to provide optimum performance.

The depth accuracy of the system for sampled data points was claimed to be 0.4mm on the basis of a 0.15 pixel error in peak location (sub-pixel interpolation of peak position was used). Image processing and surface reconstruction took about 5 minutes.

The commercial implementation of the system by Jena-Optik uses a PC platform for computing and may have made further advances in the optical and image processing systems, but little, apart from advertising literature and basic technical specifications has been published.

2.7.2.3 Discussion of Line Raster Projection

One of the most important objectives for any clinical system to measure body shape in scoliosis is that the system must be capable of reconstructing surface shape reliably. For raster contouring systems one critical reliability factor lies in the ability to identify unique markers on the back. In line raster projection, the number of structured light objects is reduced (each line is a structured light object), but each object includes many sampled data values. Within the line, individual elements, which are the pixels that

comprise the line, can be tracked continuously (because they are adjacent to each other) so that a large number of data values is achieved relatively easily.

The density of sampled data values that can be achieved is primarily limited by the optical system and the resolution of the camera/frame-store combination. In a direction approximately parallel to the lines of the grating the density of sampled data points is almost purely dependent upon resolution. In the direction perpendicular to this, the density is proportional to line spacing. The minimum line spacing will be dependent on resolution in that direction but will also depend upon optical factors such as blooming and the ability of the system to avoid flooding of the pixels (normally a CCD array) and other optical limitations.

One possible shortcoming of the Frobin system was that it did not utilise both the leading and trailing edges of raster lines in the image or use both the maxima and minima values so that a single cycle yielded only one sampled data value. This is probably because the mark/space ratio of the grating was such that the bright lines were very narrow and indeed intended only to produce a single value, but may also be due to the legacy of an earlier approach to the problem based on photographic film.

The use of a unique pattern of two different line widths will undoubtedly increase the reliability of locating fringes. However, this is at the expense of considerable additional complexity in the processing where the objective is to be able to restore any loss in fringe tracking. The system proposed in Chapter 3 suggests that a much simpler approach will provide adequate reliability.

2.7.3 Discussion of Raster Contouring

Structured light methods use relatively easily implemented optical techniques to reconstruct body shape. The key factors for the success of the methods are type of structured light pattern (the type of grating) used and the image processing to which the recorded pattern is subjected.

The number of possible structured light patterns that could be devised is large. Two patterns have been investigated here: point raster and line raster projection. Line raster

projection was generally found to be preferable because the number of structured light objects is reduced and the pattern is easily processed.

Although raster contouring requires careful optimisation of the illumination system, this is no more complex than for other methods. Often the optical system can be constructed from inexpensive "off the shelf" components and this has obvious advantages in both development and clinical systems. Higher accuracy and reliability may be achieved through increasing the quality of some system components, for example the resolution of the camera/frame-store combination, without having to significantly alter the fundamentals of the system design. Since technology in general, and especially CCD camera and frame-store technology, advances quickly with an associated decrease in cost, it is likely that a structured light system can evolve easily to embrace improved technology.

The processing of structured light patterns, whilst not always straightforward, does not usually present a heavy computational task. Non real-time processing can usually be performed using standard microcomputers with algorithms coded in high level languages. This approach adds to the flexibility of systems and allows the system to be implemented in a platform independent manner.

2.8 Comparison and Discussion of Optical Methods

2.8.1 The Optical Methods Considered

This chapter has described the investigation into some of the key methods for reconstructing human back shape, in some cases through literature review and in others through practical research. The review is by no means comprehensive in that the totality of methods for surface reconstruction would be too extensive to describe and evaluate in appropriate detail. Each of the fundamental methods may have many variants that have not been described: for example, moiré can be used with moving gratings to eliminate noise and primary patterns; phase measuring systems can process more than one phase pattern in the same image; structured light systems can use cross-raster projection patterns. The number of variations are too numerous to mention. In addition, there may be several different approaches to the image processing and surface

reconstruction algorithms. However, the approaches that were found to work to some degree have been discussed.

2.8.2 Comparison of Methods

Before progressing the research programme to an in-depth investigation of new methods, it was necessary to eliminate those basic approaches that were likely to be inappropriate for measurement of back shape. Some of the main considerations in evaluating the methods are:

i. The number of image frames required.

The images that must be acquired by the system is often an indicator of the suitability of the method. If several images must be acquired sequentially then the patient must not move during the acquisition period if good results are to be achieved. If images are acquired simultaneously then this problem does not arise. However, there may be other problems relating to the registration of images. A large number of frames may also translate to large memory requirements in the image processor.

ii. Precision or moving mechanical components.

It is desirable that the number of precision or moving mechanical components in a system should be minimised or eliminated. This requirement has already been discussed in Chapter 1.

iii. Accuracy

The accuracy of any measuring system is always an important factor. However, the primary accuracy objective for an automated system for back shape in scoliosis is that it should match or improve upon the performance of current manual methods. The accuracy of current clinical methods is low due to the limitations of conventional devices and all the methods described in this chapter should be capable of improving upon the accuracy of the conventional approach.

iv. Acquisition time

It is desirable that the acquisition time should be minimised and this has been discussed in Chapter 1. The objective is to minimise inaccuracy due to patient movement and to minimise the inconvenience to the patient.

v. Processing time

The time for processing of images should be minimised so that results are available to the clinician quickly. It is difficult and sometimes inappropriate to compare processing times because they will depend upon the computer/image processor hardware available and the budget for a system. However, it is possible to identify methods that clearly present a heavy computational task in the method.

vi. Reliability.

If a system is ever to be a useful clinical tool then reliability is of paramount importance. The overriding requirement is that surface reconstruction should be achieved regardless of the position or optical properties of the scoliosis patient and their deformity.

vii. Likely approximate cost

Some optical methods require expensive components or image processing hardware, such that their implementation for this application would be undesirable. However, a balance may be required in which other, technical, advantages are considered.

viii. Other shortcomings.

Individual methods may have shortcomings that cannot really be compared with other approaches.

Table 2.2 evaluates these criteria for each of the methods that were considered. Formulating the table is difficult because sometimes it is not really possible to compare “like with like”. However, the table gives some indications of the advantages and disadvantages of the various methods. These have generally been described in more detail in the appropriate section of this Chapter.

2.8.3 An Appropriate Method for Further Investigation

The new method that is proposed in Chapter 3 borrows from the line raster contouring principle. The research described in Chapter 3 was triggered by the limitations of current methods and their implementation by other workers. Each of the methods is discussed below, with attention being drawn to the shortcomings that either restrict or completely eliminate the method in application to back shape measurement.

i. Stereophotogrammetry

Stereophotogrammetry has one major shortcoming that severely challenges automated measurement: corresponding points in the stereo image pair must be identified. Since the human back surface is relatively structureless in intensity terms, it is difficult for even a human observer to identify corresponding points in two images with any accuracy. It would be extremely difficult, if not impossible, for an automated visual system to achieve this task. Thompson (Section 2.2.2) overcame this problem by projecting a pattern onto the object, but in this case the system really becomes a structured light system with two cameras not dissimilar to that proposed by Hill (see Section 2.7.1.2). Stereophotogrammetry was quickly eliminated as being unsuitable.

ii. Moiré

Moiré contouring is an established method for measuring three dimensional shape and has been applied extensively to back shape measurement in manual and semi-automated systems. There are a number of problems, however, with moiré that make automated analysis difficult. First, there is the “directional ambiguity” problem: it is difficult to determine whether successive fringes in the pattern represent increasing or decreasing depth. In manual and semi-automated processing this ambiguity is resolved by the operator or clinician who can intuitively judge the sign of the fringe interval. In other words, there is an input of a-priori information into the processing of the pattern. It is difficult to see how this problem could be overcome in an automated system unless a-priori knowledge of back shape (in scoliotic patients) could be incorporated.

| Method | Criteria |
|-------------------------|--|
| | Number of image frames required (simultaneously) |
| Stereophotogrammetry | |
| Moiré | |
| PSP | 4 or more. |
| FTP | |
| Optical Scanning | (Typically 1 frame per sample) |
| Point raster projection | |
| Bitwise binary encoding | (for 256 grey level encoding) |
| Line raster projection | |
| | Precision/moving mechanical components |
| Stereophotogrammetry | Metric grade cameras desirable. |
| Moiré | No. |
| PSP | Precision movement of grating. |
| FTP | No. |
| Optical Scanning | Precision galvanometer (typically). |
| Point raster projection | No. |
| Bitwise binary encoding | No. |
| Line raster projection | No. |
| | Accuracy |
| Stereophotogrammetry | $\pm 2\text{mm}$ (Thompson). |
| Moiré | Typically several mm. |
| PSP | 0.6mm RMS (Halioua). |
| FTP | Typically similar to PSP. |
| Optical Scanning | $\pm 3.5\text{mm}$ (Turner-Smith). |
| Point raster projection | $\pm 6\text{mm}$ (Hill). |
| Bitwise binary encoding | Not investigated, but $\sim 10^1\text{mm}$. |
| Line raster projection | 0.4mm on test object (Frobin). |
| | Acquisition time |
| Stereophotogrammetry | 1/25s (Frame cycle). |
| Moiré | 1/25s (Frame cycle). |
| PSP | 1s. (Halioua) |
| FTP | 1/25s (Frame cycle). |
| Optical Scanning | 1-2s (Turner-Smith) |
| Point raster projection | 1/25s (Frame cycle). |
| Bitwise binary encoding | Typically $\sim 1\text{s}$ |
| Line raster projection | 1/25s (Frame cycle). |

Table 2.2 Comparison of methods.

| Method | Criteria |
|-------------------------|---|
| | |
| | Processing time/computational task. |
| Stereophotogrammetry | Not investigated. |
| Moiré | Good for basic edge/peak method. |
| PSP | Good (simple image operations) |
| FTP | Poor, 20minutes (see Section 2.5.4). |
| Optical Scanning | Depends upon system configuration. |
| Point raster projection | Average |
| Bitwise binary encoding | Good (simple image operations). |
| Line raster projection | Average. |
| | |
| | Reliability |
| Stereophotogrammetry | Poor. |
| Moiré | Good if directional ambiguity resolved. |
| PSP | Good if unwrapping resolved. |
| FTP | Poor for this application. |
| Optical Scanning | Good. |
| Point raster projection | Average, if points unambiguously located |
| Bitwise binary encoding | Good. |
| Line raster projection | Good. |
| | |
| | Likely approximate cost (UK£) |
| Stereophotogrammetry | |
| Moiré | |
| PSP | |
| FTP | (for satisfactory processing time) |
| Optical Scanning | |
| Point raster projection | |
| Bitwise binary encoding | |
| Line raster projection | |
| | |
| | Major shortcomings |
| Stereophotogrammetry | Difficulty finding points in stereo pair. |
| Moiré | Directional ambiguity problem |
| PSP | Likely phase unwrapping problems |
| FTP | Ideally fringes full field in image, unwrapping |
| Optical Scanning | Complex control of acquisition system |
| Point raster projection | Difficulty locating points. |
| Bitwise binary encoding | Low density of sample data values. |
| Line raster projection | No <i>major</i> shortcoming. |

Table 2.2 Comparison of methods. (Continued)

Secondly, the moiré contour pattern that is selected for processing is simply the dominant component of four separate signals that are present in the image. Separating the signals is probably possible but adds an additional layer of complexity and processing to be considered in a fully automated system. For these reasons moiré contouring was rejected.

iii. Phase-stepping Profilometry (PSP)

PSP is a accurate and usable method for reconstructing back surface shape. In addition to high theoretical accuracy, the density of sampled data values is high because measured values are achieved at every pixel on the object. Unfortunately PSP requires multiple images that must be acquired sequentially with a grating or other projection component translated with precision between each image acquisition. The phase values that are calculated during processing are modulo 2π and must be unwrapped. Phase unwrapping is known, in general, to be a difficult task, but may be possible because the back surface does not usually have severe depth or intensity discontinuities (Halioua, Section 2.5.3.2, must have achieved reliable phase unwrapping). The use of multiple images and precision moving components in the method precludes its useful application here.

iv. Fourier Transform Profilometry (FTP)

FTP is an attractive alternative to PSP that has the advantage that it operates on a single image. Two major problems arise. First, the method requires the computation of the Fast Fourier Transform (FFT) twice in the processing cycle. The heavy computational task of the FFT translated to 20 minutes of processing time on the development machine. Undoubtedly this could be very much improved on more modern or specialised hardware, but it will still present a computational task considerably greater than alternative methods. Secondly, for the technique to work properly, the area for measurement should either be full field in the image or in some rectangular area within it: the processing of an irregular shape such as the human body is difficult. In addition, the phase unwrapping problem that is present in PSP also exists. For these reasons FTP was rejected.

v. Optical Scanning

Whilst optical scanning provides reliable surface reconstruction and has been successfully used for this application, a number of problems arise. First, the scanning process takes some time, up to 2 seconds in the system devised by Turner-Smith (Section 2.6.2.2). During this time it is possible that the patient might move, either accidentally or just in breathing, thus corrupting the resulting reconstruction. Secondly, a scanning system must include moving and usually precision components that, as well as wearing, must be carefully controlled and synchronised with image acquisition. These problems suggest that scanning may not be the most suitable approach to back shape measurement and the method was rejected.

vi. Point Raster Projection.

Point raster projection is a structured light method which might typically be used as a first attempt in a structured light system. The main problem is in reliably indexing points in the image. Since line raster projection appears to partially eliminate the indexing problem this method was rejected.

vii. Bitwise Binary Encoding (BBE)

BBE was not considered in detail in this investigation, partly because it was only towards the end of the investigation that the technology required (LCD panels) became widely available at a reasonable cost. Regardless of this, however, BBE provides only a very limited spatial density of sampled data: 256 points in the image for an 8 bit system. The small number of measured points limits its suitability in an effective system, and therefore it was rejected.

viii. Line Raster Projection

Line raster projection is a low cost solution that avoids precision or moving mechanical components, allows data to be acquired in a single frame cycle, presents a low computational task and yields surface reconstructions reliably with a reasonable sampled data density and accuracy. One implementation of the method was investigated by Frobin and developed to the level of a clinical tool. However, Frobin's method seems complex. The line indexing procedure is based on detecting

lines of different intensity. Given that there are local variations in intensity due to the changing reflective properties across the back surface, it might be expected that problems will arise. Frobin used sub-pixel interpolation to increase the accuracy of the method, although the value of this accuracy (to 0.4mm) for static measurements of back shape must be questionable given that the accuracy with which a patient can be positioned for measurement will be considerably poorer. In addition, the method used by Frobin only yielded one sampled data value for each peak, the notion of detecting fringe maxima and minima or leading and trailing edges (see the new method proposed in Chapter 3) was not applied. Some of the shortcomings of Frobin's method may be due to it being based upon an earlier manual method using photographic film.

2.9 Summary

This chapter has described the research into the fundamental methods of approach that might be used in a successful automated system to measure the surface shape of the back. The majority of methods were investigated through a review of the literature. However, other methods (Fourier transform profilometry and moiré) were investigated through practical research in the laboratory. For each approach, the theory underpinning the basic method was outlined and its application to back shape measurement was explored either by reporting the activity of other workers or by experimental investigation. Finally, the methods were compared and the advantages and disadvantages of each were identified.

No one existing method was completely free from limitations. The method that was both technically feasible for back surface measurement and appeared most likely to satisfy the criteria specified in Chapter 1 was line raster contouring. Although some research had been conducted by other workers, current systems based on line raster contouring had shortcomings. Further research is required to find a new method for reconstructing back shape. Chapter 3 proposes a new clinically useful method for the reconstruction of back surface shape using a method which borrows from the raster line projection principle. However, a fresh approach is taken in both the optical and image processing subsystems.

Chapter 3

A Structured Light Method for Measuring Back Shape

3. A Structured Light Method for Measuring Back Shape

3.1 Introduction

This chapter proposes a new structured light method for the reconstruction of human back shape in scoliosis. Chapter 2 outlined the literature review and research into the key measurement methods that are available for back shape reconstruction and concluded that no single method was completely suitable for measurement of back shape. Clearly there was a need for research to devise a new method that was free from the limitations of the approaches described in Chapter 2 and which satisfied the criteria specified at the end of Chapter 1.

There are two main areas of importance for the proposed method: the optical subsystem used for the acquisition and storage of structured light patterns; and the image processing subsystem used to process the structured light patterns to yield a three-dimensional reconstruction. To some extent, these two subsystems can be considered in isolation. However, it would not be appropriate to consider the optical subsystem without, for example, an appreciation of the image processing algorithms to which the pattern will eventually be submitted. Similarly, the image processing subsystem cannot be explained without a clear understanding of the underlying optical principles that produced the pattern.

The optical subsystem is described with special attention to those issues that are specific to back shape measurement. The problems encountered at the research, design and implementation stages are explained, together with their solutions and discussion of their effectiveness.

The image processing methods are explained first in overview and then by discussing each stage in the processing cycle. Wherever possible the same dataset (the same patient) has been used throughout this chapter so that the reader can see clearly the effects of the various algorithms for the *same* image. Although many of the algorithms used have a sound basis in traditional image processing techniques, the research has

indicated that they should, in almost every case, be modified considerably to suit this specific application and are therefore discussed in detail.

The image processing was implemented entirely in software on a microcomputer using the ANSI C language. The software was written in such a way that it could be easily ported to other hardware and software platforms. The software engineering and structure aspects of the system are addressed in detail in Appendix II.

3.2 Operational Considerations

The criteria for an automated back measurement system outlined in Chapter 1 specified the objectives for a measurement system. Two further areas require special consideration in the system design:

i. Area of measurement

The system should, ideally, be capable of measuring the three dimensional shape of the human back over an area from the seventh cervical vertebra (C7) to the first sacral vertebrae (S1) in the vertical direction and the full width of the back, excluding the arms, in the horizontal direction.

The area for measurement will depend upon the physical dimensions of the patient. Figure 3.1 shows this area diagrammatically. The maximum size of the area of measurement is approximately 500 mm vertically by about 400 mm horizontally. In fact, the 400 mm horizontal area size is larger than would normally be encountered, but allows for some flexibility in the positioning of patients. Since the system was to be designed for, and tested on, children in the age group 4-18 years at the Royal Liverpool Children's Hospital it was particularly important that it should be able to measure areas significantly smaller than the maximum size.




Figure 3.1 Area of the back to be measured.

ii. Ambient light conditions.

Most optical measurement systems are designed to function in circumstances where background light can be carefully controlled. Whilst scoliosis is an important condition that affects many people, most scoliosis clinics, even in large hospitals, are small and the hospital cannot normally provide dedicated space with controlled lighting conditions for an imaging system to reconstruct back shape.

The system must be able to function with background light conditions varying from blackout to normal indoor lighting levels. Whilst every effort can be made to reduce ambient light level, there are practical limitations if the system is to be operated in a hospital environment.

3.3 The Optical Subsystem

The review of optical methods for the measurement of back shape presented in Chapter 2 suggested that, as a basic approach, the line raster contouring method may be capable of producing the reliable and sufficiently accurate measurements that are required for use in a practical measurement system. Furthermore, since satisfactory structured light measurement can be achieved using white light illumination from basic projection equipment, and the resulting images can be easily acquired using standard, low cost

camera/frame-store combinations, it is an economical method for recording surface shape.

3.3.1 Basic Concept of the Optical Subsystem

The principle of the optical system is shown in Figure 3.2. Structured light surface reconstruction methods often rely, either directly or indirectly, on measuring the position of structured light markers on the object and comparing them with the position that a marker would have on some physical or conceptual reference plane. Triangulation based upon these two positions yields three-dimensional coordinates for the structured light markers on the object. The reference plane in the configuration shown in Figure 3.2 forms the xy plane of the measurement system, with z , depth, coordinates being measured as the orthogonal distance from the reference plane to the back surface.




Figure 3.2 Concept of the optical subsystem.

The camera must always be angularly offset from the projector in order for three-dimensional coordinates to be calculated. For the measurement of back shape, it is generally preferable for the camera to be vertically offset from the projector so that shadowing on the back surface is minimised.

The position of the structured light markers on the object, in this case the human back, must always be measured directly. However, two options are available for generating reference information:

- i. By considering the geometry of the optical system, it is possible to calculate the position that a structured light marker would have on a reference plane.
- ii. The position of the structured light on a reference plane may be explicitly identified by projecting the light onto a flat physical plane and recording the resulting image for processing to extract appropriate coordinates.

The second of these options was selected for this research. There are two principal reasons for explicitly recording the reference plane:

- i. Since the human back is, in optical terms, a large object it is not practically possible to illuminate it with parallel light. It is necessary to have a diverging source forming an image on the object. If the mean angle of the source to the mean plane of the back surface is significant then there is a large depth-of-field with a changing magnification of the structured light pattern across the back surface. If reference information is explicitly measured then the changing magnification will manifest in the reference image as well as in the image of the object with a minimum of net effect.
- ii. All lens systems are subject to distortions of various types. Many distortions can be corrected with the use of complex calibration objects and subsequent processing to re-map the acquired image to produce a distortion free image. The calibration process is inconvenient in a practical system and if, as may be the case, it is necessary to acquire complicated calibration data for each measurement, that is, for each scoliosis patient, the additional processing time to perform correction may be significant. However, if the reference plane is explicitly measured then, to a close approximation, any distortion to the reference plane is also manifest in the object image. Again there will be a minimum of net effect.




Figure 3.3 The CCD camera with variable focal length lens.

To ensure the flexibility of the system, a variable focal length lens, shown in Figure 3.3, was selected for imaging the structured light pattern. This allowed straight forward changing of magnification for different physical sizes of patient and allowed the patient, with structured light pattern, to be as close to full field in the image as was practically possible.

3.3.2 Choice of Structured Light Projection Pattern

The predominant challenge in selecting a structured light pattern is to devise a pattern in which individual markers can be unambiguously identified. If the approach is that the structured light pattern on the object is to be matched to the same pattern projected onto a reference plane, then the problem is further compounded because not only must markers be identified relative to each other but they must also be matched to corresponding markers on the reference plane.

A number of options are available and these have been discussed in Chapter 2. The most satisfactory arrangement that was identified in the review of methods was line raster contouring where the structured light pattern consisted of bright and dark lines. This basic principle was adopted, with considerable modification, for use in the proposed method. Because the pattern is visually similar to the fringe patterns encountered in interferometric optics, the lines will be referred to as *fringes* and the line raster pattern as a *fringe pattern*.

Having selected a fringe pattern, a further problem arises. If measurements are to be made by comparing the position of a fringe on the object and on a reference plane, it is necessary to relate a particular fringe in the object (back) image to the corresponding fringe in the reference image. Frobin solved this problem by using a unique pattern of narrow lines of slightly varying width with the slightly broader lines translating to high peaks in the intensity profile that helped distinguish individual lines (see Chapter 2). This pattern worked satisfactorily but added complexity. A much simpler alternative was used in this investigation.




Figure 3.4 Back and reference images showing reference fringe.

The solution that was employed was simply to have a fringe pattern in which one dark fringe has double the width of the others. This fringe, the reference fringe, then acts as a reference in both the object and reference images and all other fringes are indexed relative to it. Figure 3.4 shows both a back image and a reference plane recorded in this way. In subsequent image processing, the fringe below the reference fringe are assigned index values -1, -2, -3 ... and fringes above it have indices +1, +2, +3 ... and so on. In this way, the matching of fringes in the reference and back images was achieved.

A grating was prepared by reducing a black and white ruling on A4-size paper onto a 35mm lithographic slide. With the exception of the broad reference line, the effective

ruling on the grating was approximately 1 line/mm. Diffraction grating effects are not significant for this line spacing.

3.3.3 Measurement Procedure

The scoliosis patient stands on a foot-plate in a vertical frame as shown in Figure 3.5. The foot-plate was designed to place patients in a standardised position (see Chapter 4). The vertical frame is constructed from tubular iron and can be assembled or collapsed by simply inserting and removing winged nuts that attach the triangular supports. A standard 35 mm slide projector is mounted on a tripod, at a distance D_c from the frame. The projector is loaded with the grating and directed at the surface of the back. A fringe pattern such as that shown in Figure 3.4 is generated on the back. The broad reference fringe is positioned in approximately the centre of the patient's back. The fringe pattern is collected at the camera, also mounted on the tripod, at the same distance from the frame and a distance D_p vertically below the projector. The camera and projector are mounted with the optical centres of their lenses aligned vertically. The image acquired by the camera is transferred via a frame grabber into the memory of the computer. A matt black curtain is attached to the frame behind the patient to provide a dark background for the image.

The patient is removed from the image area and replaced by a flat, vertical, diffusely reflecting white screen which forms the physical reference plane from which depth coordinates are measured. The flat white screen is physically constructed as a gate that is open when a back measurement is being acquired and is closed over for acquisition of the reference plane. The camera and projector remain in the same fixed positions. A second image, of the grating projected onto the flat white plane, is acquired and stored in the computer.

The two images of fringe patterns acquired in this way contain the key data from which the three-dimensional surface is reconstructed. However, since measurements are required in absolute units (mm) rather than units based on pixels, an additional image is acquired that is used to record the information necessary to calculate the dimensions of pixels. The user places a reference object (see Section 3.3.6) in the image which is used

to calculate the effective pixel size. The resulting image is referred to here as the scaling image. The user also enters the distances D_c and D_p into the computer. Typical values of D_c and D_p are 2300mm and 1000mm respectively, giving a value of 25° for θ .




Figure 3.5 Patient positioned in frame for measurement.

Thus three images are acquired. It should be emphasised that, in principle, only one image, the image of the back with fringe pattern, is absolutely necessary for each patient in a clinical session and the system can be configured to only require the acquisition of reference and scaling images once during the session. If this is the case then the geometry of the system (the values D_c and D_p and the position of the reference fringe on the object) and the lens adjustments (the focal length of the camera and projector lenses) must remain fixed across measurements. In practice, however, the reference and scaling images must usually be collected for each patient measurement because the system geometry and focal length of the camera lens are usually adjusted for different physical sizes of patients.



Figure 3.6 Patient standing in frame.

Having collected the source data the clinician then initiates the processing of the images to reconstruct surface shape in the computer. Figure 3.6 shows a patient being measured in the frame.

3.3.4 Coordinate System

The coordinate system used for images is shown in Figure 3.7, with the image in both portrait and landscape orientation.

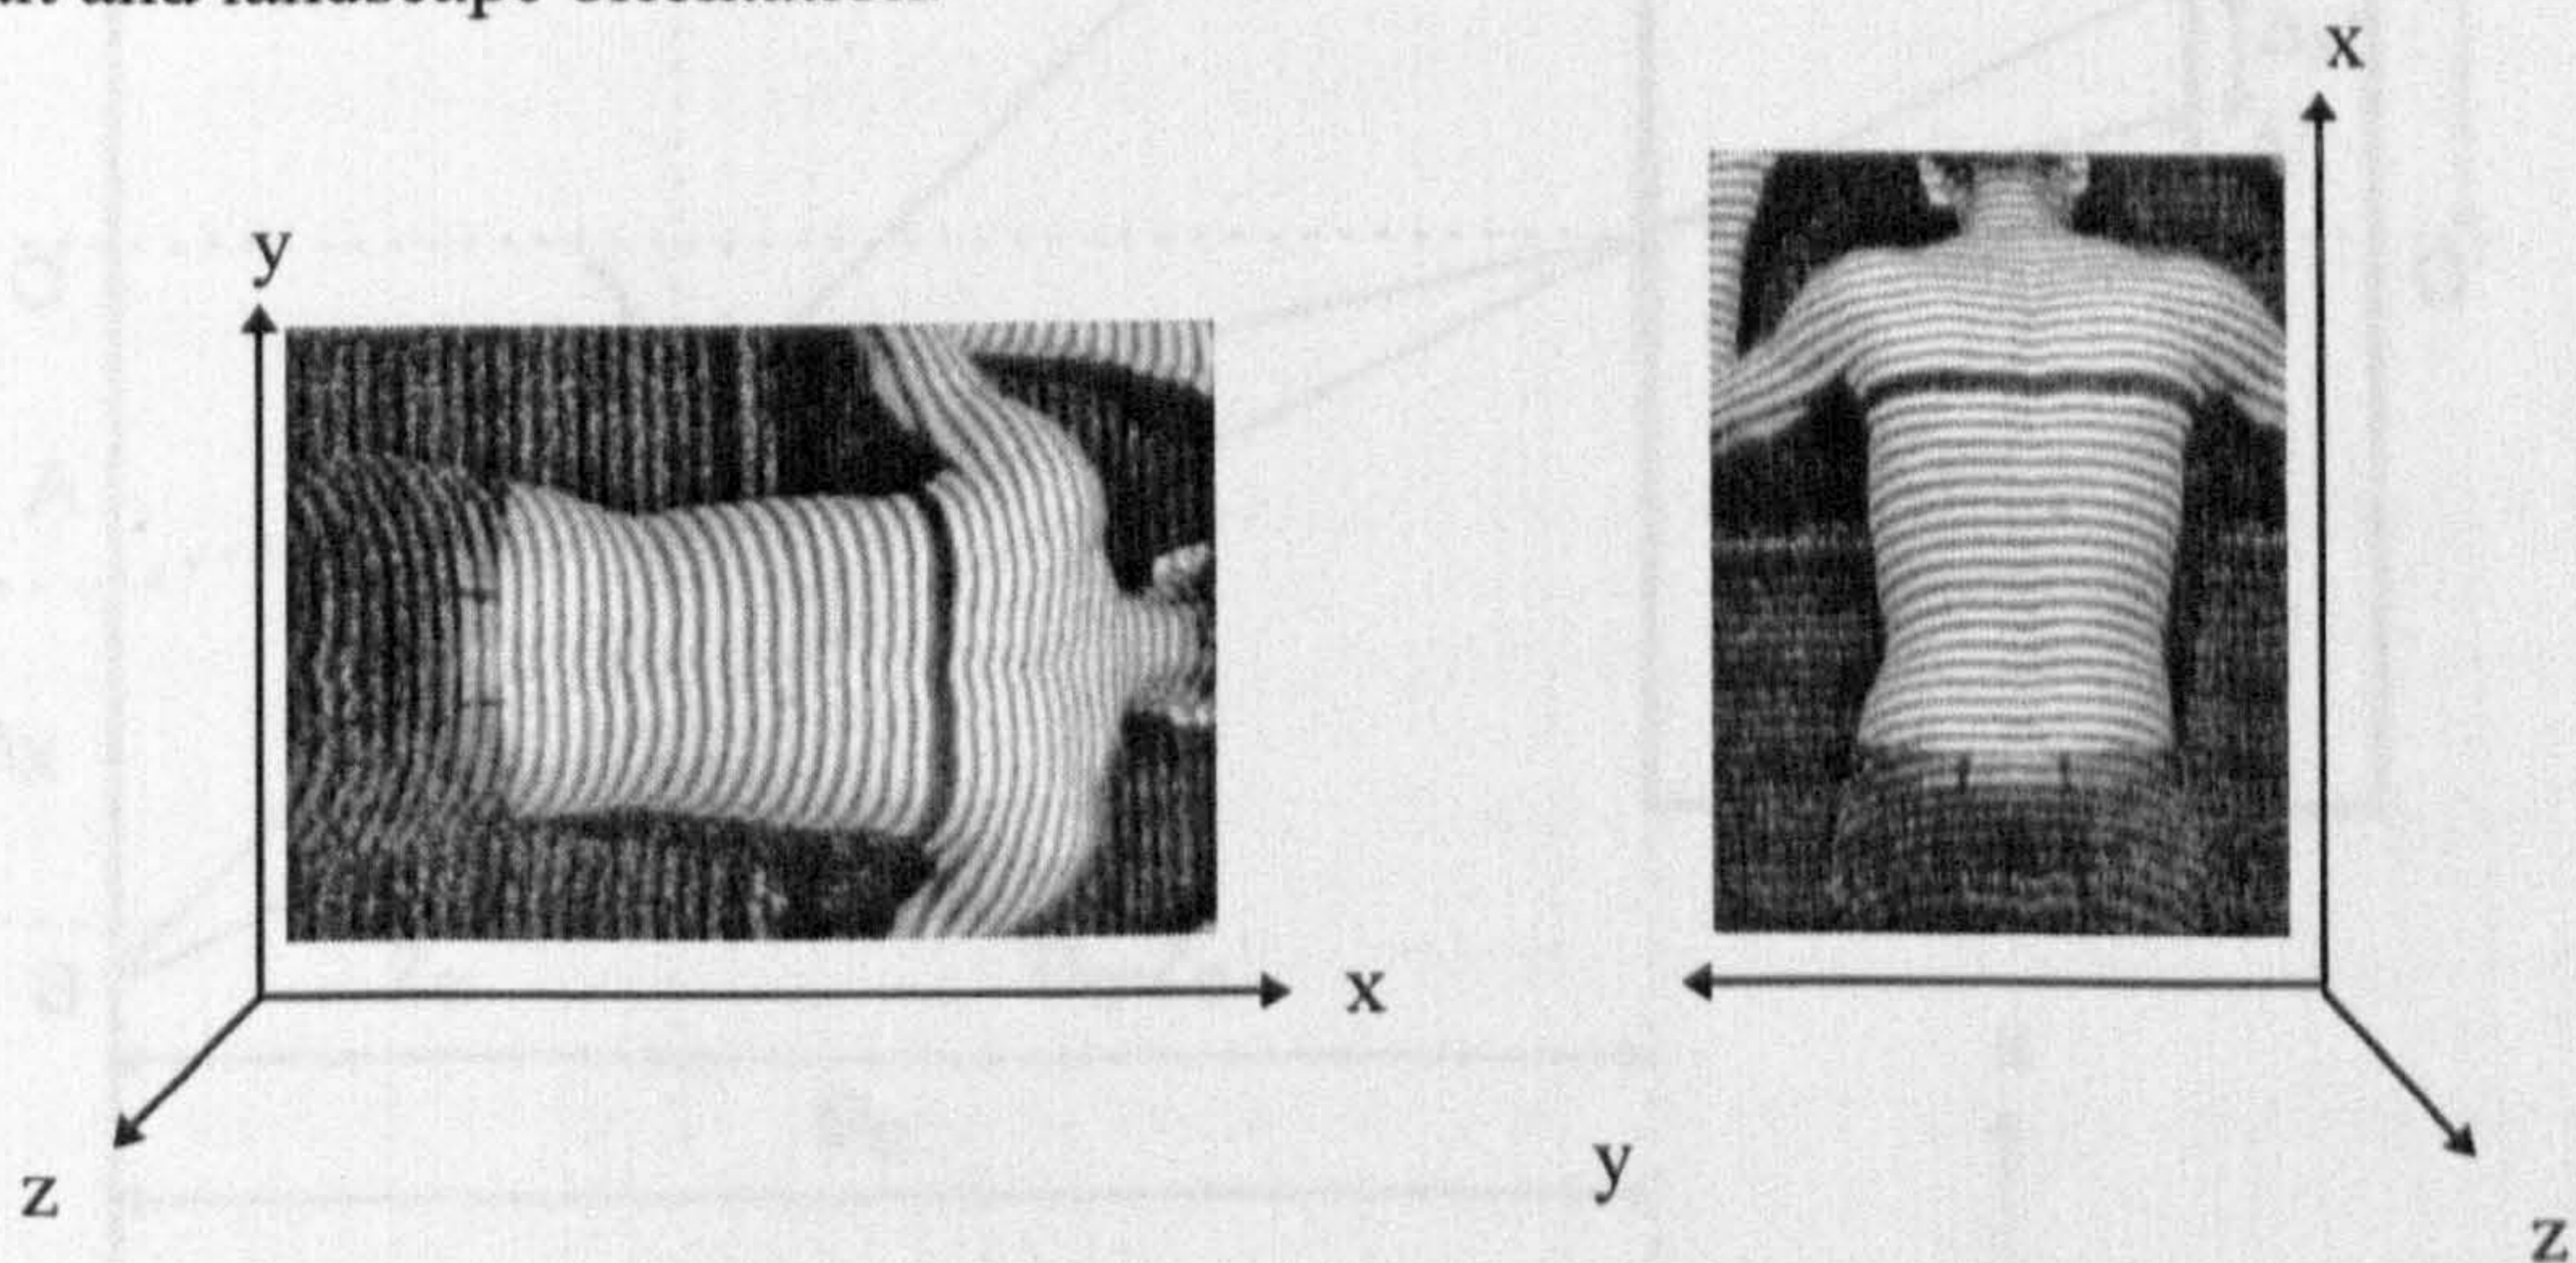


Figure 3.7 The coordinate system for images.

The choice for the coordinate system arose originally out of the desire for an image, as displayed on the monitor of the development system, to have an origin in the bottom left hand corner of the screen with intensity (for images) or depth, z , (for surface

reconstructions) to be directed out of the screen (see Section 3.4.10 for a description of the representation of the surface reconstruction).

3.3.5 Optical Geometry

Consider the optical geometry shown in Figure 3.8. A fringe from the projector at P strikes an element of the back surface at S. The pixel (x,y) position of the fringe as it falls on the back surface is recorded in the image plane of the camera at A'. The xy plane is the plane of the image and the z axis points out of the image towards the camera. The position of the fringe as it falls on the reference plane is recorded in the image plane of the camera at B'. The distance between the two fringe locations A' and B' is measured in terms of pixels to be $\Delta x'$. Using scaling information, the distance $\Delta x'$ in pixels is translated to absolute distance Δx at the reference plane.




Figure 3.8 Optical geometry.

Consideration of the similar triangles ABS and CPS gives that the z coordinate, z_s , of the surface element at S is given by

$$\frac{z_s}{\Delta x} = \frac{D_c - z_s}{D_p} \quad \text{[Equation 3.1]}$$

so that

$$z_s = \frac{D_c \Delta x}{D_p + \Delta x} \quad \text{[Equation 3.2]}$$

The z (depth) coordinate of each point lying on a fringe in the image of the back surface is thus calculated. Since D_c and D_p are known parameters of the geometry, the problem of calculating z coordinates, the main reconstruction problem, reduces to calculation of Δx . Δx is simply the vertical distance between structured light markers as they fall upon the back and upon the reference plane.

The complete surface is a set of points like S. Ultimately, the sampled data points are the pixels that lie on fringe edges (this is discussed later in this chapter). The image processing algorithms identify corresponding points in the back and reference images and apply equation 3.2 to reconstruct the surface.

3.3.6 Measurement of Pixel Size

In order to evaluate Equation 3.2 and to calculate the z coordinate for all points S on the back surface, the value of Δx in real world coordinates must be calculated from $\Delta x'$ in pixel coordinates. The relationship between Δx and $\Delta x'$ must therefore be established. It is assumed that a linear scaling factor L_x , related to magnification, which converts pixel coordinates in the x direction to coordinates in mm can be applied to $\Delta x'$ to calculate Δx using:

$$\Delta x = L_x \Delta x' \quad \text{[Equation 3.3]}$$

To calculate this scaling factor an object of known physical dimensions must be placed in the image and measured in terms of pixels by image processing algorithms. In

addition to the estimation of L_x , a similar scaling factor L_y in the y direction is also required for later calculations. L_y is not required for the calculation of z coordinates, but is used later for the calculation of aspect ratio and the parameters of clinical significance that are described in Chapter 4.

To be satisfactory in a practical clinical system, the reference object must be easily and quickly positioned. A dark, solid, circle of known radius, R , placed in the image against the white background provides a useful scaling object. The circle has the unique property of having the same dimensions in all directions so the need for precision orientation and positioning of the object in the image is largely eliminated.




Figure 3.9 Dark circle used to calculate pixel scaling factors.

In fact, the dark circle was provided by a standard 5¼" or 8" floppy diskette (obviously removed from its square casing). Not only is a floppy diskette dark and precision engineered to be circular, but it is also easily replaced if damaged or lost in a clinical environment. It was envisaged that a more "professional" dark reference object would be produced if ever the system were produced commercially! The circle is placed in approximately the centre of the image against the white reference plane and illuminated with light from the projector in the absence of a fringe pattern. The boundaries of the circle, x_1 , y_1 , x_2 and y_2 , shown in Figure 3.9, are detected using image processing algorithms. The scaling factors L_x and L_y are then calculated to be:

$$L_x = \frac{D}{x_2 - x_1} \quad \text{[Equation 3.4]}$$

$$L_y = \frac{D}{y_2 - y_1} \quad \text{[Equation 3.5]}$$

Using the value of L_x calculated in this way, Equation 3.2 can be restated in terms of pixel coordinates in the x direction to be:

$$z_s = \frac{D_c L_x \Delta x'}{D_x + L_x \Delta x'} \quad \text{[Equation 3.6]}$$

Equation 3.6 allows the calculation of z coordinates, and therefore surface reconstruction from the shift in fringe position between the back image and reference plane as measured in pixels at the camera. The principal objective of the subsequent image processing is to reliably measure the shift in fringe position, Δx , of fringes as they fall on the back and reference surfaces.

3.4 Image Processing

It has been established that reconstruction of the back surface can be achieved by measuring the shift in the position of a set of fringes as they lie on the surface of the back and on the reference plane respectively. This section discusses the methods of identifying the coordinates of fringes in the image, measuring the shift in position of the fringes and finally shows how the shift information is translated into a surface reconstruction.

An initial consideration based on the optical geometry outlined in the previous section might suggest that the detection of fringes and the ultimate reconstruction is a straight forward task. However, a number of problems arise. Besides the usual challenges of image processing algorithms, there are three special problems:

- i. The object to be measured, the back, is not full field in the image. Furthermore the object is not rectangular and therefore does not lend well to many standard image processing techniques. Many image processing algorithms need at least the *a priori*

knowledge that the object that they are designed to process does actually exist in the area of the image that they are currently processing.

- ii. Since the imaging system must perform in a variety of ambient light conditions, the contrast and mean intensity in the image may vary significantly according to the circumstances that prevail when images are acquired. In the laboratory it is possible to achieve near photographic quality blackout whilst in a clinical environment it is difficult to do more than turn off additional artificial lights. As far as possible then, the image processing algorithms should be invariant of background light and fringe visibility.
- iii. There is a significant variation in the local mean intensity across the surface. This is due to changes in the gradient of the surface and local changes in the reflective properties of skin, for example due to birth marks or scar tissue.

The following sections of this chapter outline the research into image processing algorithms to reconstruct the shape of the back from the images acquired by the optical subsystem. During the investigation each algorithm was tested and refined on a number of images of varying quality, including images acquired from patients at the Royal Liverpool Children's Hospital. Towards the end of the research programme clinical trials occasionally revealed errors in the processing of images and the algorithms were adjusted accordingly. A single dataset is followed through all the algorithms so that the results of each step in the processing cycle can be clearly seen. Although most of the earlier processing applies separately to both the reference and back images, only the back image is considered in detail because, in comparison, the processing of the flat, white, diffusely reflecting, reference plane is trivial.

A summary of the methods is presented first, followed by a detailed explanation of each step. The processing system was implemented entirely in software using the C programming language. Wherever possible the methods are described at algorithmic level. However, on some occasions it has been necessary to describe some data structures using language constructs.

3.4.1 Outline of the Image Processing Subsystem

The sequence of image processing algorithms is shown in Figure 3.10. The first task for the image processing system is to locate the coordinates of fringes. Ultimately, each fringe will be represented as a list of (x,y) coordinates which identify the set of points that represent the fringe. Ideally, for subsequent processing to be successful each image pixel should be joined to an adjacent pixel for the entire length of the fringe.

The preliminary image processing tasks can, and must, be performed on the entire image, and part of their activity is to separate those pixels that are clearly part of the back surface in the image from those pixels that are clearly the background. Later, each pixel in the image is assigned a status value that indicates whether it is a sampled data point, a background pixel or otherwise.

The back and reference images are processed to locate fringes, or more strictly the edges of fringes, and to form a binary image in which continuous lines of pixels define the location of fringes. This binary image is then the fundamental dataset from which measurements of Δx are later taken. The same algorithm is applied to both the image of the back and the reference image.

The next task is to index the fringes relative to the reference fringe. Each pixel lying on a fringe is assigned an index value that is unique to that fringe. The importance of maintaining the connectivity of the list of pixels is paramount to the indexing process because breaks in a fringe may cause the tracking of pixels and the assignment of index values to fail. In a similar manner to fringe location, the indexing algorithm is applied to both the back and reference images

The area of interest, the back, is then located in the image with a continuous boundary that encloses the area. An additional status value is used that marks a pixel as lying on the boundary of the area of interest. All subsequent operations consider only pixels that lie inside the boundary.

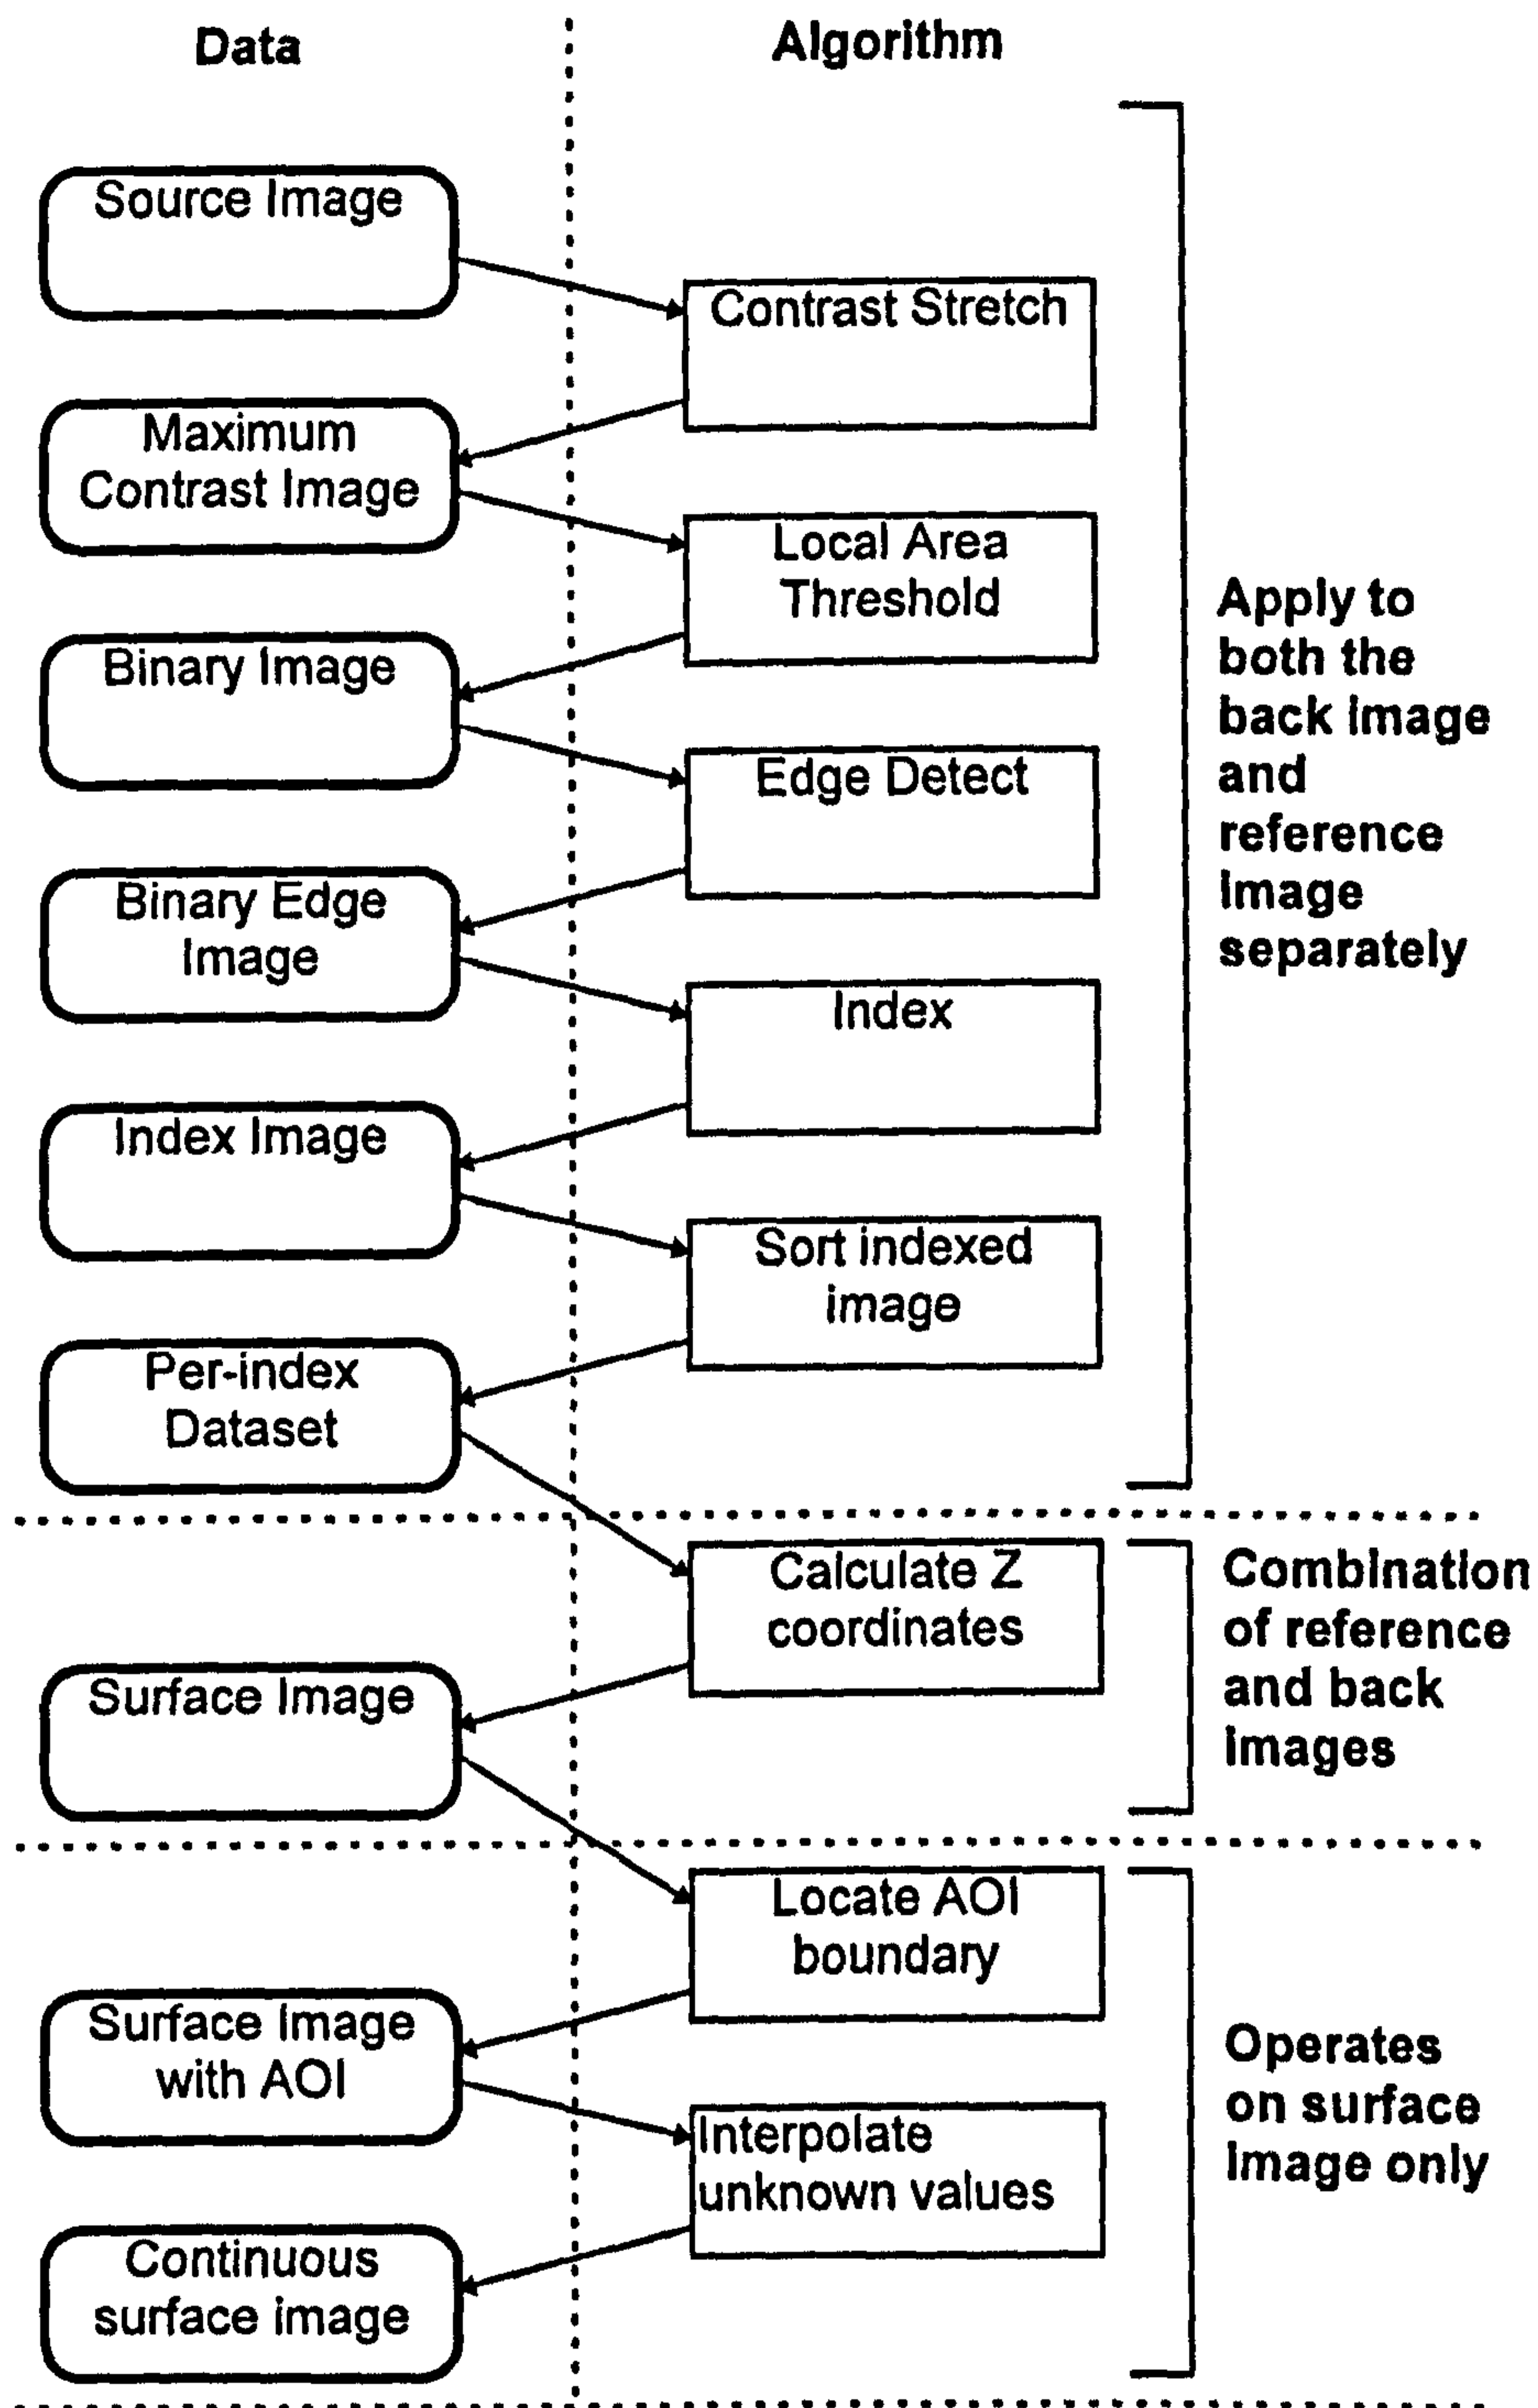


Figure 3.10 The sequence of image processing operations.

The calculation of Δx is fundamental to the measurement procedure. An image processing algorithm takes each point lying on the edge of a fringe on the back surface and searches for the corresponding point in the reference image. The search criteria is that the corresponding pixel in the reference image will lie on the same vertical line (that is, have the same y value) and have the same index number as the pixel in the back

image. In an ideal configuration, with fringes that are perfectly horizontal, it would not be necessary to specify that pixel pairs in the back and reference images must lie on the same vertical line because all pixels lying on a perfectly horizontal fringe in the reference image will have the same x coordinate. However, in practice, fringes are seldom perfectly horizontal and the location of the fringe vertically must be determined for each back/reference image pixel pair. In fact, this is an additional reason for explicitly measuring the reference plane. Clearly it is at this point in the processing of images that the back image and reference image are combined to form a new dataset of (x,y,z) coordinates for the surface.

The measurement of Δx values for all fringes yields a data set in which each point lying on a fringe edge in the back image is assigned a depth, z, coordinate. The boundary location algorithm defines the area of interest in the image. The completion of these two tasks marks the end of the actual measurement procedure. However, since the reconstruction will be later interrogated to yield clinically significant parameters that relate to scoliotic deformity, the depth coordinates of pixels lying between fringes in the image will be required. The z coordinates of these pixels are interpolated from local, known values.


3.4.2 Source Image Data

The acquired fringe pattern is digitised with a resolution of 512(x) x 256(y) pixels to 6 bits giving a range of 64 grey levels in the range 0(dark) to 63(bright). The six bits digitisation of intensity is unusual in modern grey-scale image acquisition hardware where the standard is to use 8-bit byte values and was a consequence of the desire to minimise the cost of the overall system. The software, however, is designed to cope with 8-bit images and could easily be modified to accommodate an even greater number of bits.


The projected grating has a regular rectangular, or Ronchi, profile with the exception of the dark reference fringe in approximately the centre. One might expect that the intensity profile of the light as it falls on the object and as it is recorded at the camera would have a similar, rectangular profile with a clear demarcation between bright and

dark fringes. In practice this is not the case. The intensity profile through a series of fringes more closely resembles a sinusoid. Figure 3.11 shows the intensity profile through an image with a detailed profile taken through fringes in the centre of the image. The sinusoidal nature of the profile can be clearly seen. There are several factors which may contribute to the tendency of the profile towards a sinusoid:


- i. The grating was produced using a line pattern photographed onto lithographic film. Whilst subjective inspection of the film using a travelling microscope suggested that the bright/dark transitions are good, there may be loss of quality due to slight errors in focusing which cannot be determined by the eye. A detailed analysis of the profile of the grating could easily be achieved by examining the projection grating using a microdensitometer. However, this was considered to be unnecessary because of the other overwhelming factors described here.
- ii. There is a depth-of-field problem inherent in the optical method. The angle subtended by the incident projected light to the back surface is typically 25° so that the depth of field is significant for the projection system as shown in Figure 3.12. Much of the fringe pattern will be at least slightly out of focus. In addition there is a depth of field associated with the camera system due to the depth of the patient in a direction parallel to the optic axis of the camera, but this is small in comparison to the depth of field for the projector.
- iii. No lens/aperture system can pass all the spatial frequencies required to produce a perfect Ronchi profile. The highest frequencies are necessary to reconstruct the edges of a rectangular profile and it is these frequencies that are limited by the lens system. Furthermore, there are two lens systems (the projector lens and the imaging lens) which will contribute to this effect.



(a) Source image



(b) Intensity profile



(c) Intensity profile for pixels 200-300

Figure 3.11 Intensity profile through an image.

iv. There will be diffraction effects at the edges of the bright/dark transition that will provide secondary maxima. If each fringe is considered as a single slit then there will be intensity contributions to adjacent dark areas. Note that this is a different phenomenon to the consideration of the grating as a diffraction grating, where contributions from separate slits, (separate fringes) are considered.




Figure 3.12 Depth-of-field for camera and projector.

- v. Regardless of the other effects, when fringe edges are imaged by the camera lens onto the CCD array, it is unlikely that even a close to perfect edge would lie on the exact boundary of adjacent array pixels. Typically, an edge will span at least two pixels, leading to intermediate values for intensity.
- vi. The pixels of the CCD array in the camera do not necessarily correspond to pixels in the digitised image. The CCD camera produces an analogue video signal that is then digitised at the computer. In a similar manner to that described for v. above, the new pixel may have intensity components from more than one CCD pixel, again leading to an intermediate value for intensity.

3.4.3 Contrast Stretching

Prior to the main image processing tasks, a simple contrast stretch is applied to the image. This means that after the contrast stretch the minimum intensity in the image will be zero and the maximum intensity will be one less than the number of grey levels in the image (the development system used six-bit images with 64 grey levels). The contrast stretch can be expressed as:

$$I_d(x, y) = \frac{(n_g - 1)(I_s(x, y) - I_{s_{min}})}{(I_{s_{max}} - I_{s_{min}})} \quad \text{[Equation 3.7]}$$

where n_g is the number of grey scales in the image, $I_s(x, y)$ is the intensity value of a pixel in the source image, $I_{s_{min}}$ is the minimum intensity in the source image, $I_{s_{max}}$ is the maximum intensity in the source image and $I_d(x, y)$ is the resulting intensity in the destination, contrast stretched, image.

It might at first seem that there is little value in performing this operation because it simply shifts and scales the image intensity values. However, once the contrast stretch has been performed some advantages accrue:

i. Standardised intensity range

The image, in as far as is possible, is standardised in terms of intensity. This means that it is later possible to apply algorithms that will produce similar results relatively independent of ambient light conditions.

ii. Better use of contrast range

Although the contrast stretch is simply a shift and scale operation, it will allow subsequent processing to make greater use of the intensity range. Suppose, for example, that a smoothing is to be performed by convolution of an averaging filter. The wider range of intensity values will allow for greater precision in the destination smoothed image. (A smoothing module was incorporated in system in the early phase of the development but was removed because the value of the smoothing

operation appeared to produce only a minimal improvement in the quality of the eventual surface reconstruction and presented a significant computational task).

iii. Improved visibility

The visibility of the image is improved. This is useful for clinical inspection (the image is used again later for, for example, vertebral positioning) and in testing and development.

This simple shift and scale operation is the most satisfactory way of performing a contrast improvement. Many other standard operations for contrast enhancement, for example histogram equalisation, can not be applied because they would alter the profile of the fringe intensity and the signal would be corrupted (in phase terms these operations are not phase preserving).

3.4.4 Fringe Location using Thresholding and Edge Detection

The fringe location problem is one of finding a single pixel position on a fringe that can be reliably chosen to represent the position of the fringe along a particular vertical line in the image. Common approaches to the problem include the detection of intensity maxima and minima or the detection of edges through differentiation or high-pass filtering. The direct detection of maxima and minima is known to produce problems when there is noise in the image.

The detection of edges is generally more reliable than maxima or minima detection but relies on the existence of a particular gradient over a specified number of pixels. If the fringes, as detected in the image, have varying contrast ratio and spacing then the detection of edges becomes more difficult. Again, any noise in the image will reduce the reliability of detection. Typically, these types of feature detection will yield two locations for each fringe: for example, a maxima and minima pair or leading and trailing edges.

The approach that was eventually found to be successful in this investigation was to first apply a threshold to the image to produce a binary image. The leading and trailing

edges of the binary image can then be used as the definitive locations of fringes in both the back and reference images.

3.4.5 Thresholding

The principle of binary thresholding in image processing terms is used to convert a grey-scale image of n_g grey scales into a binary image which, by definition, has two grey scales. The grey scales are normally represented as 0 (dark) and 1 (bright). The algorithm is simple. If the intensity of a pixel is greater than or equal to some threshold value then that pixel is defined to bright (1), otherwise it is dark (0). For global thresholding, that is where the same threshold is applied across the image, the operation can be expressed as:


$$I_d(x, y) = \left. \begin{array}{l} \text{if } I_s(x, y) \geq T_g \text{ then} \\ 1 \\ \text{else} \\ 0 \end{array} \right\} \quad \text{[Equation 3.8]}$$

where I_s is the source image, I_d is the destination image and T_g is the threshold value. Typically for this type of application, the value of T_g would be the mean intensity in the image or alternatively the mean of the intensity maximum and minimum.

The limitation of global thresholding is that the same threshold must apply across the whole of the image. For this application, global thresholding cannot be applied because the mean local intensity varies across the image. Consider the image set shown in Figure 3.13. Whilst the global threshold is satisfactory in some areas of the image, it is too low in others. In fact, it is impossible to select a global threshold value that can be used *reliably* for all images.

Clearly, a global threshold would not work. Furthermore, even if the failure of the algorithm as demonstrated in Figure 3.13 did not occur, a separate problem arises. It is necessary that the image processing system should identify the edge of a fringe, in other words the edge of a binary bright area, based on the indirect measurement of some phase value in the approximately sinusoidally varying intensity profile. Which value is

chosen is largely unimportant provided it is consistent across the image and consistent with measurements taken in both the back and reference images. Given that the mean intensity will vary across the surface, this could not be achieved using a single global threshold.



(a) Source image

(b) Global threshold, T_g

(c) Desired output (see later)

$$T_g = \frac{I_{s_{\max}} + I_{s_{\min}}}{2}$$

Figure 3.13 Typical failure of global thresholding.

There is, therefore, a need to adapt the threshold to account for local variations in intensity. This concept is local area thresholding. Equation 3.8 is revised to account for per-pixel changes in local mean intensity so that the thresholding principle is defined by Equation 3.9. T , the threshold, now becomes a function of pixel position (x,y) .

$$I_d(x,y) = \left. \begin{array}{l} \text{if } I_s(x,y) \geq T(x,y) \text{ then} \\ 1 \\ \text{else} \\ 0 \end{array} \right\} \text{Equation 3.9]$$

For this application, local area thresholding also presents a problem. Recall that only those areas of the image in which true fringes on the surface of the back exist should be considered for thresholding. If the threshold is determined by some local maximum and

minimum then pixels lying outside the back surface, in the background, will be determined to be bright owing to small differences in background intensity.

Other challenges to local area thresholding also exist. One method for determining a local threshold value is to divide the image into fixed rectangular units. A threshold is selected for each of the local units then becomes the local threshold for all pixels in that unit. If the threshold were selected to be the mean of the maximum and minimum values in the unit then the operation could be expressed as:

$$I_d(x, y) = \left. \begin{array}{l} \text{if } I_s(x, y) \geq T(x, y) \text{ then} \\ 1 \\ \text{else} \\ 0 \end{array} \right\}, \quad \text{[Equation 3.10]}$$

$$\text{where } T(x, y)_{x_1 < x < x_2, y_1 < y < y_2} = \frac{I_{s_{\max}} - I_{s_{\min}}}{2}$$

and x_1 and x_2 define the boundaries of the unit in the x direction and y_1 and y_2 define the boundaries of the unit in the y direction. $I_{s_{\max}}$ and $I_{s_{\min}}$ represent the local maximum and minimum values. The problem with this approach is that at the edges of a rectangular unit there will be sudden change in threshold value that would, for this application, be unacceptable.


The solution that was initially found to work for the reliable detection of fringes was:

- i. For *each pixel* in the source image, define a square local region with dimension m . In fact it would be acceptable to define the region for a small group of pixels, provided that the number of pixels in the group is small relative to the size of the local region.
- ii. Determine the maximum and minimum intensity values in the $m \times m$ local area, $I_{s_{\max}}$ and $I_{s_{\min}}$ respectively.
- iii. If the range of intensity values in the local area, $r = I_{s_{\max}} - I_{s_{\min}}$ is less than some threshold value T , then set the destination pixel to be dark (binary 0).


iv. Set the local threshold, $T(x,y)$, to be the mean of $I_{s_{max}}$ and $I_{s_{min}}$.

This series of operations can be represented as:

$$I_d(x,y) = \left\{ \begin{array}{l} r = \left(I_{s_{max}} - I_{s_{min}} \right)_{x-\frac{m}{2} < x < x+\frac{m}{2}, y-\frac{m}{2} < y < y+\frac{m}{2}} \\ T(x,y) = \frac{r}{2} \\ \text{if } r < T_r \text{ then} \\ 0 \\ \text{else} \\ \quad \text{if } I_s(x,y) \geq T(x,y) \text{ then} \\ \quad 1 \\ \quad \text{else} \\ \quad 0 \end{array} \right. \quad \text{[Equation 3.11]}$$



(a) Source Image



(b) Thresholded Image

Figure 3.14 Thresholding based on a 40 x 40 rectangular area.

Figure 3.14 shows the output of this thresholding algorithm. So, the composite thresholding mechanism used two different threshold values. The first threshold value, T_r , was used to determine whether the pixel is part of the object or part of the dark

background. The second threshold, $T(x,y)$, is the local area threshold based on local maximum and minimum values.

This type of thresholding algorithm, however, presents a significant computational task. For each pixel in the source image a total of m^2 integer comparisons must be made to calculate the local maximum and minimum values. For an image of x dimension D_x and y dimension D_y , the total number of comparisons, n_c required to process an image will be:

$$n_c = D_x D_y m^2 \quad \text{[Equation 3.12]}$$

Consider the Table 3.1 which shows values of m^2 and n_c for a 512 x 256 image:

| m | Number of Comparisons per Pixel, m^2 | Number of Comparisons per Image, n_c |
|----------|--|--|
| 10 | 1×10^2 | 1.3×10^7 |
| 20 | 4×10^2 | 5.2×10^7 |
| 40* | 1.6×10^3 | 2.1×10^8 |
| 100 | 1.0×10^4 | 1.3×10^9 |

* Eventually chosen to be the most suitable value for m


Table 3.1 Computational task for a square local area threshold.

Using a square local area threshold presented a computational task which took several minutes (typically 2.5 minutes) on the hardware used for the development system and represented a major speed bottleneck in the processing cycle.

An alternative, but similar method was devised that makes use of the fact that fringes in the image are approximately horizontal. Recall that the principal criterion for the local area threshold is that the local area must enclose at least one bright and one dark fringe. This criterion can be satisfied for approximately horizontal fringes by simply looking at a linear local area of m pixels in length and one pixel wide; in other words the local area thresholding algorithm can be defined as for the square local area by:

$$I_d(x, y) = \left. \begin{array}{l} r = \left(I_{s_{\max}} - I_{s_{\min}} \right)_{x - \frac{m}{2} < x < x + \frac{m}{2}, y \text{ constant}} \\ T(x, y) = \frac{r}{2} \\ \text{if } r < T_r \text{, then} \\ 0 \\ \text{else} \\ \text{if } I_s(x, y) \geq T(x, y) \text{ then} \\ 1 \\ \text{else} \\ 0 \end{array} \right\}, \quad \text{[Equation 3.13]}$$

So that $T(x, y)$ is calculated over a linear local area.



(a) Rectangular Area

(b) Linear Area

Figure 3.15 Comparison of local area threshold algorithms ($m=40$).

The threshold is thus calculated from local values along a vertical line in the image. Figure 3.15 shows images processed using the same local area size for both linear and square threshold areas.

The degradation in image quality achieved using a linear rather than rectangular (square) local image area was minimal and was therefore chosen for use in the system.

The processing time for the thresholding algorithm was reduced to a few seconds as might be expected from the data shown in Table 3.2.

| m | n_c for Square Local Area | n_c for Linear Local Area |
|----------|---|---|
| 10 | 1.3×10^7 | 1.3×10^6 |
| 20 | 5.2×10^7 | 2.6×10^6 |
| 40* | 2.1×10^8 | 5.4×10^6 |
| 100 | 1.3×10^9 | 1.3×10^7 |

* Eventually chosen to be the most suitable value for m

Table 3.2 Computational task for square versus linear local area thresholding.

Clearly, the linear local area is an effective solution when the degradation in thresholded image quality is minimal. The relationship between n_c for the square and linear local areas is:

$$n_{c_{\text{linear}}} = \frac{1}{m} n_{c_{\text{square}}} \quad \text{[Equation 3.14]}$$

So, for an area of size $m=40$ pixels, the linear area presents a computational task that is $1/40^{\text{th}}$ of the corresponding value using a square local area. The measured processing times associated with the images shown in Figure 3.15 are shown in Table 3.3.

| Thresholding Algorithm (m=40) | n_c | n_c Corrected* | Execution Time (seconds) |
|--------------------------------------|-------------------------|------------------------------------|---------------------------------|
| Rectangular Local Area | 2.1×10^8 | 1.2×10^8 | 157 |
| Linear Local Area | 5.4×10^6 | 4.8×10^6 | 6 |


* Corrected for limited processing near the edges of the images.

Table 3.3 Approximate number of comparisons and execution time for thresholding.

The results shown in Table 3.3 also verify that the number of comparisons required are a suitable measure of computational task, because the ratios of corrected n_c values $1.2 \times 10^8 / 4.8 \times 10^6 = 25.0$ and execution times $157 / 6 = 26.1$ are in good agreement. The linear local area threshold algorithm was selected for continued use.

There are two parameters that must be optimised for a linear local area thresholding algorithm. These are T_r , the range threshold, and m , the size of the local area. The range threshold, T_r , in Equation 3.13, needs to be adjusted carefully. A value of T_r that is

too low will cause background areas in the image to be marked as bright which is clearly undesirable. The correct value of T_r will mark bright pixels lying on fringes to be bright, dark fringes to be dark, and the background to be dark. A value of T_r that is too high will cause loss of valid bright fringe data, particularly at the edges of the back image. The effect on the thresholded image of changes to value of T_r is shown in Figure 3.16.



(a) Too Low ($T_r=10$) (b) Correct ($T_r=15$) (c) Too High ($T_r=25$)

Figure 3.16 Effects of altering the range threshold (for $m=40$, $n_g=64$).

Optimisation of the m parameter is less critical. However, there will still be a degradation of thresholded image quality for extreme values of m . The main criterion for the value of m is that it should be at least large enough to ensure that, when processing the back surface area, the local area encompasses at least one bright/dark fringe pair.

However, a value of m that is too large will no longer reflect the local area objective of the algorithm: areas become so big that maximum and minimum values remote from the pixel of interest will influence the threshold value. Figure 3.17 shows the effect of changing the value of m . A value of $m=40$ was found to be robust in a clinical environment.

The final solution used values $m=40$ and $T_r=15$ (for image of $n_g=64$ grey levels) as being stable, conservative values that allowed the system to perform adequately in a range of circumstances including a clinical environment.

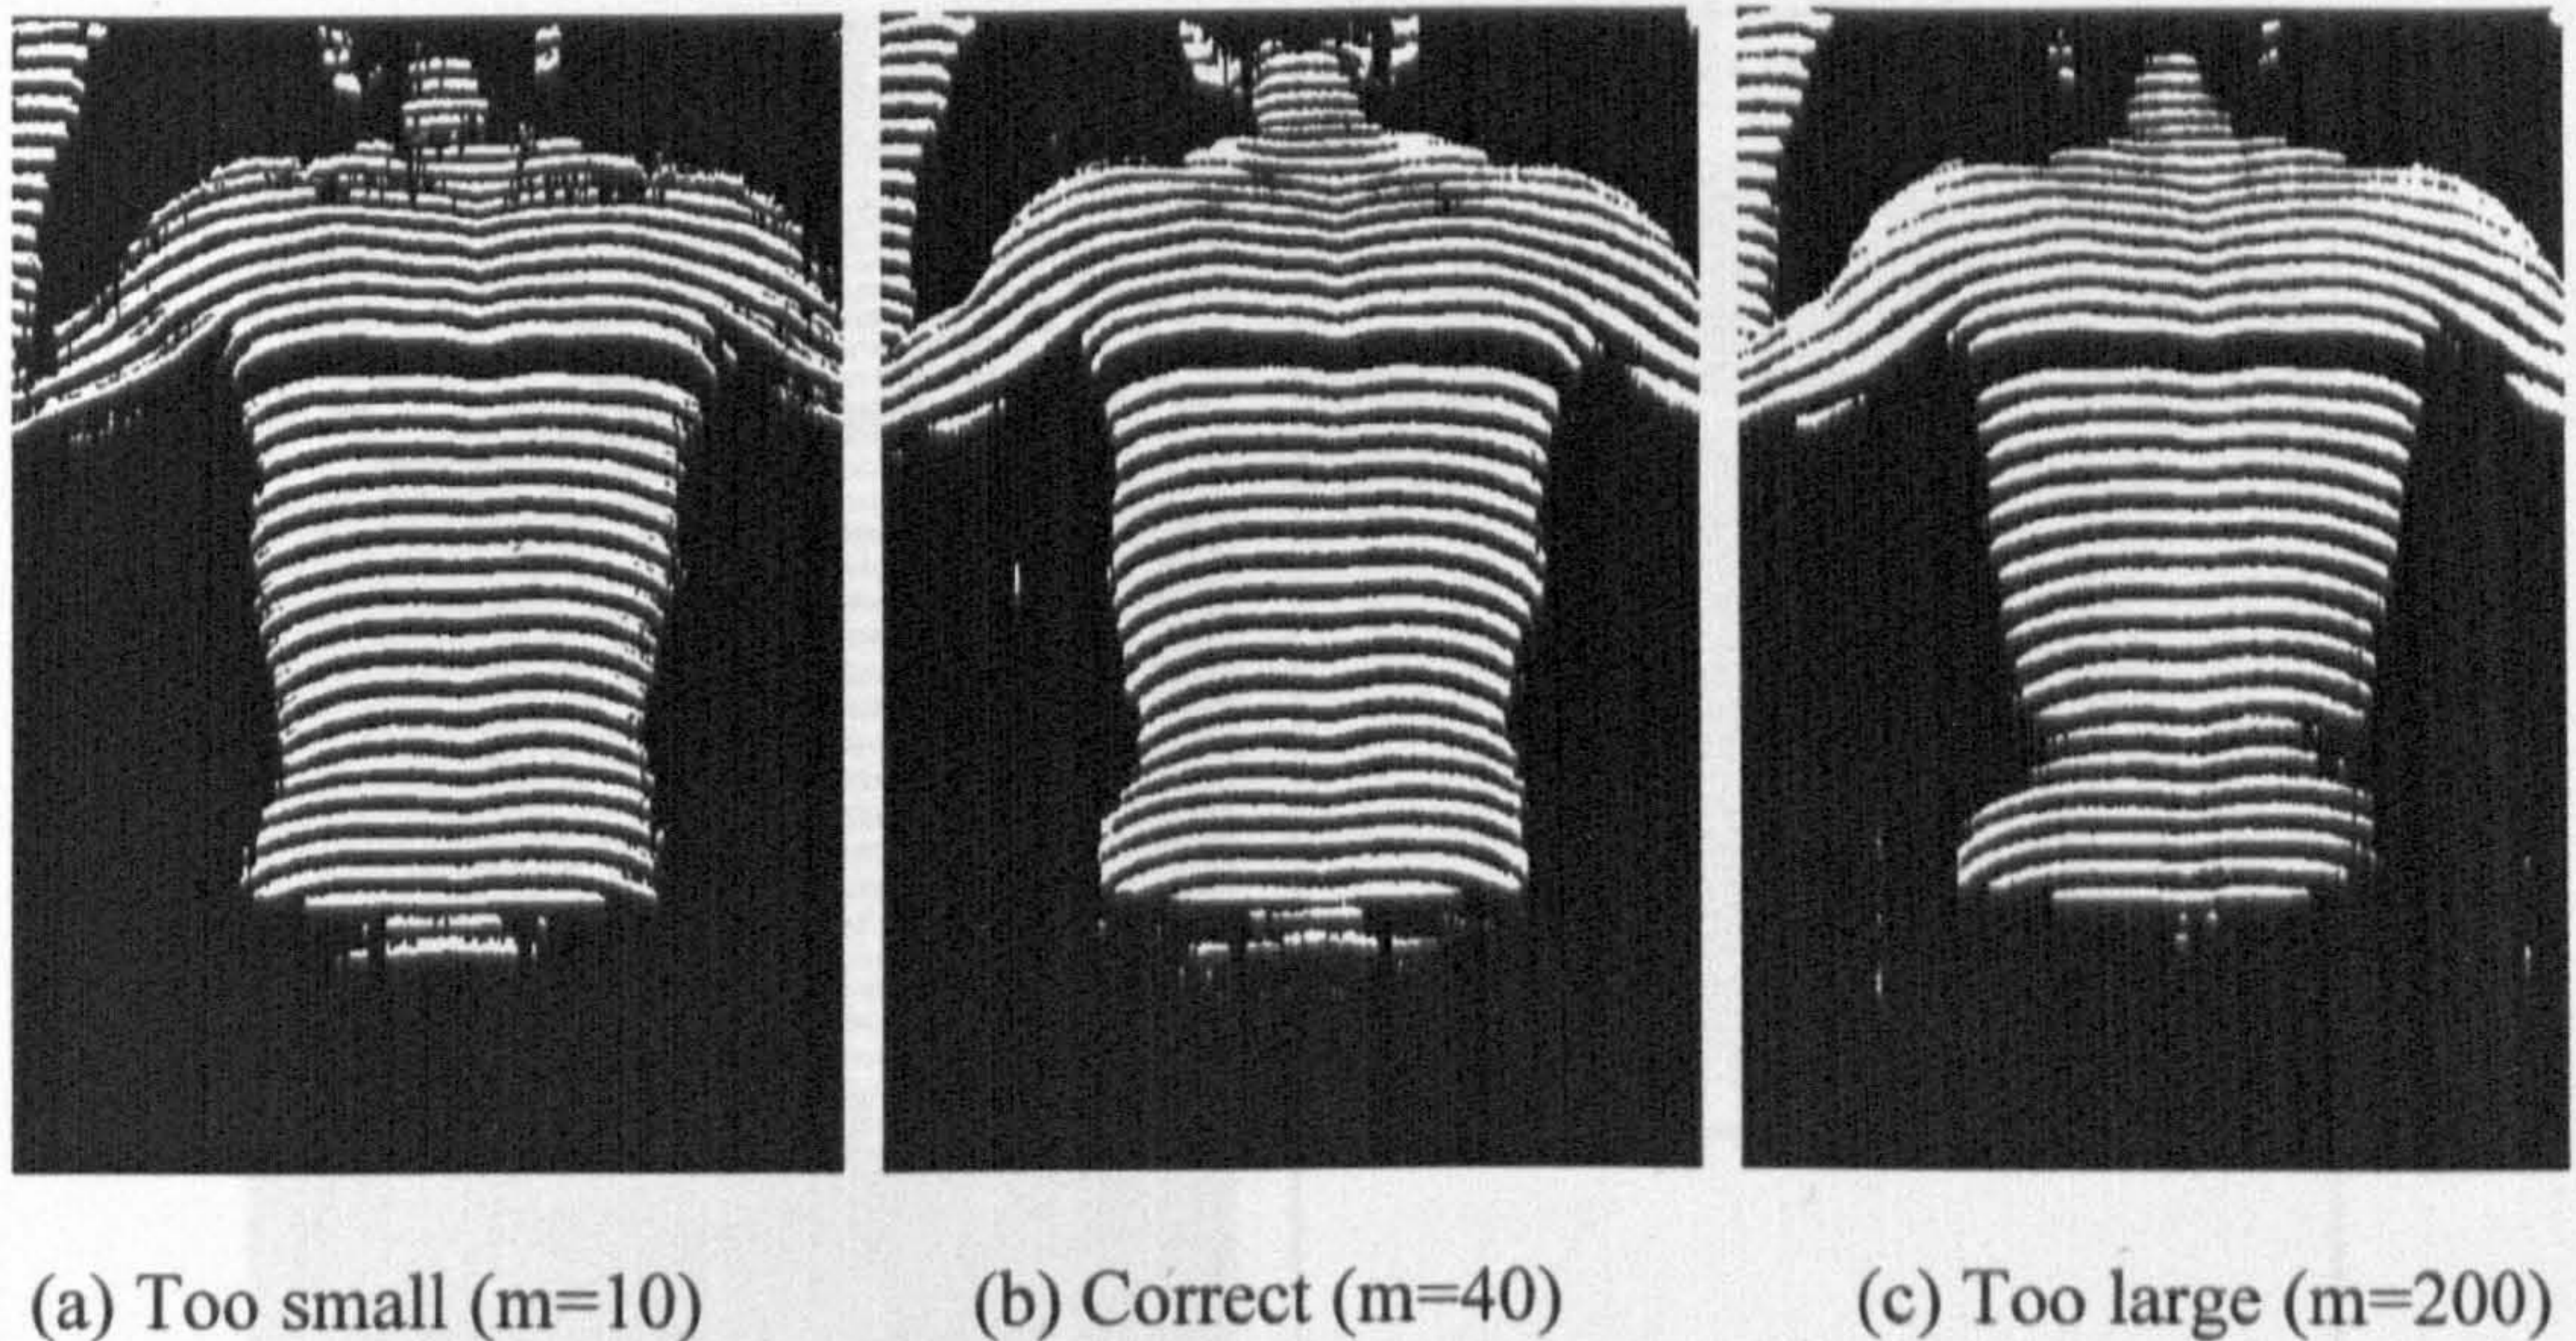


Figure 3.17 The effects of altering the size of the local area (for $T_r=15$, $n_g=64$).

3.4.6 Edge Detection

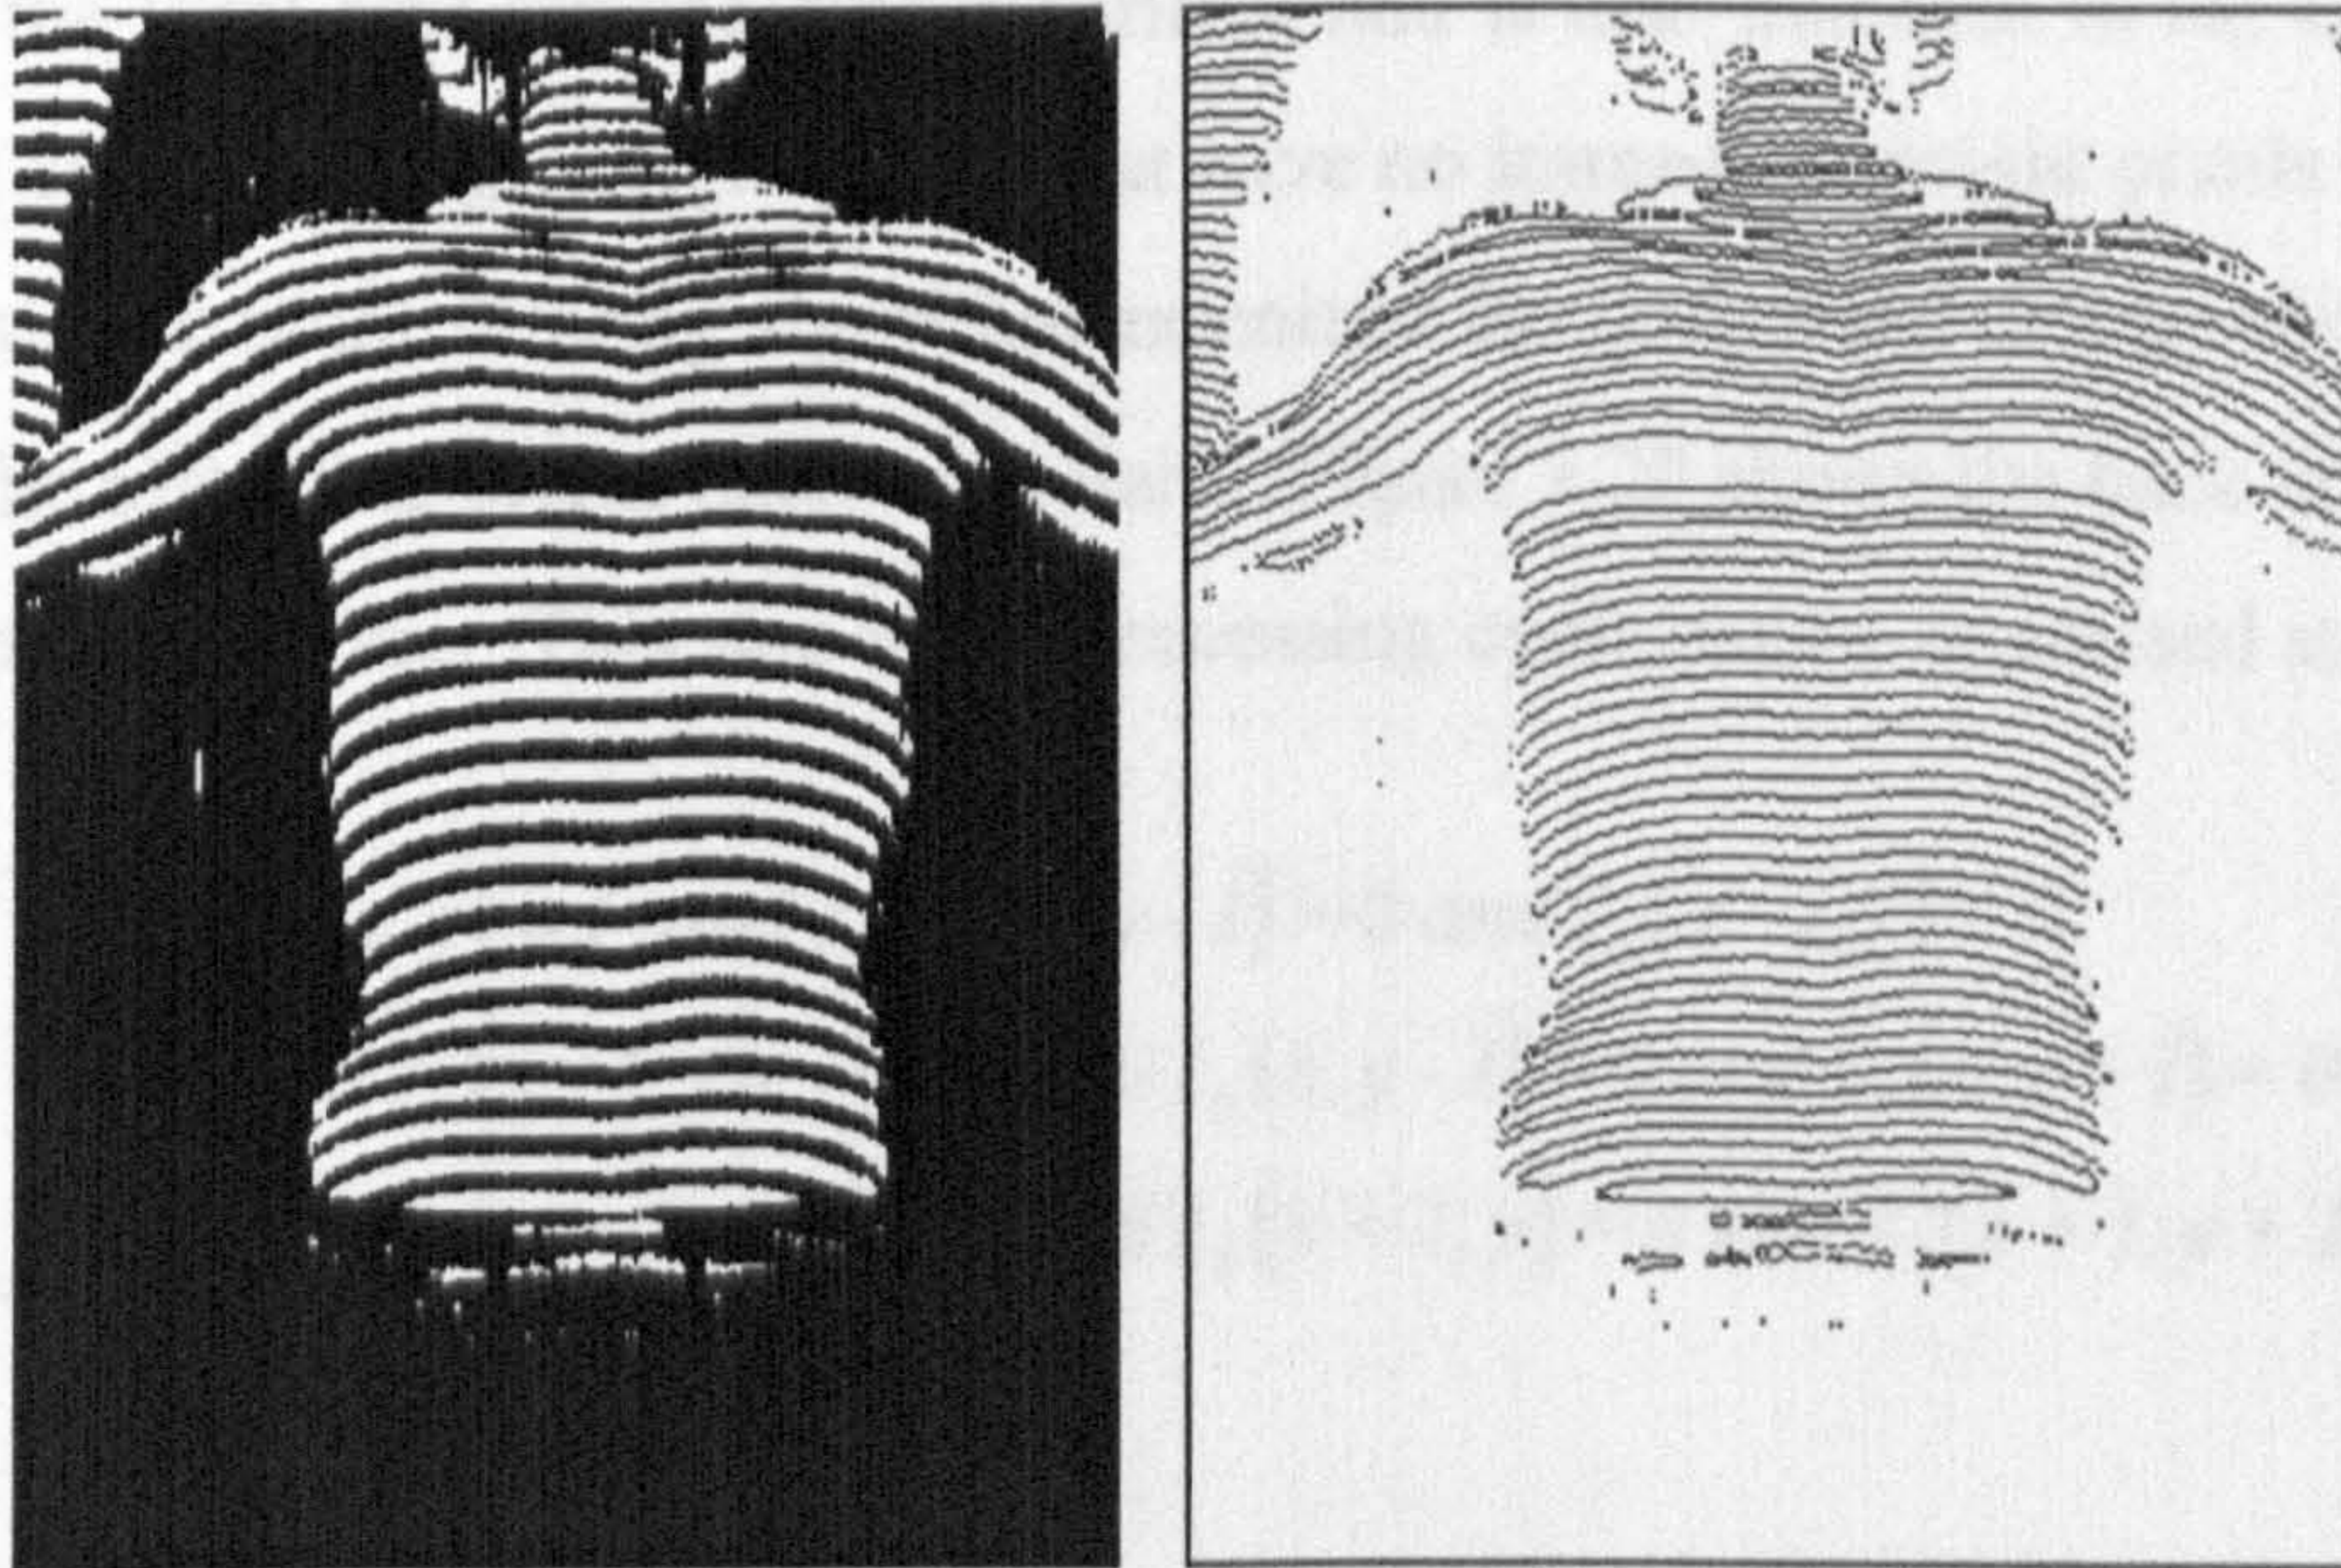
Ultimately, there must be a single point (pixel) for any y value that is chosen to be the x coordinate of the fringe. Processed fringes must therefore be represented as binary fringes that are a single pixel in width as shown in Figure 3.18. The production of the binary edge image from the thresholded image is straightforward. A simple algorithm takes each vertical line in the image and detects a binary step in adjacent pixels. The operation can be expressed as:

$$I_d = \left. \begin{cases} \text{if } \left([I_s(x, y) = 0 \text{ and } I_s(x + 1, y) = 1] \text{ or } [I_s(x, y) = 1 \text{ and } I_s(x + 1, y) = 0] \right) \text{ then} \\ 1 \\ \text{else} \\ 0 \end{cases} \right\}$$

[Equation 3.15]

The operation does not present a significant computational task. The combination of the thresholding algorithm with edge detection forms the binary dataset from which

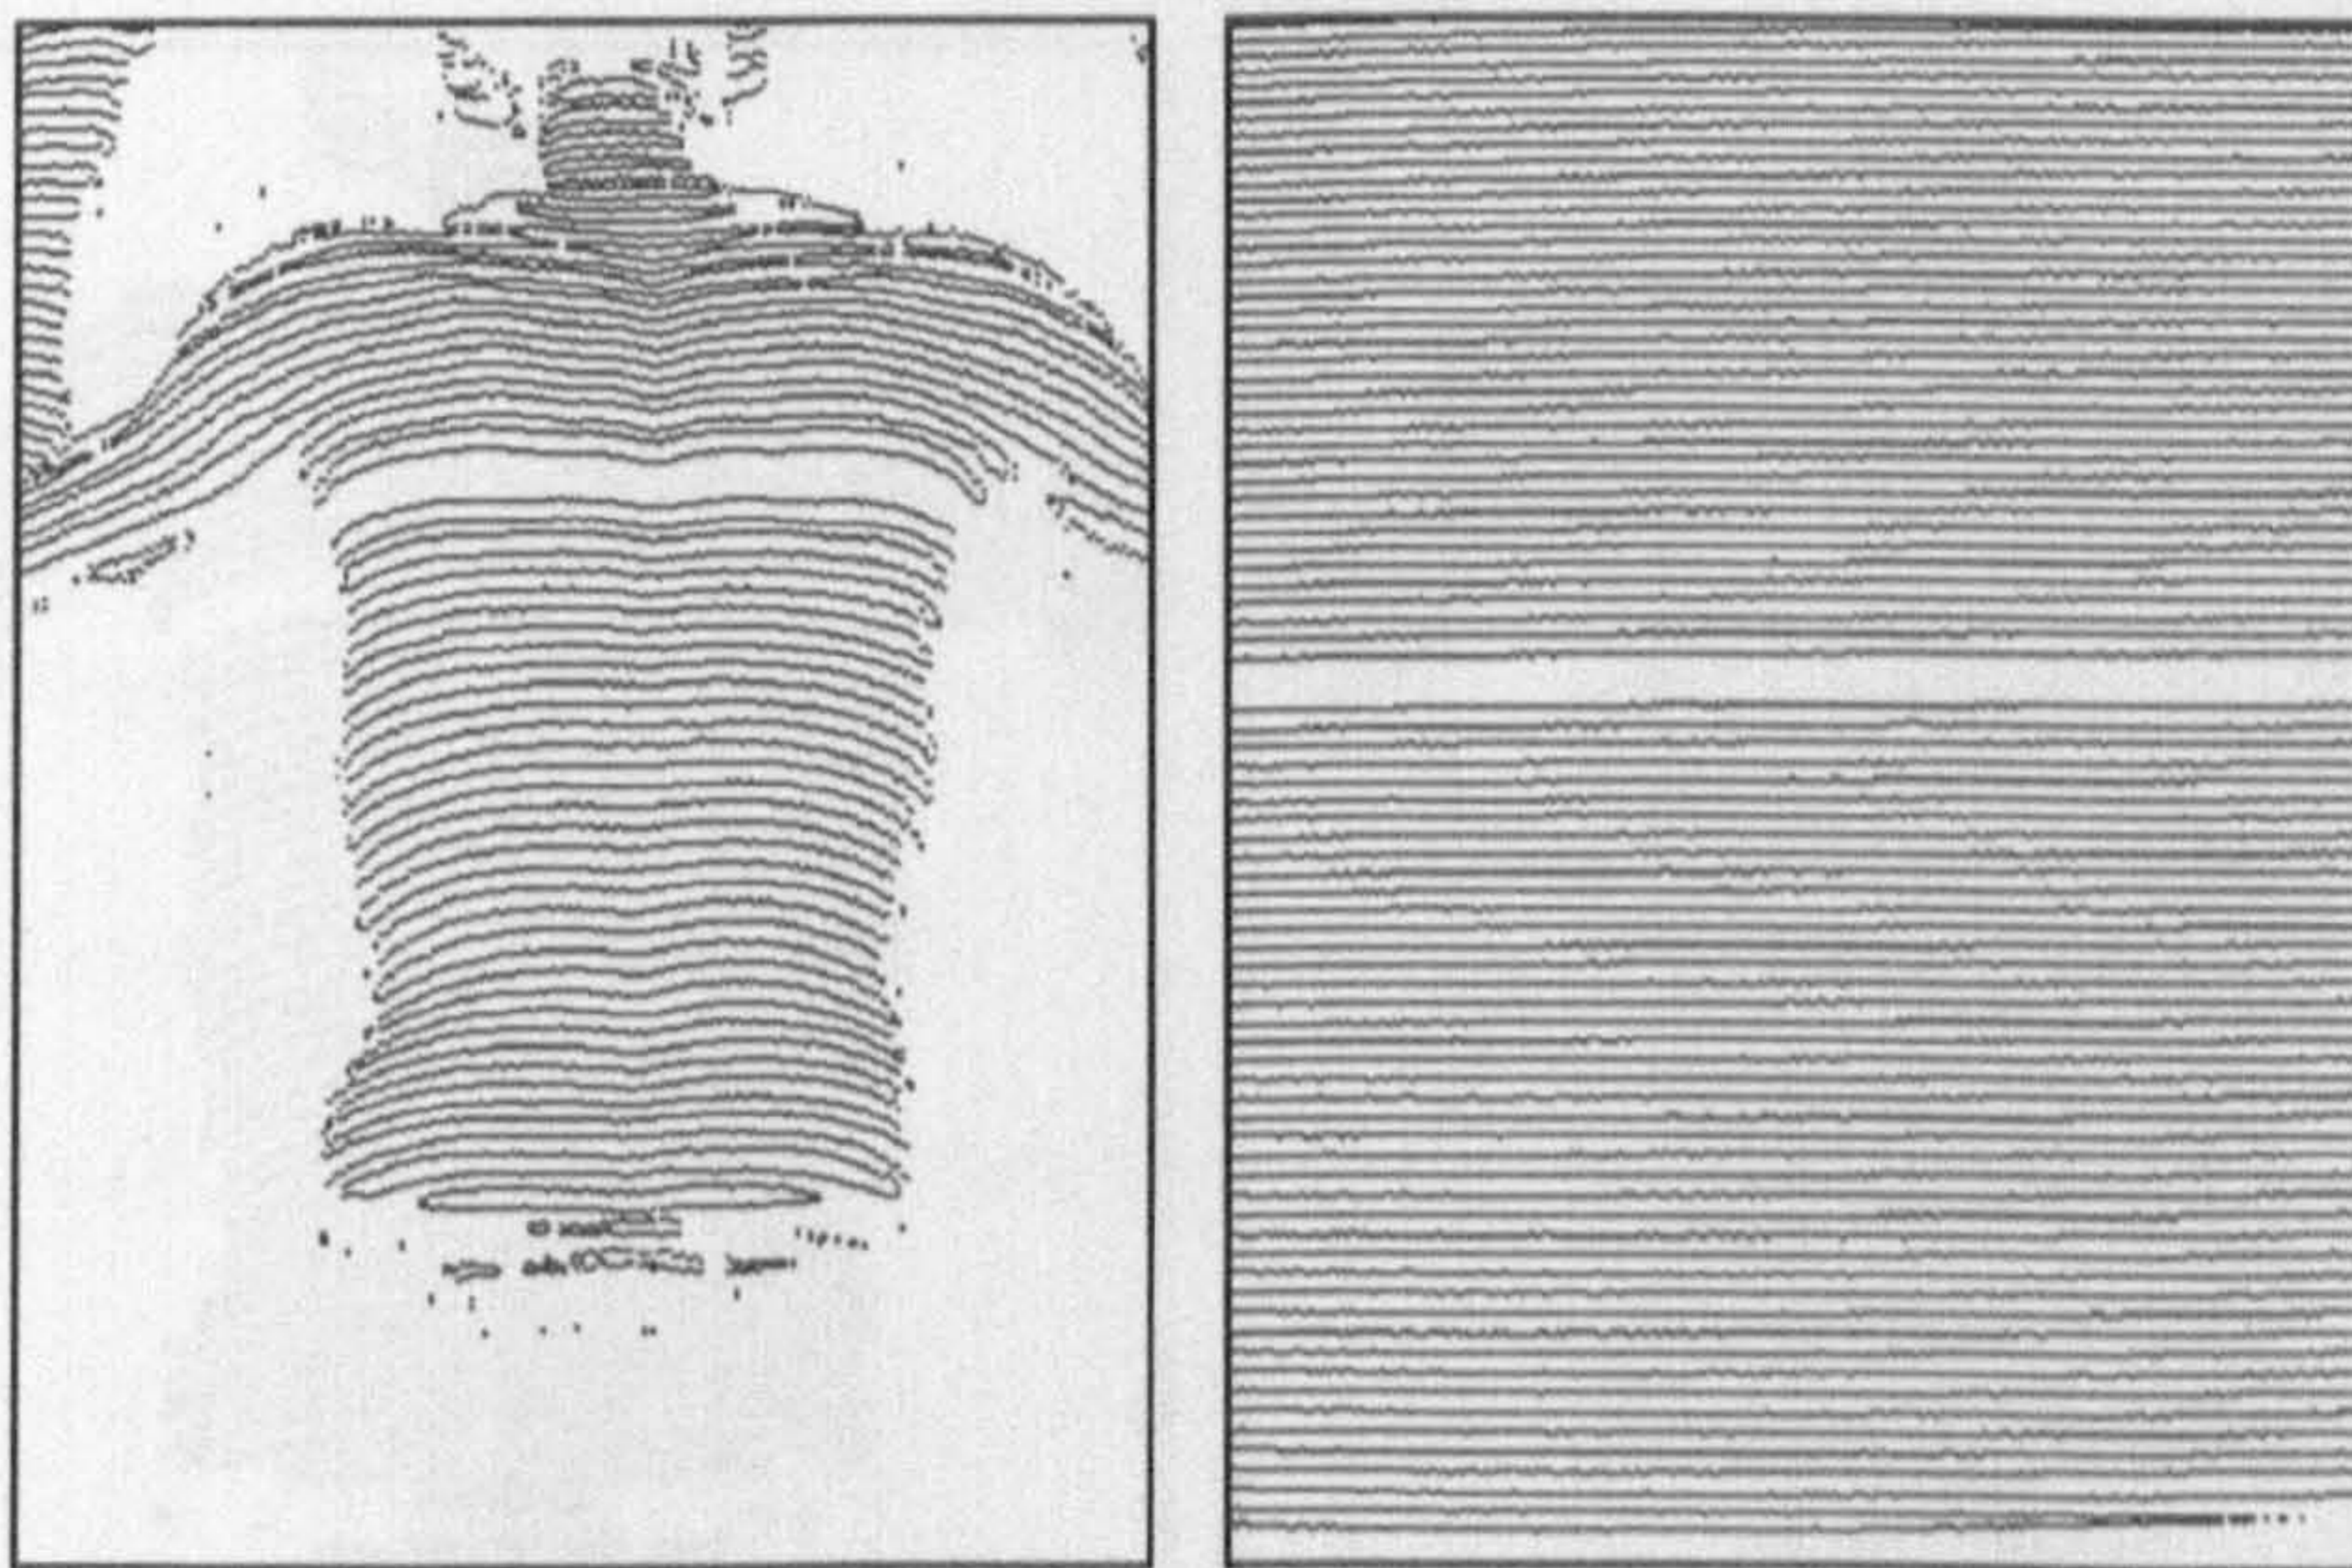
measurements will be taken. The algorithms are applied to both the back and reference images and the two images are shown side-by-side for comparison in Figure 3.19. The edge detected images form the images from which measurements are ultimately made.



(a) Binary source image

(b) After edge detection.

Figure 3.18 The binary edge image (inverted to improve visibility on printed page).



(a) Back edge image

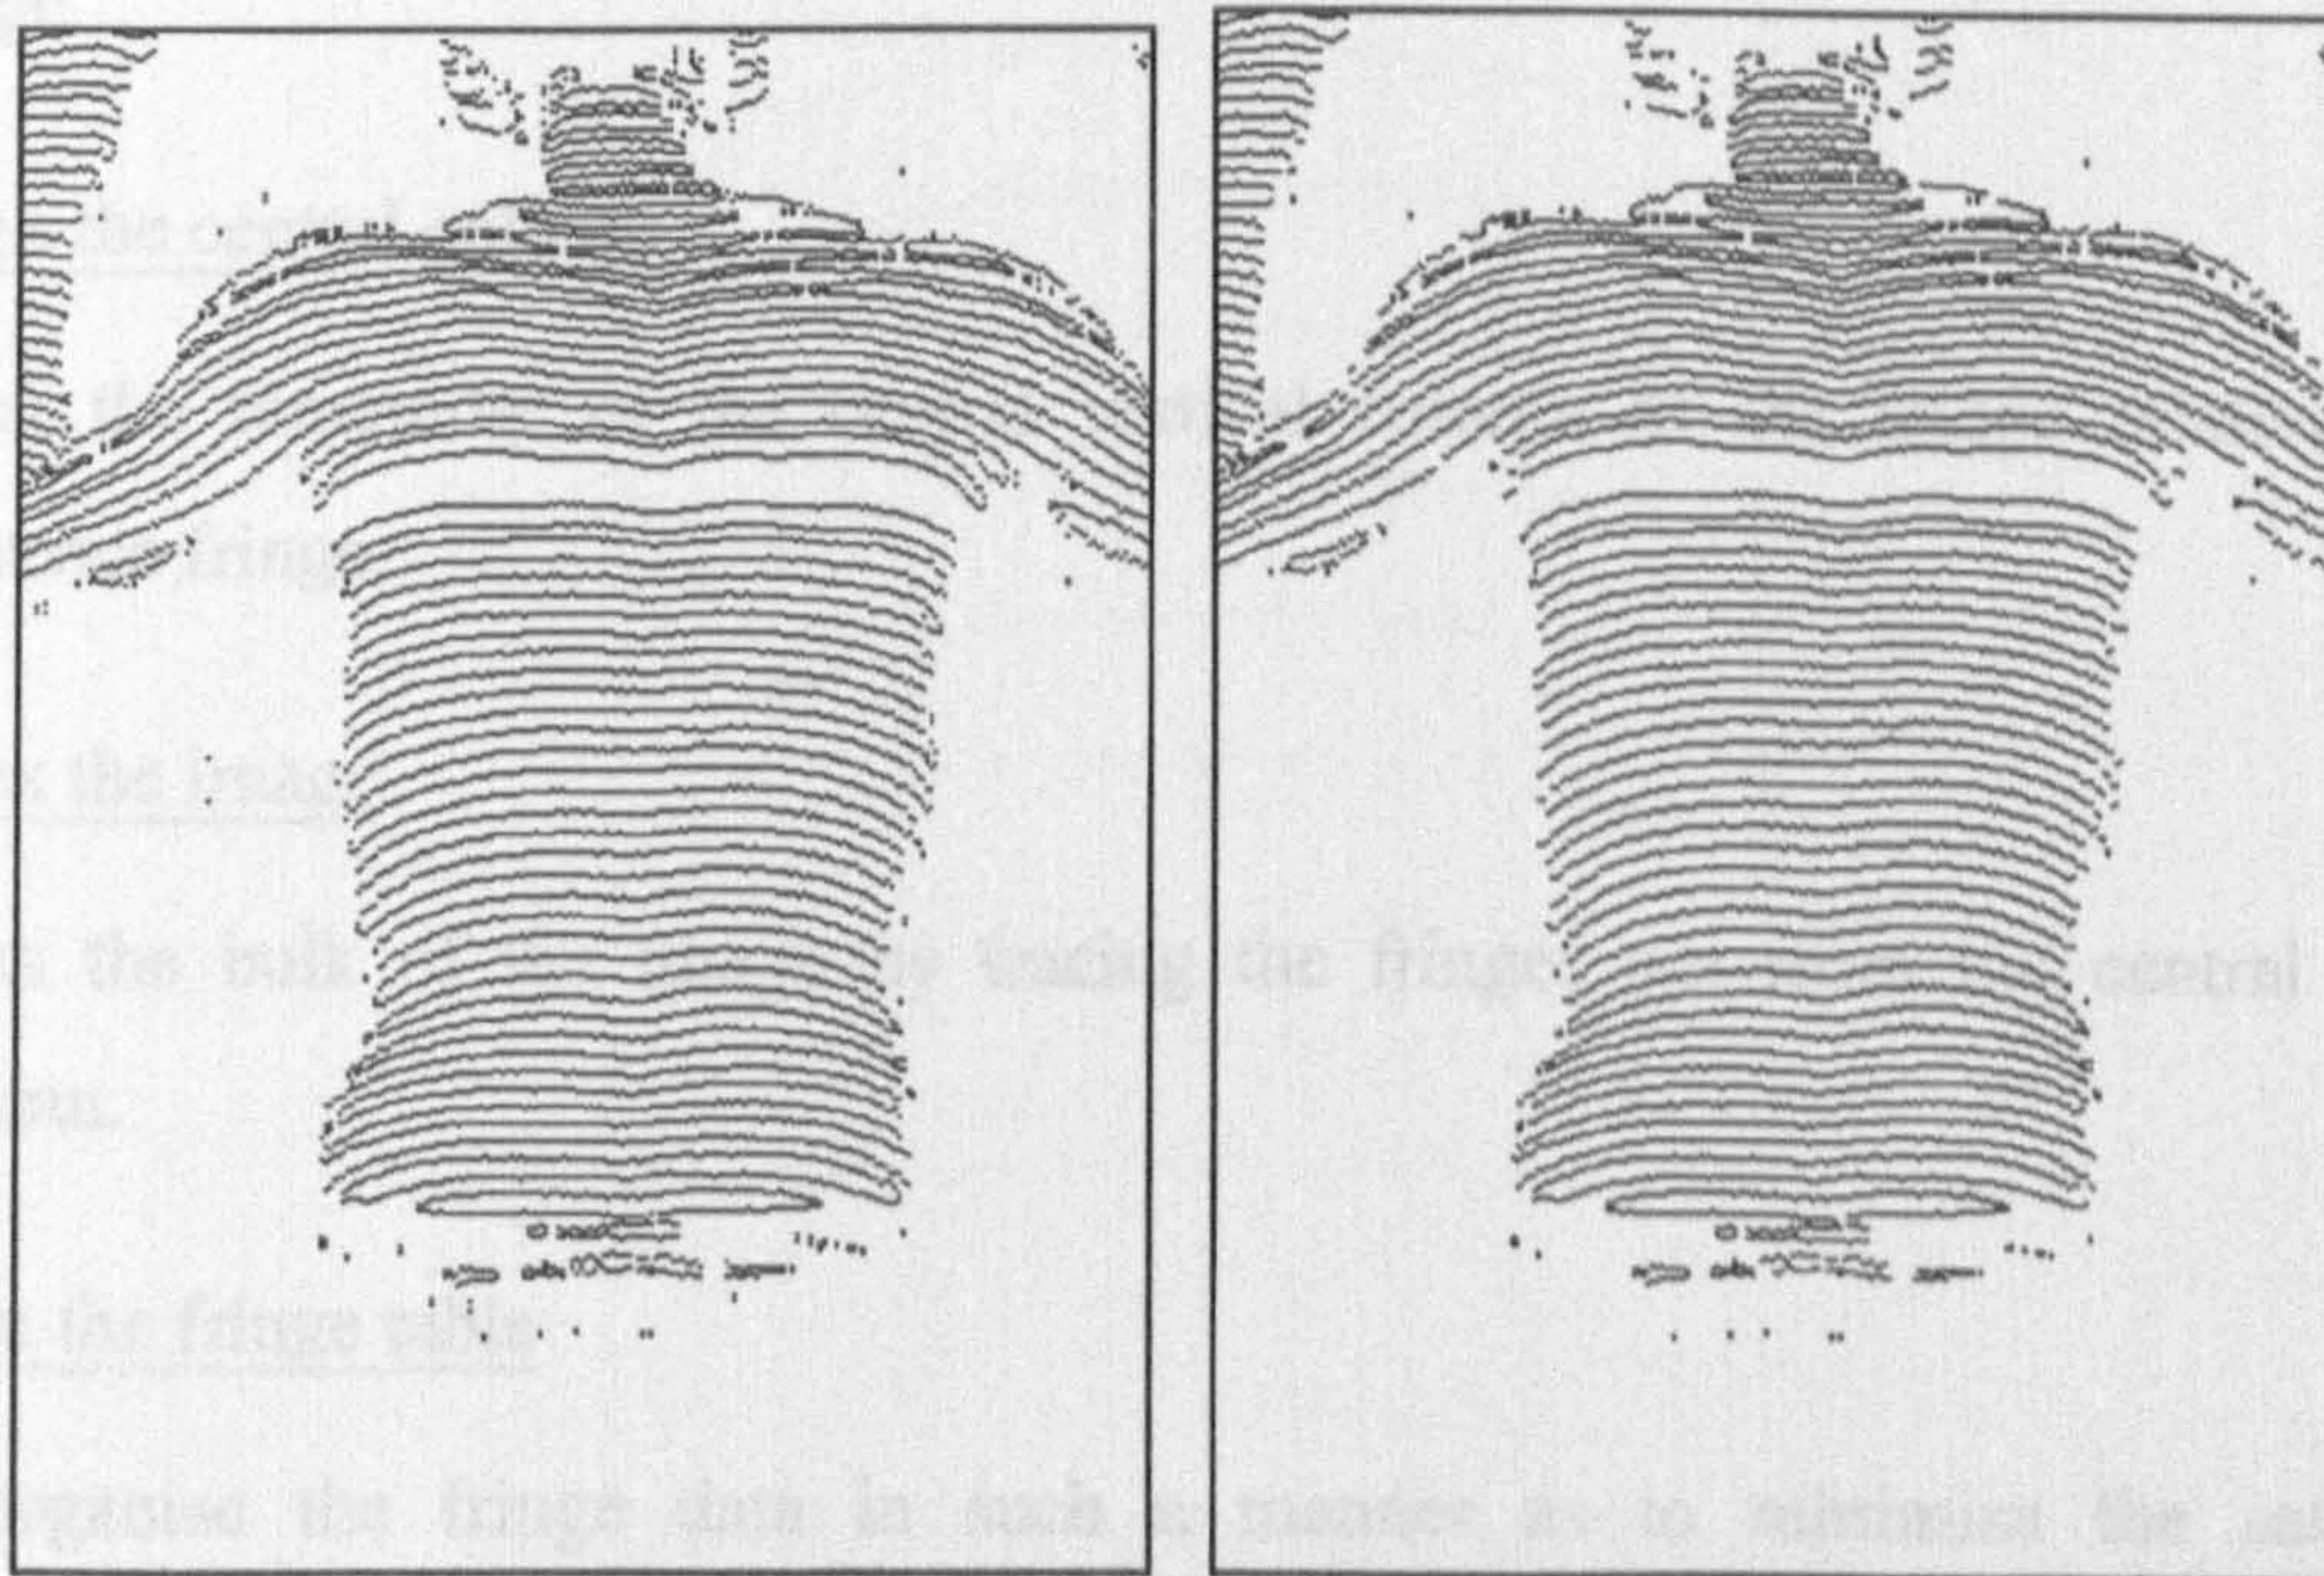
(b) Reference edge image

Figure 3.19 Back and reference images after edge detection.

3.4.7 Point Filtering

At the edges of the back the data quality becomes poor and this will ultimately translate to spurious depth measurements in the reconstruction. The poor data results from failure of the local area thresholding algorithm and is also manifest in the edge image. The worst cases are single bright pixels that have no immediate bright pixels adjacent to them. A simple point filtering algorithm examines the image and detects this condition then sets the offending bright pixel to be dark. Figure 3.20 shows the back image before and after point filtering. This step in the processing cycle can be expressed as:

$$I_d(x,y) = \left. \begin{array}{l} \text{if } I_s(x,y) = 1 \text{ and } I_s(x-1,y-1) = 0 \text{ and } I_s(x-1,y) = 0 \\ \text{and } I_s(x-1,y+1) = 0 \text{ and } I_s(x,y-1) = 0 \text{ and } I_s(x,y+1) = 0 \\ \text{and } I_s(x+1,y-1) = 0 \text{ and } I_s(x+1,y) = 0 \text{ and } I_s(x+1,y+1) = 0 \text{ then} \\ 0 \\ \text{else} \\ I_s(x,y) \end{array} \right\} \quad \text{[Equation 3.16]}$$



(a) Before point filtering

(b) After point filtering

Figure 3.20 The back image before and after point filtering.

The improvement is subtle and is more pronounced in some images than others. The sample image shown in Figure 3.20 does not show significant improvement: nevertheless, examination of pixels down the left of the trunk does show that some spurious pixels have been eliminated.

3.4.8 Fringe Ordering

The optical method depends intrinsically upon the ability of the system to match fringes in the back image with their corresponding fringes in the reference image. To do this, it is necessary to assign each pixel lying on a fringe an index value that is unique to the fringe. The process of doing this is *fringe ordering*. It should be noted that so far the reference fringe has been considered to be the central broad dark fringe. Now that the image has been changed to be a binary edge image, there are two reference fringes that represent the leading and trailing edges of the original dark reference fringe. There are four principal steps in the ordering of fringes. These are:

i. Locate the reference fringes

Search for and locate the reference fringes along the central vertical column of the image.

ii. Index the central column

Index the remainder of the central vertical column of the image relative to the reference fringe.

iii. Index the image

Index the bulk of the image by tracing the fringes out from the central vertical column.

iv. Form the fringe table

Re-organise the fringe data in such a manner as to minimise the search for corresponding pixels in the reference and back image. The data is organised into a data structure called the *fringe table*.

3.4.8.1 Searching for the Reference Fringes

The search for the reference fringes relies on two assumptions

- i. That all fringes intersect the central vertical column of the image. For an image of dimensions $512(x) \times 256(y)$ the central vertical column will be $y_c=127$.
- ii. That the surface does not contain severe local changes in gradient that could lead to the reference fringe not have a spacing clearly greater than the others.

These are reasonable assumptions that have been found to be valid for almost every fringe pattern encountered during the investigation.

The central vertical column of the image is considered in isolation. The column is examined, starting at the bottom of the image. Fringes are examined in groups of six, that is five inter-fringe spaces as shown in Figure 3.21.




Figure 3.21 Finding the reference fringes.

The intersections of six adjacent fringes with the central vertical column of the image are located to be x_1 , x_2 , x_3 , x_4 , x_5 and x_6 . The separations between the fringes are

calculated in terms of pixels to be $\Delta x_1 = x_2 - x_1$, $\Delta x_2 = x_3 - x_2$... and so on. Five inter-fringe spacings are determined in this way. The middle of these five spacings is considered as a candidate to be the reference fringe pair. In principle 3 fringe spacings could be used rather than 5. However, using 5 spacings reduces the capacity for error because the top and bottom two are averaged to decrease random errors in the fringe position.

The algorithm to detect the reference fringes depends upon identifying a fringe spacing that is significantly greater than for adjacent fringes. The reference fringes will have a spacing that, on a flat plane, will be three times the spacing of adjacent fringes (the absence of a fringe in the original image translates to the absence of two fringe edges and hence three inter-fringe spacings). The upper reference fringe will ultimately be assigned an index +1 and the lower assigned an index -1.

The spacing of the middle fringes is compared with the mean spacing of the two fringe spacings immediately above and below it. If the spacing of the middle fringe is Δx_m , the mean spacing of fringes above it Δx_{hi} and the mean spacing of the fringes below it Δx_{lo} , then the following defines the operation performed:

$$R_s = \left\{ \begin{array}{l} \Delta x_{lo} = \frac{\Delta x_1 + \Delta x_2}{2} \\ \Delta x_m = \Delta x_3 \\ \Delta x_{hi} = \frac{\Delta x_4 + \Delta x_5}{2} \\ \text{if } (\Delta x_m > 1.5\Delta x_{lo} \text{ and } \Delta x_m > 1.5\Delta x_{hi}) \text{ then} \\ \quad \frac{\Delta x_m}{\Delta x_{hi} + \Delta x_{lo}} \\ \text{else} \\ 0 \end{array} \right. \quad \text{[Equation 3.17]}$$

R_s is then the ratio of the middle fringe spacing, Δx_m , to the sum of the mean spacings of the fringes immediately above and below it, $\Delta x_{lo} + \Delta x_{hi}$. However, a meaningful ratio is only calculated in circumstances where the middle fringe spacing is at least 1.5 x the

mean spacing above and below it, otherwise R_s is set to be zero to indicate that this middle fringe is not a candidate for the reference fringe.

The ratio R_s is *determined for every set of five fringe spacings* along the central vertical column of the image. In fact, each value is recorded in a data structure of the type shown in Figure 3.22.

```
typedef struct
{
    float Rs ;
    int xNegative ;
    int xPositive ;
} Fringe Ratio ;
```

Figure 3.22 Data structure for defining R_s .

In the data structure, `xNegative` is the x coordinate of the fringe to be assigned an index -1 and `xPositive` is the x coordinate of the fringe to be assigned an index +1. Also, in the actual implementation of the algorithm, values of $R_s=0$ need not be recorded since the fringe is not being considered as a candidate to be the reference fringe.

At the end of this operation, there exists a set of values for R_s for every fringe except the top and bottom two fringes in the image which are not processed (it is not possible to designate these to be middle fringes). The fringe with the highest value of R_s is then determined to be the reference fringe and its edges, on the central column of the image are assigned indices accordingly.

A new ordered image is constructed to hold the indexed fringe pattern. Pixels in this image are either zero to indicate that they do not lie on fringes, or hold an index value that indicates the fringe with which they are associated. The first additions to the ordered image are the two pixels that represent the intersection of the upper and lower edges of the reference fringe with the central vertical column of the image. As more indexed pixels are identified they are added to the ordered image.

then the bright pixel in the edge image is mapped into the indexed image and assigned the same index as the pixel that was found in the indexed image.

The pattern of searching is crucial to the success of the algorithm. Both the shape of the search matrix and the order in which pixels in the indexed image are examined is critical. The search matrix, together with the order in which it is searched is shown in Figure 3.23. The bold column represents the central vertical column. The black pixels represent pixels that lie on fringes in the indexed image that have already been indexed. Each of the black pixels in the diagram will have the same index value. The grey pixels are bright pixels in the edge image. The algorithm has encountered a bright pixel, P, in the edge image. In fact the diagram depicts an undesirable situation: there has been a loss of connectivity along the fringe with a break in the fringe between P and the nearest indexed pixel.

The fringe tracking algorithm uses the search matrix to attempt to index the fringe by searching first for an indexed pixel at pixel a. If pixel a is an indexed pixel then P inherits its index and the index image is updated. If pixel a is not an indexed pixel then the process is repeated for pixels b, c, d and e. If an indexed pixel has still not been found then the search is cast wider to include pixels f, g and h.

Notice the order in which the search occurs. Since in the left hand side of the image the fringes will tend to slope down to the left, P is more likely to find an indexed pixel to its north-east than to its south east. For this reason the algorithm examines pixel b before pixel c, pixel d before pixel e and pixel g before pixel h. The increase in efficiency is negligible. The real value arises when, say, both pixels d and e are indexed pixels. The algorithm will cause P to inherit the index of d rather than e, which is generally desirable.

The general downward trend of fringes is not always observed in all areas of the image. For example, for some patients with a severely protruding scapula the trend will be in the opposite direction for that area of the image. Such conditions may cause fringe tracking to be lost. However, the downward trend is almost always present at the edges of the back where image quality is degraded and where the directional bias to fringe

tracking is of most value. Pixels to the right of the central column are processed in the same manner but the search matrix is vertically reflected as shown in Figure 3.24.




Figure 3.24 The search matrix for the right half of the image.

In this way, a complete image of indexed fringes is built up. When the indexing algorithm has been completed, every pixel in the edge image has either been linked into a fringe or has been discarded because it cannot be linked to a fringe. Figure 3.25 shows the result of the indexing algorithm for back and reference images.




Figure 3.25 Result of the indexing algorithm.

The fringes in Figure 3.25 are colour coded to indicate the index numbers. The colours are reused for different indices because of the limited number of colours that can be meaningfully printed and interpreted.

| Colour Code | Index |
|-------------|-----------------------------------|
| Black | +1 , -1 (Reference fringes) |
| Yellow | $\pm (2 + 6n)$, $n=0,1,2 \dots$ |
| Blue | $\pm (3 + 6n)$, $n=0,1,2 \dots$ |
| Magenta | $\pm (4 + 6n)$, $n=0,1,2, \dots$ |
| Cyan | $\pm (5 + 6n)$, $n=0,1,2 \dots$ |
| Red | $\pm (6 + 6n)$, $n=0,1,2 \dots$ |
| Green | $\pm (7 + 6n)$, $n=0,1,2 \dots$ |


Table 3.4 Fringe colour coding for Figure 3.25.

3.4.8.4 Indexed Image Editing

The entire image processing system is designed to function completely automatically without the need for intervention by the user, who is usually a clinician. However, the image processing can benefit from manual intervention. Consider the back image shown in Figure 3.25. In the areas around the neck and at the bottom of the image spurious data does exist. To detect and correct this spurious data automatically would be complex because it is difficult to discriminate between spurious and valid pixel values. These data points do not have any detrimental effect on the correct data which occupies the vast majority of the surface. In order to improve the quality of the image, the user has the option to edit the image data manually if he so desires. Two image editors are provided.

The first image editor, the trim editor, provides the user with the ability to trim the edges of the back surface to eliminate spurious edge effects. For both trim and pixel editors, the binary edge image, not the indexed image, is updated. After this has been done the indexing procedure is invoked to produce a new indexed image. Although re-indexing the image is not strictly necessary for the trim editor because pixel interconnections are not changed, re-indexing provides a convenient method for producing the updated indexed image and allows the same algorithm to be used whether

the image is trim edited or pixel edited. Figure 3.26 shows the back surface before and after trim editing.




(a) Before Trim Editing (b) After Trim Editing

Figure 3.26 Back image before and after trim editing.

Clearly the spurious data at the boundary of the surface have been eliminated. The user is only presented with the edit option for the back image: the reference image is not subject to the same problems.

The second image editor available to the user is a pixel editor that allows the knowledgeable user to repair any breaks in the continuity of fringes. To use the editor, the user must have a basic understanding of the way that the indexing works. The data displayed is the binary edge image overlaid with the indexed image. Again the edge image is modified as pixels are added and deleted, and the indexing algorithm is invoked to produce a new indexed image. The pixel editor has been a valuable tool for repairing poor quality images when good reconstruction is required. In early clinical trials its use has avoided the necessity of recalling the patient for a second imaging session.



(a) Before Pixel Editing

(b) After Pixel Editing

Figure 3.27 The back image before and after pixel editing.

Figure 3.27 shows the back image prior to and after pixel editing. There is a clear improvement in the data quality around the neck area. It should be emphasised that this type of editing is not used in the majority of cases and that when it is used, care is taken only to repair existing fringes and not to add new data.

3.4.9 The Fringe Table

After fringe ordering, both the back and reference images hold pixel values that are index numbers. Eventually, the system will need to compare the x coordinates for corresponding pixels in the back and reference images.

The simplest algorithm to identify pixel pairs in the back and reference images might be to scan the back image and when a fringe pixel is encountered to scan that column in the reference image to find a pixel with the same index. This solution would be very inefficient because a search would be required for every fringe pixel in the back image.

A better approach is to reorganise the back and reference images into a new two dimensional data structure, the *fringe table*. What is required is a table that holds, for a particular index and column (y), the x coordinate of the pixel with that index and

column number. The operation to convert an indexed image to a fringe table can be expressed as:

$$X[I(x, y), y] = \begin{cases} \text{if } I(x, y) \neq 0 \text{ then} \\ x \end{cases} \quad \text{[Equation 3.18]}$$

where X is the fringe table and $I(x, y)$ is the indexed image. The condition $I(x, y) \neq 0$ indicates that only valid indexed pixels are entered ($I(x, y) = 0$ indicates the background or a pixel lying between fringes). So the x coordinate of a pixel with index i along column y will be $X[i, y]$.

It may also be noted that although this reformulation of the dataset is designed for the speed of subsequent processing it also represents a significant data reduction. Usually there will not be more than 50 fringes in the image so that the data-set is reduced from being a 512 x 256 pixel image to a 50 x 256 element data structure.

There are two further problems:

- i. For a given index I , there may not be an valid entry in both the reference index table and the back index table for some or all values of y . A restatement of this would be to say that it is possible that all or part of an indexed fringe that exists in the back image may not exist at all in the reference image and vice versa. This situation usually arises for fringes at the bottom of the back image or at the top of the reference image. The angle of the projector to the object (both back and reference plane) is such that for a fringe at the bottom of the back image, the corresponding fringe in the reference image may be beyond the field of view of the camera. Similarly, for a fringe at the top of the reference image, the corresponding fringe in the back image may also be beyond the field of view.

The only optical solution would be to ensure that this event cannot occur by inspecting the back and reference images, counting fringes, at the time of acquisition. This is not a practical solution for a real clinical system because the complexity of set-up and measurement procedure needs to be minimised.

- ii. For a given column y , there may not be valid entries in the fringe tables for some or all values of i . This condition most commonly occurs at the left and right sides of the back image where few fringes exist. If the trim editor has been used to clip the arms out of the image then there will not be *any* fringes at the left and right edges of the back image.

Equation 3.18 did not specify values in the fringe table for these conditions. A solution was devised to overcome both of these problems. Once the indexed images have been calculated, the significance of the reference fringes is eliminated: each image is simply a set of indexed fringes. It is therefore, no longer necessary to maintain +1 and -1 fringe indices for the reference fringes. It is, however, important to ensure that any alterations to index numbers in the back image are also applied to the reference image and vice versa.

First, the indices of all fringes in both indexed images are shifted to have positive values. This is achieved by finding the minimum index in the combined images and adding it, plus 1, to each index. After the shift all indices in both images are positive. When the conversion of the images to fringe tables occurs, both fringe tables will have the same dimensions, but will not have valid entries for some indices. An incidental value to shifting the indices to be positive is that the indices in the fringe table will also be positive and this is of value if the tables are implemented as arrays in languages such as C where negative array indices are not permitted: the code is simplified.

Secondly, it is necessary to provide an indicator for each fringe table entry that verifies that an entry is valid. One solution would be to use a negative value for the x coordinate to indicate an invalid entry, effectively encoding a status value into the entry. An alternative solution is to have, for each entry in the fringe table, a separate status value. Each entry in the fringe table becomes a data structure that has an x coordinate field and a status field. The status field can be a single bit set to 1 to indicate that the x value is valid, 0 to indicate otherwise. Another advantage of implementing the fringe table entries as data structures is that an additional field can be added to the data structure later if required. For example, it might be that additional fields could be added that

would be used in the handling of the (rare) condition where indexing has failed and more than one fringe in, say, the back image has the same index: some arbitration could be implemented based on other local fringe positions.

After the indexed images have been converted to fringe tables, they are discarded because they are not required for subsequent processing.

3.4.10 Representing Three-dimensional Coordinates

The image processing system must ultimately yield a set of (x,y,z) coordinates that represent the surface shape of the human back. Given that the source data is a digital image, a two dimensional array of pixels, and that sampled surface values correspond to pixels in the image, it was reasonable that the surface should be represented in a similar fashion.

The surface is represented as a two dimensional array of floating point pixels. The first dimension of the array represents the x direction and the second dimension represents the y direction. If a pixel at (x,y) has a known z (depth) coordinate associated with it, then the value of the array element at that pixel is the z value. Another way of stating this would be to say that the three-dimensional coordinate set is a floating point digital image that represents surface depth, rather than intensity. This convention is in widespread use in surface reconstruction and is referred to here as the *surface array*. The x and y coordinates are simply pixel values. The scaling factors L_x and L_y (see Section 3.3.6) are applied to measurements in the x and y directions to translate pixel units into actual real-world coordinates in millimetres when required. The z values are stored in millimetres.

To simply record three dimensional coordinates in this manner is not adequate for representing all the information about a surface. For each pixel, it is necessary to record a status value to indicate other information about the surface at that point. Status values are used, for example, to determine whether a pixel lies within the area of interest in the image (in other words whether a pixel is a valid measured point or whether it is associated with the background).




Figure 3.28 Surface representation.

The array of status values is referred to as the *status array* and is shown with the surface array in Figure 3.28. The complete set of status values used are shown in Table 3.5. When processing is complete, all status values should be greater than or equal to 2, and the status values for pixels lying on the surface should be greater than or equal to 3.

| Status Value | Mnemonic | Meaning |
|--------------|-----------------|--|
| 0 | ZS_Test | Reserved for testing. |
| 1 | ZS_Unassigned | Pixel not assigned a valid status value. |
| 2 | ZS_Out | Pixel lies outside the area of interest. |
| 3 | ZS_In | Pixel lies inside the area of interest. |
| 4 | ZS_Hard | Pixel is a true sampled data value. |
| 5 | ZS_Interpolated | Pixel is an interpolated data value. |
| 6 | ZS_Boundary | Pixel is a boundary pixel. |
| 7-255 | not defined | Unused/reserved for future use |

Table 3.5. Status values associated with surface coordinates.

To manipulate two arrays of values such as these in a computer language is unwieldy and may sometimes be computationally inefficient. In fact each pixel in the dataset is constructed as a small data structure in C shown in Figure 3.29.

```
typedef struct
{
    float z ; /* The Z coordinate */
    int s ; /* The Status Indicator */
} t_ZCoords ;
```

Figure 3.29 The C data structure defining a z coordinate.

The composite array which holds both the surface array and the status array is constructed as shown in Figure 3.30.

```
t_zCoords surface[512][256] ;
```

Figure 3.30 The C declaration of the surface array for a 512 x 256 pixel image.

The array constructed in this way is referred to as the *surface image*. Status values change during the processing cycle. For example, as soon as it is known that a pixel lies inside the boundary of the area of interest, it can be marked as *ZS_In*. However, if the pixel is later assigned interpolated z value then its status will be updated to be *ZS_Interpolated*. Clearly it is implicit that an interpolated value must lie inside the area of interest. Part of the merit in this scheme of status values is that at any stage in the processing, progress can be monitored and a surface image or map of status values displayed.

3.4.11 Calculation of z Coordinates

The equation for calculating the z (depth) coordinates was given in Section 3.3.5 to be:

$$z_S = \frac{D_C L_X \Delta x'}{D_P + L_X \Delta x'} \quad \text{[Equation 3.19]}$$

The equation is both simple and computationally desirable. The $D_C L_X$ multiplication can be performed in advance of the calculation of z so that for each measurement of z there are two floating point multiplications ($L_X \Delta x'$ and $(D_C L_X) \Delta x'$), one floating point addition ($D_P + (L_X \Delta x')$) and a floating point division of the denominator into the numerator. Moreover, there is no requirement for the calculation of trigonometric functions as is often the case when working with optical geometries. The calculation of z_S does not, however, present a heavy computational task in comparison with, say, the thresholding algorithm.

The calculation of z coordinates from the fringe tables is straightforward. Each (x,y) pixel in the surface image has two fields, .s the status field and .z, the actual z


coordinate for the point. The operation to calculate the surface image can be expressed by Equations 3.20 and 3.21. For the $.z$ field:

$$z_S(x, y).z = \left. \begin{array}{l} \text{if } X_{back}[i, y].s = 1 \text{ and } X_{reference}[i, y].s = 1 \text{ then} \\ \frac{D_C L_X(X_{back}[i, y].x - X_{reference}[i, y].x)}{D_P + L_X(X_{back}[i, y].x - X_{reference}[i, y].x)} \\ \text{else} \\ \text{undefined} \end{array} \right\} x=X_{back}[i, y]$$

[Equation 3.20]

$$z_S(x, y).s = \left. \begin{array}{l} \text{if } X_{back}[i, y].s = 1 \text{ and } X_{reference}[i, y].s = 1 \\ ZS_Hard \\ \text{else} \\ \text{unmodified (initialised to } ZS_Unassigned) \end{array} \right\} x=X_{back}[i, y]$$

[Equation 3.21]



(Red= $ZS_Unassigned$, Cyan= ZS_Hard)


Figure 3.31 The status array after z calculation

After this operation has been performed, the surface image holds status values that are $ZS_Unassigned$ for the bulk of the image, with status values ZS_Hard for those pixels

that have valid depth measurements. Those pixels that have ZS_Hard status values also have valid z measurements in the z field. Figure 3.31 shows the status array after the z calculation.

3.4.12 Defining the Boundary of the Back

Some of the subsequent algorithms, notably the interpolation, require that the boundary of the back surface in the image is clearly identified. The top and bottom boundaries of the image can be defined well by the top and bottom fringes. However, the detection of the left and right boundaries is more difficult.




Black=ZS_Boundary

Figure 3.32 The status array after boundary location.

A simple algorithm was devised that takes the endpoints of each fringe and links the sets of left and right endpoints together using straight lines in the status array. In order to do this, the fringes are quickly traced to find their endpoints in a manner similar to the that described for the fringe indexing except that it is very much simplified because the indexing has already been performed. Figure 3.32 shows the status array after boundary location.

The back surface is completely enclosed by the top and bottom fringes and the boundary pixels. Those pixels lying outside the surface are marked as such in the status array by setting the status value to *ZS_Outside*. This is achieved by counting pixels in from the edges of the image in both the x and y directions until a boundary pixel is encountered. More complex methods than this can be devised that will also cope with the surface being re-entrant in both the x and y directions. However, this was not necessary because the simple technique worked reliably. Figure 3.33 shows the status array after the assignment of *ZS_Outside* values.



Green=*ZS_Outside*

Figure 3.33 Assignment of pixels outside the back surface

At the end of this operation the surface image holds valid z coordinates for pixels lying on fringes, a marked area of those pixels that are outside the surface, the boundary pixels for the back and unassigned pixels that lie between fringes.

3.4.13 Interpolation

The object of the surface reconstruction is to yield a continuous surface for which estimates of z coordinate can be extracted for all points lying on the surface. It is therefore necessary to assign z coordinates to pixels lying between fringes on the surface. An approximation of the z values can be achieved by performing a linear interpolation, in the x direction, based upon the z values of the surface at fringes immediately above and below a pixel. Two special problems arise:

- i. The interpolation should be tolerant of occasional breaks in the continuity of the fringes where pixels that should be part of a fringe are not marked as such. For the system as described in this chapter, this phenomenon will not occur, because the indexing algorithm will fill in estimated positions for such pixels. However, more complex indexing algorithms might yield an indexed image which is more complete but has breaks in continuity. It is desirable that the interpolation should be able to cope with this condition by interpolating between any two measured z coordinates that lie on fringes.
- ii. The interpolation algorithm must not attempt to interpolate the z coordinates for pixels lying on or outside the boundary.

The interpolation operation can be expressed for the surface array and status arrays as:

$$z_{SD}(x, y).z = \left\{ \begin{array}{l} \text{if } z_{SS}(x, y).s = ZS_Hard \text{ then} \\ z_{SS}(x, y).z \\ \text{else if } z_{SS}(x_1, y).s = ZS_Hard \text{ and } z_{SS}(x_2, y).s = ZS_Hard \\ \text{and } z_{SS}(x, y) \neq ZS_Boundary \text{ and } z_{SS}(x, y) \neq ZS_Outside \\ z_{SS}(x_1, y).z + \frac{x - x_1}{x_2 - x_1} (z_{SS}(x_2, y).z - z_{SS}(x_1, y).z) \\ \text{else} \\ \text{not defined} \end{array} \right.$$

[Equation 3.22]

$$z_{SD}(x, y).s = \left\{ \begin{array}{l} \text{if } z_{SS}(x, y).s = ZS_Hard \text{ then} \\ ZS_Hard \\ \text{else if } z_{SS}(x_1, y).s = ZS_Hard \text{ and } z_{SS}(x_2, y).s = ZS_Hard \\ \text{and } z_{SS}(x, y) \neq ZS_Boundary \text{ and } z_{SS}(x, y) \neq ZS_Outside \\ ZS_Interpolated \\ \text{else} \\ ZS_Unassigned \end{array} \right.$$

Equation 3.23]

where x_1 is the x coordinate of the nearest fringe (i.e. pixel with status ZS_Hard) below the current pixel and x_2 is the x coordinate of the nearest fringe above the current pixel. Figure 3.34 shows the values in the status array after the interpolation.




Figure 3.34 The status array after interpolation.

The reconstruction is now complete. The surface image holds the z coordinates for each pixel lying on the surface and status values for all pixels in the image have been assigned, with the exception of those few pixels near the edge of the surface where interpolation was not possible. Figure 3.35 shows a projection of the surface.




Figure 3.35 A projection of the surface

3.4.14 Simplifying the Status Values

Once the reconstruction has been achieved it will be subjected to further processing to extract clinically meaningful parameters that relate to scoliotic deformity, and this processing is described in Chapter 4. However, some of the status values such as ZS_Hard and $ZS_Boundary$ are not required for subsequent processing. All that is required is that for each pixel in the image there should be a status value that indicates whether the pixel is part of the surface or not. The status array is therefore updated so that it hold values either ZS_In or ZS_Out . The re-mapping of status values can be expressed as:

$$z_{SD}(x, y) = \left. \begin{array}{l} \text{if } z_{SS}(x, y).s = ZS_Out \\ \text{or } z_{SS}(x, y).s = ZS_Unassigned \\ \text{or } z_{SS}(x, y).s = ZS_Boundary \text{ then} \\ ZS_Out \\ \text{else} \\ ZS_In \end{array} \right\} \quad [\text{Equation 3.24}]$$

Figure 3.36 Shows the new status map.



Figure 3.36 Status array after simplification.

3.4.15 Processing the Scaling Image

The method for calculating pixel size was outlined in Section 3.3.6. The scaling procedure simply involves placing a dark circle of known radius in the image area and measuring its diameter in the x and y directions in terms of pixels.

The image processing to extract the dimensions of the circle is straightforward. A global threshold based on the mean value of the maximum and minimum intensity values in the scaling image is applied to produce a binary image. The coordinates of the edges of the circle are then determined by counting in columns and rows from the edges of the image until a column or row is found that has a dark pixel. This row (or column) is then chosen to have the x (or y) coordinate of the edge of the circle. The scaling factors relating a pixel to real world coordinates in mm is then calculated using Equation 3.4.

In order to conserve secondary memory space, the scaling images are not stored permanently by the system. The image is processed online at the time of data acquisition and only the scaling factors L_X and L_Y are recorded permanently.

3.5 System Implementation

The platform selected for implementation of the system was the Acorn Archimedes Microcomputer (25MHz Acorn RISC Processor) running a proprietary operating system called Arthur and using the Watford Electronics Video digitiser as a plug-in frame-store.

This hardware selection was a compromise between processing power, cost and ease of programming. The back image processing system was implemented completely in software using the C language.

The system was implemented as a layered architecture of software modules. Figure 3.37 shows a simplified diagram of the structure.

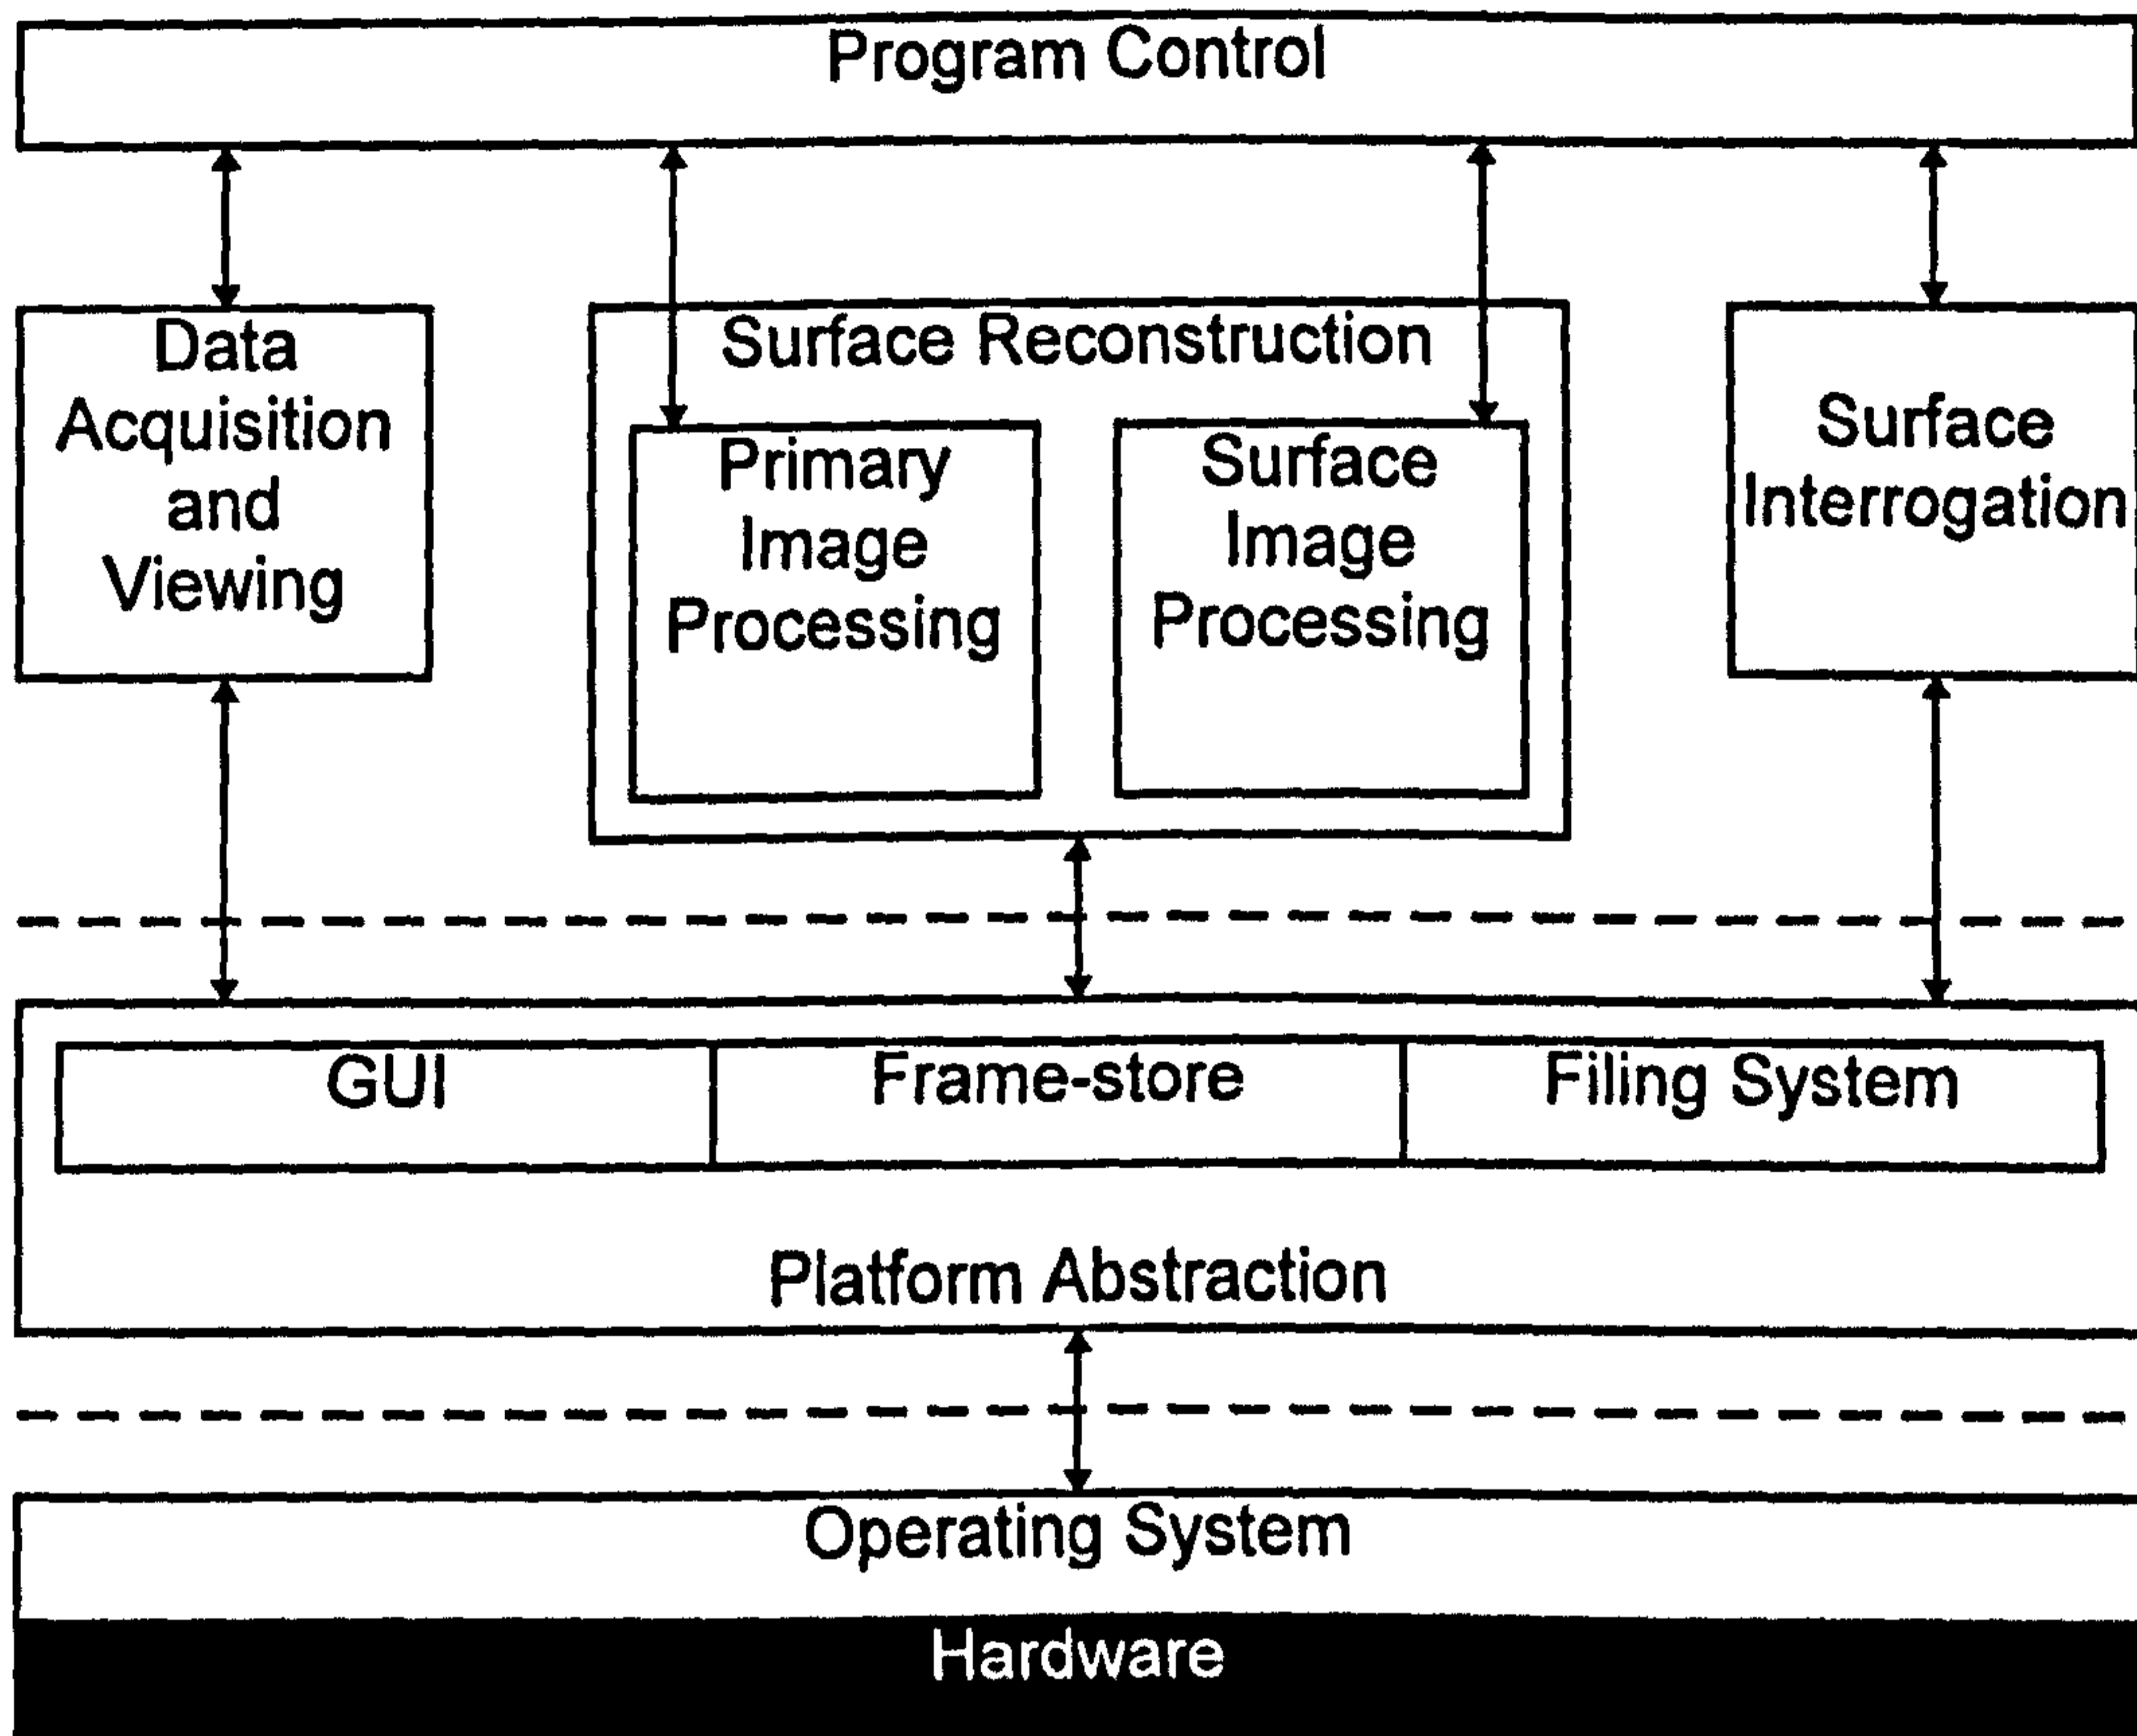


Figure 3.37 Simplified layered system structure.

A complete review of the system implementation is presented in Appendix II.

3.6 Discussion of the Proposed Method

For the proposed method there were two main areas of research: to investigate a new optical approach to back surface reconstruction; and to investigate new image processing algorithms that were specialised for the application and that would cope with the wide range of back surfaces and optical conditions that may be encountered in a clinical environment. Each of these is discussed in the sections that immediately follow. Inevitably, attention is drawn to the known shortcomings of the current method.

3.6.1 The Optical Subsystem

The optical method that was used dictated the approach to surface reconstruction and the measurement procedure. The algorithms for processing images are highly specific to the images produced by this particular optical subsystem and to the measurement of back shape. For example: the processing system is unable to deal with fringe patterns that do not have a reference fringe; also, the fringe tracking algorithm used in the ordering of fringes assumes that fringes will tend to curve downwards in the left and right sides of the image. As such, the method is specialised for the measurement of back shape. However, with a little modification, it could in principle be adapted for the measurement of other surfaces.

The optical geometry is particularly elegant. The projector produces diverging light so that the angle of illumination changes across the object (the back) and the plane of the grating is not parallel to the imaging plane of the camera or the reference plane. It might have been expected that the angle of illumination would have to be explicitly calculated for each fringe in the image and that the changing angle would then be incorporated into the calculation of z (depth) coordinates. In fact this is not necessary because the adjustment is inherent in the triangulation process.

The main practical objective for the optical system was that it should be easy to set up and use. In as far as is possible this objective has been realised. There are no precision or complex procedures in the set-up (except that it is desirable that the camera should be horizontal and this is currently achieved using a small spirit level). There is no explicit calibration of z coordinate values. The accuracy results for in-plane measurements presented in Chapter 4 suggest that depth measurements have the required accuracy. Furthermore, the lack of systematic errors in depth measurements suggests that there may be little value in a calibration exercise.

3.6.1.1 The Data Acquisition Problem

The main shortcoming of the measurement system as a clinical tool is the time taken to acquire the dataset for a measurement. This is identified in Chapter 4 to be the only serious area in which the system fails to meet the criteria specified in Chapter 1. Three

images must be acquired: the back surface image, the reference image and the scaling image for each patient. The criteria was that a measurement session should be possible in less than a minute and although the system is easily capable of coping with this acquisition speed, it cannot be realised because of the time that is spent by the operator positioning the patient, reference plane and scaling disc.

3.6.1.2 The Framework and Mounting of Components

The frame that houses the reference plane, and in which the patient stands, ensures that the patient can be placed in a reasonably standard position for image acquisition, and that the depth of field for measurement rarely exceeds 20cm and includes the reference plane. The camera and projector are mounted on a tripod some distance in front of the frame. The mounting of the camera and projector on a tripod is less than ideal. Firstly the tripod, when fully extended for measurement is bulky and inconvenient in a clinical environment. It can be easily knocked and its position may alter slightly during an extended session as the legs of the tripod settle. To some extent this can be eliminated by making sure that all adjustable components are tightly locked into position but some drift is inevitable. Secondly, the projector is a heavy object that is mounted on a platen at the top of the vertical arm. The slightest knock of the tripod causes the arm to sway and the projector with it. Even the slide exchange mechanism of the projector which is used to remove the grating for acquisition of the scaling image will send the projector into a brief harmonic motion which lasts several seconds before the system is stable enough for measurements to be taken. A more rigid mounting arrangement is required.

3.6.1.3 The Projector and Gratings

The 35mm slide projector provided a simple and cost effective means of projecting the fringe pattern onto the object. The cost was at least an order of magnitude less than most industrial or commercial metrology specification white light projectors. The provision of an auto-focus lens on the projector was beneficial because it improved the time taken for acquisition. One key advantage of using a standard 35mm slide projector in the development programme was that, because 35mm slides can be easily produced by photographing A4 sized patterns on paper, a variety of slides with different patterns

(mainly different line spacings) could be produced at negligible cost. However, there are shortcomings of using a standard slide projector. The intensity distribution across the projection area is not uniform, being brighter in the centre than at the edges, although this should not be a problem because the local area nature of the image processing should eliminate the effects of this feature. Also, the quality of gratings produced photographically can never be as good as for the gratings ruled on glass that are often used in optical applications: the translucent areas of the grating are not completely bright and the opaque lines are not completely dark.

3.6.1.4 The Camera and Frame-store (Digitiser)

The camera was a standard CCTV monochrome camera that produced an analogue signal that was subsequently digitised and the image placed in the conventional memory of the microcomputer. The digitiser limited the system in two ways. Firstly, the resolution of the digitiser was 512 pixels in the x direction by 256 pixels in the y direction. If a camera with a higher resolution were used then it would be possible to project and resolve a greater number of fringes on the object and thus have a higher density of sampled data. Secondly, the digitiser was capable of only six bits of intensity resolution. This limits the accuracy with which fringe edges can be located (recall that the actual observed profile is similar to a sinusoid). A higher specification framestores would be desirable.

Setting the focusing, focal length and aperture of the camera is currently performed manually by the operator who adjusts the setting by looking at the image produced on the computer screen prior to grabbing the first of the set of images. Fortunately, the ability of the system to produce a surface reconstruction is not *critically* dependent upon the optimisation of these settings, however, performance (accuracy and reliability of the thresholding and ordering algorithms) will be degraded as image quality decreases. Automated or semi-automated support for the optimisation of the camera settings would be desirable.

3.6.2 The Image Processing Subsystem

The overriding requirement for the image processing subsystem was that it should be reliable in that it should be capable of producing a surface reconstruction in a broad range of circumstances including the ambient lighting levels in the environment, unusual body shapes and different reflective properties of the back. Investigating a method that could be reliable in a clinical environment is significantly more challenging than for a laboratory investigation in which the object and environment can be carefully controlled. In almost every case where there was a conflict between providing increased reliability or some other improved system feature or function, the matter was resolved in favour of increased reliability. The approach to reliability may be unusual in a measurement system: often the overriding requirement is for high accuracy. Reliability was not suggested as a specific criteria for the system in Chapter 1 because it was assumed to be axiomatic that any clinical tool must have this property.

Failure of the system to produce a surface reconstruction is rare. Of over two hundred measurement sessions (patients) the system has only failed in a handful of cases. In the majority of these the failure was due to operator error, for example, failure to place the scaling disc in the scaling image or using an aperture setting at the camera which produced heavy flooding of the CCD array. There are also other conditions that may arise which whilst they do not actually cause system failure, do limit the usefulness of the system. For example, occasionally the patient will have a winged scapula, that is a scapula that protrudes unusually severely from the back: shadowing occurs and the tracking of fringes fails. However, this will simply mean that the surface data for a region of the back is not available, and does not really constitute system failure.

3.7 Summary

Chapter 2 concluded that, of the fundamental approaches that were considered, no single method was completely satisfactory and therefore new methods should be investigated. This chapter has described research into a new method for reconstructing the surface shape of the back.

A new structured light method was proposed that overcomes many of the problems inherent in existing techniques. A structured light pattern was devised that could be easily and unambiguously interpreted. The optical principle was described. Image processing algorithms, specific to the optical method and the application, to reconstruct surface shape from the structured light pattern were investigated and optimised to provide satisfactory results in the wide range of conditions that would be expected in a clinical environment. Finally, the proposed system was discussed and its shortcomings were identified.

The surface reconstruction generated by the proposed method is suitable for further processing to extract clinically significant parameters that relate to scoliotic deformity. The research into the extraction of these parameters is included in Chapter 4, which also evaluates the method in terms of the criteria for a successful system that were specified in Chapter 1.

Chapter 4

System Evaluation and Clinical Application

4. System Evaluation and Clinical Application

4.1 Introduction

Chapter 3 outlined the optical and image processing principles of the new method for reconstructing back shape that is proposed in this work. The method was implemented in a working clinical system. This chapter evaluates the clinical system and describes how clinical information can be obtained from the reconstruction.

Chapter 1 concluded with a set of criteria for a practical automated system. To evaluate the system as a whole it is appropriate to return to these criteria and to determine whether they have been fulfilled by the proposed system. To summarise briefly, the proposed criteria for a system were that it should:

- i. measure the patient in the standing position.
- ii. be non-contact.
- iii. have a target accuracy of 5mm.
- iv. measure parameters of clinical significance.
- v. produce repeatable measurements.
- vi. acquire and process data in less than one minute.
- vii. provide long term storage of body shape.
- viii. be simple to use by a clinician.
- ix. cost less than £2000.
- x. be portable or collapsible.

Each of these criteria is addressed in the immediately following sections. The summary at the end of the chapter returns to the list of criteria and determines whether each of the criteria has been satisfied.

4.2 Measurement of the Patient in the Standing position

Clearly the method proposed in Chapter 3 satisfies the requirement that the patient be measured in the standing position. The challenge was to place the patient in a position that was as standard as possible so that the results of successive imaging sessions over perhaps a series of months or years could be validly compared. In the earlier research,

the patient simply stood upright and was positioned according to the clinician's directions. It became evident that this arrangement was unsatisfactory because it was difficult to place the patient in the same or similar position for a separate imaging session at a later date.

A number of options were considered to address the positioning problem:

i. Rotating and/or shifting the dataset.

One option to help standardise patient position might have been, after surface reconstruction, to have modified the reconstructed data so that all depth (z) measurements are relative to some plane defined by the data set itself. For example, a least-squares fit of the surface to a plane from which depth values could subsequently be re-calculated would be a convenient option. This option was rejected on the grounds that depth would no longer be true coordinates measured from a known and standard reference plane.

ii. Fixing skeletal landmarks.

It would have been possible to devise an extension to the frame in which the patient adopts a particular position by contact with protruding supports. For example, it would be possible to place the shoulders in contact with an adjustable (for height and width of the patient) mount so that all patients are measured in a similar position. This option was rejected on the grounds that it may force the patient to adopt a position that may be different from their normal standing position.

iii. Opaque markers: a body orientated reference plane.

By placing a minimum of three markers on the skin surface, perhaps at the C7 vertebrae and at the "dimples of Venus" to the left and right of the S1 vertebrae, a body centred reference plane can be defined from which depth values can be calculated. Image processing algorithms could be used to identify the markers and construct the reference plane. This option was rejected on two grounds. First, as for (i) the measurements are no longer taken from a known and standard reference plane. Secondly, the presence of markers on the skin will effect the projected fringe pattern.

Since the markers are likely to be of approximately the same size as the fringe spacing it would be extremely difficult to remove their effects on the intensity profile through local area methods. The alternative would be to acquire a separate image of the patient with markers on the skin and then to remove the markers for fringe pattern imaging: the problem of patient movement between the two images would then arise.

iv. A foot-plate.

The position in which a patient stands will be primarily governed by the position of their feet. A simple foot-plate, shown in Figure 4.1 was designed to help the patient stand in a standard position. The foot-plate fitted into the frame.




Figure 4.1 Foot-plate design.

This simple foot-plate allowed the patient to be positioned in a relatively standard position and avoided the problems of alternative options. It also offers the patient a stable standing position with feet apart which reduces sway.

There are disadvantages in measuring the patient in the standing position. Conventional methods such as the measurement of ATI and STAS (see Chapter 1) measure the patient in the forward bending or prone position and it is not possible to reproduce clinical parameters of scoliotic deformity that rely on this position. However, it is possible to measure equivalent parameters for the standing position.

Clearly, to compare parameters measured in the standing position with those measured in, say, the forward bending position is not appropriate, but if *similar* parameters can be measured then they are more likely to be accepted by the clinical community than if completely new parameters are devised (such as the volumetric asymmetry of the back)

that have no record of clinical use in either the routine clinic or research programme. The measurement of clinical parameters is addressed later in Section 4.5.

4.3 A Non-Contact System

The optical system that has been devised is completely non-contact. The patient simply steps onto the foot-plate and stands erect while the operator adjusts the focal length of the camera lens and presses a mouse button to grab the fringe image. No markers of any kind are placed on the back and the body is completely free of constraint.

4.4 Accuracy

It is difficult to measure the accuracy of a system for the measurement of human body shape using a patient or other individual. To do so would require that the true body shape be measured first using some other method with known and high accuracy, and then measured using the system *without any movement of the patient between the two measurements*, and finally the two sets of measurements compared to produce error values from which accuracy can be inferred. This is not possible because a patient cannot remain stationary for the required time. However, it is possible to measure the fundamental accuracy of the measurement technique used in the system.

4.4.1 An Experiment to Measure Accuracy

The accuracy of the system was measured using a series of plane measurements at different depth (z) positions. Usually the range of depth values that will be measured by the system, the depth of the area of measurement of the back, will not exceed 200mm. The objective of the accuracy experiment was to investigate the accuracy of the system in measuring depth coordinates in the depth range that would normally be encountered in the measurement of patients.

The system was set-up as shown in Figure 4.2. The frame, housing the reference plane, was positioned at a horizontal distance of 2340 mm from the camera and projector. The projector was mounted at a distance 1090mm above the camera. This gave an angle of 25° between the optic axes of the camera and projector at the object. These settings are typical for the system as it is used in a clinical environment. The experiment was

conducted in a laboratory under the type of background lighting conditions that are typically available in a clinical environment. The experiment was conducted in daylight and the curtains were drawn to exclude excessive light.




Figure 4.2 Configuration for accuracy measurement.

The apparatus was placed on a polished floor and a line was drawn with a sharp pencil on the floor parallel to the bottom of the reference plane. Lines parallel to the reference plane were also drawn at 10mm intervals from the reference plane in the direction of the camera, to a depth of 200mm. The reference plane was translated 10mm forward (by physically shifting the frame) and aligned with the 10mm line drawn on the floor and recorded by the system. In principal this measurement is equivalent to the measurement of the patient except, of course, here a flat plane is being measured. The plane was then moved back to the normal (0mm) position for the reference plane and a normal reference image was acquired. A 200mm dark disk was then placed against the

reference plane in the absence of fringes and a scaling image was acquired and processed in the usual way to calculate the pixel scaling factor.

This procedure was repeated with the plane being placed at 20mm, 30mm, ... , 190mm, 200mm. Separate reference plane (0mm) data were acquired for each measurement. There are two reasons for conducting the measurements in this way. First, the system is configured to operate in precisely this way for a patient (acquire patient surface image, then acquire the reference image and then acquire scaling disk image: in that order). Secondly, the system may alter slightly between measurements if any part of the apparatus is accidentally knocked so that it is desirable to collect reference data soon after each measurement.

4.4.2 Results of the Accuracy Experiment

The images acquired in the experiment were processed by the system in fully automatic mode: the image editors were not used to correct any fringe tracking errors or to manually intervene in any way. The pixel scaling factors measured by the system were $L_x=1.2$ and $L_y=1.8$ so that, for example, in the y direction 1 pixel represented 1.8mm at the reference plane. Figure 4.3 is a graph of the true, positioned, z value for each plane measurement versus the mean measured value for the plane. The error bars show the worst case deviations of each measurement from the mean.

The graph shows the approximately linear relationship between the positioned values and measured values that would be expected from a measurement system. It should be noted that although the frame was positioned as accurately as possible, the frame is a large object, approximately 2m in height, it was not possible to guarantee that the frame did not flex, perhaps by a millimetre or more, during positioning or measurement.




Figure 4.3 Graph of true versus measured values for the accuracy experiment.




Figure 4.4 Error in the mean measurement versus z position.

The error in the mean measurement, that is the difference between the mean measurement of the plane and the true position of the plane is shown in Figure 4.4 for each of the plane positions. The graph shows no evidence of a systematic error in the measurements. The mean error in the measured position of the plane was 1.01mm. The mean standard deviation from the mean measured value for measured points (pixels) in the image across all measurements was calculated to be 0.87mm.

4.4.3 Comments on Accuracy

The automated system investigated in this work is capable of measuring depth coordinates for flat surfaces with a mean error of approximately 1 mm. Given that the mean standard deviation of all measurement sets was 0.87 mm then, *if* the distribution of errors is assumed to be normal it may be said that 68% of the measured values for z coordinates will lie within $1.01\text{mm} + 0.87\text{mm} = 1.88$ mm of their true values [similarly 95% of measured z values will lie within 2.75 mm of the true value, 99.7% within 3.62 mm and so on].

Obviously, these results are for measurements taken from a flat plane. It should be expected that the accuracy with which a body surface can be measured will be poorer because the possibility of errors in thresholding and indexing, and spurious errors at the edges of the back surface. However, it is not uncommon to quote accuracy based on measurement of a flat plane for this type of measurement system (Frobin 1988, Turner-Smith 1988).

4.5 Measurement of Parameters of Clinical Significance

The two principal clinical parameters for measuring surface shape that were described in Chapter 1 were the Angle of Trunk Inclination (ATI) and Trunk Asymmetry Score (TAS). It was observed that these parameters for measuring the extent of scoliotic deformity are usually measured with the patient in the forward bending position, and sometimes with the patient in the prone position.

The reasons for conventionally measuring the patient in the forward bending or prone positions are twofold. First, the forward bending position accentuates the scoliotic deformity in many patients, making the back hump more obvious. It is also observed

that whilst the forward bending position may be a useful diagnostic and monitoring position, it may not reflect the scoliosis in the normal standing position, of which the patient is most aware. The second reason why the forward bending position is used is because of the limitations of conventional measuring instruments. Any measure of angle or asymmetry must have a datum line or plane from which measurements are taken. In general, to form such a line or plane is difficult without careful set-up and positioning to the patient. A special case, however, exists when measurements are to be made relative to a horizontal reference line or plane: the pull of gravity can be used directly in devices such as the Scoliometer or with the aid of a spirit level coupled to the Formulator Body Contour Tracer.

In the automated system that has been described in this investigation, the constraints of these devices are largely removed and it is possible to measure the patient in the standing position. The datum line or plane is provided through the reference plane of the system.

4.5.1 Vertebral Positioning

Both the ATI and STAS parameters are required at each vertebral level. A problem arises in the automated system because the position of the individual vertebrae, or even the line of the spine in the image, are not known.

Before ATI or STAS can be measured from the surface reconstruction, the positions of the vertebrae must be known. There is no obvious method for doing this based on the surface reconstruction itself. In some cases the spinous processes of the vertebrae protrude slightly from the back and it is conceivable that an automated system might be capable of detecting these small bumps and automatically indexing them as vertebrae with little or no manual intervention. However, such a strategy would work for only a small subset of patients and would be difficult to implement reliably.

The approach taken in the system described here is to request manual intervention from the user. Initially, the source image of the patient was displayed on the screen and the clinician was requested to identify estimated positions of vertebrae in the image. Vertebral identification was performed after surface reconstruction. The clinician

responded to a series of requests from the system to click the mouse cursor over the appropriate position in the image for each vertebrae. Whilst this option for vertebral location was satisfactory, it was time consuming for the clinician.




Figure 4.5 Vertebral positions as displayed by the system.

A variation on this technique was implemented to reduce the time taken in identifying vertebrae. The clinician is requested to place the mouse cursor and click over first the S1 (Lowest) vertebrae and then the C7 (highest) vertebrae. He then draws a line, also using the mouse, joining S1 to C7 along the estimated line of the spine. The system then estimates the position of the vertebrae along the line of the spine automatically. The relative sizes of vertebrae are included in the calculation of vertebral positions: the data for relative vertebral sizes was acquired by measuring size of vertebrae on a dry bone specimen at the University of Liverpool.

Figure 4.5 shows the vertebral locations identified on a scoliosis patient. The positions are stored by the system for future use and subsequent data interrogation sessions give the user the option of using old vertebral positions or identifying a new set.

4.5.2 Measurement of Angle of Trunk Inclination

The objective for measuring the standing ATI was to simulate the use of the Scoliometer, except of course in the standing position. The Scoliometer is 200mm in

length. A simple algorithm performs a least squares fit of a 200mm long line to the surface shape at a particular vertebral level as shown in Figure 4.6.




Figure 4.6 Measurement of ATI.

The midpoint of the line coincided with the position of the appropriate vertebra. The angle between the line of best fit and the reference plane, θ_{ATI} , was calculated from the gradient of the line and taken to be the ATI value at that spinal level. Figure 4.7 shows the multi-level ATI map for the patient and vertebral positions shown in Figure 4.5.

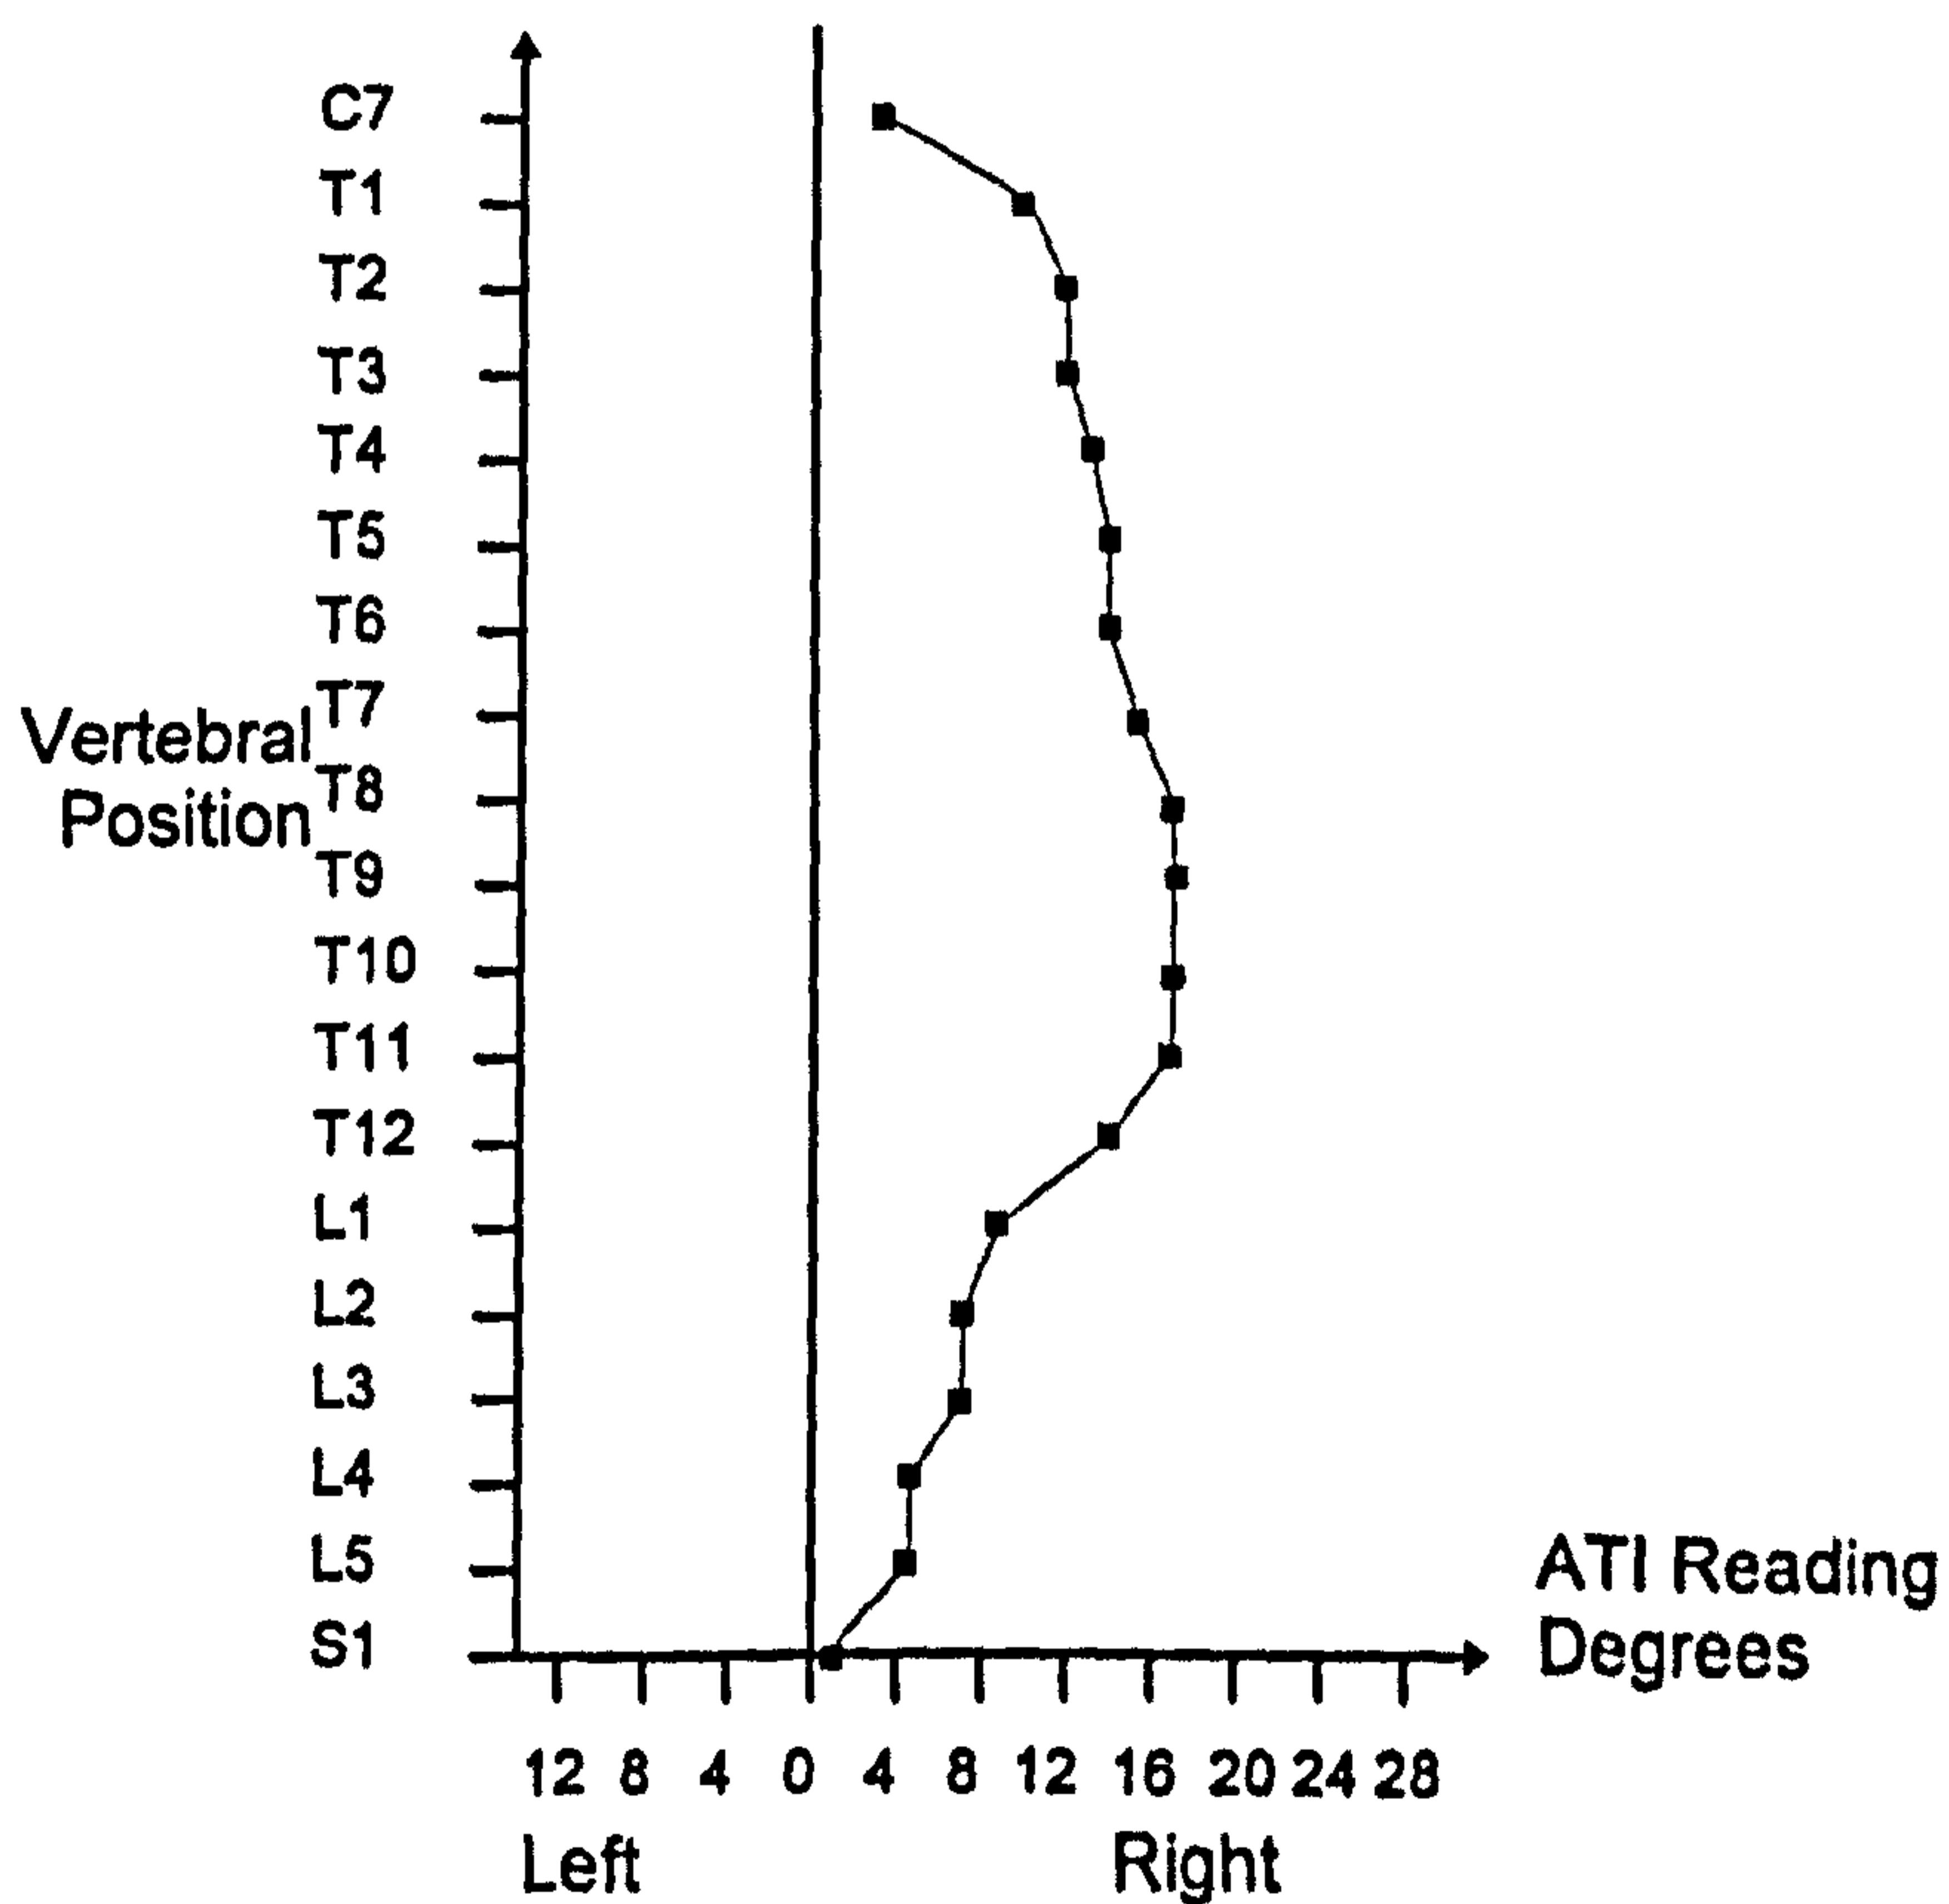


Figure 4.7 ATI map for the patient shown in Figure 4.5.

When interpreted in this manner, in addition to being measured in the standing rather than forward bending position, the ATI is measured slightly differently than might be expected from the strict definition of ATI given in Chapter 1, where the angle measured is the angle of the tangent to the rib hump on the left and right of the back to the datum line. In this case a line of best fit is used because it more closely simulates the activity of placing the Scoliometer onto the back surface.




Figure 4.8 ATI values obtained for the three sets of readings compared.

In order to compare, to some extent, the measurements obtained using the automated system with those obtained using conventional methods, a plaster cast was taken of a patient at the Royal Liverpool Children's Hospital. The patient had been admitted for spinal surgery with a right sided adolescent idiopathic scoliosis. The apex of the curvature, measured on a radiograph was at T8 and the Cobb angle (see Chapter 1) was 75°. The patient was placed in the prone position and multiple vertebral level ATI measurements were taken in the conventional manner. Without moving the patient, a plaster mould of the patient was made. A thin plaster shell was formed using plaster

impregnated bandage. The shell was then removed and the patient released. The mould was then used to make a plaster cast of the patient's back shape.

After setting and drying the plaster-cast, it was placed on a horizontal surface and multiple vertebral level measurements were again taken with the clinician estimating the vertebral positions on the plaster cast. Lastly, the cast was mounted in the system, the surface was reconstructed and measured using the semi-automated method outlined in the previous section. Figure 4.8 shows the three sets of results obtained, and they appear to be in good agreement. The comparison of results simply verifies the measurement of a clinical parameter for the same surface.

4.5.3 Measurement of Standardised Trunk Asymmetry Score

The other parameter of clinical interest in this investigation was STAS. The general method for measuring the parameter was outlined in Chapter 1. The method relies upon subtracting paired surface depth values on either side of the vertebrae (or midline) and processing the sum of the differences using a simple formula. The separation of the divisions for depth coordinate pairs is calculated on the basis that they occur at intervals of 5% of the lateral chest diameter.

An estimate of the lateral chest diameter for each vertebral level in the image can be obtained from the status array described in Chapter 3. The transition from the surface to the background is represented by a change of status value. The distance across the image is measured in this way.

The depth coordinates at either side of the vertebrae (the midline) are sampled from the surface. For each sampled depth coordinate, the depth values of the appropriate pixel and its eight neighbours are averaged to reduce noise in the result.

Figure 4.9 shows the STAS readings produced for the patient and vertebral positions shown in Figure 4.5.

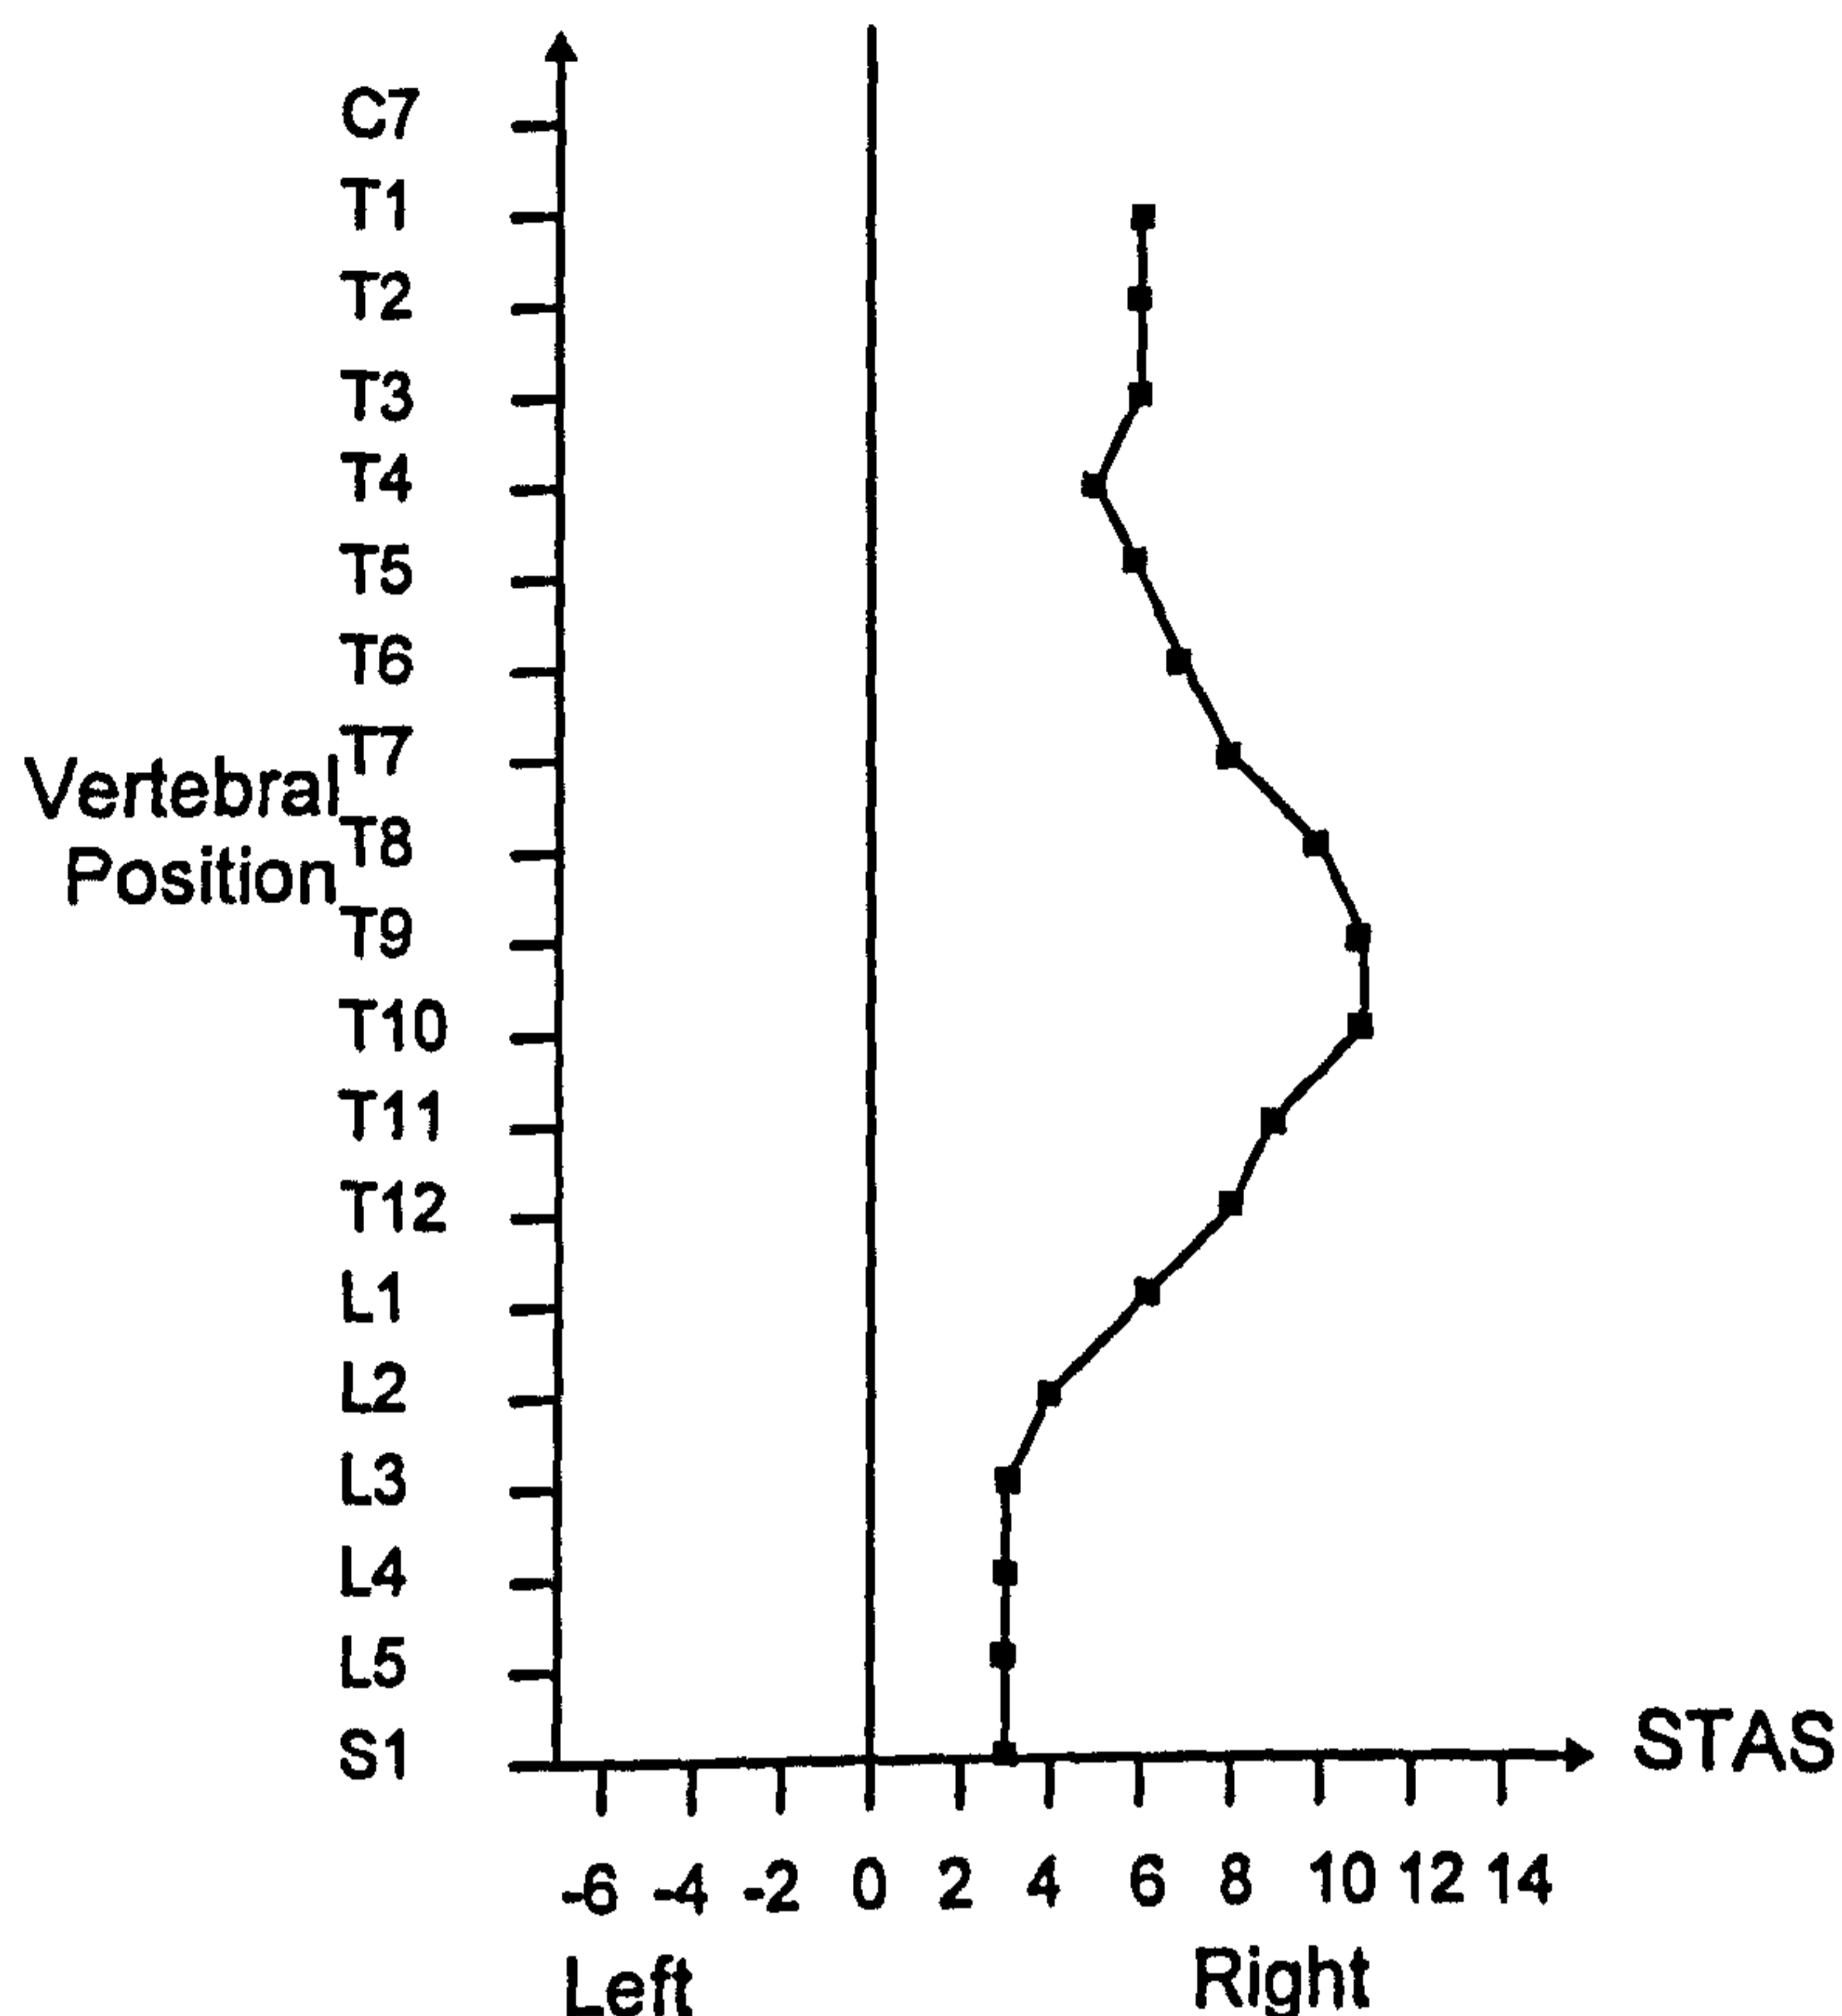


Figure 4.9 STAS readings produced by the system for the patient from Figure 4.5

4.5.4 Provision of Other Clinically Useful Information

In addition to the clinical parameters ATI and STAS the system is capable of producing other clinical information, that whilst not a measurable parameter, is of interest to the clinician. Projections of the three-dimensional dataset such as that shown in Figure 4.10 provide a useful overview of the surface and verify that the three-dimensional data at least looks correct (and increases user confidence in the system).

An alternative method of viewing the surface is provided by contour maps such as those shown in Figure 4.11. The contour maps are of special interest to some clinicians because they are similar to the moiré images that are sometimes used in scoliosis clinics and research. Figure 4.11(a) is the normal back that was used in Chapter 3. Figure 4.11(b) is the contour map of the image used in Figure 4.5 in the discussion of vertebral positioning. The hump on the scoliotic back is clearly visible in the contour map.





Figure 4.10 Projection of a back surface.

The system is also capable of drawing coronal or sagittal sections of the surface, selected using the mouse on the source image, so that the clinician can see surface profiles along horizontal or vertical lines in the image. These additional surface representation tools are a useful, but not essential, addition to the system. All graphical data can be printed out and stored in patient records for later reference.



(a) A normal back

(b) A scoliotic back

Figure 4.11 Contour maps of produced by the system.

4.6 Repeatability

If the repeatability of the system were assessed on the basis of comparing the surface reconstruction that is achieved in successive, independent measurements of a patient then a useful measure of repeatability would be difficult. It is impossible to place the patient in the same position more than once. A volumetric difference between surface reconstruction measurements would have little value even if the distribution of the difference could be plotted. In addition, two measurements with a significant volumetric difference could produce the same, or very close, values for parameters of scoliotic deformity.




Figure 4.12 ATI values for two independent measurements.

A more useful measure of repeatability is to compare the values of ATI and STAS between successive measurements. These measurements are relatively independent of the swaying of the patient in the sagittal plane because they are measures of the asymmetry of the back in the coronal plane. Similarly, swaying of the patient or other different positioning in the transverse plane will simply affect the position of the patient

in the image. The motion that affects ATI and STAS most significantly is the rotation of the body about a vertical axis. This will produce the coronal asymmetry that is measured by these parameters.

Figure 4.12 shows a typical difference between two multiple vertebral level ATI measurements of the same patient taken a few minutes apart. A systematic deviation of 3-4 degrees is evident. This is because the patient will have been positioned at a slightly different angle for the second measurement. A similar systematic error is usually observed in the multiple vertebral level STAS curve.

Whether this type of repeatability error is acceptable clinically in patient management has yet to be determined. It may be that there is value in standardising the standing ATI and STAS measurements so that readings represent the distribution of the curve about the mean. A comprehensive statistical study of repeatability of the measurement of clinical parameters has not yet been undertaken.

4.7 Speed of Data Acquisition and Processing

The time taken for the acquisition and processing of shape data is an important factor in evaluating the system as a clinical tool. There are two considerations: first, the speed with which the system can theoretically acquire and process data; and second, the effective speed with which the system can provide results which includes the time taken for the clinician to position the patient, reference plane and scaling object, and the time for interactive processing to extract clinically significant parameters.

4.7.1 Speed of Data Acquisition

Because the frame-grabber acquires an image in a single frame cycle (1/25s), the speed of actually acquiring the image of fringes on the patient is negligible. Of greater importance, however, is the time taken for the clinician to position the patient for measurement, acquire the image of the patient and also to acquire an image of the reference plane and the scaling image.

In the original system design, the acquisition of the reference and scaling images was separated completely from the acquisition of the back surface image. The collection of

the reference and scaling images was applied after the collection of the first back image of an imaging session and thereafter was only applied if the geometry of the system changed (e.g. if the tripod holding the projector and camera was readjusted for height or after accidentally being knocked) or the focal length of the camera lens was changed (e.g. increased for a small child) or if the lens of the projector was re-focused. For each back image, the clinician had to select the reference/scaling dataset which was identified by filename.

The separate acquisition of back and reference/scaling images was unsatisfactory because:

- i. In practice, either the geometry of the system or the focal length of the imaging lens was altered for almost every patient so the advantage of not having to acquire the reference/scaling images for each patient was minimal.
- ii. The clinician had to actively recall the appropriate reference dataset and this was confusing in a busy clinical environment, especially when the processing may be performed later in the day after all acquisitions had been made.

In the current system, the clinician is prompted to position the patient, then the reference plane and finally the scaling disk for *each measurement*. The system then handles the dataset as a complete unit and the capacity for human error is much reduced and the handling of data by the system is much simplified.

The complete acquisition of data from the moment of arrival of the patient (that is the adjustment of the system and the collection of the back surface, reference and scaling images) typically takes 5 minutes per patient, but this is reduced if the clinician is an experienced operator or if little adjustment is required.

4.7.2 Processing Time

The complete processing cycle to reconstruct the back surface shape typically takes between 20 and 30 seconds. The processing time is not constant because it depends upon the number of fringes in the images.

Table 4.1 shows an approximate breakdown of the times for the major processing tasks in the system for the image that was used in the description of the processing for the method proposed in Chapter 3. Some algorithms are invoked twice, once for the back image and then again for the reference image. The values are rounded to the nearest second: this was a limitation of the function library used to time the processing. Whilst these values are not very accurate, they do provide an indicator of the distribution of computational task and might be used to identify which algorithms would benefit most from improved efficiency or alternative approaches.

| Algorithm | Time (s) | Invoked | Total time (s) |
|-------------------------|----------|---------|----------------|
| Contrast stretch | 3 | 2 | 6 |
| Thresholding | 6 | 2 | 12 |
| Edge extraction | 1 | 2 | 2 |
| Point filter | 1 | 2 | 2 |
| Fringe ordering | 1 | 2 | 2 |
| Construct fringe table | 1 | 1 | 1 |
| Calculate z coordinates | 1 | 1 | 1 |
| Interpolate | 2 | 1 | 2 |
| Status organisation | 1 | 1 | 1 |
| Housekeeping | 2 | 1 | 2 |
| Total | | | 31 |

Table 4.1 Approximate breakdown of processing time.

After surface reconstruction, the clinician must identify vertebral positions before clinical parameters can be calculated. Typically this will take an additional two minutes, but does not have to be performed immediately: it can be done later when the data is recalled.

4.8 Provision of Long Term Storage of Data

The system records almost all of the acquired data onto disk for future retrieval. The only data that is not recorded is the scaling image that is used to measure the x and y pixel scaling factors. Instead of recording the image, the scaling factors are calculated immediately after acquisition. To the user, this presents a small time delay after the acquisition of the scaling image whilst the image is processed to extract the scaling factors.

The total data size for a single measurement set is typically 230 kilobytes for the system as described in Chapter 3. The distribution of file size is shown in Table 4.2. The table represents the data sizes for the image used throughout Chapter 3 for describing the system. Other images will have similar file sizes.

| File | Size in bytes |
|------------------------------|----------------------|
| Back image file | 108,974 |
| Reference image file | 127,569 |
| Geometric configuration file | 32 |
| Vertebral locations file | 212 |
| Patient data | 34 |
| Edge location file | 30,262 |
| Total | 267,083 |

Table 4.2 File sizes used by the system for a typical dataset.

The reason that the image files will vary in size is that they are compressed. The compression factor for the files is not constant. The compression algorithm is supplied in executable form by the manufacturer of the frame grabber and details were not available on its methodology. However, the files could easily be stored uncompressed or using a different compression algorithm, and conversion of existing data is straight forward.

Although disk space on modern microcomputers is relatively cheap, the standard disk size for the microcomputers used in this investigation was only 40 megabytes. Of this, about 30Mb was available for storage of system data. This allowed space for 112 image data sets. In the short term space was not a problem, but would eventually become so. An archiving utility was incorporated into the system to allow the transfer of datasets to floppy disks should the need arise. The archiving utility also served as a general purpose backup tool and backups are occasionally performed in this way. The clinician needs no understanding of operating system utilities to perform the backup; a series of prompts instructs the user to enter the dataset name (hospital number) and then insert blank floppies as required. An additional facility of incidental value in the archiving procedure stores the data onto floppy disks formatted for the Microsoft MSDOS

operating system so that it can subsequently be transferred to almost any computer system.

It is conceivable that the size of the dataset could be significantly reduced if the reference image and/or the back surface image were discarded after reconstruction and the surface reconstruction stored instead. This option was not considered desirable because continued research to improve image processing or other system functions could not be applied retrospectively.

During processing, the system may produce two additional files that are included in the dataset. If the back image is edited using either the trim editor or the pixel editor (or both) then the edited version of the binary image is stored for subsequent retrieval. When the data set is processed later and the user opts to intervene in the processing cycle using an image editor, then the option of loading a previously edited image is available. Also, once the user has identified vertebral positions then these are stored in the dataset (see Section 4.5.1). In subsequent processing sessions the user is similarly presented with the option of either specifying new vertebral positions or using stored values.

4.9 Simplicity of Clinical Use

The simplicity of use of the system for the clinician was a basic design axiom that has heavily governed the investigation. The key aspects of the system that have made it simple to use for the clinician are:

i. System setup

The system can be easily set up in under half an hour. The details of the setup and measurement procedure are outlined in Chapter 3. The system does not need to be set up by clinical staff: set-up can be performed by almost anyone with only a minimum of training.

ii. Simple measurement procedure.

The measurement procedure has already been discussed in Chapter 3. In essence all that is required is the capture of three images: of the patient, the reference plane and

the scaling disk. These can be performed quickly and the absence of a complex calibration procedure such as is common in automated vision systems has been avoided.

iii. Provision of a graphical user interface.

The provision of a usable interactive GUI was a time consuming task, but one which has very much simplified the use of the system in comparison to the early laboratory development system. The user navigates a series of simple screen menus that both simplify the acquisition, reconstruction and interrogation of the surface and prevent the user from performing illegal operations such as attempting to extract parameters from a surface before it has been reconstructed. All the menu selections and interactive processing is performed using a mouse and pointer. The menu options available are outlined in Appendix II.

These three features have enabled the measurement system to be easily used in a clinical environment as a tool in the assessment of scoliosis.

4.10 Cost

The desire to minimise cost was a major contributing factor to the system design. The approximate cost of the current system at the time of purchase and construction is shown in Table 4.3.

| Item | Cost £ |
|----------------------|-------------|
| Computer system | 1200 |
| Camera | 400 |
| Camera lens | 200 |
| Projector | 100 |
| Tripod | 150 |
| Frame | 100 |
| 35mm Slide Projector | 100 |
| Total | 2250 |

Table 4.3 Breakdown of approximate cost of the system.

The total cost of the system, therefore, was approximately £2250 in 1990. However, since then the cost of an equivalent microcomputer and camera will have fallen.

4.11 Portable and Collapsible

Because the scoliosis clinic is not usually able to assign permanent space for an automated system to measure body shape the system needs to be collapsible. There are three large components in the system: the microcomputer and screen, the tripod with camera and projector and the frame with reference plane. The computer can easily be dismantled and stored away. The tripod can be collapsed and the telescopic legs withdrawn. The frame was constructed to be easily collapsed to lie flat. The triangular supports on the frame are secured using wing nuts that can easily be tightened and removed without any special tools. The delicate components of the system (camera, projector, grating etc.) are placed in a suitable carrying case modified for the purpose.

Whether the system is portable is debatable. The large frame and computer are the most bulky items and would be difficult to transport in all but the largest of vehicles. It is easily conceivable that the system could be implemented on a laptop computer in the future and this would eliminate the bulk of the computer. In addition, a lightweight, perhaps aluminium, frame with the required strength characteristics could probably be constructed which could be more easily collapsed and would fit into a smaller space.

4.12 Summary of the Evaluation Exercise

To evaluate the suitability of the system for the purpose for which it is intended it is appropriate to return again to the criteria for an automated system that were discussed in Chapter 1 and listed at the beginning of this chapter. Table 4.4 compares the criteria with actual system functionality.

Table 4.4 suggests that in general the system is satisfactory. There are two areas in which it has not satisfied the criteria. The most important of these is the data acquisition time. It should be noted that data acquisition time is the time required for the clinician to complete a measurement session for a patient as opposed to the time taken to acquire a fringe pattern from a patient (which is performed in a single frame cycle, 1/25s). With the system as it is currently configured it is difficult to see how this problem might be

overcome. However, if the system were configured to use a single reference and scaling image over an extended session (several back measurements) then the acquisition time would be very much reduced.

| Criteria | Desired Functionality | Actual Functionality | Criteria satisfied? |
|---------------------------------|--------------------------|---|--------------------------------------|
| Standing position | Standing | Standing | Yes |
| Non-contact | Non-contact | Non-contact | Yes |
| Accuracy | 5mm | 1mm mean deviation. | Yes |
| Data acquisition and processing | 1 minute | 5 minutes acquisition 20-35 s processing | No (acquisition) Yes (processing) |
| Long-term Storage | Capable of storage | Capable of storage | Yes |
| Simple to use by clinician. | Simple to use | Simple to use | Yes |
| Cost | < £2000 | £2250 (1990) | Yes (current cost) |
| Portable and Collapsible | Portable, Collapsible | collapsible | Yes (Collapsible) No (Portable) |

Table 4.4 Comparison of desired functionality with actual functionality.

The other criteria that has not really been satisfied is the portability issue. It was suggested that the system could be implemented on a laptop computer and that the frame could be constructed differently. These solutions would increase the portability of the system, again with some increase in cost.

4.13 Summary

This chapter has evaluated the new method for surface reconstruction that was proposed in Chapter 3 in terms of the criteria that were identified in Chapter 1. In addition, the methods that were used to calculate parameters of clinical significance for the working clinical system have been described. The method satisfies the criteria, with the exception that data acquisition is more time consuming than desired.

Chapter 5

Conclusions and Recommendations for Further Work

5. Conclusions and Recommendations for Further Work

5.1 Introduction

The measurement of back shape using optical and image processing methods presents a significant challenge to the investigator. Unlike many industrial objects, very little accurate *a priori* knowledge about shape can be assumed. Although the back has a basic form, there are many variations in surface shape and reflective properties. The problem is further compounded because the patient cannot stand perfectly still so it must be assumed that a slowly moving object is being measured.

This thesis has described the content and progression of a research programme into methods for automating the back surface measurement process and providing the scoliosis clinician with useful information about back shape. Chapter 1 specified the objectives for the investigation, explained the need for an automated system and identified a set of criteria for such a system.

Chapter 2 described the research into the initial methods of approach that were considered likely to be successful. The methods were investigated either through practical research or through a review of literature and consideration of theory. Some methods presented clear restrictions on the measurement process so that it would be difficult or impossible to satisfy the criteria identified in Chapter 1. Other methods were, perhaps unnecessarily, complex or demanding of the optical or computer hardware, and the complexity or demands were not always justified in terms of the performance of the method.

Chapter 3 proposed a new method for measuring the shape of the back. The optical method is certainly novel in its application to the measurement of back surface shape and may also be new in the more general area of surface reconstruction (a review of literature has not identified its use elsewhere). Many of the image processing techniques are standard image processing operations, but in almost every case they have been adapted for this specific application and as such represent new work.

Chapter 4 returned to the criteria for a successful method that were proposed in Chapter 1. The final phase of the research programme was to determine whether the criteria were satisfied by the proposed method. The majority of the criteria were satisfied and potential solutions to those which were not are discussed in the recommendations at the end of this chapter.

5.2 Conclusions

The primary purpose of this research programme was to investigate methods for reconstructing the three-dimensional shape of the human back and to propose a new and effective solution to the measurement problem. A number of initial approaches were investigated and although each of the methods, which were described in Chapter 2, had merits, no single approach was completely satisfactory. The research progressed to propose a new method reconstructing back shape.

The review of literature revealed serious shortcomings in stereophotogrammetry and phase stepping profilometry which, although both methods have been used for the measurement of back shape, limit their useful application. The principal problem with stereophotogrammetry was that, because the back is an apparently structureless object, it is almost impossible to locate corresponding points in the stereo image pair.

Phase stepping profilometry (PSP) requires three or more images to be acquired in succession, with a grating being shifted with precision between each image. The main shortcoming is that the patient must remain perfectly still between the acquisition of each of the three images. It would be difficult, if not impossible, for the patient to maintain a perfectly still standing position for the required time, and for this reason PSP was not investigated beyond the level of a literature survey.

Moiré contouring was of special interest in the research programme because non-automated moiré techniques are already in use in the clinical environment. Currently the clinician is able to simply observe a moiré topogram or photograph and analyse it, often subjectively, as a contour map of the surface. Preliminary research considered the theory of the technique and investigated a semi-automated method for reconstructing surface shape. Two principal problems with moiré contouring were identified. First,

although the contour fringes produced by moiré are lines of equal depth from a grating, the absolute fringe order is not known: a direction ambiguity, sometimes called the “Hills or valleys” problem exists. Secondly, moiré fringes have poor visibility and suffer from systematic noise and second order effects. For these reasons, the research into moiré contouring did not progress beyond a feasibility study. A publication (see Appendix III) arose from the moiré work.

The research into Fourier transform profilometry (FTP) was conducted in the laboratory. The application of FTP to back shape measurement had not been reported in the literature. Unlike PSP, FTP requires only a single image. However, FTP presented a heavy computational task that took approximately 20 minutes for phase demodulation. The long processing time rendered the technique unsuitable for inclusion in a system that must provide a surface reconstruction in less than a minute. The research into the application of FTP to measurement of back shape yielded a publication (see Appendix III).

The line raster contouring work proposed by Frobin (see Chapter 2) used a complex pattern of light lines of varying widths to identify key fringes that were used to index the fringe pattern. The slightly broader lines were observed in the resulting image to be brighter (higher intensity) than the narrower lines and were identified on this basis. Both the key fringes and other fringes were located by detecting the maxima of a Gaussian curve fitted to the intensity profile. Whilst this approach by Frobin produced a working clinical system, it may be more complex than is necessary for the back surface reconstruction problem.

The main novel contribution of this work lies in the research into a new method for reconstructing back surface shape. The optical arrangement borrows from the basic line raster contouring principle, but is modified for the method proposed in Chapter 3. Rather than projecting a pattern of narrow bright lines on a dark background (as proposed by Frobin) a pattern of light and dark lines of equal spacing was projected. This allowed both the leading and trailing edges of fringes in the intensity profile to be utilised and simplified the processing of the fringe pattern. The exception to this was that one of the fringes, the reference fringe, was significantly broader than the others

and acted as the key fringe from which all others could be indexed. In this way the need to identify key fringes based on their intensity was eliminated: instead the *separation* of fringes was used. Since the early dissemination of this research, the method appears to have been adopted by other workers (Sciandra 1995), although it is not clear whether the principle was conceived independently or as a result of this work, (the publications arising from the work in this thesis were not referenced by Sciandra).

The optical approach used in the proposed method required a specialised sequence of image processing operations to reconstruct surface shape. The use of thresholding to generate a binary image is a well established technique in image processing. The thresholding algorithm that was used in this work, however, was specially adapted for fringe patterns and further for patterns that are not full field in the image. The broad reference fringe in the source images required special treatment and led to a new fringe indexing technique that was specific to the new optical method. The new method was reliable and the research programme progressed to produce a working clinical system. Four publications arose from the technical aspects of the research into the proposed method (See Appendix III).

The secondary objective of this work was to demonstrate that the method was capable of producing parameters of clinical significance that relate to scoliotic deformity. In principle, it should be possible to measure almost any parameter that is conventionally measured from the surface shape of the back with the patient in the standing position. Two parameters were selected for this aspect of the work: Angle of Trunk Inclination (ATI) and Standardised Trunk Asymmetry Score (STAS) and the capability of the system to measure both was demonstrated. Both parameters have conventionally been measured with the patient in the forward bending position, although some workers have, for research purposes, measured the same or similar parameters with the patient in the standing position. Clearly, the proposed method was only capable of measuring these parameters in the standing position.

One important aspect of the measurement of parameters of scoliosis was that the measurements should be repeatable. The repeatability of the measurements is limited by the accuracy with which the patient can be repositioned for successive measurement

sessions. It was almost impossible to place the patient in exactly the same position twice, so that slightly different results are only to be expected. However, in the majority of cases it is found that re-positioning the patient leads to a systematic error in the resulting multi-level ATI or STAS curve. Both curves have the same shape, but are typically offset by a few units. This suggests that the patient is rotated slightly from the hips. Four publications arose from the clinical aspects of the research (see Appendix III).

The last objective was to devise an automated system based on the proposed method and to verify that the system was a usable clinical tool. Such a system was constructed and was evaluated in terms of the specific criteria that were identified for a useful clinical measurement tool in Chapter 1. Most of the criteria for a successful system have been satisfied. The patient is measured in the standing position, the system is non-contact, the accuracy (based on in-plane measurements) is satisfactory and subject to patient positioning problems, the measurements are repeatable. The system is easy to use by a clinician or other operator and provides long-term storage of data from which back shape can be reconstructed. The system could currently be assembled for less than £2000 and a portable version incorporating a laptop computer should be possible. The main shortcoming of the method is that the acquisition of data for a patient is longer than the target acquisition time, although this is a limitation of the speed with which the operator can position the patient and reference plane rather than a limitation of the method. The clinical system was installed at the Royal Liverpool Children's Hospital where it was used in routine clinical sessions.

5.3 Recommendations for Further Work

5.3.1 Fundamental Approaches

Several years have elapsed since the beginning of the research programme, the practical research work, and the production of this thesis. In that time there have been developments in the technology that is available to the researcher. There are two developments of importance. Firstly, the speed of affordable microcomputers has increased significantly so that the processing of images can be performed perhaps an order of magnitude more quickly. Secondly, the performance and range of features of

low-cost framestores and frame-grabbing devices has improved. It would be appropriate that any future work should first re-evaluate the range of methods that are available.

The possibility of using colour pattern projection, acquisition and image processing was not considered to be an option for this research because of the cost of colour acquisition hardware. However, colour image acquisition is now easily affordable. For a few hundreds of pounds, true colour cameras and frame-grabbers can be purchased. Such hardware could benefit the type of structured light system proposed in this work. One of the main challenges in the proposed system, that heavily influenced the research into the optical and image processing methods, was the need to compare the position of fringes on the back with corresponding fringes on a reference plane. This gave rise to the reference fringe from which all other fringes on both back surface and reference plane were indexed. The use of colour would eliminate the need for this kind of reference fringe. Each fringe could be allocated a unique colour and indexing would proceed on the basis of fringe colour. The need for fringe tracking would be eliminated and breaks in fringes due to shadowing would effect only the shadowed area and not the remainder of the length of the fringe.

Fourier transform profilometry was not pursued beyond initial feasibility research partly because of the significant processing time that is required for the calculation of the FFT, which must be performed twice during the processing cycle. The increase in performance of microcomputers in recent years would significantly increase the speed with which the FFT can be calculated and may well make the method suitable for inclusion in practical automated systems with the increased accuracy and sampling density that is associated with phase sensitive methods.

5.3.2 A Composite Tool for Assessing Scoliosis

Currently, the Royal Liverpool Children's Hospital use three objective methods for assessing the extent of scoliosis in a patient. First, radiographs are used in some circumstances to examine the underlying deformity, although they are not used for routine monitoring of the progression of the disease, and parameters such as Cobb angle

are measured. Second, the surface shape of the back with the body in several types of forward bending position is assessed and parameters such as ATI and STAS are measured. Lastly, of course, the patient is measured in the standing position using the new method proposed in this work. A composite clinical tool for the assessment of scoliosis would embrace all three methods.

To incorporate radiographs into the system is relatively straight forward. Indeed, this facility already exists in the development system, but no attempt has been made to explore the relationship between the radiograph and the surface information. Currently, a radiograph can be collected by simply directing the system camera at a radiograph mounted on a light-box and acquiring an image. If digital X-rays were available in electronic or magnetic form then incorporation into the system may be simpler still. A more advanced system would not only allow the storage of the radiograph but would also integrate the radiograph with the surface data, perhaps by overlaying the radiograph onto a surface reconstruction or contour map. In this way the clinician could see both the surface reconstruction and the underlying deformity in the same visual image. The poor quality of conventional radiographs means that automated processing to extract, for example, Cobb angle would be difficult, but semi automated processing, perhaps with the clinician marking appropriate points on the screen using a mouse cursor should be possible.

If the patient is imaged from behind in the forward bending position as per Adam's test (see Chapter 1) then the automatic measurement of ATI and STAS in the forward bending position should be easily possible. The acquired image would simply have the outline of the back. A simple edge detection algorithm could locate the back profile and the measurement of ATI reduces to the task of measuring the mean gradient of the surface over the midline of the spine. The measurement of STAS in the forward bending position is also relatively trivial.

In addition, other more complex imaging facilities could be incorporated into the system. For example, some scoliosis workers have noted a tendency for scoliosis patients to be tall and slender. It may be that this body shape is related to the condition and that there is value in collecting data regarding the entire shape of the body. It might

not be necessary to use three-dimensional reconstruction but instead to process the two-dimensional image to measure, for example, the ratio of the length of the arms and legs to the height of the patient and so on. Such information may be useful for research purposes.

5.3.3 Clinical Research

That the patient should be measured in the standing rather than forward bending position is a hypothesis rather than a statement of fact. The rationale behind measuring the patient in the standing position is that it is the shape of the back in this position that is of most concern to the patient. The debate on the value of standing versus forward bending measurement lies in the clinical domain and is therefore outside the scope of this investigation. However, the method proposed in this research, which measures the patient in the standing position may be a useful tool in collecting data that can be used to contribute to a debate over the issue.

It is doubtful whether the relationship between the surface shape of the back and the underlying structural scoliosis can ever be reliably determined. Indeed the subjective experience gained during the investigation suggests that no *universally applicable* relationship exists. Whilst the surface shape of the back is affected in almost all cases of scoliosis, not all underlying structural deformities exhibit the same surface deformity. Two children with very similar underlying deformities may exhibit very different surface shape deformities. Perhaps the solution may be simply to look for any abnormality in surface shape without attempting to categorise the underlying deformity on the basis of surface shape. The data acquired using the system resulting from this research could contribute to an investigation of the relationship between the surface deformity and the underlying structural deformity of the spine.

The parameters of clinical significance that were investigated in this investigation were, in engineering terms, very simple. The ATI is simply the mean angle between the surface and a reference plane measured over a distance of 200mm (the length of the Scoliometer). STAS is a simple measure of the asymmetry in a cross section of the back. However, these parameters have enabled objective assessment of scoliosis from

back surface shape. Other more ambitious parameters and asymmetry techniques could be investigated. For example, a composite parameter with weighted contributions from ATI-type gradients, STAS-type asymmetry scores, the position of the vertebrae and the overall volumetric asymmetry of back could be devised. It should, however, be noted that any new parameter of scoliotic deformity must be proven to be clinically useful if it is ever to find widespread application.

5.3.4 Improving the Set-up and Acquisition Time

One important shortcoming of the current configuration of the clinical system is the time taken for each individual patient measurement. The positioning of the patient for measurement and the acquisition of the back image does not present a major problem, and anyway it is difficult to envisage how this might be improved. The problem is that a fresh reference image and scaling image must be acquired for each patient measurement. Although the system can be configured to relate a series of back images to a single reference plane, this is not implemented in the current clinical system because, for the majority of back images, the magnification of the camera lens is altered. The simplest solution might be to adopt three or four standard magnifications and to acquire reference and scaling information for each at the start of a series of measurements. For each patient, one of the standard settings would be selected and the focal length of the lens adjusted accordingly. Whether the focal length could be set with sufficient accuracy, however, would have to be determined experimentally.

The camera has three settings which must be optimised for good system performance: focal length (magnification), aperture and focus. These are currently adjusted manually by the operator who views the image on the computer monitor. The settings for the focus and aperture are difficult. First, the pixelisation effect on the screen means that accurate focusing is difficult. Second, the operator sees an image that is modified in intensity by the monitor settings: changing the brightness or contrast of the image on the monitor may give the illusion that the aperture setting is acceptable when it is not. One solution to both of these problems would be to implement a focusing and exposure program. A small area of the image, which is continually being updated, perhaps the central vertical column, could be extracted and processed to produce data to support the

focus and aperture optimisation. In the case of focusing, out of focus fringes in the image will exhibit a lower contrast ratio (the ratio of the fringe intensity maximum to the intensity minimum over a complete cycle). In the case of aperture setting, the objective is to avoid clipping of the intensity profile by having the aperture too wide and to ensure that a reasonable proportion of the grey scale range is utilised. In a semi-automated system both of these could be implemented by requesting the operator to make adjustments, whilst a fully automated system would control the lens directly.

5.3.5 Improving the Fringe Ordering

Shadowing on the back surface due, for example, to a winged scapula can cause loss of data in the fringe ordering algorithm. The algorithm relies upon being able to track a fringe continuously from the central vertical column of the image to end of the fringe. Data is not only lost in the shadowed area but also for the remainder of the fringe length.

It should be possible to detect, recover and order fringes for which tracking has been lost by locating a shadowed area and attempting to work around it. For example, if a fringe appears to be significantly shorter than an adjacent fringe then the adjacent fringe could be tracked for some distance and then used as the starting point for a search algorithm to find a local binary edge that has not yet been allocated an index. If one is found then the tracking process can continue along that fringe using the index value of the fringe for which tracking failed. Status values would be used as indicators of the confidence in the indexing. For example, the same fringe is tracked as a result of conventional ordering then the index in use would have precedence over the estimated index. Alternatively, if two estimated indices are applied to a fringe then the algorithm could either arbitrate using other information or discard the remainder of the fringe.

5.3.6 Image Processing Performance

The time taken for the reconstruction of surfaces is small in comparison with the time taken for patient positioning and data acquisition, there is always value increasing processing speed. The micro-computer used during the investigation is now outdated and a more modern machine would improve the speed of processing significantly. The clock speed of the current development system is 25MHz. Currently (1996) the clock

speed of moderately priced microcomputers is 133MHz. Although the clock speed of processors alone is not necessarily a valid performance indicator for specific algorithms, an increased clock speed of this proportion, together with other improvements in system performance will significantly improve performance of the processing system. No specific claim for the increased performance using a modern microprocessor is claimed here, but it is realistic to expect that the current processing time for reconstruction would be reduced from 30 seconds down to a few seconds and perhaps less than a second.

Because the reference and back images are processed separately there is an inherent high level potential for parallelism in the image processing system. In simple parallel architecture with two processors, one processor could be handling the back image whilst another handled the reference image, and this would almost halve the processing time.

Many of the image processing algorithms would also benefit from parallel implementation at a lower level. At the very least, it is often possible to break the image into several portions with each portion being processed by a different processor. The thresholding algorithm is one example of a process that would benefit from this type of parallel processing.

The cost of a specialised image processing system to perform parallel processing would preclude its use in a system of the kind envisaged in this work. However, it is now possible to purchase a standard PC-type computer (typically using the Intel Pentium Processor) that has several processors. Such a machine might be suitable for use as the hardware platform for the system. Multiprocessing operating systems such as Microsoft Windows NT would ease the development of the software and provide inter-process communication facilities.

The objective of a real-time or near real-time implementation of the proposed method would be to assess the deformity dynamically. For example, the effects of breathing on back shape could be investigated. Alternatively, dynamic measurement could be used to adjust the positioning of the patient, perhaps to exhibit the worst case, best case or mean

value of some parameter. Early consultations with clinicians have suggested that this might be beneficial.

5.4 Other Applications

The automated system that was investigated for this project should have a number of applications either in its current or in modified form. In principle the system or a similar system is suitable for the measurement of many types of industrial or biomedical objects. Two offshoots of this work that are currently being explored are:

i. Patient positioning during radiotherapy setup and treatment.

The position of a patient in a radiotherapy machine must be carefully measured so that the machine can be configured to deliver the correct radiation dose to a localised internal volume. A system to measure body shape is being developed to tackle this problem. In addition to the initial patient setup, it is desirable to monitor the movement of the patient during treatment and a real-time or near real-time option is also being investigated.

ii. Wheelchair seat design for the disabled.

Many wheelchair users spend considerable time sitting in the chair and it is desirable that individual wheelchairs should be customised for the individual. The objective of the research will be to design a body shape measurement system that provide the data necessary to design the shape of the wheelchair seat.

5.5 Dissemination of this Research

A number of papers on this work have been published in appropriate journals and conference proceedings. The subjects covered by the papers include moiré contouring (1 paper), Fourier transform profilometry (1 paper), the proposed method for back surface reconstruction (4 papers), clinical application and experience of using the system based on the new method (4 papers). In addition there have been 13 short or abstract-only publications. A list of the research publications arising from this work is included as Appendix III.

References

References

- Adams W. (1865). "Lectures on the pathology and treatment of lateral and other forms of curvature of the spine." London: Churchill.
- Adler N.S., Csongradi J. and Bleck E.E. (1984). "School Screening for Scoliosis." Western Journal of Medicine, 141(5). p. 631-633.
- Atkinson J.T., Groves D.G. and Lalor M.J. (1982). "The measurement of wear in dental restorations using dual-source contouring." Wear, 76. p. 91-104.
- Bangemann M. (1993). "Europe and the Global Information Society: Recommendations to the European Council". Report to the European Council of Ministers in Corfu 1994, Brussels: European Commission DGXIII.
- Bremberg S. and Nilsson-Berggren B. (1986). "School screening for adolescent idiopathic scoliosis." Journal of Pediatric Orthopedics, 6(5). p.564-467.
- Brunie L. (1992). "Accurate 3-D modelization of the scoliosis by matching 3-D images (MR,CT) with 2-D Xrays." In: Dansereau J. (Editor), Proceedings of the International Symposium on 3-D Scoliotic Deformities (Montreal), Stuttgart: Gustav Fischer Verlag. p. 11-17.
- Bunnell W.P. (1984). "An objective criterion for scoliosis screening." Journal of Bone and Joint Surgery, 66A(9). p. 1381-1387.
- Burwell (1989). The 1989 Report on School Screening for Scoliosis: The Relation to Aetiology. Proceedings of the British Scoliosis Society.
- Burwell R.G., James N.J., Johnson F., Webb J.K. and Wilson Y.G. (1983). "Standardised trunk asymmetry scores." Journal of Bone and Joint Surgery, 65-B(4). p. 452-463.
- Burwell R.G., Patterson J.F., Webb J.K. and Wojcik A.S. (1990). "School screening for scoliosis - the multiple ATI back shape appraisal using the Scoliometer with observations on the sagittal declive angle." In: Neugerbauer H. and Windischbauer G.,

(Editors) Surface Topography and Body Deformity. Stuttgart: Gustav Fischer Verlag, p. 17-23.

Byrd J.A. (1988). "Current theories on the etiology of idiopathic scoliosis." Clinical Orthopedics, 229, p. 114-119.

Cobb J.R. (1948). "Outline for the study of scoliosis." In: American Academy of Orthopaedic Surgeons, Instructional Course Lectures. St. Louis: C.V. Mosby. p. 261-275.

Cooley J.W. and Tukey J.W (1965). "An algorithm for the machine calculation of complex Fourier series." Mathematics of Computation 19(90). p. 297-301.

Dangerfield P.H., Williams L., Dorgan J.C. and Rajasekeran S. (1992a). "Multiple spinal level measurement of body surface deformity in scoliosis". In: Alberti A., Drerup B. and Hierholzer E. (Editors), Surface Topography and Spinal Deformity. Stuttgart: Gustav Fischer Verlag. p. 212-216.

Dangerfield P.H., Scutt D., Dorgan J.C., Li Y. Pearson J.D. and Brodie D.A (1992b). "The effect of body position on the three-dimensional surface deformity of scoliosis." In: Dansereau J. (Editor), Proceedings of the International Symposium on 3-D Scoliotic Deformities (Montreal), Stuttgart: Gustav Fischer Verlag. p. 139-147.

Dansereau J. (Editor) (1992). Proceedings of the International Symposium on 3D Scoliotic Deformities (Montreal). Stuttgart: Gustav Fischer Verlag.

Darawulla J.S. and Rajan U. (1992). "Role of moiré topography in diagnosis and follow-up for scoliosis.", In: Alberti A., Drerup B. and Hierholzer E. (Editors), Surface Topography and Spinal Deformity. Stuttgart: Gustav Fischer Verlag. p. 209-211.

Debrunner H.U. (1972). "Das Kyphometer", Z. Orthop, 110. p. 389-392.

Dickson R.A., Stamper P., Sharp A and Harker P. (1980). "School screening for scoliosis: cohort study of clinical course". British Medical Journal, July 1980. p. 265-267.

Frobin W. and Hierholzer E. (1981). "Rasterstereography: a photogrammetric method for measurement of body surfaces." Photogrammetric Engineering and Remote Sensing 47(12). p. 1717-1724.

Frobin W. and Hierholzer E. (1983). "Automatic Measurement of body surfaces using rasterstereography." Photogrammetric Engineering and Remote Sensing 49(3). p. 377-384.

Frobin W. and Hierholzer E. (1988). "Real-time rasterstereography using a solid state camera." Proceedings of SPIE 1030. p. 28-34.

Goetze H.G. (1973). "Der rotationsindex bei idiopathischen thorakalsskoliosen." Z. Orthop., 111. p. 737-743.

Goodall A.J.E., Burton D.R. and Lalor M.J. (1994). "The future of range image measurement using binary-encoded pattern projection." Optics and Lasers in Engineering, 21. p. 99-110.

Halioua M., Krishnamurthy R.S., Liu H. and Chiang F.P. (1983). "Projection moiré with moving gratings for automated 3D topography." Applied Optics, 22(6). p. 850-855.

Halioua M. and Liu H.C. (1989) "Optical three-dimensional sensing by phase measuring profilometry." Optics and Lasers in Engineering 2. p. 182-215.

Halioua (1990). Commercial literature supplied by Dimensional Systems Inc., New York.

Hill D.L., Raso V.J., Durdle N.G. and Peterson A.E. (1992). "A video-based technique for trunk measurement." In: Alberti A., Drerup B. and Hierholzer E. (Editors), Surface Topography and Spinal Deformity. Stuttgart: Gustav Fischer Verlag. p. 52-56.

Ichioka Y. and Inuiya M. (1972). "Direct phase detecting system." Applied Optics, 11(7). p. 1507-1514.

Idesawa M. (1982). "Scanning moiré method and automatic measurement of 3D shapes." Applied Optics, 16(8). p. 181-186.

James J.I.P. (1976). Scoliosis. London: Churchill Livingstone.

Kawara, T (1979). "Corneal topography using moire contour fringes." Applied Optics, 18(21). p. 3675-3678.

Kshirsagar S. (1994). Ph.D. Thesis. Liverpool John Moores University.

Meadows D.M. (1970). "Generation of surface contours by moiré patterns." Applied Optics, 9(4). p. 942-947.

Moe J. (1978). Scoliosis and other spinal deformities. Philadelphia: Saunders.

Moreau, M.J., Hongxing J., Raso V.J. and Hill D.L. (1992). "The effects of TRSH instrumentation on scoliotic deformities." In: Dansereau J. (Editor), Proceedings of the International Symposium on 3-D Scoliotic Deformities (Montreal). Stuttgart: Gustav Fischer Verlag. p. 171-175.

Nash C. and Moe J. (1969). "A study of vertebral rotation." Journal of Bone and Joint Surgery, 51-A. p. 223.

Neugerbauer H. and Windischbauer G. (1995) "Survey on three years of practical work using moiré-topography for the early diagnosis of scoliosis, hump back, flat back and funnel chest." In: D'Amico M., Merolli A. and Santambrogio G.C. (Editors) Three-dimensional Analysis of Spinal Deformities: Technology and Informatics 15, Oxford: IOS Press. p. 303-306.

Pearson J.D., Atkinson J.T., Lilley F., Burton D.R., Goodall A.J., Hobson C.A., Kshirsagar S. and Search D.J. (1994). "Optical Measurement of Solder Bonds on Printed Circuit Boards." Proceedings of SPIE, 2340. p. 231-240.

Pelsue K. (1990). "Precision, post-objective, two-axis, galvanometer scanning.", Literature supplied by General Scanning Incorporated, Massachusetts.

- Perdriolle R. (1979). La Scoliose. Son etude tridimensionnelle. Paris: Maloine S.A.
- Pin L., Mo L.Y. Lin L., Hua L., Changsa, Hui H.P. Hui D.S., Chang B.D., Chang Y.Y., Yuan L. (1985). "Early diagnosis of scoliosis based on school screening." The Journal of Bone and Joint Surgery, 67-A(8). p.1202-1205.
- Pirodda L. (1982). "Shadow and projection moiré techniques for absolute or relative mapping of surface shapes." Optical Engineering 21(4). p. 640-649.
- Pun W.K., Luk K.D.K, Lee W. and Leong J.C.Y (1987). "A simple method to estimate the rib hump in scoliosis." Spine, 12(4). p. 342-345.
- Sciandra J., De Mauroy J.C., Rolet G., Kohler R. and Creach J.P. (1995). "Accurate and fast non-contact 3-D acquisition of the whole trunk." In: D'Amico M., Merolli A. and Santambrogio G.C. (Editors) Three-dimensional Analysis of Spinal Deformities: Technology and Informatics 15. Oxford: IOS Press. p. 81-85.
- SRS (Scoliosis Research Society) (1980). Spinal screening program handbook. Chicago: Scoliosis Research Society.
- Suzuki S., Yamamuro T., Shikita J., Shimozo K. and Iida H. (1989). "Ultrasound measurement of vertebral rotation in idiopathic scoliosis." Journal of Bone and Joint Surgery, 71(B). 252-255.
- Swain I.D., Daunt S.O'N., Robertson J.C. and Isherwood P.A. (1986). "Dynamic moiré topography for the determination of postural asymmetry." In: Harris J.D. and Turner-Smith A. (Editors), Surface Topography and Spinal Deformity. Stuttgart: Gustav Fischer Verlag. p. 77-83.
- Swain I.D., McKee V., Borelli P. and Robertson J.C. (1992). "The use of moiré fringe topography in the assessment and treatment of chronic back pain." In: Alberti A., Drerup B. and Hierholzer E. (Editors), Surface Topography and Spinal Deformity. Stuttgart: Gustav Fischer Verlag. p. 199-201.
- Takasaki M. (1970). "Moiré Topography", Applied Optics, 9(6), p.1457-1472.

Takeda M., Ina H. and Seji K. (1982). "Fourier transform method of fringe pattern analysis for computer based topography and interferometry." Journal of the Optical Society of America, 72(1). p. 156-160.

Thompson F. (1985). Stereophotogrammetry: a method of evaluating scoliosis and kyphosis, M. Ch. Orth. Thesis, University of Liverpool.

Turner-Smith A.R. (1988). "A television/computer three-dimensional surface shape measurement system." Journal of Biomechanics, 21(6). p. 515-529.

Willner S. (1984a). "Development of trunk asymmetries and structural scoliosis in prepuberal school children in Malmö: follow up study of children 10-14 years of age." Journal of Pediatric Orthopedics, (4). p. 452-455.

Willner S. (1984b) "Prevalence study of trunk asymmetries and structural scoliosis in 10-year-old school children." Spine, 9(6). p. 644-647.

Womack K.H. (1984). "Interferometric phase measurement using spatial synchronous detection." Optical Engineering, 23(4). p. 391-395.

Acknowledgements

Acknowledgements

The author would like to thank Professor C.A. Hobson (Director of Studies) for his support throughout the research programme and relentless effort to progress this work to a thesis. Without the enthusiasm and clinical support of Dr. P.H. Dangerfield (Supervisor, University of Liverpool) the application of the work to scoliosis, and the funding of the project, would not have been possible.

The academic and personal support of Dr. J.T. Atkinson (late), Professor D.R. Burton, Professor M.J. Lalor, Prof. R. Morgan, Dr. D.M Harvey, Dr. D.A. Hartley, Mr. F.Lilley and Dr. D. Groves is gratefully acknowledged..

I would like especially to thank the staff and patients of Alder Hey Children's Hospital under Mr. J.C. Dorgan for their support and willingness to participate in clinical experiments and trials, sometimes at great personal inconvenience.

This work was partly funded by the Arthritis and Rheumatism Council.

Lastly, I must acknowledge the support of my wife, Vivienne, for her patience during the writing of this thesis.

Appendix I
Glossary of Medical Terms

Appendix I

Glossary of Medical Terms

Adam's test. A method for observing the asymmetry of the back by viewing the patient from behind in the forward bending position.

Aetiology (Etiology). The science of the causes of a disease or condition.

Angle of Trunk Inclination (ATI). The angle between the horizontal plane and a plane across the posterior aspect of the trunk at the point or points of maximum deformity with the patient in the forward bending position. The ATI can also be interpreted to refer to a similar angle measure in the standing position.

Anterior. The front (of the body).

ATI. See *angle of trunk inclination*.

C1...C7. See *cervical vertebrae*.

Cervical vertebrae. A collection of seven vertebrae at the top of the spine. Usually indexed from the top using indices C1...C7. C7 has special interest because it occurs at the top of the back and can easily be located by palpation due to the existence of a large spinous process.

Cobb angle. A measure of the lateral curvature of the spine. Measured on spinal radiograph by drawing lines parallel to the upper border of the upper vertebral body and the lower border of the lowest vertebral body of a structural curve and measuring the angle between the perpendiculars to these lines.

Coccygeal vertebrae. See *coccyx*.

Coccyx. A group of four small vertebrae, the coccygeal vertebrae, at the lower end of the spine (below the sacrum) that are fused.

Coronal plane. A plane bisecting the body from front to back. See Chapter 1, Section 1.2.1.

Coronal section. A section of the body taken in the coronal plane.

Cranium. A large bone structure that form the bulk of the skull and accomodates and protects the brain.

Debrunner kyphometer. A device for assessing the extent of a kyphosis or lordosis from the surface of the back.

Formulator body contour tracer. A device for recording the asymmetry of the back from which TAS and STAS can be measured.

Forward bending position. A position in which the scoliosis patient is often placed for measurement. The patient bends forward from the waist and the back curves. Scoliotic deformity is often more obvious in this position. Variations in the forward bending position exist, for example sitting forward bending in which the patient bends forward from a sitting position.

Idiopathic. Any disease or condition of unknown cause or that arises spontaneously.

Kyphosis. An abnormal curvature of the spine in the sagittal plane in which the effect is to deform the spine to be convex i.e. the body hunches forward.

L1...L5. See *lumbar vertebrae*.

Lateral plane. See *transverse plane*.

Lordosis. An abnormal curvature of the spine in the sagittal plane in which the effect is to deform the spine to be concave. The opposite of *kyphosis*.

Lumbar vertebrae. A group of vertebrae in the lower spine. Usually indexed from the top as L1...L5.

Muscular dystrophy. A hereditary progressive weakening and wasting of the muscles.

Palpation. Examination by touch.

Posterior. The back (of the body).

Radiograph. An X-ray image formed by the attenuation of X-rays that pass through the body. Bony structures attenuate the x-rays more severely than soft structures so that the image is predominantly one of the bones.

Sacral vertebrae. See *sacrum*.

Sacrum. A large triangular bone at the base of the spine, formed from five vertebrae, the sacral vertebrae, that are fused in the adult. The process of fusion begins in the womb and is not complete until approximately the twenty-fifth year of life. The sacrum mounts the spine in the pelvis.

Sagittal plane. A plane bisecting the body from left to right. May also refer to a plane parallel to this. See Chapter 1, Section 1.2.1.

Scoliometer. An inclinometer used to measure ATI.

Scoliosis. (A disease causing) a lateral curvature of the spine. The curvature is therefore in the transverse plane.

Spinal cord. The elongated collection of nerves attached to the brain (arguably a part of the brain) extending down the length of the spine and contained by the vertebral canal.

Spinal deformity. Any abnormality in the structure of the spine. Scoliosis is one type of spinal deformity.

Spinal radiograph. An x-ray image of the spine. A full length spinal radiograph includes all the vertebrae of the spine from C7 to the pelvis and may include part of the pelvis itself.

Spinous process. A bony structure, part of a vertebrae, that protrudes from the vertebral body toward the back surface and serves for the attachment of muscles and ligaments. Can often be located through palpation of the back surface.

Surface deformity. An abnormality in the surface shape of the back resulting from an underlying spinal deformity.

T1-T12. See *thoracic vertebrae*.

TAS. See *trunk asymmetry score*.

Thoracic vertebrae. A group of twelve vertebrae in the area of the spine corresponding to the thorax. These vertebrae are attached to the ribs. Usually indexed from the top as T1...T12.

Transverse(or lateral) plane. A plane bisecting the body horizontally (with the individual in the standing position). May also refer to a plane parallel to this. See Chapter 1, Section 1.2.1 .

Trunk Asymmetry Score (TAS). A measure of the asymmetry of a coronal cross-section of the back, calculated by measuring the difference in depth of the back at five positions to the left and right of the spine.

Standardised trunk asymmetry score (STAS). A modified version of TAS that is standardised to have equivalent meaning for patients of all sizes.

Vertebral rotation. The rotation of vertebrae about the axis of the spine.

Appendix II
System Implementation Details

Appendix II

System Implementation Details

A2.1 Introduction

This appendix describes and discusses the implementation of the new method proposed in Chapter 3. The hardware and operating system of the machine used for the investigation are reviewed and their selection is justified. The software design and implementation are discussed in more detail than was possible in Chapter 3. Lastly, the system as it is presented to the user (as a series of interactive screens and menu options) is described.

A2.2 The Hardware Platform

A2.2.1 The Computer

The Acorn Archimedes 540 is based upon a 32-bit Reduced Instruction Set Computer (RISC) processor, the Acorn RISC Machine ARM Chip, operating at 25MHz over a flat memory architecture. It uses 27 32-bit registers, and a small instruction set of 44 simple instructions. Each instruction executes in a single clock cycle so that the processor yields roughly 25 Million Instructions Per Second (MIPS). The development machine was purchased in 1988 at a cost of approximately £2000 with 4Mb of dynamic RAM. At the time of purchase Acorn *claimed* that the processing power of the *earlier* 4MHz version of the processor machine was *equivalent to* (Source: Acorn OEM Products Catalogue):

- 10 x IBM PC AT running BASIC benchmarks.
- a 16.67 MHz Motorola 68020
- 2 to 4 x DEC VAX 11/780 running high level benchmarks.

The criterion for measuring processing power was not specified any more precisely.

These benchmarks suggested that the Archimedes would give reasonable performance and, coupled with consideration of cost of the various machines available at the time, provided an economical performance/cost ratio. Furthermore, it seemed likely that the

system would be used heavily for the processing of integer data (e.g. binary images) which are known to lend well to RISC type processing.

A2.2.2 The Video Digitiser

The Watford Electronics Video Digitiser was capable of capturing 512 x 256 resolution images to 6 intensity bits in a single frame cycle (0.04 seconds). It should be noted that the device was not strictly a frame-store, but a frame digitiser. The difference between the two is that a frame-store will digitise the incoming video signal and place the resulting image into its own local memory which can then be accessed over a bus from the host computer. A video digitiser, however, has no memory on the device itself. Images are digitised and placed directly in to the main memory of the machine. The lack of memory on the digitiser was, presumably, the key contributor to the low price of the device, £147, when true frame-stores typically cost ten times that figure.

A2.2.3 The Operating System

Initially, the operating system provided for the Archimedes was Arthur, a character based operating system with some limited graphics ability. All system services, including graphics functionality are implemented as software interrupts (SWI's). Low level graphics were available for the drawing of simple lines, triangles and circles through calling SWI's. During the investigation, Acorn released the successor to Arthur: RISCOS, a full multitasking operating system with a modern Graphical User Interface (GUI) and Application Programming Interface (API). Although RISCOS was installed on the development machine, its advanced features were not used. There were two reasons for this:

- i. One of the objectives was to produce a software system which was, as far as possible, independent of platform. Using the RISCOS API would have resulted in significant portions of the code for the system being difficult to port to other platforms if and when required.

- ii. By the time RISCOS became available a sizeable portion of the code was already written and would have required significant revision using time that was better spent on research.

The video digitiser had a built in ROM which held the code for the SWI's to perform basic operations such as grabbing and displaying images, as well as optimised routines to save and load images to disk. From the programmer's perspective the device was very simple to use and minimised the time overhead in writing low level code to control image acquisition.

The Acorn Archimedes computer with the Watford Electronics Video Digitiser therefore provided a simple, cost effective hardware solution for the system. Wherever possible, however, the software system was constructed to be independent of the hardware and operating system.

A2.3 Software Design and Implementation

The complete system comprises some 10,500 lines of source code (excluding code comments, and other code that was written to investigate moiré and phase measuring methods) that are distributed among 65 files of source code. Clearly, an implementation of this size required careful planning and strict adherence to a software design in order to function reliably.

A2.3.1 Choice of Programming Language

The criteria for selecting a programming language to use for the implementation of the system were:

- i. That the language should have high level primitives and structures so that image processing algorithms could be quickly investigated without the unnecessary time and complexity overhead of controlling the machine at a low level.
- ii. The language should have at least some low level functionality so that the program can, for example, call system services, manipulate pointers to memory and control hardware.

- iii. The language should conform to a major international standard and be in widespread use so that it would be portable to other platforms and was likely to be supported by manufacturers for the foreseeable future.
- iv. The language should support user definable data structures so that large and complex data structures can be easily implemented and modified.

The programming languages available for the platform were LISP, Prolog, FORTRAN 77, ANSI C and Acorn Basic. C was chosen on basis that it most closely satisfied the criteria.

A2.3.2 The Software Design

The system was implemented as a layered architecture of software modules. Figure A2.1 shows a simplified diagram of the structure. This figure was included in Chapter 5 and is repeated here as an aid to understanding the text.

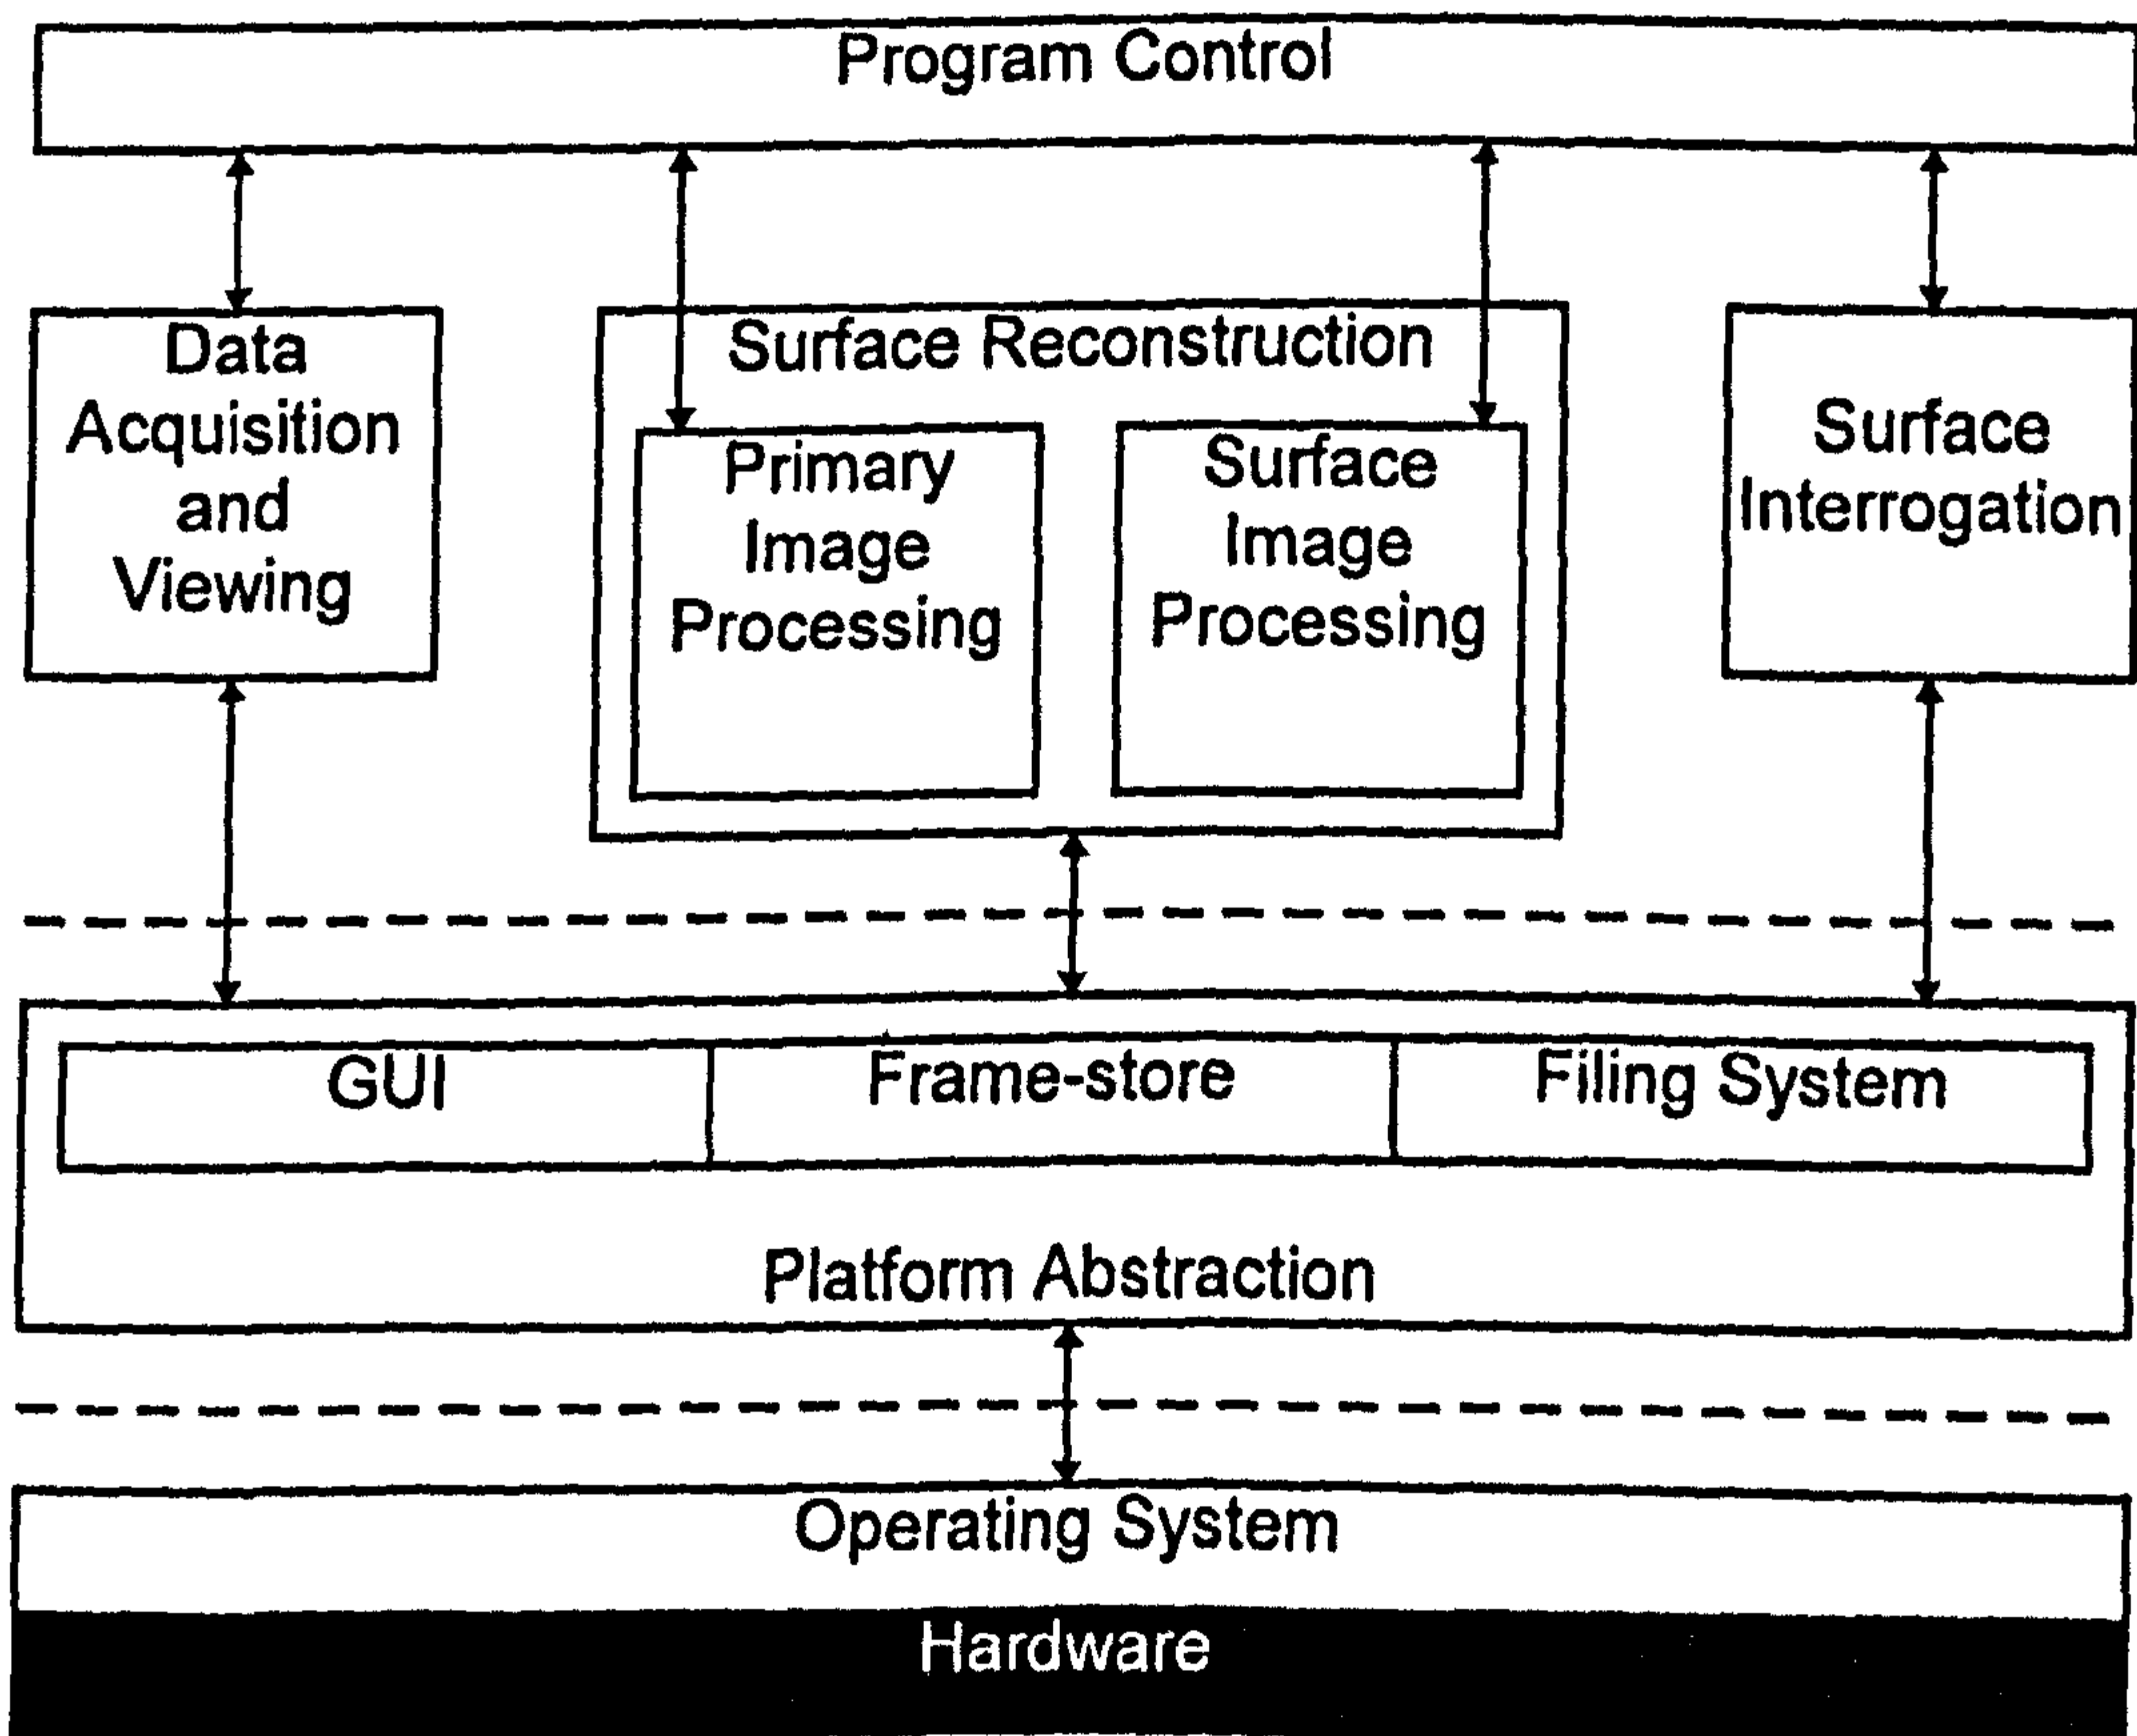


Figure A2.1 Simplified layered system structure.

At the lowest level is the operating system and hardware. The operating system provides the software with a suite of software interrupts that can be called, in principle from anywhere in the system. However, this application is designed to limit the platform specific code (that is, any code which makes direct calls to the OS) to a layer of code called the *Platform Abstraction Layer*.

There are three main modules that comprise the platform abstraction layer: The GUI, the Frame-store Interface and the Filing System Interface.

The graphical user interface handles all the input and output of data from and to the user and calls are made to it by higher layers. The GUI, which is normally mouse driven but also provides keyboard support, provides routines for the construction on the screen of menus and graphical input and output. Example of function calls to the GUI are shown in Figure A2.2.

```
GUI_Construct_Menu( "Interrogate Surface" ) ;  
GUI_Add_Menu_Item("A. View 3d Reconstruction" ) ;  
GUI_DrawLine( x1 , y1 , x2 , y2 ) ;
```

Figure A2.2 Examples of function calls to the GUI.

The Frame-store Interface provides a set of functions for controlling the frame-store, or digitiser, in a manner independent of the underlying hardware. Images are represented in the system as integer arrays held in the main memory and are independent of the underlying hardware. Examples of calls to the frame-store interface are shown in Figure A2.3.

```
ImageIO_GrabImage( backImage ) ;  
ImageIO_DisplayImage( refImage ) ;  
ImageIO_InitialiseDevice() ;
```

Figure A2.3 Examples of calls to the Frame-store Interface

There is some overlap between the responsibility for graphics between the frame-store interface and the GUI. For example, the on-board ROM on the video digitiser used for development held optimised routine for displaying images that performed significantly better than the corresponding GUI routines. Also, it is possible that some frame-stores

will display images on separate monitors (not the host monitor) so that the Frame-store Interface had to collaborate closely with the GUI.

| File Type (Mnemonic) | Status | Description |
|-------------------------|------------------|---|
| BAC | Always required | Original source image of the back |
| EDG | Optional | Used for temporary storage of edge image during index editing of indexed image, and for later recall. |
| GXY | Always required | Storage of system geometry and pixel scaling factors. |
| MVL | Optional | Storage of vertebral positions for later recall. |
| PDT | Always required. | Patient details (name, date-of-birth, date-of-observation etc.) |
| REF | Always Required | Original source image of the reference plane. |
| XRY | Optional | X-ray image of the patient, if exists. |

Table A2.1 File types used by the system.

The filing system interface handles the saving and loading of all data to and from the disk system. It supports different filing systems. Currently, there are two filing systems supported: the Acorn Advanced Disk Filing System (used on machines at the Liverpool Alder Hey Children's Hospital) and the Acorn SCSI filing system (used on the development machine), but other filing systems could easily be incorporated.

A dataset is identified by the hospital number and filing system interface provides routines that save images and other data independently of the underlying file-system. The complete set of files used for a dataset are shown in Table A2.1. Examples of calls to the filing system interface are shown in Figure A2.4.

```
FileIO_Save_Back_Image( imagel ) ;
FileIO_Load_Patient_Data( pd ) ;
FileIO_Get_File_List( listString ) ;
```

Figure A2.4 Examples of Calls to the Filing System Interface.

The next layer of the software is the main application layer that performs the main activity of the system. The Data Acquisition and Viewing module supervises the acquisition of source data. The user is first prompted for information about the patient:

Name, Hospital number and so on). The module then prompts the user to position the patient for measurement, to image the reference plane, to enter system some geometric parameters (e.g. D_c and D_p) and finally to position the circle. The image processing to extract the pixel scaling factors from the scaling image is also performed by this module. The user can also view the source dataset quickly without performing any processing. This is necessary so that the clinician can recall the patient for another measurement session if any data was entered incorrectly or images are clearly inadequate.

The Surface Reconstruction Module is the core software that performs the operations that are described at algorithmic level earlier in this chapter. The algorithms are implemented in two separate modules. The primary image processing module performs all the operations on the source images up to and including the indexing of fringes. These are the operations that apply to the back and reference images independently. The surface image processing module take the two sets of indexed fringes and combines them to reconstruct the surface image.

The surface interrogation module provides the user output on back shape and the methods it uses are described in Chapter 4. This module is principally concerned with the display of the surface and the extraction of clinically significant parameters that relate to back shape and spinal deformity.

Lastly, the highest level module, the Program Control module, coordinates the modules below it and passes data between them. Some application level validation is performed, for example, when a user selects a patient dataset for processing, the program control module will check which files are available for that patient and will prevent further processing if any essential file is missing or cannot be accessed. The Program Control Module also initialises the hardware and software for the system and shuts down the program elegantly (and also the whole computer system if required).

A2.4 The User Perspective

The user navigates the software using a series of menu's for which only valid options can be selected (illegal operations such as attempting to display a surface before it has

been reconstructed are impossible). All user input and output, including the selection of files for processing is performed interactively using a mouse and pointer. A complete display of all the various user screens and data displays is beyond the scope of this appendix, but Figure A2.5 shows the principle set of menu options.

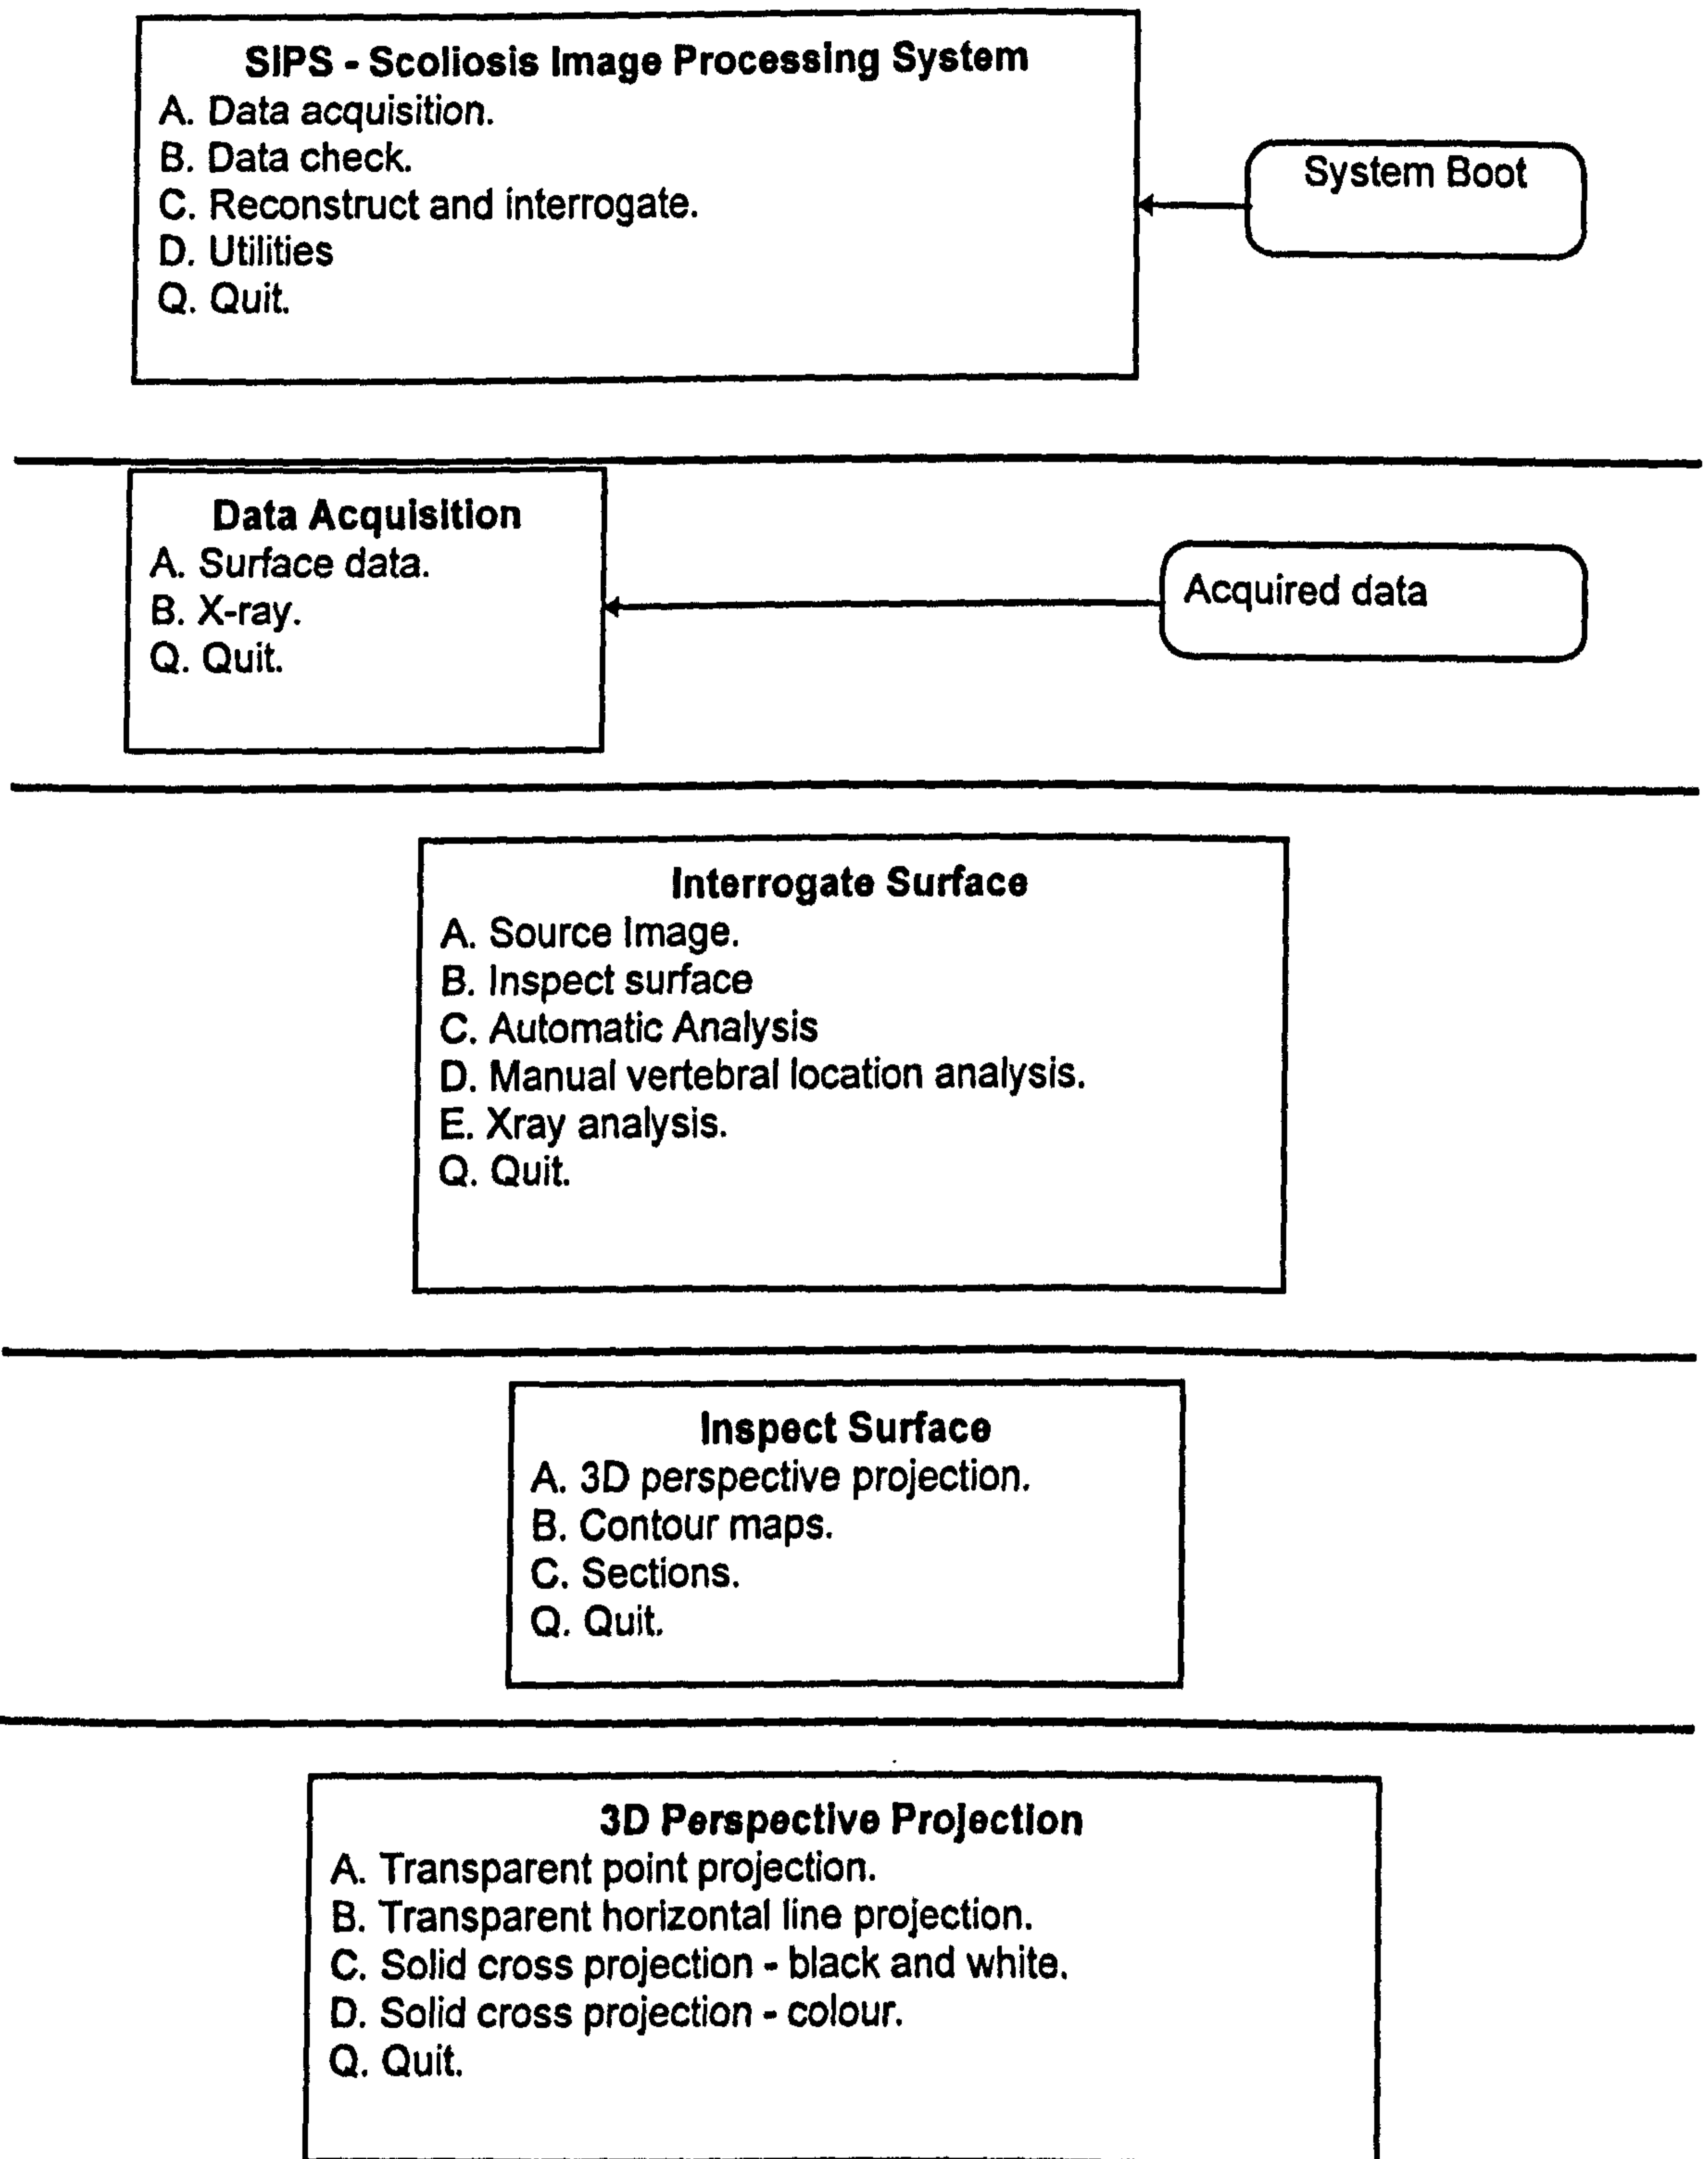


Figure A2.5 The menu system (continued overleaf).

Contour Maps
A. Colour-coded solid.
B. Colour-code line.
C. Black and white.
D. Continuous grey scale.
Q. Quit

Surface Sections
A. Coronal section.
B. Sagittal section.
Q. Quit.

Automatic Analysis
A. ATI analysis.
B. TAS analysis.
C. Surface statistics.
Q. Quit.

Manual Vertebral Location Analysis
A. Inspect/select vertebral positions.
B. Spinal tracing.
C. ATI analysis.
D. STAS analysis.
Q. Quit.

Vertebral
Positions

Utilities
A. Archive data to floppy diskette.
B. Retrieve archived data.
C. Delete data set.
D. Create ASCII image file.
Q. Quit.

Figure A2.5 continued.

**Appendix III
Publications**

Appendix III Publications

1. Fourier Transform Profilometry

Pearson J.D., Dangerfield P.H., Burton D.R., Hobson C.A. and Dorgan J.C. (1995). "Fourier Transform Profilometry for Measurement of the Three-dimensional Shape of the Human Back in Scoliosis." In: D'Amico M., Merolli A. and Santambrogio G.C. (Editors), Three-dimensional Analysis of Spinal Deformities: Health Technology and Informatics (15). Oxford: IOS Press. p. 9-14.

Abstract

This paper describes preliminary work in the application of Fourier Transform Profilometry (FTP) to the measurement of human back shape. Optical fringe patterns are generated on the surface of the back and the FTP method, which is phase sensitive, is used to reconstruct the surface shape of the back. The method is relatively resilient to changes in the reflective properties of the skin and overcomes the problem of low sampling density which is inherent in edge or peak extracting methods.

2. Moiré Contouring

Groves D., Dangerfield P. and Pearson J.D.(1990). "Advanced computer analysis of moiré contour images of the human back." In: Neugebauer H. and Windischbauer G., (Editors), Surface Topography and Body Deformity. Stuttgart: Gustav Fischer Verlag. p. 107-114.

Abstract

A method is described and demonstrated which enables the accurate and detailed analysis of moiré contour images of the human back. The method will enable the measurement of clinically accepted parameters of scoliosis that are measured by the Scoliometer, kyphometer and Formulator Body Contour Tracer. In addition the method of analysis has enabled a new single coefficient of spinal deformity to be proposed which is a normalised measure of three-dimensional asymmetry about the spine. The

authors propose its use: (i) in following the progression of the spinal curvature in individual patients (ii) as the main parameter for scoliosis detection in school screening. The results obtained from following individual patients will enable a value of the parameter to be determined above which an individual will be considered to have scoliosis and require further investigation. The method of analysis is demonstrated using non-automated contour image acquisition. An automatic high resolution system incorporating the technique is under development.

3. Publication on the Technical Aspects of the Proposed Method

Pearson J.D., Hobson C.A. and Dangerfield P.H. (1993). "Automated measurement of human body shape and curvature using computer vision." In: Wilson D.C. and Wilson J.N. (Editors), Proceedings of SPIE 2035. p. 45-57.

Abstract

A system to measure the surface shape of the human body has been constructed. The system uses a fringe pattern generated by the projection of multi-stripe structured light. The optical methodology used is fully described and the algorithms used to process acquired digital images are outlined. The system has been applied to the measurement of the shape of the human back in scoliosis.

Other papers on the method include:

Pearson J.D., Dangerfield P.H., Atkinson J.T., Gomm J.B., Dorgan J.C., Hobson C.A. and Harvey D.M. (1992). "Measurement of body surface topography using an automated imaging system." Acta Orthopædica Belgica, 58 Supplement I, p. 73-79.

Pearson J.D., Dangerfield P.H., Hobson C.A. and Li Y. (1992). "An automated visual system for the measurement of the three-dimensional deformity of scoliosis." In: Dansereau J. (Editor), Proceedings of the International Symposium on 3-D Scoliotic Deformities. Stuttgart: Gustav Fischer Verlag. p. 50-56.

Hobson C.A., Pearson J.D. and Dangerfield P.H. (1992). "An automated visual system for the measurement of human spinal deformities." Proceedings of the Institute of Electrical Engineers 354, p. 389-392.

4. Publications on the Clinical Application of the Proposed System

Dangerfield P.H., Pearson J.D., Dorgan J.C. and Hobson C.A. (1995). "Clinical Experience of Employing Computerised Shape Analysis to measure the 3D deformity of scoliosis." In: D'Amico M., Merolli A. and Santambrogio G.C. (Editors), Health Technology and Informatics (15), Oxford: IOS Press. p. 231-238.

Dangerfield P.H., Scutt D., Dorgan J.C., Li Y. and Pearson J.D. (1992). "The effect of body position on the three-dimensional deformity of scoliosis." In: Dansereau J. (Editor) Proceedings of the International Symposium on 3-D Scoliotic Deformities. Stuttgart: Gustav Fischer Verlag. p. 139-147.

Dangerfield P.H., Pearson J.D., Hobson C.A., Li Y., Brodie D.A., Scutt D. and Dorgan J.C. (1992). "The cost benefits of an automated imaging system for topography of the back in scoliosis." Healthcare and Accountancy. p. 319-328.

Dangerfield P.H., Pearson J.D., Nunn N., Dorgan J.C., Klenerman L., Hobson C.A. and Harvey D.M. (1992). "Measurement of the angle of trunk inclination using a computerised imaging system." In: Alberti A., Drerup B and Hierholzer E. (Editors) Surface Topography and Spinal Deformity. Stuttgart: Gustav Fischer Verlag. p. 17-21.

5. Additional Short Papers and Abstract-only Publications

Dangerfield P.H., Pearson J.D., Li Y., Dorgan J.C., Hobson C.A. (1993). "Three-dimensional analysis of back shape topography in scoliosis using an advanced computer imaging system." Revue de Chirurgie Orthopedique 1993, (Masson, Paris), EFORT Supplement 432. (Abstract only).

Dangerfield P.H., Scutt D., Dorgan J.C., Li Y., and Pearson J.D. (1993). "The effect of body position on the three-dimensional surface deformity of scoliosis." Clinical Anatomy, 6(1). p. 61. (Abstract only).

Dangerfield P.H., Dorgan J.C., Pearson J.D., Hobson C.A. and Yiang J. (1993). "Computer image processing in scoliosis: Clinical experiences." Proceedings of the British Scoliosis Society, 18, p. 22. *This paper won the British Scoliosis Society Esola Prize (and cheque for £250) for the best paper at the 1993 conference. (Abstract only).*

Dangerfield P.H., Pearson J.D., Hobson C.A. and Dorgan J.C. (1992). "Measurements of back surface topography in scoliosis using an automated imaging system." Clinical Anatomy, 5(3). p. 238-239. (Abstract only).

Dangerfield P.H., Pearson J.D., Hobson C.A., Dorgan J.C. and Klenerman L. (1992). "Measurement of back surface topography in scoliosis using an automated imaging system." Journal of Bone and Joint Surgery, 74-B, Supplement I. p. 95-96. (Abstract only).

Dangerfield P.H., Pearson J.D., Li Y., Dorgan J.C. and Hobson C.A. (1992). "Three-dimensional analysis of back shape topography using an advanced computer imaging system." Journal of Bone and Joint Surgery, 75-B, Supplement II. p. 174. (Abstract only).

Dangerfield P.H., Pearson J.D., Li Y., Hobson C.A. and Dorgan J.C. (1992). "A low cost automated imaging system (SIPS) for scoliosis." Journal of Bone and Joint Surgery, 74-B Supplement III. p. 310. (Abstract only).

Dangerfield P.H., Scutt D., Dorgan J.C., Li Y., Pearson J.D., Rajasekeran S. and Brodie D.A. (1992). "The effect of body position on the three-dimensional surface deformity of scoliosis." Journal of Bone and Joint Surgery, 74-B, Supplement III, p. 307. (Abstract only).

Pearson J.D., Dangerfield P.H., Li Y., Hobson C.A. and Dorgan J.C. (1992). "Three-dimensional body shape measurement of scoliosis using an automated imaging system." Proceedings of the European Spinal Deformities Society, (Lyon, France), p. 59-60. (Short paper).

Pearson J.D., Dangerfield P.H., Li Y., Hobson C.A. and Dorgan J.C. (1992). "Multiple level measurement of the three-dimensional surface deformity of scoliosis." Proceedings of the European Spinal Deformities Society, (Lyon, France), p. 255.

(Abstract only).

Dangerfield P.H., Groves P.H. and Pearson J.D. (1989). "Advanced computer analysis of moiré contour images and the assessment of ribcage deformity in scoliosis." Proceedings of the British Scoliosis Society 14th Annual General Meeting, p. 31.

(Abstract only).

Dangerfield P.H., Groves D. and Pearson J.D. (1989). "Advanced computer analysis of moiré contour images and the assessment of ribcage deformity in scoliosis." Journal of Bone and Joint Surgery, 72-B(2), p. 338. (Abstract only)

Dangerfield P.H., Groves D. and Pearson J.D. (1989). "Computer recording and analysis of surface topography for the assesment of rib-cage deformity in scoliosis." Clinical Anatomy, 3 (3), p. 64-65. (Abstract only).

Dangerfield P.H., Groves D. and Pearson J.D. (1989). "The Natural History of Scoliosis and Imaging of the spine." Proceedings of GES (French Scoliosis Society), (23). p. 56-58. (Short paper).