

Ultrasound Imaging for Robotic Knee Arthroplasty



William Kerr

A thesis submitted for the degree of

Doctor of Engineering

August 2016

Copyright Declaration

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

The copyright of this thesis belongs to the author under the terms of United Kingdom copyright acts as qualified by University of Strathclyde Regulation 3.50. Due

acknowledgement must always be made of the use of any material in, or derived from this thesis.

Signed:

Date:

Acknowledgements

I would like to thank my supervisors, Prof. Philip Rowe and Prof. Gareth Pierce, for the opportunity to conduct this research. Their expertise, guidance and willingness to explore new ideas have been a much needed source of inspiration and support.

I would also like to convey my gratitude to Grant Smilie and Alec Ward for their effort and helpful ideas in constructing a myriad of awkward parts. For their constant willingness to help in all technical areas, I would like to show my appreciation for Rahul Summan, Joseph Jackson, Charles MacLeod and Maxim Morozov. I would also like to thank Nicholas Smith for answering a daily barrage of questions without judgement.

Furthermore, I would like to express my deepest gratitude to my parents and sister for providing unwavering support and understanding throughout this project and beyond. Finally, I would like to show my appreciation for Monica, who made any hardship worthwhile.

Abstract

Osteoarthritis is a degenerative disease affecting the articulating joints of the body and is particularly prevalent in the knee. The degradation of the joint can cause disabling symptoms in the sufferer, the current remedy to which is total or unicompartmental knee arthroplasty. While traditional knee arthroplasty procedures are well established through decades of success, the method by which the unwanted bone is removed is outdated and inaccurate. In answer to this, robotically guided knee arthroplasty systems provide a higher level of accuracy when reshaping the bone, which is responsible for an improved clinical outcome. Many of these systems require a preoperative 3D model of the joint, with which the surgery can be planned and

implemented. Currently, Computed Tomography is the chosen method to achieve this, providing the required level of accuracy consistently. However, this method is costly and applies potentially dangerous doses of ionising radiation to the patient. Ultrasound imaging has the potential to provide an alternative to this by offering comparable accuracies, while reducing cost and eliminating the risk of ionising radiation.

Ultrasound imaging has become commonplace in biomedicine and in NonDestructive Testing (NDT). Despite their common ancestry, the two fields of research have bifurcated, with recent developments finding little crossover. One of the recent research topics in NDT has been Synthetic Aperture (SA) imaging methods. Despite the possibility of improved imaging capabilities, these methods have found little uptake in biomedicine. This thesis, therefore, applies SA methods to the task of reconstructing the bony surfaces of the knee, in an effort to apply the advances within NDT to biomedical imaging. Employing an industrial robot to provide automated and accurate probe positioning, 3D reconstruction of complex bony surfaces was found to be possible at a fundamental level to the required submillimetre accuracy. However, it was found that only certain SA widths would allow for accurate surface reconstruction. Additionally, a number of inaccuracies were found to be caused by 3D effects with 1D arrays. Finally, the effects of soft tissue inclusion were investigated, with results suggesting that soft tissue would not seriously compromise the ability to accurately reconstruct surfaces.

Contents

Copyright Declaration	i
Acknowledgements	ii
Abstract	ii
Contents	iii
List of Figures	viii
List of Tables	xxv
List of Abbreviations	xxvii
Chapter 1 Introduction	1
1.1 Project Background	1
1.1.1 Anatomy of the Knee	1
1.1.2 Osteoarthritis	6
1.1.3 Robotic Knee Arthroplasty.....	9
1.1.4 Preoperative Imaging Methods	13

1.2	Ultrasound	17
1.3	Aims of the Thesis	26
1.4	Contributions to Knowledge	28
1.5	Thesis Structure	29
1.6	Publications Arising from this Thesis	30
1.6.1	Journal Papers	30
1.6.2	Conference Papers	31
Chapter 2	Ultrasound Data Capture and Processing	31
2.1	Introduction	31
2.2	Data Capture and Processing Techniques	32
2.2.1	A Scan	32
2.2.2	B-Scan	35
2.2.3	C-Scan	36
2.2.4	Plane B-Scan	37
2.2.5	Phased Arrays	39
2.2.6	Linear B-Scan	43
2.2.7	Dynamic Depth Focussing	44
2.2.8	Sector B-Scan	45
2.2.9	Full Matrix Capture	48
2.2.10	Total Focussing Method	50
2.2.11	Common Source Method	57
2.2.12	Synthetic Aperture Focussing Method	58
2.2.13	Biomedical Synthetic Aperture Techniques	61
2.2.14	Bone Surface Contour Extraction	65
2.2.15	Bone Surface Reconstruction	67
2.3	Conclusion	68
Chapter 3	System Component Design Considerations and Testing	69
3.1	Introduction	69

3.2 Coupling Media.....	70
3.3 Probe Positioning and Tracking	71
3.3.1 Probe Manipulation	71
3.3.1.1 Probe Manipulation Methods	72
3.3.1.2 Probe Mount Design and TCP Calibration	75
3.3.1.3 Path Planning	81
3.3.2 Path Accuracy	88
3.3.2.1 Methods and Materials	89
3.3.2.2 Results.....	94
3.3.2.3 Discussion.....	96
3.3.2.4 Conclusion	100
3.4 Conclusion.....	100
Chapter 4 Imaging Known Geometries	102
4.1 Materials and Methods	102
4.2 Results	112
4.2.1 Cuboid.....	116
4.2.2 Sphere.....	125
4.2.3 Cylinder.....	134
4.3 Discussion	142
4.4 Conclusion.....	156
Chapter 5 Composite Human Distal Femur Surface Reconstruction	156
5.1 Introduction	157
5.2 Methods and Materials	157
5.3 Results	170
5.3.1 Point Cloud Comparison	170
5.3.2 Mesh Comaprison	181
5.4 Discussion	190
5.5 Conclusion.....	199

Chapter 6	Bovine Distal Femur Reconstruction	202
6.1	Introduction	202
6.2	Bovine Sample Preparation	202
6.3	Bovine Distal Femur Surface Imaging and Reconstruction.....	205
6.3.1	Materials and Methods	205
6.3.2	Results	210
6.3.2.1	Point Cloud Comparison.....	210
6.3.2.2	Mesh Comparison	215
6.3.3	Discussion	222
6.4	Conclusion.....	231
Chapter 7	Bone Surface Imaging With Soft Tissue – a Modelling Approach....	232
7.1	Introduction	232
7.2	Finite Element Modelling of Soft Tissue Interfaces	233
7.3	Materials and Methods	234
7.4	Results	246
7.5	Discussion	254
7.6	Conclusion.....	260
Chapter 8	Bone Surface Imaging With Soft Tissue – an Experimental Approach	261
8.1	Introduction	261
8.2	Materials and Methods	262
8.3	Results	267
8.4	Discussion	276
8.5	Conclusion.....	279
Chapter 9	Future Work, Practical Implementation and Conclusions.....	280
9.1	Reconstruction of Complex Bony Surfaces	281
9.2	Reconstruction of Complex Bony Surfaces with Preceding Soft Tissue Interfaces.....	283
9.3	Reconstruction of the Intact Knee Joint	287

9.4 Practical Implementation	290
9.5 Conclusions	300
References	304
Appendix A.....	331
A1. Calibration Part Design.....	331
Appendix B.....	341
B1. Results Tables	341
Appendix C.....	346
C1. Results Tables	347

List of Figures

Figure 1.1: A sagittal section of the knee joint, showing the two bone types in the femur, tibia and patella, as well as the location of the bones relative to one another (Gray, 1918).	2
Figure 1.2: Anterior view of the right knee, with patella and associated ligaments omitted (Gray, 1918).....	3
Figure 1.3: Posterior view of the left knee, showing the four main ligaments – ACL, PCL, MCL, LCL (Gray, 1918).	5
Figure 1.4: Part (a) shows the anterior femoral musculature, showing three of the four muscles termed the quadriceps: rectus femoris, vastus lateralis and vastus medialis, with vastus intermedius being covered by rectus femoris. Part (b) shows the posterior muscles, including the hamstring group, where semimembranosus is labelled ‘1’, semitendinosus ‘2’ and biceps femoris ‘3’ (Gray, 1918).	6
Figure 1.5: Visual description of the progression of OA at the microscopic level, with ovals representing chondrocytes, dots representing GAGs and curved lines collagen fibrils (Arokoski et al., 2000). Three stages of OA progression are illustrated, showing degradation of the superficial layer, loss of GAGs and a further loss of chondrocytes.	8
Figure 1.6: the components of the MAKO RIO system employed in partial knee MAKOPlasty, showing (from left to right) the robotic arm with burr, the stereoscopic infrared optical tracking system and the control panel.	12
Figure 1.7: Flow diagram of the processes involved in robotic knee arthroplasty, from diagnosis to surgical procedure.	13
Figure 1.8: The system described by (Salcudean et al., 1999) for imaging of the carotid artery © [2013] IEEE. The probe maintains good contact with the skin using movement of the manipulation system, which is controlled by an external expert.	22
Figure 1.9: Three different ultrasonic inspection methods using robotic arms. Part (a) shows contact testing, with the top image showing testing with a single element probe, while the bottom shows that with a phased array. Part (b) shows immersion testing, again with the top showing a single element, while the	

bottom shows an array. In both (a) and (b), the green area represents the region imaged. Part (c) shows through transmission testing, with a transmitter on the left and a receiver on the right.	25
Figure 2.1: An example of showing both pitch-catch and pulse-echo modes. The top half of the image depicts the physical setup, while the bottom displays the A scans.	34
Figure 2.2: Construction of a B-scan image, showing mechanical movement of a single element probe. The time of flight data is displayed vertically, while the horizontal axis indicates the position of the transducer.	36
Figure 2.3: Raster scan path for probe and raised circular area (left) and the resultant C-scan image (right), with a difference in height being represented by a change in colour.	38
Figure 2.4: Configuration of a conventional one dimensional array, showing the pitch, P , the element width, W , and the kerf, K	39
Figure 2.5: Data capture for a plane B-scan for a 16 element probe, using a 4 element sub aperture. The sub aperture moves electronically along the length of the array, producing a different region of focus at every transmission. ..	40
Figure 2.6: Construction of parameters required for phase delays for the i th element of a N element array.	41
Figure 2.7: Field characteristics of a single element transducer (a-d) and a 32 element array (e-h), both operating at 5MHz in water.....	43
Figure 2.8: Sound fields of a 32 element array operating at 5MHz in water, focussed at (from left to right) 3cm, 6cm and 9cm. Also shown are associated time delays for the individual array elements.	46
Figure 2.9: Sound fields of a 32 element array operating at 5MHz in water, steering the beam to angles of 15, 0 and -15 degrees (left to right). Also shown are associated time delays for the individual array elements.....	47
Figure 2.10: The FMC procedure showing transmission and reception cycles, with grey representing an inactive element, red an element being fired and blue a receiving element.	50
Figure 2.11: Geometric configuration for one iteration of the TFM algorithm. The red line represents the path of the ray from the transmitting element to the point under reconstruction (x_p, z_p) . The blue line is the path from that point to the receiving element.	53

Figure 2.12: From a simulated FMC of a point reflector. The TFM image using only time domain signals (a), the TFM image using Hilbert-transformed signals (b), a typical unaltered A-scan (blue) with the envelope (red) as a consequence of the Hilbert Transformation (c).	55
Figure 2.13: Transmit-receive matrices of a four element array for TFM (a), CSM (b), single element receive SAFT (c) and multiple element receive SAFT (d).	59
Figure 3.1: A pair of Kuka KR5 Arc HW robots within the robot cell in the FIRST laboratory in CUE.	75
Figure 3.2: Exploded view of probe mount and the probe. The unlabelled white lines illustrate how the three parts come together.	79
Figure 3.3: The probe mount showing the extra screw added for stability (a) and the probe and probe mount attached to the robot via the tool changer (b). ...	80
Figure 3.4: Possible collision scenario with first probe mount design, showing the robot colliding with the tank. With the dotted line representing the water line, collisions could be possible when attempting to achieve a top-down perspective of the sample, while keeping the probe face fully submerged.	81
Figure 3.5: Exploded view of the second probe mount for Tool 2. The stabiliser, which ensures the probe maintains the appropriate placement relative to the robot, is marked.	82
Figure 3.6: The probe mount (a) and the probe in place on the robot (b). The base was remade in plastic to minimise electrical interference from the robot.	83
Figure 3.7: A representation of the augmented hemispherical path followed by the probe, with an example of a sample at the centre. The probe face would be tangential to the surface of the hemisphere, with the TCP pointing at the centre of the hemisphere at all times.	84
Figure 3.8: Graphical representation of three motion types defined in KRL: linear (a), point to point (b) and circular (c) (KUKA, 2010).	86
Figure 3.9: formulation of the coordinates of start, end and auxiliary points for the hemispherical path, with a sphere shown for reference.	89
Figure 3.10: The different components of the hemispherical scan represented with different colours, with the part of the scan completed by Tool 2 in black.	91

Figure 3.11: The camera setup relative to the robot position, with all cameras on one side of the robot, due to spatial constrictions in the robot cell. Also detailed is the face of one of the cameras, showing the LED strobe ring and the lens in the centre.93

Figure 3.12: The positions of the six cameras relative to each other. Also shown is the marker set and the TCP (white box).95

Figure 3.13: The marker set employed, with the sphere representing the TCP attached to the gold coloured screw (a) and the calibration wand (b). The silver screw in part (a) was employed to adjust the position of the TCP sphere.96

Figure 3.14: An example of the results of ICP, taking unmatched data (a) and transforming one set to minimise error between the sets (b).97

Figure 3.15: Absolute error between paired points of Vicon and Kuka robot data (a, b), difference between distance of consecutive points in each data set (c, d) and the previous value normalised by the corresponding distance between consecutive points provided by Vicon (e, f).99

Figure 3.16: The position of the robot at a point on one of the two curves identified as having high error. Most of the cameras are positioned to the left of the robot, meaning the markers are, for the most part, blocked by the robot itself. With fewer rays intersecting the spheres because of line of sight blockages, the accuracy of the reconstruction degraded. 102

Figure 4.1: The three test pieces used: a sphere, cuboid and flattened off cylinder. Also shown is the mount, to which each was attached to allow full access. 107

Figure 4.2: The probe being manoeuvred around the sphere sample, ensuring immersion of the probe and a lack of contact between robot and water at all times. 109

Figure 4.3: The cuboidal path employed in the cylindrical sample, showing two passes over the top to extend coverage. Also shown are the probe orientations relative to the surface. 109

Figure 4.4: The three additional scan paths employed for inspection of the cylindrical sample. 110

Figure 4.5: Graphical representation of the TFM and SAFT synthetic aperture definitions, with the elements constituting the SAFT aperture shown in a lighter

shade. Defining θ as a relatively small angle (eg. $<40^\circ$) will ensure that $\beta-\alpha$ is less than N , so that the reconstructing aperture will always be smaller than with TFM.	114
Figure 4.6: A typical SAFT image from the cuboid test piece, showing the element positions as red circles. The surface profile is seen at the top of the image, while the return signal reflection can be seen as a strong reflection towards the bottom of the image. Also present is an obvious backwall reflection, shown as the second surface representation.	115
Figure 4.7: The result of coordinate extraction using only global thresholding, with extracted points represented as yellow circles (a). All unwanted reflections were included in the extraction. Also shown is the final contour extraction method (b), where all the unwanted reflections have been eliminated from the extraction.	116
Figure 4.8: the resultant point cloud from the scan, of which the image in Figure 4.6 was a part. The coordinates are relative to the KUKA World axes.....	118
Figure 4.9: A TFM image produced from a data capture when the probe was operating into water, with no reflective materials present. Shown in detail is the result of ring down.....	119
Figure 4.10: The Faro-derived reference models of the sphere (a), the cylinder (b) and the cuboid (c), each of which shows the support onto which they were each screwed.	120
Figure 4.11: A typical TFM image from the cuboidal path inspection of the cuboid test piece (a). Also displayed is the contour extraction for the image (b), where unwanted regions have been extracted.	123
Figure 4.12: The resultant point cloud from the cuboidal path inspection of the cuboid test piece using SAFT. A side on view (a) and a top down view (b) are displayed. While the cuboid shape is obvious, the sides have extended past the edge of the sample.....	124
Figure 4.13: The resultant point cloud from the cuboidal path inspection of the cuboid test piece using TFM. A side on view (a) and a top down view (b) are displayed. The side lobes have dominated the surface extraction, resulting in a noisy surface representation.	125

Figure 4.14: The resultant point cloud from the cuboidal path inspection of the cuboid test piece using TFM, with a -5dB threshold limit. Showing a side on view (a) and a top down view (b), it can be seen that multiple surface representations exist.	126
Figure 4.15: Typical TFM (a) and SAFT (b) images from the hemispherical scan of the cuboidal sample. Sidelobes are present in the TFM image, while none can be seen in the SAFT image.	127
Figure 4.16: The TFM (a) and SAFT (b) derived point clouds, resulting from the hemispherical scan of the cuboidal sample. Side lobes have resulted in the noisy TFM-based point cloud in (a), with the poor surface representation in (b) resulting from the inaccuracy associated with the robot path.	128
Figure 4.17: The result of the full hemispherical scan of the cuboid (processed with TFM) with a threshold of -5dB rather than -10dB (a), resulting in a cleaner point cloud. Also shown is the comparison with the reference model.....	129
Figure 4.18: Results of SAFT processing from individual parts of the hemispherical scan that show positional errors (a), (b) associated with specific parts of the probe path. Also shown are two parts of the scan combined which show more accurate point placement (c) and that combined cloud when compared to the reference model (d).	130
Figure 4.19: Typical TFM (a) and SAFT (b) images from the cuboidal scan of the spherical sample. While side lobes are present in (a), a thicker surface representation has been reconstructed in (b).	132
Figure 4.20: An example of a poor surface profile from SAFT processing found with the sphere sample, showing a smearing effect with the surface.	133
Figure 4.21: The TFM (a) and SAFT (b) derived point clouds, resulting from the cuboidal scan of the spherical sample.....	134
Figure 4.22: Resultant SAFT point clouds from a single pass along one side of the sample (a) and a top side pass (b), both showing a lack of curvature expected from the arc of the sphere.	134
Figure 4.23: The comparison between the topside scan of the sphere processed with SAFT and the reference model, matched manually. As the probe moves from the centre of the sphere, the error increases.	135

Figure 4.24: The comparison between both the topside scan of the sphere processed with SAFT, with a -6dB threshold (a) and that with TFM, with a -5dB threshold (b) and the reference model. Less side lobe contribution is obvious in (b), due to a lower threshold level.	136
Figure 4.25: The TFM (a) and SAFT (b) derived point clouds, resulting from the hemispherical scan of the spherical sample.	137
Figure 4.26: The comparison between the full hemispherical scan of the sphere, processed with SAFT and the reference model, matched using ICP. The top view of the point cloud is detailed, showing the high levels of error found at axis crossings.	138
Figure 4.27: The resulting point cloud from the full hemispherical scan of the sphere (processed with TFM) (a) and the comparison with the reference model, matched using ICP.	138
Figure 4.28: The TFM (a) and SAFT (b) derived point clouds, resulting from the cuboidal scan of the cylindrical sample. Part (a) shows high levels of side lobe extraction, while (b) displays a cleaner representation, but with inaccurate placement.....	140
Figure 4.29: The comparison between both cuboidal scan of the cylinder processed with SAFT, with a -12dB threshold (a) and that with TFM, with a -5dB threshold (b) and the reference model. Side extension is present in part (a), while (b) shows a high side lobe contribution.	140
Figure 4.30: The TFM (a) and SAFT (b) derived point clouds, resulting from the hemispherical scan of the cylindrical sample. Side lobes have resulted in the noisy TFM-based point cloud in (a), with the poor surface representation in (b) resulting from the inaccuracy associated with the robot path.	141
Figure 4.31: The TFM-derived point cloud resulting from the hemispherical scan. The threshold limit employed was -5dB. The increase in threshold has resulted in lower side lobe extraction, but has also resulted in true surface extraction also.	142
Figure 4.32: The comparison between both the TFM (a) and SAFT (b) derived point clouds and the reference model for Pass 1. TFM reconstruction resulted in an accurate surface representation, while SAFT displayed side lobe extraction and backwall reflections.	143

Figure 4.33: The comparison between both the TFM (a) and SAFT (b) derived point clouds and the reference model for Pass 2. TFM produced an accurate surface representation, while SAFT resulted in backwall extraction. ...	144
Figure 4.34: The comparison between both the TFM (a) and SAFT (b) derived point clouds and the reference model for Pass 3. Both point clouds show a flattening effect towards the centre of the arc.	145
Figure 4.35: The dimensions of the individual elements from the 5MHz linear array (a), the 3D beam shape from a single element in water (b), a two dimensional cut of the beam intensity at $y = 0$ and a further two dimensional cut of the beam intensity at $x = 0.08$ (d). The X dimension relates to the distance from the probe face, while the Z dimension relates to the height of the individual elements, meaning the Y dimension is where the full probe aperture is located.	148
Figure 4.36: A single ultrasonic transceiver (grey) emits a wide beam shape, the outer parts of which do not return to the transducer, due to specular reflection on a planar surface. The centre of the beam is reflected back to the transducer. The beam is depicted in the height (ie. elevation) dimension of the transducer and the blue represents where the resultant image plane would be located.	149
Figure 4.37: The effect of diffuse reflection from a sharpened corner, with rays travelling back from outside the image plane due to the height of the element (a) and the diffuse reflections caused by the edge of the sample.	150
Figure 4.38: The effect of specular reflection from a curved surface, with rays travelling back from outside the image plane due to the angle of reflection from the surface.	152
Figure 4.39: Reconstruction scenarios from a number of pixels on a shape similar to the cylindrical sample. The blue elements represent the associated receiving apertures, the arrows show the main direction of the specular reflection, the green points denote the pixel positions and the triangles show how the receiving aperture is determined.	154
Figure 4.40: TFM (a) and SAFT (b) images of the cylindrical sample when located beyond the edge of the probe. TFM reconstructed the surface, while SAFT was incapable of doing so.	156

Figure 4.41: TFM reconstruction of a simulated FMC of a point reflector. Element positions are displayed as red circles, while the paths of travel for two pixels at (a, b) and (c, d) are shown for the same transmission and reception. It is clear that the path to the centre of the side lobe is the same distance as the path to the point reflector.	158
Figure 4.42: An A-scan from the point reflector FMC, showing the time of travel for the point reflector, D, and the range in which pixel contributions will be non-zero, P.	159
Figure 4.43: Examples of ‘ghost’ surfaces where the whole surface has been mirrored either side of the true surface (a) and one particular strong surface has been mirrored, while the weaker has not (b).	160
Figure 5.1: The composite human distal femur model, showing drilled hole through the centre and metallic base for stability (a). Four screws have secured the base to the table to eliminate movement and the sample has been spray painted to account for laser absorption (b).	164
Figure 5.2: Reference model captured with Faro laser scanner (a) and the representative mesh supplied by Sawbones, with no drilled hole (b)....	165
Figure 5.3: The second path used for the inspection of the composite human distal femur, with separate segments denoted by different colours.	166
Figure 5.4: The physical setup of the experiment, showing the robot with probe attached, manoeuvring around the composite bone. The bone was completely submerged in tap water, as was the face of the probe.	167
Figure 5.5: A typical 20° SAFT-derived image (left), with the extracted surface (right), with extracted points displayed as black dots. The detail of the extracted points shows the jagged nature of the extracted contour.	169
Figure 5.6: The extracted contour (in black) after the application of a smoothing function, eliminating the jagged appearance seen in Figure 5.5.	170
Figure 5.7: The point cloud resulting from a cuboidal path scan, with no smoothing applied. The blue points represent the final point cloud after point elimination, while the red points represent those points that have been eliminated. The blue panel represents one of the boundaries past which points were eliminated.	171
Figure 5.8: The result of the nearest neighbour point elimination algorithm on a cuboidal scan, processed using SAFT, with no smoothing applied. The final points	

are shown in blue, while the eliminated points are displayed in red. While outliers have been eliminated, so have genuine surface points. 172

Figure 5.9: The mesh reconstruction process carried out in Geomagic Wrap. The original point cloud (a) was wrapped, resulting in a jagged surface mesh with a number of holes (b). The holes were then filled (c), after which smoothing and polygon relaxation were applied (d). This resulted in a continuous surface mesh, which was representative neither of the physical bone or the original point cloud. As such, a planar trimming tool was employed (e) to remove the unwanted part of the mesh (f). 174

Figure 5.10: Idealised example illustrating the weakness in mesh to mesh comparisons. The red lines represent the surface of the reference mesh, while the blue lines represent that of the compared mesh. The red points signify the vertices of the reference mesh, while the black points in (a) represent the vertices defining that part of the surface and the green points in (b) denote the points sampled across the whole surface. 175

Figure 5.11: Final surface mesh of an unsmoothed SAFT data set (a) and the densely sampled point cloud (b). Also shown is the detail of the point cloud, illustrating the transparency and, therefore, discontinuous nature of the model..... 176

Figure 5.12: Four viewing angles of the compared point cloud produced using 20° aperture SAFT processing, with no smoothing function. 178

Figure 5.13: Four viewing angles of the compared point cloud produced using 20° aperture SAFT processing, with smoothing function. While the individual lines are cleaner than in Figure 5.12, a number lines have been located inside the sample. 179

Figure 5.14: Four viewing angles of the compared point cloud produced using TFM processing, with smoothing function. The point clouds are significantly more sparse than with SAFT reconstruction. 180

Figure 5.15: A typical FMC processed using 30° aperture SAFT (a) and TFM (b), with part (b) displaying side lobes. 181

Figure 5.16: Four viewing angles of the compared point cloud produced using 20° aperture SAFT processing and the hemispherical scan path, with no smoothing function. The intercondylar notch has not been reconstructed as well as was the case with the cuboidal scan path. 183

Figure 5.17: Four viewing angles of the compared point cloud produced using 20° aperture SAFT processing and the hemispherical scan path, with smoothing function.	184
Figure 5.18: The point cloud produced using 40° aperture SAFT processing and the hemispherical scan path, with no smoothing function. Large errors caused by side lobes are present on the far left and right of the image, as well as at the top.	185
Figure 5.19: Four viewing angles of the compared point cloud produced using TFM processing and the hemispherical scan path, with no smoothing function. The higher thresholding associated with TFM processing resulted in a sparser point cloud.	186
Figure 5.20: Mean errors for the point clouds associated with the two scan paths, with smoothed and unsmoothed examples. The most important factor in mean error was the probe path, with the cuboidal path producing the smallest errors.	187
Figure 5.21: Result of error reduction procedure on unsmoothed, TFM point cloud. The procedure has not eliminated the outliers, instead removing the true surface.	188
Figure 5.22: The mean error achieved for data compared with the laser-scanned reference mesh for the cuboidal scan path. ‘Smoothing’ refers to the application of a smoothing function to the 2D points forming contours in each image, ‘nearest neighbour’ denotes the application of the outlier elimination algorithm and ‘sampling’ indicates whether or not the mesh itself or a densely sampled point cloud representing the mesh was employed	189
Figure 5.23: Two viewing angles of errors resulting from matching TFM-derived (a) and 20° SAFT-derived (b) meshes with no smoothing and no outlier elimination applied that have been densely sampled and compared with the laser scan data.	190
Figure 5.24: Two viewing angles of errors resulting from matching 20° SAFT-derived meshes with no smoothing and no outlier elimination applied. One has been compared directly with the laser scan data (a), while the second (b) has been densely sampled then compared. Part (b) offers greater detail, as (a) has a low vertex density.	191

Figure 5.25: Two viewing angles of errors resulting from matching 20° SAFT-derived meshes with no smoothing and no outlier elimination applied. One set has been sampled and compared with the laser scan data (a), while the second (b) was compared with the CAD reference model. ...	193
Figure 5.26: Two viewing angles of errors resulting from matching 30° SAFT-derived sampled meshes with no smoothing and no outlier elimination applied. One set has been produced using the hemispherical scan path (a), while the second (b) was generated using the cuboidal scan path. The inaccurate surface point placement caused by use of the hemispherical path caused the construction of erroneous surfaces.	195
Figure 5.27: Errors resulting from matching 40° SAFT-derived (a) and TFM (b) sampled meshes with no smoothing and no outlier elimination applied with a hemispherical path. The meshes produced with the hemispherical path showed no pattern in the placement of inaccurate surfaces.	196
Figure 5.28: Errors resulting from matching 20° SAFT-derived sampled meshes with no smoothing and no outlier elimination applied. The colour limit has been set to 3mm to better expose the regions containing the highest error. Regions with a high gradient tend to have higher levels of error.	197
Figure 5.29: Representation of a possible contour extraction scenario, displaying the difference between points with smoothing applied (a) and those without (b), showing how smoothing can increase mean error.	198
Figure 5.30: The line of sight (blue line) of the probe relative to the sample at a number of positions along the circular path, illustrating how the hemispherical path allowed no access to the intercondylar surfaces.	200
Figure 5.31: The result of the wrapping procedure on the 20° SAFT-derived point cloud with no smoothing and no outlier elimination applied showing the incomplete nature of some of the surface meshes produced.	203
Figure 5.32: The incident (blue) and specularly reflected (green) rays, showing how certain regions can be rendered inaccessible due to specular reflection and the probe angle.	206
Figure 5.33: 20° SAFT images from the region described in Figure 5.32, showing the image threshold limits of -18dB (a) and -8dB (b), so that the angled surface was below the extraction limit.	206

Figure 6.1: Front view of the bovine sample (a), showing a tendon attachment site with non-osseous tissue still intact. Also shown is a side view of the sample (b), showing both small holes and the separation of the shaft and the epiphysis.	211
Figure 6.2: The separation of the femoral shaft and the distal epiphysis, with the remains of the connective cartilage visible on both portions.	212
Figure 6.3: The path employed for inspection of the bovine distal femur, showing the six divisions of code execution.	213
Figure 6.4: A typical TFM image from the data set (a), and the result of surface extraction using both the method described in Chapter 5 (b) - the ‘Top’ method - and that herein (c) – the ‘Bottom’ method.	216
Figure 6.5: Comparison of the mean error attained using different contour extraction methods. ‘Top’ refers to the extraction method described in Chapter 5, due to the extracted points being closer to the top of the image. ‘Bottom’ refers to the method described in Chapter 6 for the same reason. The compared point clouds employed no smoothing function and had the nearest neighbour point elimination algorithm applied.	218
Figure 6.6: Comparison of mean error, standard deviation and maximum error for the four processing methods. The values were calculated as a mean of every parameter permutation, with a general deterioration visible as aperture increases.	219
Figure 6.7: Two viewing angles of the compared point cloud produced using 30° aperture SAFT processing, with no smoothing function at 1% velocity.	220
Figure 6.8: Two viewing angles of the compared point cloud produced using 30° aperture SAFT processing, with no smoothing function at 3% velocity. The density of the point cloud was decreased significantly relative to that in Figure 6.7.	221
Figure 6.9: Two viewing angles of the compared point cloud produced using TFM processing, with no smoothing function at 1% velocity. This shows an even greater decrease in density than that seen in Figure 6.8	222
Figure 6.10: Comparison of the mean error attained using different contour extraction methods. The compared meshes employed no smoothing function, did not have the nearest neighbour point elimination algorithm applied and employed a 1% TCP velocity.	223

Figure 6.11: Comparison of mean error, standard deviation and maximum error for the four processing methods. The values were calculated as a mean of every parameter permutation.	224
Figure 6.12: Four viewing angles of absolute errors resulting from matching 20° SAFT-derived mesh with no smoothing, no outlier elimination and using a single top side pass.	225
Figure 6.13: Two viewing angles of errors resulting from matching 30° SAFT-derived sampled meshes with no smoothing, no outlier elimination applied and a single top side pass. One mesh has been produced using a 1% TCP velocity (a), while the other has been created using a 3% velocity (b), showing how the decrease in point cloud density can be detrimental to the mesh reconstruction.	226
Figure 6.14: Four viewing angles of the 20° SAFT-derived mesh with no smoothing, no outlier elimination, using a single top side pass and using the ‘Bottom’ contour extraction method. The errors displayed are from comparison with the mesh of the same parameters, but created using the ‘Top’ extraction method. Negative errors denote that the reference mesh (the ‘Top’ mesh) was outside the surface, while positive numbers imply that the reference mesh was inside the compared mesh.	228
Figure 6.15: Two viewing angles of errors resulting from matching sampled meshes with no smoothing, with no outlier elimination applied and a single top side pass, created using a 1% TCP velocity. One mesh has been produced using 40° SAFT (a), while the other has been created using TFM (b). .	229
Figure 6.16: A typical 20° SAFT image from the bovine scan. The red points show the points extracted using the ‘Bottom’ method, while the yellow show that for the ‘Top’ method. As the ‘Top’ points are higher, this can result in a larger overall surface reconstruction.	231
Figure 6.17: Two point clouds resulting from 20° SAFT processing, 1% TCP velocity, no smoothing, with one using a single top side pass (a) and the other using a double (b), showing an increase in density in regions of high error.	232
Figure 6.18: Typical regions of high error displayed in red (a, c, e, g) and the accompanying incomplete area in the reference laser scan data.	235

Figure 6.19: Causes of incompleteness in the laser scan data denoted with a red arrow. One section shows general degradation of the osseous tissue caused by boiling (a), while the other shows remnants of soft tissue (b).	236
Figure 7.1: Custom coloured MRI images showing a coronal (a) and sagittal (b) cross section.	243
Figure 7.2: Segmented images of both samples, with materials separated by lines.	245
Figure 7.3: Final segmented images before input to PZFlex (a, b) and after material assignment in the image importation tool (c, d).	246
Figure 7.4: The dimensions used to scale the pixel coordinates in the materials table; the maximum width of the femoral condyle (a) and that of the medial condyle (b).	249
Figure 7.5: Every probe position (grey rectangle) and tissue map employed for the 2.25MHz probe simulations. Parts (a) and (b) are not of the same scale as (c-h).	251
Figure 7.6: Probe position and tissue map employed for the 5MHz and 1MHz probe simulations.	252
Figure 7.7: Part (a) shows the TFM image from part (e) of Figure 7.5 and part (b) shows the two contour extraction methods, with black representing the new extraction method.	254
Figure 7.8: Upsampling of the reference model, with the original points in blue and the new, upsampled points in red.	256
Figure 7.9: SAFT (a) and TFM (b) images from the FMC resulting from the image shown in Figure 7.5 (b), showing a more accurate surface representation with TFM.	257
Figure 7.10: Composite images of TFM images of knee 1 (a), SAFT images of knee 1 (b), TFM images of knee 2 (c) and SAFT images of knee 2 (d).	258
Figure 7.11: Errors resulting from the comparison of the data and the reference model form knee 1, showing the SAFT-derived points obtained using the ‘New’ (a) and ‘Top’ (b) contour extraction methods. Also shown are the TFM-derived points, found using the ‘New’ (c) and ‘Top’ (d) extraction methods.	260
Figure 7.12: Errors resulting from the comparison of the data and the reference model from knee 2, showing the SAFT-derived points obtained using the	

‘New’ (a) and ‘Top’ (b) contour extraction methods. Also shown are the TFM-derived points, found using the ‘New’ (c) and ‘Top’ (d) extraction methods.	261
Figure 7.13: SAFT images created using 1MHz (a), 2.25MHz (b) and 5MHz (c) probes. Also shown are TFM images produced using 1MHz (b), 2.25MHz (d) and 5MHz (f) probes.	262
Figure 7.14: Errors resulting from the comparison of the data and the reference model from knee 1, with the columns showing 1MHz, 2.25Mhz and 5MHz probes, from left to right. The rows display SAFT processing with the ‘Top’ extraction method, SAFT processing with the ‘New’ extraction method, TFM processing with the ‘Top’ extraction method and the TFM processing method with the ‘New’ extraction method, from top to bottom.	264
Figure 7.15: Summation of the pressures at all time steps for the 2.25MHz probe on knee 2, showing how ultrasonic energy inaccurately penetrated to the back of the patella.	267
Figure 7.16: Snapshots of the travelling pressure wave and the associated A-scans for the 1MHz (a, c) and 5MHz (b, d) probes, firing and receiving on element 64. There is a clear decrease in the width of the wave packet with the 5MHz transducer.	269
Figure 7.17: Simulated surface reflections from a sphere using 2.25MHz (a) and 5MHz (b) 64 element probes, showing more concentrated and intense side lobes in part (b).....	271
Figure 8.1: The apparatus used to retrieve reference FMC data from the bare condylar surface.	275
Figure 8.2: The apparatus used to retrieve FMC data from the condylar surface whilst situated in soft tissue mimicking material.....	276
Figure 8.3: Formation of the first two triangles during the contour extraction, with the first being solid blue and the second being dashed red.	277
Figure 8.4: Graphical representation of the ‘Radial’ extraction method. From the original image (a), a centre point is created. From this, a triangle is produced (b) and any pixels above the threshold (yellow points) within the triangle are identified (c). The distance of each of these pixels to the centre point is calculated and all the points are discarded, with the exception of that with the	

smallest distance (shown in red). This process is repeated over 180°, creating the final contour.	278
Figure 8.5: SAFT (a, c, e) and TFM (b, d, f) images resulting from FMCs captured with bare bone (a, b), non-contact tissue mimic (c, d) and contact tissue mimic (e, f).	280
Figure 8.6: SAFT (a, c, e) and TFM (b, d, f) images resulting from FMCs captured with bare bone (a, b), non-contact tissue mimic (c, d) and contact tissue mimic (e, f).	282
Figure 8.7: SAFT (a, c, e) and TFM (b, d, f) images resulting from FMCs captured with bare bone (a, b), non-contact tissue mimic (c, d) and contact tissue mimic (e, f). Also shown are the surface representations, extracted using the ‘Top’ (yellow), ‘Radial’ (black) and ‘New’ (red) methods.	284
Figure 8.8: SAFT (a, c, e) and TFM (b, d, f) images resulting from FMCs captured with bare bone (a, b), non-contact tissue mimic (c, d) and contact tissue mimic (e, f). Also shown are the surface representations, extracted using the ‘Top’ (yellow), ‘New’ (black) and ‘Radial’ (red) methods.	286
Figure 8.9: A quarter of the tissue mimic material, cut from the container using a scalpel, displaying an isotropic dispersion of powder.	287
Figure 8.10: The partially exposed sample after removal of a quarter of the tissue mimic material, showing small gaps between the osseous tissue and the mimic.	288
Figure 8.11: Proposed setting procedure for achieving multiple layers of soft tissue mimic, each with different scatterer densities.	291
Figure 9.1: The KUKA LBR iiwa 14 R820.	299
Figure 9.2: Colour slices of the sagittal (a) and coronal (b) planes of the knee joint. Also shown are corresponding maximum pressure maps (c, d).	301
Figure 9.3: Knee in extension (a) and flexion (b), showing the possible regions for inspection for movement analysis in red.	302
Figure 9.4: Possible ‘box’ design for ultrasound CT, with four probes (a), each with two axis of motion (b), allowing for full coverage using the cuboidal path employed herein.	305
Figure 9.5: Method of employing two probes at once, with the leg ensuring transmissions from one probe do not interfere with those of the other.	306
Figure 9.6: Ring configuration of ultrasound computed tomography, with two probes with linked rotational and linear movement.	307

Figure 9.7: Action of balloon coupling, with deformation allowing for continuous contact.	308
Figure 9.8: Possible double robot configuration allowing for full access to the knee joint.	309
Figure 9.9: Possible position of immobilisation straps that would ensure no movement of the joint during ultrasonic bone surface capture.	312

List of Tables

Table 3.2: Specifications of the Kuka KR5 Arc HW 6 axis robots in CUE (C Mineo et al., 2012).	74
Table 3.3: Properties of the four available probe types.	78
Table 3.4: summary of the results of the comparison between the Vicon positional data and the robot positional data.	100
Table 4.1: The error results from the comparison between the TFM-derived point clouds and the reference model.	143
Table 4.2: The error results from the comparison between the SAFT-derived point clouds and the reference model.	143
Table 4.3: Summary of the comparisons for the different samples and parameters.	146
Table 5.1: Results of the original point cloud data from the cuboidal scan, before the application of surface meshing. The bracketed angle next to all SAFT results denotes the angular aperture employed.....	177
Table 5.2: Results of the original point cloud data from the hemispherical scan, before the application of surface meshing. The bracketed angle next to all SAFT results denotes the angular aperture employed.	182
Table 6.1: The parameters and results from the five smallest mean errors for the composite sample.	238
Table 6.2: The parameters and results from the five smallest mean errors for the bovine sample. All values were achieved using a 1% TCP velocity and the ‘Top’ contour extraction method.	238
Table 7.1: Material properties and the associated sources as employed in the PZFlex model.....	247

Table 7.2: Results of the comparison between the resultant 2D point clouds and the reference model.	259
Table 7.3: Results of the comparison between the resultant 2D point clouds from different frequencies and the reference model.	263

List of Abbreviations

ACL	Anterior Cruciate Ligament
API	Array Performance Indicator
AUT	Automatic
AUT	Automatic External
EXT	
CPU	Central Processing Unit
CUE	Centre for Ultrasonic Engineering
CIRC	Circular
CUDA	Compute Unified Direct Architecture
CT	Computed Tomography
CAD	Computer Aided Design
CAM	Computer Aided Manufacture
CAOS	Computer Assisted Orthopaedic Surgeries
CMM	Coordinate Measuring Machine
DSL	Diagnostic Sonar Ltd.
FIRST	Facility of Innovation and Research in Structural Testing
FPGA	Field Programmable Gate Arrays
FE	Finite Element
FMC	Full Matrix Capture
FRD	Full Raw Data
FRD	Full Raw Data
GPGPU	General Purpose Graphics Processing Unit
GAG	Glycosaminoglycans
GPU	Graphics Processing Unit
HDL	Hardware Description Language

ICP	Iterative Closest Point
KRL	Kuka Robot Language
PZT	Lead Zirconate Titanate
LED	Light Emitting Diode
LIN	Linear
MRI	Magnetic Resonance Imaging
T2	Manual High Velocity
T1	Manual Reduced Velocity
MEX	Matlab Executable
M-SAF	Multi-Element Synthetic Aperture Focussing
NDT	Non-Destructive Testing
OA	Osteoarthritis
PDM	Point Distribution Model
PSF	Point Spread Function
PTP	Point to Point
PVA	Polyvinyl Alcohol
PCL	Posterior Cruciate Ligament
PG	Proteoglycans
RSI	Robot Sensor Interface
RIO	Robotic Arm Interactive
SNR	Signal to Noise Ratio
SPL	Spline
SSM	Statistical Shape Model
SA	Synthetic Aperture
SAFT	Synthetic Aperture Focussing Technique
SF	Synthetic Focussing
SRA	Synthetic Receive Aperture

STA	Synthetic Transmit Aperture
TCP	Tool Centre Point
TFM	Total Focussing Method
TFM	Total Focussing Method
UKA	Unicompartmental Knee Arthroplasty

Chapter 1 Introduction

1.1 Project Background

Osteoarthritis (OA) of the knee is a disease affecting almost five million people in the UK (ArthritisResearchUK, 2013). When the symptoms associated with the condition become disabling, knee replacement is usually chosen as the method of treatment. While traditional surgeries are relatively successful, the cutting methods employed are inaccurate and outdated. Robotic surgery offers an alternative to this, providing greater cutting accuracy and, arguably, a better patient outcome. A patientspecific model of the knee is required for such operations and currently Computed Tomography (CT) is the preferred method of preoperative imaging. This method, however, is costly and applies ionising radiation. As such, ultrasound is proposed as an alternative imaging modality. Ultrasound imaging is well used in both medical imaging and Non-Destructive Testing (NDT), but the two fields have developed in relative isolation. Due to this, techniques common in one field are rarely used in the other. Therefore, there is the potential to draw knowledge from both areas to improve imaging outcomes. This was the rationale for this thesis.

This chapter presents the background to the project and demonstrates the need for the proposed ultrasound system. Osteoarthritis and the associated pathology are discussed, as is robotic knee arthroplasty. Additionally, preoperative imaging methods, their capabilities and weaknesses are reviewed. Further to this, the current state of ultrasound imaging in both NDT and biomedicine is presented. Finally, the aims of the project are introduced, based on areas where current research is lacking.

1.1.1 Anatomy of the Knee

The knee is a large, load bearing joint which of the lower limbs which allows for smooth locomotion. Often considered a simple hinge joint, the knee is in fact more complex, with variable axes of flexion and extension due to the allowance of some

medial and lateral rotation (Standring, 2008). The joint consists of three bones: the femur (thigh bone), the tibia (shin bone) and the patella (kneecap). Each of these can be seen in Figure 1.1, which shows a sagittal section of the knee joint with the distal femur and proximal tibia. In these regions, the bones are comprised of cortical and cancellous bone. Cortical bone has a highly compact structure and forms a dense shell around the cancellous bone, which has a lower density and has a sponge-like structure. As the whole bone is not dense, it can dampen the forces applied to the joint during locomotion (J. Smith, 2011).




Figure 1.1: A sagittal section of the knee joint, showing the two bone types in the femur, tibia and patella, as well as the location of the bones relative to one another (Gray, 1918).

Both the femur and tibia possess two condyles which act as an articulating interface during flexion and extension, as can be seen in Figure 1.2. All of these condyles are covered in articulating cartilage, a low friction, wear resistant soft tissue which allows for smooth articulation and provides load bearing (Pearle, Warren, & Rodeo, 2005). The femoral condyles are rounded, while the tibial condyles form a plateau, as can be

seen in Figure 1.1. These shapes are complementary, allowing the femoral condyles to articulate smoothly with the tibial condyles during flexion and extension.




Figure 1.2: Anterior view of the right knee, with patella and associated ligaments omitted (Gray, 1918).

While the condylar interfaces provide the main articulation of the knee joint, the patellofemoral articulation also contributes (Standring, 2008). The patella operates as a region of protection for the front of the knee and can increase the leverage of the quadriceps (Scuderi & Tria, 2010). The group of bony structures comprising the knee also includes the fibula. However, this bone has no load bearing function.

One of the soft tissue structures within the knee joint is the meniscus: a crescent-shaped region of fibrocartilaginous tissue which rests on the outer parts of the tibial plateau, as can be seen in Figure 1.2. Towards the centre of the condyles, the edges of the meniscus are thin, while the outer edges are thick. This structure serves a number of purposes

including spreading applied forces by increasing the region of contact between the tibia and femur, distribution of synovial fluid, inhibition of soft tissue becoming trapped in the joint and an increase of the conformity between the tibial and femoral condylar surfaces.

The femur, tibia and fibula are held together by ligaments, which are connective tissues that specialise in connecting bones (Motesharei, 2014). The four main ligaments of the knee joint are the Anterior Cruciate Ligament (ACL), the Posterior Cruciate Ligament (PCL), the Medial Collateral Ligament (MCL) and the Lateral Collateral Ligament (LCL). These are displayed in Figure 1.3, though it should be mentioned that the MCM is labelled the tibial collateral ligament and the LCL is termed the fibular collateral ligament. Working to hold the bones together and maintain stability, the ACL and PCL work together to maintain the correct placement of the femoral and tibial condyles. The PCL prevents anterior movement of the femur, while the PCL prevents posterior movement. The MCL and PCL tighten when the leg is extended, preventing valgus and varus forces respectively.




Figure 1.3: Posterior view of the left knee, showing the four main ligaments – ACL, PCL, MCL, LCL (Gray, 1918).

The control of the motion of the joint is administered by the muscles. Extension is controlled by the quadriceps, a group of four muscles, as shown in Figure 1.4 (a). The hamstrings are a group of three muscles (semitendinosus, semimembranosus and biceps femoris) which control knee flexion. These are located posteriorly, while the major extensors are positioned anteriorly. While the main role of these muscle groups is to produce motion, they also serve to stabilise the leg during locomotion by the counteraction of each group against the other (J. Smith, 2011).




Figure 1.4: Part (a) shows the anterior femoral musculature, showing three of the four muscles termed the quadriceps: rectus femoris, vastus lateralis and vastus medialis, with vastus intermedius being covered by rectus femoris. Part (b) shows the posterior muscles, including the hamstring group, where semimembranosus is labelled '1', semitendinosus '2' and biceps femoris '3' (Gray, 1918).

1.1.2 Osteoarthritis

Osteoarthritis is a degenerative disease that causes pain to the sufferer and can eventually lead to complete loss of mobility in the affected joint (J A Buckwalter & Mankin, 1997). It is thought to be the most common cause of disability amongst the elderly in the UK (Peat, McCarney, & Croft, 2001), with as much as 60% of over 65s suffering some form of the disease (Croft, 1990), 18.4% of over 70s suffering from OA of the knee (Anderson et al., 1993) and 10% of those over the age of 55 suffering disabling symptoms caused by knee OA (Jordan, 2003). The strong association with age suggests that not only the number, but also the proportion of the population suffering from the disease is set to rise, with a predicted continuation in growth of the percentage of the population over the age of 65 in the UK and across the world (Lee, 2007). This suggests that the cost to the NHS and other health services of treating the disease will continue to grow.

In addition to this, the occurrence of OA also increases with obesity (Wright, Riggs, Lisse, & Chen, 2008). Suspected causes are believed to be either increased mechanical burden or metabolic effects on cartilage. Both are due to excess fat in the tissue (Manek, Hart, Spector, & MacGregor, 2003). Labelled a pandemic, obesity is becoming more prevalent, with a predicted 11 million more people likely to be categorised as obese in the UK by 2030 (Y. C. Wang, McPherson, Marsh, Gortmaker, & Brown, 2011). This will add further to the number of OA sufferers and to the costs of caring for them. Ethnicity also plays a role in the chances of developing of the disease, with variation between different racial groups being observed (Croft, 1990). However, the differences between ethnicities may also be affected by the likelihood of obesity in these groups (Wright et al., 2008), which emphasises the effect of obesity on the likelihood of suffering from the condition. While the disease presents itself almost entirely in the elderly, it is also found occasionally in younger subjects. These people have usually participated heavily in sports, with the likelihood of the disease greatly increasing in those who sustain injuries (Caine & Golightly, 2011). Levels of physical exercise also have a bearing on the progression of OA, with former top level athletes at risk of developing relatively premature OA (Kujala et al., 1995).

It is at the microscopic level that changes due to OA are first encountered (National Collaborating Centre for Chronic Conditions, 2008). Healthy cartilage can be observed in Figure 1.5, showing high levels of glycosaminoglycans (GAG) in all areas, intact cartilage surface and collagen fibrils closely packed and oriented parallel (if in the superficial zone) and perpendicular (if in the deep zone) to the cartilage surface. The early stages of OA see depletions in the concentration of proteoglycans (PG). These are molecules found within cartilage made up of a core protein and a GAG chain, responsible for the control of the flow of water in and out of the cartilage (Muir, 1977). This loss of PGs causes an increase in the water content of the cartilage and is accountable for exposure of the superficial layer of the cartilage, through disorganisation of the fibril network. This causes the cartilage to soften and lose compressive strength (J A Buckwalter & Mankin, 1997). In addition to this, the subchondral bone begins to thicken, reducing the volume of compressible cartilage. In the latter stages of the disease, the superficial layer begins to erode completely, as can be seen in part three of Figure 1.5. This exposes the cartilage of the intermediate and deep zones to large stresses and strains, causing further damage (Arokoski, Jurvelin, Väättäinen, & Helminen, 2000). These processes continue to progress, with further PG depletion, fibrillation of the cartilage surface and worsening osteosclerosis of the subchondral bone.




Figure 1.5: Visual description of the progression of OA at the microscopic level, with ovals representing chondrocytes, dots representing GAGs and curved lines collagen fibrils (Arokoski et al., 2000). Three stages of OA progression are illustrated, showing degradation of the superficial layer, loss of GAGs and a further loss of chondrocytes.

When the cartilage matrix is first disrupted, chondrocytes respond by increasing synthesis and degradation of the matrix. In the latter stages of the disease, however, the chondrocyte response is non-existent and the cartilage continues to wear away until gone. This is represented in Figure 1.5 by a decrease in the density of the chondrocytes as the disease progresses. Eventually, this will lead to bone on bone interactions in joints, which will cause significant pain to the sufferer. The subchondral bone sees changes including increased density and formation of osteophytes. Such alterations ultimately lead to changes in the shape of the entire joint, which, in turn, causes instability and, more importantly, pain (J A Buckwalter & Mankin, 1997).

1.1.3 Robotic Knee Arthroplasty

A number of surgical procedures exist which are intended to alleviate the pain of OA and to return the patient as close to normal function as possible. These include total knee arthroplasty (TKA), unicondylar knee arthroplasty (UKA), osteotomy, muscle release, debridement and allotransplantation and autotransplantation of grafts (Joseph A Buckwalter & Lohmander, 1994). Each of these techniques has associated benefits and weaknesses that suit individual patients. For example, proximal tibial osteotomy is a procedure which aims to shift weight from the damaged part of the knee to a healthy part by cutting away a wedge of or implanting a wedge into the bone from underneath the healthy condyle. This is considered appropriate for younger, healthier patients (Callahan, Drake, Heck, & Dittus, 1995). Such cases, however, are in the minority, with the majority of those with severe knee disease (ie. OA progression that warrants surgery) being over 65 (Juni, 2003). As such, pain relief and the ability to perform simple tasks, such as climbing the stairs, becomes paramount for the patient (Dieppe et al., 1999; Williams & Garbuz, 2010). With these goals in mind, the literature suggests that total knee arthroplasty is the most reliable option for most patients (Bachmeier et al., 2001; Rand & Ilstrup, 1991; Woolhead, Donovan, & Dieppe, 2005).

TKA is considered to be the most trustworthy of the possible surgical treatments of OA. It has earned this trust through decades of reliability, with one study reporting survivorship of 93% of all types of TKA after five years (Rand & Ilstrup, 1991), another giving 92% at ten years (D. J. Weir, Moran, & Pinder, 1996) and a third reporting 96% also at ten years (Malkani, Rand, Bryan, & Wallriches, 1995). This procedure involves removing a layer of bone from the surface of both the femoral and tibial condyles, which is then replaced by a prosthesis. In traditional surgeries, the bone is removed to a specific thickness and shape using an extramedullary or intramedullary jig, which allows for a saw to cut away the unwanted bone in the correct positions through guides (Williams & Garbuz, 2010). In addition to this procedure, surgeons may also opt for unicompartmental knee arthroplasty (UKA), in which only one of the condylar compartments undergoes surgery. This allows the healthy bone and cartilage of the other joint compartments to remain intact and for the anterior cruciate ligament (ACL) and posterior cruciate ligament (PCL) to be maintained. Employing the most suitable

procedure could help decrease the likelihood of complications during surgery and increase functionality (Callahan et al., 1995). It should be noted, however, that UKA has suffered from poor survival rates relative to TKA (Herring, Laskin, & D, 1978; Insall & Aglietti, 1980), meaning the implants, rather than the patients have not failed. However, UKA has had a smaller test group, with only 8% of the knee arthroplasties in England and Wales being UKAs (Ellams, Swanson, & Young, 2011).

The alignment of the knee achieved during surgery is an important factor in determining the success of the procedure (Denham, Jeffery, & Morris, 1991; Fang, Ritter, & Davis, 2009; Lotke & L, 1977). While the traditional cutting methods yield high levels of success, they can only achieve alignment within two or three degrees of the 'perfect' alignment (Reed & Gollish, 1997). This is because these procedures require the use of anatomical land marks and human decision making during the cutting procedure, which can lead to poor alignment and inaccurate placement (Citak et al., 2012; Ill et al., 1995). These are factors thought to be related to poor patient outcome (Bäthis et al., 2004; Fang et al., 2009; Huang et al., 2012). As such, improving the process by which the bone is cut should, in turn, improve the accuracy of the final alignment and hence the clinical outcome.

Robotic surgical systems offer a more accurate method of removing unwanted bone. This produces significantly higher cutting accuracies than traditional procedures (Citak et al., 2012; Harvie, Sloan, & Beaver, 2012). Robotic systems are, therefore, associated with improved patient outcome (Choong, Dowsey, & Stoney, 2009; J. Decking et al., 2004; R. Decking, Markmann, Fuchs, Puhl, & Scharf, 2005). There are a number of robotic knee arthroplasty systems available (Netravali, Shen, Park, & Bargar, 2013) and, therefore, a number of different procedures. However, the partial knee MAKOPlasty procedure, using MAKO's RIO (Robotic Arm Interactive orthopaedic) system (Mako Surgical Corp., Fort Lauderdale, FL, USA) is the world leader and will be considered herein (Hagag, Abovitz, Kang, Schmitz, & Conditt, 2011). The system is displayed in Figure 1.6. The preoperative procedure begins by obtaining a three dimensional, patient-specific bone model of the knee joint by way of CT. This can then be used to determine the optimal size of implant and the desired placement thereof. The actual surgical procedure begins by inserting bone pins into both the femur and tibia, to which

optical tracking arrays are attached. These arrays are tracked by an infrared stereoscopic optical tracking system. Landmarks on the surface of the skin are then registered using a blunt probe which also has an optical tracking array attached. At this point, the skin is incised, allowing access to the bony surfaces of the knee and checkpoint pins are inserted into the femur and tibia (Roche, O’Loughlin, Kendoff, Musahl, & Pearle, 2009). At these points, the same probe is employed to register bony landmarks and the surface within the incision. Doing this matches the positional data gathered via the optical tracking system to the CT data set, meaning that movement of the optical markers can be translated to movement of the three dimensional CT model of the knee of the patient (Cobb & Pearle, 2013). The position of the robot relative to the knee is achieved by way of fixed optical tracking arrays. The dynamic positioning of the tool tip is achieved by way of reconstruction using the angles read on each of the six rotary joints.



Figure1.6: the components of the MAKO RIO system employed in partial knee MAKOPlasty, showing (from left to right) the robotic arm with burr, the stereoscopic infrared optical tracking system and the control panel.

At this point, the unwanted bone can be milled away, which is achieved by a high speed cutting burr which is attached to the end effector of the robotic arm. The robotic arm has six degrees of freedom, allowing the surgeon to have full control of both the movement of the burr and the operation thereof (Hagag et al., 2011). However, when the burr enters a region where cutting is not wanted, haptic feedback is employed so that the movement of the arm is actively restrained. While the surgeon may override this by force, it usually serves to restrict the cutting of bone to those volumes that need to be removed. Trial prosthetic condylar heads are then fitted to ensure the implant sizing and placement has been achieved as planned. If this is the case, matching real parts are implanted and the wound is closed. The entire process involved in a MAKOPlasty partial knee replacement is described diagrammatically in Figure1.7.

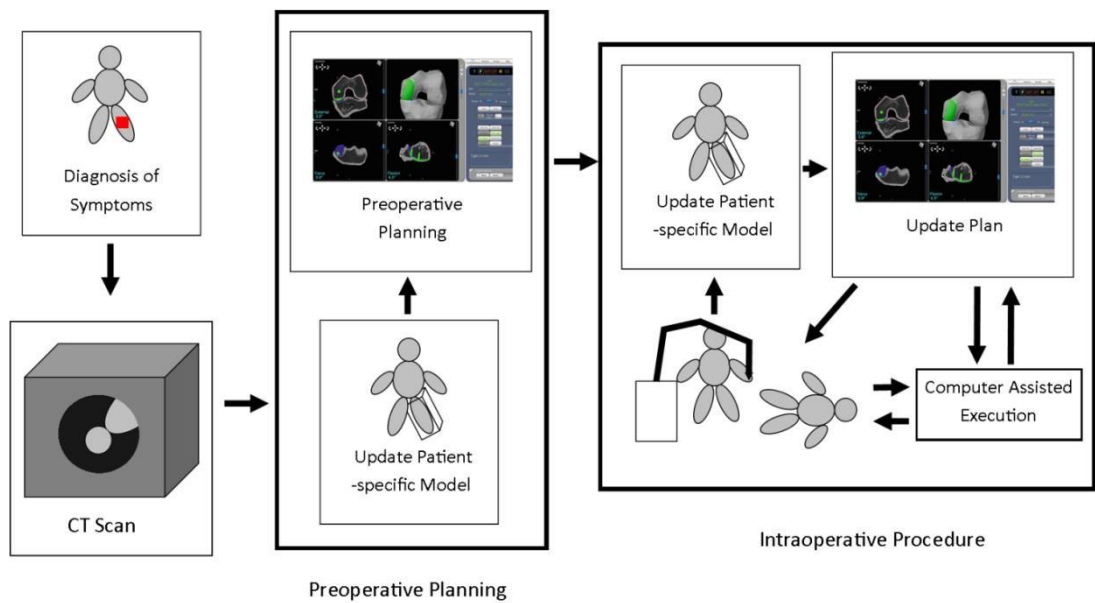


Figure1.7: Flow diagram of the processes involved in robotic knee arthroplasty, from diagnosis to surgical procedure.

1.1.4 Preoperative Imaging Methods

At the heart of this system and all image-based Computer Assisted Orthopaedic Surgeries (CAOS) is preoperative imaging. These systems typically employ CT to produce the patient specific three dimensional models (Citak et al., 2012; Netravali et al., 2013), with MAKOPlasty procedures using CT exclusively (Cobb & Pearle, 2013; Hagag et al., 2011; Pearle, Kendoff, Stueber, Musahl, & Repicci, 2009; Roche et al., 2009). For imaging bone and, more specifically, building three dimensional likenesses of the region of interest for the purposes of surgery planning, CT has been well used since the 1980s. It is a useful tool, as it allows relatively accurate visualisation of sections of bone and hence the possibility of simulating the effects of a biomedical implant and the positioning thereof (Fishman et al., 1992). It has been shown to be capable of depicting the geometry of bony surfaces, such as the femoral condyles, with one study achieving an impressive accuracy of 0.3mm (Siu, Rudan, Wevers, & Griffiths, 1996). However, in normal clinical settings, using commercially available CT scanners, the accuracies attained vary between 0.5mm and 1mm (Viceconti, Zannoni, Testi, & Cappello, 1999), with a typical resolution of $0.52mm \times 0.52mm \times 2mm$ (W. Yao, Abolmaesumi, Greenspan, & Ellis, 2005).

The geometries of the volumes under investigation are found using segmentation, which is a method of delineation of certain structures within an area of interest. This is performed slice by slice on 2D images, which can then be combined to create a 3D model. The accuracies stated are a measure of the accuracy of the segmentation, which is influenced not only by the properties of the CT machine itself, but also by the method of segmentation. These measurements of accuracy presented herein are based on the mean distance between the automatically and manually segmented bone (Hu, Hoffman, & Reinhardt, 2001; W. Yao et al., 2005). Although manual segmentation is seen here as being 'perfect', it, too, would be affected by poor image quality. Automated segmentation, however, is advantageous, in that the reproducibility between two operators will not be an issue and the segmentation process will take less time than would be required for the manual equivalent. The accuracy of these processes is of vital importance to the success of the procedure.

Imaging using CT has a number of disadvantages, including the fact that the contrast between soft tissues is low. Consequently, preoperative planning cannot take into account alignment of the soft tissues, but only that of the bony structures (Roche et al., 2009). Another prominent disadvantage of CT is that patients must receive a dose of ionising radiation during the imaging. A standard abdominal scan will apply a dose of 10mSv to the subject, compared to the national yearly average dose from natural sources of 2.23mSv (Watson, Jones, Oatway, & Hughes, 2005). Knee arthroplasty procedures would require full leg scans which give a standard dose of 2.7mSv, but have been reduced to as little as 0.53mSv (Henckel et al., 2006). The radiation dose required to preoperatively image for arthroplasty is, then, lower than that needed for other forms of CT imaging. However, the effect of this should not be underestimated, with such scans adding to cumulative radiation exposure which can increase the risk of cancer in some individuals (Sodickson, Andriole, Prevedello, & Nawfel, 2009). While diagnostic X-rays alone are thought to be responsible for 0.6% of the risk of cancer in the UK and higher in other parts of the world (de González & Darby, 2004), the use of CT (which employs many X-rays in one scan) is thought to have increased this to as much as 2% in the US (Albert, 2013; Brenner & Hall, 2007).

In addition to this, there is, importantly, a significant cost associated with CT. The cost per scan varies greatly from country to country but one study in the states found scans of single regions of the body ranging from \$1,400 - \$2,500 (Fred, 2004), which would have to be paid by the patient or his insurance company. In those countries that offer universal health care, the cost will not be charged to the individual but will put strain on the budget of the governing health board. The cost of the CT scanners bought for use in England and Wales was £579,000 each. The average cost of a single scan ranges from £54-268, with the average at £121 (M. Davies et al., 2011). In addition to this, these machines have a life time of between 7 and 10 years, meaning purchase costs will remain high for the NHS.

While CT is the most typical form of preoperative imaging in CAOS, Magnetic Resonance Imaging (MRI) is also a possible method of imaging bony surfaces, although it is used mostly to image soft tissues. MRI has also been employed in evaluating macro-scale bone geometry, with one study achieving an accuracy of

1.44±1.1mm, with a voxel size of $1.37\text{mm} \times 1.37\text{mm} \times 5\text{mm}$ (Magnenat-thalmann & Schmid, 2008). While the results show that this method is capable of producing useful results, there is a drawback, in that MRI acquires the intensities not from the bone surface, as is the case with CT, but from water molecules within the internal structure. As bone is non-uniform, having both cancellous and compact parts, the returned intensities can cause complications in the segmentation procedure (Magnenat-thalmann & Schmid, 2008).

There are a number of other disadvantages associated with MRI. One of these is the use of a powerful electromagnet. This causes many problems, such as the need for a specially designed room to house the equipment safely, so that no metallic objects are in the vicinity of the machine. This also increases the cost of these systems, due to the associated expenditure of building a dedicated room. The magnet also determines the dimensions of the machine, which includes the size of the aperture into which the patient must fit. As this space is usually quite small, those suffering from claustrophobia might have to be sedated to remain still. The small aperture size also restricts morbidly obese patients from using the device. Open MRI, in which the patient is placed between two plates, has, however, helped to solve some of these problems (Hailey, 2006). The superconducting electromagnet is required to switch gradients quickly, which produces a very loud hammering noise of around 115 dB (Counter, Olofsson, Grahn, & Borg, 2005). With 90dB often considered the cut off between safe and dangerous acoustic exposure (Clark, 1991) and scan times lasting from 15 to 90 minutes, the risk of hearing damage is such that headphones must be worn. Additionally, the fast changing gradients that cause huge changes in the magnetic field that can induce oscillating currents in the patient's body. This may cause involuntary stimulation of the nerves and muscles. These effects can, however, be minimised by ensuring that the patient does not move during the procedure (Silva, Silva, Egito, & Carriço, 2006).

Another safety concern with MRI use is the effect the magnetic field can have on the function of biomedical implants. The magnetic qualities of such implants might make them malfunction or become displaced, causing harm or even death (Frank G Shellock, 2002). Due to the fact that orthopaedic implants tend not to be made with ferromagnetic

materials, they have been found to be of no risk to the patient and cause no artefact in the returned images (Mechlin, Thickman, Kressel, Gefter, & Joseph, 1984; F G Shellock & Crues, 1988). New implant materials should, however, be tested to ensure their safety. Due to the risk associated with this, clinicians must take great care in screening patients for contraindicating implants and should keep up to date with the growing number of devices and their safety (Sawyer-Glover & Shellock, 2000). This presents a continually expanding screening process and the difficulty of risk assessment when a patient has an implant.

The most debilitating drawback associated with MRI is the cost. The MRI scanners used in England and Wales cost £895,000, with the cost of a single scan ranging from £84-£472, with an average of £184 (M. Davies et al., 2011). As with CT, these machines have a predicted lifetime of from 7 to 10 years. This implies that the equipment will have to be continually renewed, adding constantly to the associated cost.

From the information presented, it is clear that the two currently available imaging methods that could be employed in preoperative imaging have some drawbacks. Any other imaging method that would wish to replace these methods would need to achieve similar levels of accuracy (ie. attain mean errors in the range of 1mm). In addition, the alternative method must also offer some advantage over the existing methods, as they have been proven to produce results of an acceptable standard. To improve on CT, any alternative method should eliminate as far as possible any danger to the patient. Further to this, a possible replacement technique should be available at a lower cost. Both CT and MRI suffer from the requirement of dedicated rooms, purchasing of expensive scanners and specially trained staff to administer the scans, making the cost of each scan very high. As such, an alternative system should be relatively portable, allowing for placement in non-dedicated rooms and for movement between different institutions. Finally, the systems should not be wholly reliant on specialist staff for operation. If safety requirements were lowered relative to both CT and MRI, as would be the case with an ultrasound system, this could be achieved.

An ultrasound system could offer the solution. Diagnostic ultrasound is known to be safe (Barratt et al., 2006b; McDicken, 1981), requiring no use of ionising radiation.

Due to the relative lack of safety concerns, an ultrasound system would not, in principle, require a dedicated area of operation. In addition, ultrasound systems are relatively inexpensive (Barratt et al., 2006b; Krysztoforski, Krowicki, SwiatekNajwer, Bedzinski, & Keppler, 2011). Ultrasound also shares with CT and MRI the advantage that it is noninvasive, eliminating the need for arthroscopic inspection and, in turn, the risk of further damage by inspection (Bae, Schumacher, & Sah, 2007).

1.2 Ultrasound

Ultrasound is often defined as sound with a pitch too high to be perceived with human hearing. The frequency at which audible sound becomes ultrasound is usually defined as 20kHz (Hughes, 2001). While humans are incapable of hearing such frequencies, other species, such as bats, employ it not only to hear, but to ‘see’ (Moir, Jackson, & Windmill, 2013). The ability to see using non-visual means was recognised as early as 1794 by Lazzaro Spallanzani and Charles Jurine who blindfolded bats and found they still maintained the ability to navigate successfully during flight (Griffin & Galambos, 1941; Kane, Grassi, Sturrock, & Balint, 2004). The first step towards realising the ability to generate and measure ultrasound came with the discovery of the piezoelectric effect in 1880 by Jacques Curie and Pierre Curie (Erikson, Fry, & Jones, 1974; Kholkin, Pertsev, & Goltsev, 2008). They discovered that electrical charges accumulate in certain materials when mechanical force is applied (Arnau & Soares, 2008). In 1881, it was predicted mathematically by Gabriel Lipmann that the converse would occur, such that when a voltage is applied to such a material, mechanical deformations occur. This was then proven experimentally by the Curies (Arnau & Soares, 2008).

The ability to ‘see’ using sound was first realised by Reginald Fressenden in 1912, with the invention of the Fressenden Oscillator (Frost, 2001). This device was essentially a loudspeaker operating in water, using electro-acoustic transduction. While it was used primarily for communication, it was also employed for echo ranging. Driven by the sinking of the Titanic (Kane, Grassi, et al., 2004), it gave ships the ability to see underwater obstacles. The onset of the First World War brought the ability to see underwater to the fore, with the capacity to detect submarines becoming paramount.

Paul Langevin and Constantin Chilowsky designed and built an echolocation system, capable of detecting submerged submarines. With Langevin being a former student of Pierre Curie, the system employed the piezoelectric effect via the excitation of quartz crystals. This system is considered to be the predecessor to all modern ultrasound probes (Kane, Grassi, et al., 2004).

Following the end of the Second World War, single element ultrasound systems had been developed for use in industrial non-destructive testing as flaw detectors for metals (Firestone, 1946). In the 40s and 50s, ultrasound had been investigated for biomedical uses, such as brain tumour detection (Dussik, 1941) and soft tissue characterisation (Kane, Grassi, et al., 2004). However, it was through methods in non-destructive testing that biomedical imaging became commonplace. Professor Ian

Donald of the University of Glasgow believed that the traditional clinical methods of diagnosis of conditions showing grossly indented abdomen were inadequate. Stemming from his interest in the work of John Wild, Donald postured that ultrasound could be used to determine whether the cause of the indented abdomen was a cyst or a myoma. This was put to the test when, in 1955, Donald travelled to Kelvin Hughes Ltd. in Glasgow with a range of samples. Using one of their single element flaw detectors, he was able to clearly see a difference in the readings between the two tissue types. This ability was actually used to correctly diagnose a patient who had been wrongly diagnosed using the traditional methods, which led to a landmark paper within diagnostic ultrasound (Donald, MacVicar, & Brown, 1958).

From here, Donald's research continued with the invention of automatic two dimensional scanning, the first ante-partum diagnosis of placenta previa using ultrasound, the development of a method of measuring the biparietal diameter of the foetal head and the first detection of early pregnancy using a full bladder to image the pelvic viscera (Kane, Grassi, et al., 2004; Kurjak, 2000). Equally as important as his own research, the work of Donald spawned an entire field of research within biomedical imaging, leading eventually to Doppler imaging, three dimensional imaging and duplex imaging (H. Azhari, 2010). With biomedical ultrasound now being commonplace in imaging of the heart, brain, liver kidney, foetal and reproductive

systems, it has taken its place alongside X-ray, CT, MRI and other nuclear methods as a standard form of medical imaging (Erikson et al., 1974).

While the origins of biomedical ultrasound imaging lie in ultrasonic non-destructive testing, research in the two fields has diverged from that point. Since its inception, biomedical ultrasound has grown at a fast rate, outpacing NDT quite substantially in terms of technological advancement. There are two reasons for this: the market size of biomedical ultrasound is larger than that of NDT and the imaging needs of the two fields are vastly different (Lines, 2003). The demand for ultrasound biomedical imaging is great, which has produced a large mass market. The market for NDT is, on the other hand, more niche, due to the smaller consumer group and the huge variety of imaging needs. A general dissimilarity that defines the imaging needs of the two fields is the variety of the materials and objects being imaged. Materials in NDT tend to have a wide range of impedances and velocities, while those in biomedical imaging are usually quite similar, seeing as soft tissue imaging is dominant. Additionally, the materials in biomedical imaging are usually isotropic with similar levels of attenuation, while in NDT anisotropy is possible, as are a wide range of attenuations. As well as this, NDT employs a number of wave modes, while biomedical imaging is usually limited to longitudinal waves. Finally, the subjects in NDT are usually static and remain so permanently. Those in biomedical ultrasound, on the other hand, are prone to movement, which can lead to mechanical deformation of the materials and objects under investigation (Lines, 2003). This, in turn, implies that while NDT can acquire data over a long period of time, biomedical imaging requires real time data visualisation.

The last difference in particular was fundamental in the early uptake of phased array transducers in biomedical imaging (Lines, 2003). Arrays allow for two dimensional images to be produced in real time (Drinkwater & Wilcox, 2006). While biomedical imaging came to require this, the associated cost outweighed the benefit of real time imaging in NDT, with real time inspection being preferred but ultimately considered unnecessary. However, NDT came to employ phased arrays as standard (Wilcox, Holmes, & Drinkwater, 2007) due to the ability to carry out a large number of inspection types with one probe (McNab & Campbell, 1987).

The two fields have swapped ideas (Jensen, Nikolov, Gammelmark, & Pedersen, 2006; Sutcliffe, Weston, Dutton, Charlton, & Donne, 2012) and NDT has started to catch up with biomedical ultrasound in terms of development. However, the fields still remain largely isolated (Lines, 2003).

1.3 Current State of the Field

Ultrasound imaging is routinely employed in the areas of cardiology, radiology, obstetrics, gynaecology; employing endoscopic, transesophageal, transrectal, laparoscopic, neurosurgical, vascular and intravascular imaging (Szabo, 2014). What is common to all these areas and imaging methods is that they are designed to image and characterise soft tissue. While this has led to great advancements, providing Doppler and three dimensional imaging, skeletal imaging has remained relatively untouched. While there are a few commercial systems available for niche applications, the market is very small (H. Azhari, 2010). This reason for this is that X-rays and CT are proven methods of producing an accurate depiction of bony surfaces. In saying this, there is an area of research dedicated to musculoskeletal ultrasound imaging, which evaluates both bone and the attached tissues. While much of the research in this field is in the characterisation of the tissues attached to bone (Klauser & Peetrons, 2009), a large proportion of the research is centred on rheumatology (Kane, Balint, Sturrock, & Grassi, 2004), much of which requires imaging of bony surfaces. This includes identification of bone erosion, measurement of deformation of the profile of the bone and joint spacing (Kane, Grassi, et al., 2004). However, each of these is merely an identification of the existence conditions or the progressions thereof. As such, highly accurate depiction is neither employed nor needed.

One of the major problems with most biomedical ultrasound systems is that they are hand-held, requiring a trained operator. The potential for large intra- and interoperator variability is, therefore, great (Vieyres et al., 2006), as evidenced by the need for high teaching standards (Klauser & Peetrons, 2009). The most obvious way to rid any ultrasound system of the dependence on the skill of the individual operator is to automate the process of probe movement and data capture.

This has been investigated in the form of robotic ultrasound systems. The need for these stemmed from the lack of trained operators in certain regions. It was suggested that jointed robots could be controlled by some external operator in an area where such an operator was not available (Priester, Natarajan, & Culjat, 2013). In this vein, an early robotic ultrasound system with six degrees of freedom was designed for examination of the carotid artery (Salcudean, Bell, & Bachmann, 1999), the operation of which can be seen in Figure 1.8. This system used contact probes, employed operator intervention to maintain good imaging parameters and had a robot design based on parallelogram linkage. A number of other robotic arm systems were developed for telemedicine. One of these had the purpose of diagnosing problems in the shoulder (Norihiko Koizumi, Nagoshi, Hashizume, & Mitsubishi, 2009) and allowed for seven degrees of freedom. Another group created a robot that allowed for six degrees of freedom, but did not employ the usual ‘arm’ design. This, again, was intended for telemedicine and was employed in general ultrasound investigations (Vieyres et al., 2006).



Figure 1.8: The system described by (Abolmaesumi et al., 2002) for imaging of the carotid artery ©[2002] IEEE. The probe maintains good contact with the skin using movement of the manipulation system, which is controlled by an external expert.

Other systems were not intended for telemedicine, with one being designed for inspection of the carotid artery (Nakadate et al., 2010). Importantly, this study showed that an unskilled operator could achieve higher levels of accuracy using an automated ultrasound device than could be attained by a skilled operator operating manually. The same group also developed methods for movement compensation (Nakadate et al., 2011), which can be a problem in ultrasound imaging. There was also a further attempt to image the carotid artery using a probe controlled by a robotic arm (Ciompi, Pujol, Ferr, & Radeva, 2012). While those systems designed for telemedicine were created to allow ultrasound operators to work on patients beyond their reach, the other systems described in this paragraph were intended to improve upon the performance of operators. This was, for imaging of the carotid artery at least, achieved. One thing that united both types was that they all employed contact probing. Additionally, with the exception of one, all these methods were intended for soft tissue imaging.

There were other robotic systems which employed more unique robot designs, such as a robot wrist design, which allowed for only three degrees of freedom (Arbeille et al., 2003; Gourdon, Poignet, Poisson, Vieyres, & Marche, 1999). This was designed for telemedicine, with abdominal imaging in mind. One team used a platform with jointed legs to image the abdomen (Masuda, Kimura, Tateishi, & Ishihara, 2001) and another system which involved pulling belts attached to an ultrasound probe (Masuda, Vilchis, Troccaz, & Cinquin, 2003). Both of these systems employed contact probes and were intended for abdominal imaging. As mentioned before, all of these systems were designed with only soft tissue imaging in mind.

A number of biomedical systems have employed “servoing” as a means of robot control. This method uses the data captured by the ultrasound probe to inform the movement of the robot in a closed loop, in real time (El-Hawary, Roth, King, Chess, & Johnson, 2006). One system allowed a human operator to control the robot along one axis, while the robot automatically moved along the others based on the images received from the ultrasound probe (Abolmaesumi, Salcudean, Zhu, Sirouspour, & DiMaio, 2002). This provided images of the carotid artery, with operator independent results. Another group has employed servoing as a method of generally improving data capture, providing optimisation of probe position (Chatelain et al., 2015), while a

further group used it to ensure that a target structure is kept in the image region (Krupa & Chaumette, 2005). Ultrasound based servoing has been employed in tracking needles, tool tips, fiducials, tumours and organs (Novotny, Stoll, Dupont, & Howe, 2007; Priester et al., 2013) and has been used to compensate for patient movement (Chanel, Nageotte, Vappou, Luo, & Cuvillon, 2014; N Koizumi et al., 2012; Nakadate et al., 2011).

Robotically guided ultrasound probes have also been researched and employed in NDT. The inspection of parts is of vital importance within NDT, allowing for the evaluation of material properties and the discovery and characterisation of defects. This process covers a range of objects, ranging from small, complex parts, up to very large structures (C Mineo, Herbert, Morozov, & Pierce, 2012). As such, automated inspection can be incredibly beneficial, allowing for a reduction in cost, a faster inspection time and a higher quality examination. In a number of industries, custombuilt inspection systems have been created, which are expensive and take time to be produced. Standard, commercially available robots are an alternative to this, allowing for a less expensive solution that can be delivered faster and have greater flexibility (Schwabe, Maurer, & Koch, 2010).

Because of the nature of the structures under investigation, many of which are too dangerous for humans to access, much of the research into robotic ultrasound centres on mobile robots (Gordon Dobie, Summan, Pierce, Galbraith, & Hayward, 2011). Many of these are capable of climbing (Kang, Kim, Son, & Choi, 2003; Menegaldo, Santos, Ferreira, Siqueira, & Moscato, 2008; Sattar, Rodriguez, Shang, & Bridge, 2006; Shen, Gu, & Shen, 2005; White, 2005), are miniaturised (Gordon Dobie, Galbraith, Macleod, Summan, & Pierce, 2012; Gordon Dobie et al., 2011) and some even work in groups, providing multiple sensing techniques (G Dobie, Summan, MacLeod, Pierce, & Galbraith, 2014). Three axis scanning systems have also been employed to inspect small parts (Hillger, Bühling, & Ilse, 2014). In a similar vein, scanning gantries have also been employed to scan large components (Cuevas, López, & García, 2012; Hasse & Maurer, 2004).

More flexible 6-axis robotic arms are, however, employed within NDT for component inspection. A number of inspection techniques are employed in cooperation with

robots, including ultrasound, eddy currents (Sattar & Brenner, 2006), thermography (Schmidt & Dutta, 2012), visual tests (C Mineo et al., 2012) and other electromagnetic techniques (Mohr, Haase, & Ungerer, 2014). These can be employed either as single robot configuration or with multiple robots working in cooperation. Focussing on ultrasound, inspection can be achieved in a number of ways, including single robot contact inspection of one side of a structure, immersion testing or through transmission testing (Mohr et al., 2014). As is shown in Figure

1.9, these can be carried out with either single element probes or with phased arrays. While single element probes only allow for a one dimensional inspection, phased arrays allow for an entire area or volume to be inspected in one scan. These inspections can be carried out in contact with the specimen, in which case a gel or grease contact medium would be placed between the probe and the surface of the specimen. This can also be achieved with immersion testing, in which both the probe and the piece are submerged, with water acting as the contact medium. Finally, through transmission can be employed, which uses jets of water as a means of matching the impedance of the ultrasound elements to the piece under inspection.

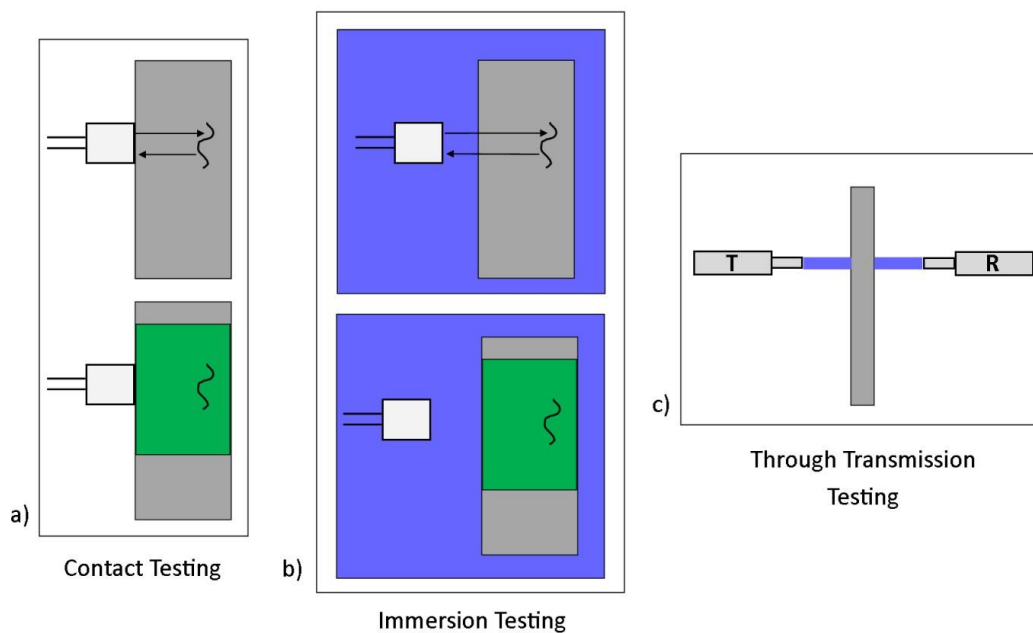


Figure 1.9: Three different ultrasonic inspection methods using robotic arms. Part (a) shows contact testing, with the top image showing testing with a single element probe, while the bottom shows that with a phased array. Part (b) shows immersion testing,

again with the top showing a single element, while the bottom shows an array. In both (a) and (b), the green area represents the region imaged. Part (c) shows through transmission testing, with a transmitter on the left and a receiver on the right.

These types of robotically guided ultrasound systems are employed in a variety of fields, such as the automotive, aerospace, petrochemical and nuclear industries (C Mineo et al., 2012). Because of the diverse needs of these industries, ultrasonic inspection is applied using a number of different approaches. In the aerospace, automobile and other sectors for example, the use of composites has become common due to the need for lightweight, high strength materials (Cuevas et al., 2012;

H. Wang, Xu, & Xu, 2013). As with all materials employed in such industries, it is important to determine whether or not flaws are present in the structure, both in the manufacturing stage and once in use. Due to the levels of ultrasonic attenuation in composite materials, the through transmission technique is usually employed (Z. Lu et al., 2013). While this practice may be necessary for this field, it may not be suitable for others. For example, when inspecting welds with complex geometries, a single robot configuration might be better suited (Nilsson, Appelgren, Henrikson, & Runnemalm, 2012). Other constraints, such as those found in the nuclear industry, might also determine the robot configuration, with size and safety concerns often restricting the number and type of robot (Mohr et al., 2014).

Ultrasonic inspection requires a normal incidence to be maintained between the probe and the surface of the piece. Robotic arms are well placed to meet this need, in that they allow for complex paths to be followed with a high degree of precision (H. Wang et al., 2013). The robot path can be taught manually, with an operator jogging the robot and saving each position of the path. This is, however, very time consuming and requires that the operator gain access to all necessary areas of inspection (C Mineo et al., 2012). This can be alleviated by the use of Computer Aided Design (CAD) models, which allow paths to be planned specific to the part. This model can either be the original Computer Aided Manufacture (CAM) model or can be generated using a volumetric scanning system, such as a laser line probe. A number of systems have been

developed to combine the processes of capturing surface geometry, tool path creation, robot control and ultrasound data capture and processing (Cooper, Nicholson, Liaptsis, Wright, & Mineo, 2013; Cuevas & Hernández, 2014).

This type of ultrasound imaging centres on the detection and characterisation of flaws within a material. While there are a number of inspection methods which employ surface reconstruction (Fradkin et al., 2013; Grotenhuis & Hong, 2012), surface geometry remains relatively unimportant in this field of research. There are, however, a number of parallels between robotically guided ultrasound in NDT and that within bone surface imaging. For example, there is a need for accurate and precise positioning. In NDT, this is necessary to find an accurate position for any defect, while in bone imaging, it is vital to correctly reconstruct the surface. There is also a similarity in the importance of path planning. As mentioned, maintaining perpendicularity is of vital importance. However, while in NDT there are methods to find accurate surface geometries, this is unavailable in bone imaging, as the purpose itself is to acquire the surface geometry.

1.3 Aims of the Thesis

As it stands, only CT is employed in the acquisition of patient-specific models for preoperative planning in robotic knee arthroplasties. The body of work presented in this thesis is an attempt to provide a proof of concept that demonstrates the capability of ultrasound imaging to offer a legitimate alternative to the current standard. As such, the focus of the work was to create a system capable of reconstructing bony surfaces with accuracies comparable with those achieved using CT. The research focus of the Centre for Ultrasonic Engineering (CUE) is in industrial NDT. This provided a unique opportunity to explore methods usually limited to NDT for use in biomedical imaging. In doing so, a number of methods and outcomes novel to the fields of both medical ultrasound imaging and ultrasonic NDT have been necessary.

Medical ultrasound imaging has advanced greatly over the last few decades, spreading into standard practice in a number of areas within medicine. However, most of the

research within the field has centred on soft tissue imaging. While a number of research efforts have been made into imaging the surface of bone, these have largely been limited to the registration of landmarks to models acquired using CT. The ability to accurately reconstruct bony surfaces using ultrasound, independent of any other imaging modality has yet to be achieved. This work attempts to provide evidence related to the feasibility of such an approach. Additionally, while a number of systems have employed bespoke robotic manipulators, very few studies have employed commercially available robots. Nor did most of these studies employ robots specifically for their ability to accurately position probes, but rather to allow remote control. To address this, commercially available industrial robots have been employed for probe manipulation, providing accurate and precise probe positional data. Furthermore, synthetic aperture data processing techniques have not been used extensively within medical imaging. Hence in this thesis, the Total Focussing Method (TFM) and variations of the Synthetic Aperture Focussing Technique (SAFT), employed in NDT for their superior imaging capabilities, have been utilised.

Complex, three dimensional surface reconstruction by way of ultrasound imaging has, to the best of the author's knowledge, yet to be investigated within NDT. While a number of studies have employed two dimensional surface reconstruction, the purpose of this has been to account for refractive effects while evaluating structural defects within the part. In contrast, the surface of the knee is evaluated, returning a surface profile for a complex, three dimensional volume.

With these points in mind, the work presented herein is an attempt to bridge the gap between ultrasound imaging in NDT and biomedicine. While the problem is located firmly in biomedicine, surface geometry measurement of the knee for arthroplasty surgery is, in essence, a form of NDT. As such, methods and techniques usually limited to employment in one or the other field have been investigated and employed in tandem, providing previously unused abilities and advantages.

1.4 Contributions to Knowledge

A number of novel avenues of research have been investigated herein which have an impact on ultrasonic imaging in both biomedicine and NDT. These include:

- The development of an ultrasound system which is capable of fully reconstructing complex, bony surfaces with sub-millimetre accuracy. This has included novel combinations of high level imaging methods, probe manipulation and data capture. Crucially, this was achieved without a dependence on *a priori* data, as has been the case in other ultrasound based bone surface geometry reconstruction systems.
- The use of synthetic aperture imaging methods within a biomedical application. These techniques allow for higher resolving power than B-scan methods, which can improve accuracy and can increase coverage using a single data capture. They have been employed within biomedicine before, but have not made use of high resolution images, created using standard probes with high element counts – the scenario usually found in NDT research. As such, imaging methods usually confined to NDT have been transferred to biomedicine and exploited therein.
- The employment of synthetic aperture techniques for surface imaging and reconstruction. The advantages of these methods have never been exploited in bone surface imaging, with traditional B-scan methods being employed universally. It was found that surface profiles of high accuracy can be produced using a narrow aperture SAFT. Further, the use of synthetic aperture methods for surface imaging has only been investigated for dualmedia refraction compensation research. TFM has been employed in this capacity exclusively as yet, but was found to be the inferior imaging method for NDT samples, with a narrow aperture SAFT providing the most accurate surface profiles.
- Investigations into the performance of FMC-based imaging methods when imaging bone with intact soft tissue. Given that such techniques had never been employed in bone imaging in any capacity and that FMC was known to have low SNR, it was unknown if the bony surface could be isolated amongst noise

and surrounding tissues. Using a soft tissue phantom and FE simulation, FMC datasets were produced and processed using synthetic aperture methods, with the results indicating that FMC-based imaging could in fact produce accurate surface representations.

- Study of the effect of imaging the surfaces of three dimensional parts with a one dimensional array, both in terms of imaging effects and of 3D reconstruction using 2D images. It was found that edges can be responsible for reflections when the surface is outwith the imaging plane, causing inaccurate 3D reconstruction. Further, curved surfaces can be responsible for a distortion effect in the resulting images, due, once again, to out of plane reflections. This is of particular importance to NDT inspections, as the effects were particularly pronounced in simple, relatively polished surfaces.

1.5 Thesis Structure

The remainder of this thesis begins in Chapter 2 **Error! Reference source not found.** with a description of standard data capture and processing techniques available within both NDT and biomedicine. Examples of general uses, strengths, weaknesses and relevance to the prescribed application are presented and examined.

In Chapter 3, descriptions of the individual components of the probe manipulation system are presented. A robotic arm was employed for this purpose, requiring design of probe mounts and tool path planning. In addition to this, the accuracy of the robotic positioning system was assessed using an optical tracking system, confirming the viability of the robotic arm in both probe manipulation and position recording.

The ability of the ultrasound imaging system to accurately recreate the surface profiles of a number of simple geometries is investigated in Chapter 4. The use of synthetic aperture methods in surface imaging is rare and, as a consequence, the understanding of the strengths and limitations of such methods was unknown. The effects of flat, edged and curved surfaces were studied, with a number of effects being encountered when employing 1D arrays to image 3D surfaces. The effects of varying aperture width

are also presented in an attempt to determine whether SAFT or TFM should be employed in the final application.

Chapter 5 and Chapter 6 present the results of reconstruction of the surface of a human-shaped, composite distal femur and a real bovine distal femur using the ultrasound imaging system. Using robotic probe positioning, FMC was employed as the data capture method, which allowed for a variety of synthetic aperture widths to be tested. Surface point clouds were created by extraction of surface profiles from each image. These were then employed to create the full surface reconstructions, which allowed for comparison to reference models. This was, therefore, a new way of producing accurate surface representation using ultrasound, as the dependence on *a priori* data was eliminated.

An investigation into penetration and noise issues in soft tissues when using FMCbased imaging methods is presented in Chapter 7. In the absence of cadaveric experimentation, both soft tissue phantoms and FE simulation were employed, which required more advanced surface profile extraction methods. Finally, Chapter 9 presents future work, ideas for practical implementation and conclusions.

1.6 Publications Arising from this Thesis

1.6.1 Journal Papers

Kerr W, Pierce SG, Rowe P, “Investigation of Synthetic Aperture Methods in Ultrasound Surface Imaging Using Elementary Surface Types”, *Ultrasonics*, 72, 165-176

Kerr W, Rowe P, Pierce SG, “Accurate 3D Reconstruction of Bony Surfaces Using Ultrasonic Synthetic Aperture Techniques for Robotic Knee Arthroplasty”, *Computerized Medical Imaging and Graphics* (under review)

1.6.2 Conference Papers

Kerr W, Rowe P, Pierce SG, “Measurement of the Geometry of the Distal Femur Using Robotic 3D Ultrasound”, *Computer Aided Orthopaedic Surgery CAOS*, Vancouver, Canada, June 2015

Kerr W, Rowe P, Pierce SG, “Measurement and Reconstruction of the Geometry of the Distal Femur Using Robotic 3D Ultrasound”, *25th Congress of the International Society of Biomechanics ISB*, Glasgow, United Kingdom , July 2015

Chapter 2 Ultrasound Data Capture and Processing

2.1 Introduction

There are a number of ultrasonic imaging methods available for use, each of which employs different data capture and processing techniques. Of vital importance to the choice of imaging method is an understanding of how the characteristics of each technique determines the ability to accurately depict the object in question. When considering this choice, factors such as financial cost, computational expense, physical limitations, coupling and reflecting material properties and object geometry must all be taken into account. The imaging methods of interest within NDT research tend to differ from those within biomedical ultrasound. Firstly, the two fields seek to depict materials and objects possessing completely different acoustic properties. In addition, each field is driven by varying user demands and expectations. Accuracy is often paramount within NDT, while biomedical ultrasound focusses more on human friendly visualisation for the identification of problems, rather than accurate metrology. As such, each field provides a number of imaging methods, with different strengths and weaknesses.

The purpose of this chapter is to describe relevant ultrasonic imaging methods by detailing the associated operational principles and assess performance based on those

principles. This evaluation is then employed to determine which of these techniques is most suited to imaging the bony surface of the knee. Using this as a foundation, a number of imaging methods are described and reviewed. Additionally, some of the imaging methods described within this chapter are used extensively in proceeding chapters. As such, this review is presented as a reference for understanding these techniques.

2.2 Data Capture and Processing Techniques

There are a number of methods for acquiring ultrasonic data and then for processing, all of which have distinctive strengths and weaknesses that make them particularly suitable for certain tasks. The purpose of this section is not only to detail these strengths and weaknesses, but to, at least in part, justify the choices of imaging methods that have been made and to do so based on the described qualities.

2.2.1 A Scan

The first, most common and basic method of data capture and display is the amplitude modulation scan, more popularly known as the A scan. The concept of the A scan is that ultrasonic energy is displayed as a function of time. With knowledge of the speed of sound, c , within the medium, the distance from the transducer at which a reflector is positioned can be determined simply using equation 2.1, where d is distance and t is time. The reflector will present itself as values of higher amplitude than found elsewhere.

$$d = ct \tag{2.1}$$

There are two main methods by which A scans are performed: pulse-echo and pitchcatch. The latter operates by a transducer transmitting a pulse, which is then gathered by an ultrasound receiver. The recorded signal is the waveform obtained from the receiving transducer. Pulse-echo, on the other hand, employs the same transducer as both the transmitter and receiver.

Both methods are exemplified in Figure 2.1, which shows detection of a simple reflector by way of an A scan. On the right hand side, the pitch-catch mode shows one spike in the associated waveform, which represents the presence of the reflector. On the left hand side of Figure 2.1, the pulse-echo mode shows two spikes in the waveform representing, firstly, the reflector and, secondly, the back wall of the piece. To determine the location of the reflector, it must be taken into account that the waves have to travel to the reflector, then back. Ascertaining the distance from the transmitter to the reflector requires, then, that the distance found using equation 2.1 is divided by two, to account for the return leg of the journey.

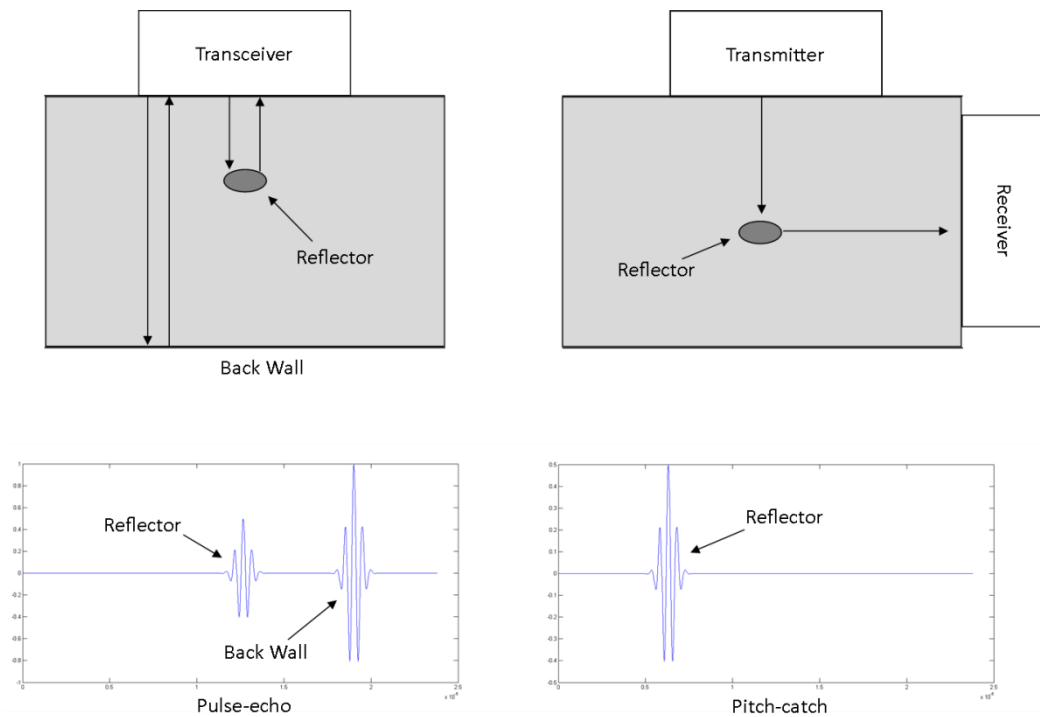


Figure 2.1: An example of showing both pitch-catch and pulse-echo modes. The top half of the image depicts the physical setup, while the bottom displays the A scans.

A-scans are the simplest means of data interpretation and, as such, there exist many examples of their use within both NDT and biomedicine. They are used in metallic, concrete and composite structures (McCann & Forde, 2001), single element probes providing simple A-scans are capable of delivering information on grain size (Holler & Schmitz, 1983; Palani, Joseph, Jayakumar, & Raj, 1995), characterising adhesive bonds (Adams & Drinkwater, 1997; Seydel, 1975), measuring thickness (Krause et al., 1997) and characterising composites (Adams & Cawley, 1988), as well as many other procedures. It is the use of A-scans in flaw detection, however, that presents the greatest similarity to the imaging bony surfaces. Flaws, such as cracks, can be described as discontinuities in the material being inspected. These discontinuities must have an acoustic impedance mismatch that allows reflection to be returned, just as there is between soft tissue and bone.

The use of A-scans within the diagnostic medical ultrasound is extensive, finding use in cartilage thickness measurement (Jurvelin, Lyyra, Rasanen, & Kolmonen, 1995; Suh, Youn, & Fu, 2001; Töyräs et al., 2001; J. Q. Yao & Seedhom, 1999), assessment of biomechanical properties of soft tissues (C. P. Brown et al., 2012; Saarakkala, Laasanen, Jurvelin, & Töyräs, 2006; Tatarinov, Sarvazyan, & Sarvazyan, 2005; Toyras, Rieppo, Nieminen, Helminen, & Jurvelin, 1999; Töyräs et al., 2001), mechanical properties of bone (Gregg et al., 1997; Hoffmeister, Jones, Caldwell, & Kaste, 2006; Kotha, DePaula, Mann, & Guzelsu, 2008) and in ophthalmology (Fledelius, 1997). Related more closely to the imaging of bone surfaces, A-scans have also been employed as a means of intra-operatively registering the bones of the patient in the coordinate system of the navigation system to that of the preoperative CT image. This was achieved by palpating landmarks via ultrasound rather than by gaining physical access (Heger, Portheine, Ohnsorge, Schkommodau, & Radermacher, 2005; Mozes & Chang, 2010).

A-scans are of vital importance to high level imaging techniques, in that they allow tools to be calibrated and problems to be diagnosed. The use of A-scans as a method

of imaging, while useful in some cases, is simply not practical for creating a three dimensional image of a knee joint independent of other imaging methods. The one dimensional nature of the display of data allows for no meaningful imaging to be performed, particularly for a task that is three dimensional in nature.

2.2.2 B-Scan

Brightness scans, usually shortened to B-scans, are simply a way to display a group of A-scans. A single element probe is moved mechanically along one axis of the piece. Each A-scan is then stacked along the horizontal axis, with the waveforms being represented by colour, rather than the vertical height of a line. In the example shown in Figure 2.2, a number of defects and a back wall have created peaks in the waveforms, the amplitude of which has been represented by a redness in colour. While the single value B-scan described here provides a two dimensional representation of the area of interest, the quality of the returned images is usually low, failing to provide an accurate depiction of the desired reflectors. The prime cause of this lack of clarity is the fact that only a single element transducer is employed. These are afflicted by near field fluctuation, travel-dependant decay, fixed beam width and invariable focus.




Figure 2.2: Construction of a B-scan image, showing mechanical movement of a single element probe. The time of flight data is displayed vertically, while the horizontal axis indicates the position of the transducer.

This type of scan is rarely seen in biomedical imaging, but is quite standard in NDT. It sees use in imaging in concrete, with the testing of the structural integrity of bridges (Christoph Kohl & Streicher, 2006) and that for faults in concrete members and structures in a number of applications (C. Kohl, Krause, Maierhofer, & Wostmann, 2005; Krause et al., 2001). It also finds use in the detection of defects in wood (Hasenstab & Osterloh, 2009) and metallic structures, such as nuclear reactor components (Merola et al., 2002) and welds (Blomme, Bulcaen, & Declercq, 2002), as well as many others. Despite the high level of use within NDT, this technique would be difficult to implement in this project, as it would require huge numbers of data captures to acquire all the necessary data. As such, it would require that the specimen remain static for a long period of time. While this might be acceptable in a certain situations, such inspection of bridges, it would not be possible to subject a patient to such a protocol.

2.2.3 C-Scan

C-scans, otherwise known as constant depth scans, are, again, a way of displaying a large number of A-scans. In this method, the transducer is moved mechanically in two dimensions – usually following a raster pattern. Alternatively, instead of the transducer being moved, the specimen might be moved, while the probe remains stationary. While B-scans display each full waveform, in C-scans, each pixel represents a depth. The depth is determined by where on the time axis the peak is, with the peak simply being a point where the amplitude of the signal has passed a given limit. This is exemplified in Figure 2.3, where raised areas are red, while those lower down are blue.




Figure 2.3: Raster scan path for probe and raised circular area (left) and the resultant C-scan image (right), with a difference in height being represented by a change in colour.

This technique is widely employed in acoustic microscopy (Lemons, 1974), finding use in both biomedical (Saijo et al., 1997) and more industrial settings (Schwarzschild, 1979). These transducers work in the GHz range of frequencies, allowing for very high resolutions. These qualities are not only unnecessary for imaging large scale structures, like knee joints, but they have associated with them impractical physical restrictions, such as a requirement for the probe to be very close to the specimen, due to high levels of attenuation. Speaking of C-scans more generally, it might be said that they appear to be well suited for imaging of the knee, in that they essentially image surfaces. However, the two dimensional nature of the probe movement is restrictive and due to the use of single element probes, they suffer from the same problems associated with B-scans.

C-scans are also employed in the same ways B-scans are above, owing to the similarity in their method of data capture and presentation. For the reason explained in Section 2.2.2, C-scans would be very difficult to implement in this project.

2.2.4 Plane B-Scan

This method mimics the B-scan described in Section 2.2.2. However, while the Bscans described previously employ a single element transducer, plane B-scans use arrays of

transducers, which are groups of transducers positioned close together. The most common of these is a one dimensional linear array. The geometric configuration of a standard one dimensional array is shown in Figure 2.4. Displayed in the illustration is the distance between the centre of one transducer and the centre of the nearest to it, known as the pitch. Also exhibited are the width of one element and the distance between the edge of one transducer and the next, commonly termed the kerf.




Figure 2.4: Configuration of a conventional one dimensional array, showing the pitch, P , the element width, W , and the kerf, K .

Arrays such as that shown in Figure 2.4 allow for electrical scanning, rather than mechanical. Instead of moving the probe along an axis of the piece, the position of the probe is fixed and a group of elements is fired. This subset of the full number of elements is known as a sub aperture. Once the first sub aperture has been fired, the position of the sub-aperture changes and fires again. This is repeated until all the elements have fired at least once. The way in which the sub aperture moves is variable, with it being possible that sub apertures may overlap. This is explained in Figure 2.5, which shows a 16 element array performing a plane B-scan with a subaperture 4 elements wide. The sub aperture moves right along the horizontal axis, with N sub apertures being fired in total. Return signals are recorded then recorded on using the elements of the sub-aperture.




Figure 2.5: Data capture for a plane B-scan for a 16 element probe, using a 4 element sub aperture. The sub aperture moves electronically along the length of the array, producing a different region of focus at every transmission.

2.2.5 Phased Arrays

Thus far, only single element or unfocussed arrays have been considered. This section discusses transducer arrays capable of focussing the sound beam to a particular point. These are known as phased arrays and achieve this by delaying the point at which individual elements are excited. Each of the elements produces an associated wavefront when fired, each of which interacts such that a natural focus is produced, as can be seen in Figure 2.5. It is possible to apply delays such that these wavefronts, through constructive interference, combine at a given location. Considering the example depicted in Figure 2.6, for an array consisting of N elements, the delay for the i th element can be determined. The first step in this is to calculate the distance from the centre of that element to the focus point, r_i . This can be done simply using Pythagoras, as shown in equation 2.2.

$$r_i = \sqrt{\text{focus } z^2 + (\text{focus } x - \text{centre } i)^2} = \sqrt{r_{iz}^2 + r_{ix}^2} \quad 2.2$$

With both this distance and the speed of sound within the medium, the time delay for the i th element, τ_i , can then be found simply by using equation 2.1, replacing d with r_i . This ensures that the wavefronts from each element arrive at the focus point at the same time. The same technique can also be applied to the signals in reception, with delays being applied to the resultant A-scans.




Figure 2.6: Construction of parameters required for phase delays for the i th element of a N element array.

One of the main problems associated with focussed transducer arrays is the creation of secondary lobes. These are unwanted sound beams radiating from the probe in a direction that differs from the primary beam. These present themselves in two forms: side lobes and grating lobes. The former is found in both phased arrays and single element transducers. The number of side lobes and their magnitude is determined by the ratio of the aperture and the wavelength (Shung & Zipparo, 1996). As the size of the ratio increases, the number and magnitude of the side lobes also increases. In the case of arrays, the aperture is considered to be the distance from the start of element 1 to the end of element N .

The effect of this can be seen in Figure 2.7, where parts (a) and (b) show a single element 5MHz transducer operating into water, with an aperture of 0.5mm. These

sound fields were produced using a method based on Huygens' Principle which treated each element as an infinite number (made finite to allow for calculation) of point radiators. This was bespoke code written in Matlab script. In parts (a) and (b), the ratio of the aperture and the wavelength is 1.69. Only two side lobes can be seen in part (a) and, 1.5mm away from the transducer face, no side lobes are present, as can be seen in part (b). Parts (c) and (d) show a transducer with the same properties as before, but with an aperture of 1mm, implying a ratio of 3.38. It is clear from parts (c) and (d) that there is both an increase in the number of side lobes and their intensity.

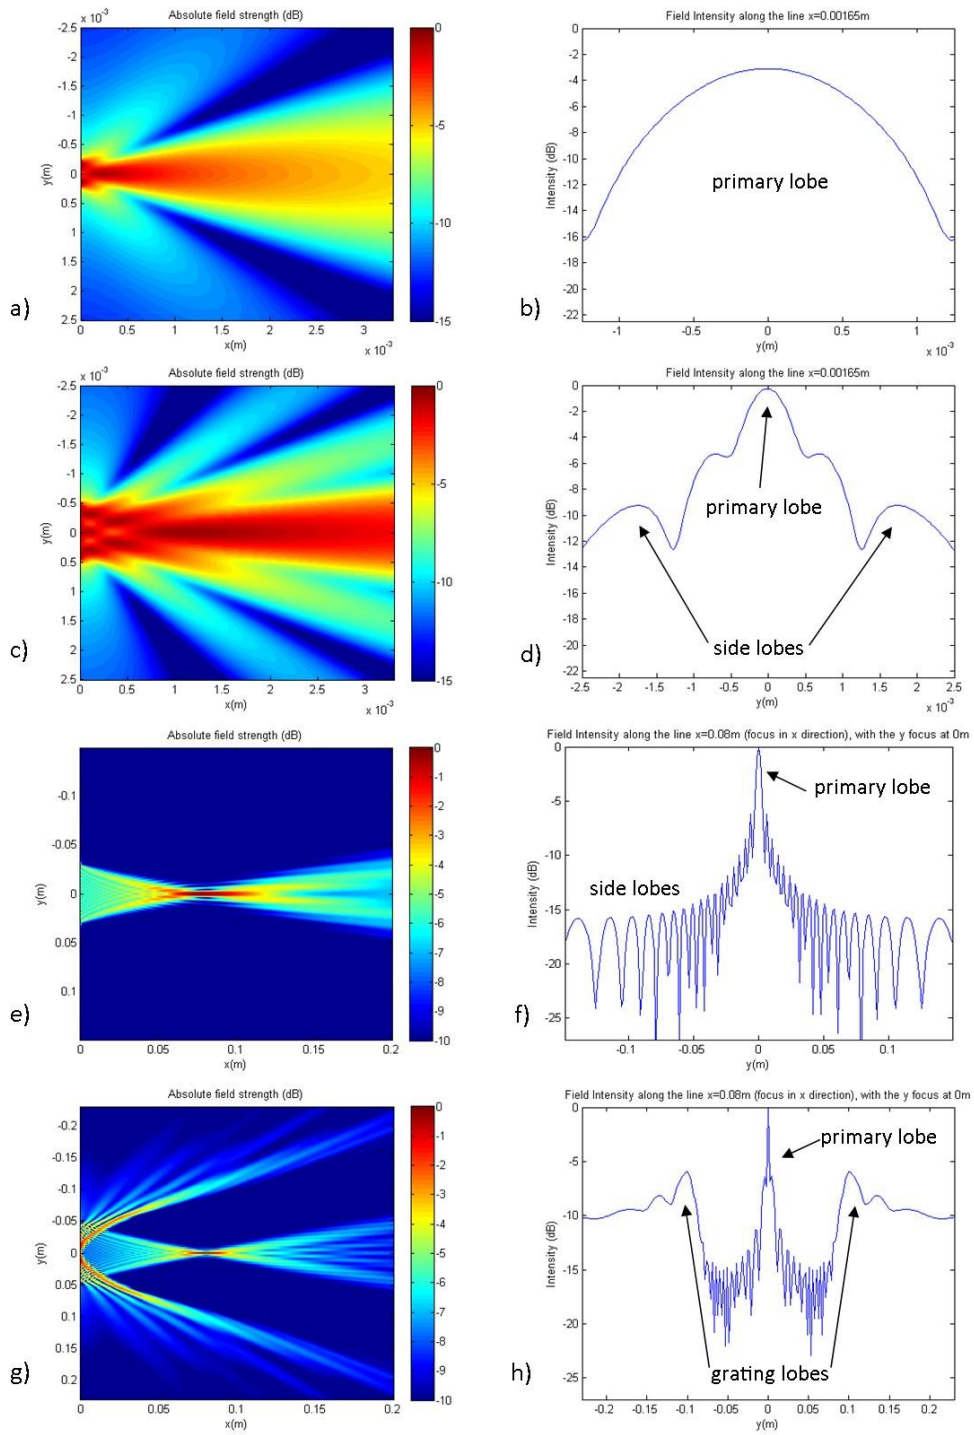


Figure 2.7: Field characteristics of a single element transducer (a-d) and a 32 element array (e-h), both operating at 5MHz in water. In parts (a), (c), (e) and (g), the full radiation field is shown, with the transducer on the left of the image. Parts (b), (d), (f) and (h) show a vertical section through the associated field.

The effect of side lobes is important within methods which employ a single element in transmission. In particular, imaging methods based on Full Matrix Capture (FMC) – a phased array method which employs only one element in transmission – requires a spherical wave, such that the element acts like a point source. As can be seen in parts (a-d) in Figure 2.7, it is obvious that the larger aperture width does not provide the kind of uniform radiation that would be seen with a point source. The smaller aperture element, on the other hand, while still possessing an obvious focus, does not display the same lack of uniformity, as can be seen when comparing parts (b) and (d) of Figure 2.7.

In addition to side lobes, grating lobes are also a possibility. These occur only in arrays and are a consequence of the periodic spacing between elements (Hindi, Peterson, & Barr, 2013). These differ from side lobes in that they tend not to occur in as close proximity to the primary beam and are of greater intensity. These can, however, be avoided by maintaining a spacing of less than $\lambda/2$ or by eliminating the periodicity of the element spacing (Gavrilov, Hand, Abel, & Cain, 1997). This is demonstrated in Figure 2.7 (e) and (h), where a 32 element array is operating at 5MHz into water, with an element spacing less than $\lambda/2$. In this case, only side lobes are present. Parts (g) and (h), on the other hand, show the same array, but with an element spacing greater than $\lambda/2$. In this instance, grating lobes are obviously present, showing a more significant intensity than the side lobes.

The reason secondary lobes are of importance is that they influence the imaging capability of probes. Both side lobes and grating lobes carry the danger of causing an echo from a strong reflector. It is difficult to differentiate between reflections caused by the primary beam and any secondary beams. As such, it is impossible to determine from the resultant image what is simply an artefact and what is a true representation of the object under inspection.

2.2.6 Linear B-Scan

One way in which phased arrays can be applied is in focussed linear B-scans. These are the same as plane B-scans in every way, but for the fact that each sub aperture is

focussed to a fixed depth. The action of a linear B-scan would be very similar to that seen in Figure 2.5. However, arrays with higher element counts might be used, which would, by virtue of being capable of having more elements in sub apertures, allow for a smaller and more intense area of focus.

This method of imaging has found some success in registering bones in space to MRI and CT preoperative images intra-operatively by finding the correct transformation to apply to the preoperative model. This was achieved by (Barratt et al., 2006b), using freehand ultrasound and optical tracking to return the three dimensional position of the probe. This study achieved an RMS error as low as 0.06mm for positioning of three femoral fiducial markers.

2.2.7 Dynamic Depth Focussing

Phased arrays are also used to focus at multiple depths in a single scan. Termed dynamic depth focussing, the technique applies a series of delays that allow the wavefronts to combine constructively at focus points along an axis perpendicular to the transducer face. Figure 2.8 exhibits this, showing the same probe focussing at a range of depths. The lengths of the lines above the sound field images represent the relative delay for each of the 32 elements in the array. While this is demonstrated using a transmit beam, dynamic depth focussing is usually carried out using one transmit signal that has not been delayed. The received signal is then delayed for the required number of depths in post processing. The method, however, remains the same. This technique allows for the entire depth of a volume of interest to be inspected in one data capture.




Figure 2.8: Sound fields of a 32 element array operating at 5MHz in water, focussed at (from left to right) 3cm, 6cm and 9cm. Also shown are associated time delays for the individual array elements.

2.2.8 Sector B-Scan

Otherwise known as a sectorial, azimuthal or angular scan, a sector scan moves the beam through a number of angles, making a sweeping motion, as shown in Figure 2.9. Unlike those described above, sector scans do not focus the beam, but rather steer it. The time delay for the i th element can be calculated using equation 2.3, where p is the distance between the centre of the i th transducer to the $(i+1)$ th transducer, θ is the steering angle and c is the speed of sound within the medium.

$$\tau_i = \frac{p(i-1) \tan(\theta)}{c} \quad 2.3$$

This method of scanning usually employs every element in the array in both transmit and receive modes, which provides a means of imaging over a range of angles quickly.

The displayed image resulting from such a scan usually combines the data from the full sweep, producing an image in the form of an acute sector.




Figure 2.9: Sound fields of a 32 element array operating at 5MHz in water, steering the beam to angles of 15, 0 and -15 degrees (left to right). Also shown are associated time delays for the individual array elements.

Sector scans are employed extensively in biomedical sonography, with the most recognisable ultrasound imaging method being sector scans used in obstetrics. While these scans are used in everyday practice, they are also heavily employed in biomedical research. They have, for example, been used in research into the improvement of transcranial imaging (Gilmour & Subramanian, 2011). In this study, a beam was mechanically swept through small angular increments, with the purpose of directing the ultrasound energy towards a particular area of the brain through a narrow cranial window.

This method can also be extended to three dimensions using the same principles. Research in biomedical imaging has tended towards high element counts to achieve high levels of resolution, accuracy and coverage (Yuling Chen, Nguyen, & Yen, 2011; Seo & Yen, 2007; S. W. Smith et al., 2002; Yen, Steinberg, & Smith, 2000). Sector scans have been employed to help guide surgeons intra-operatively (Welch, Johnson,

Bax, Badr, & Shahidi, 2000), with one study assessing pelvic orientation during surgery, with errors of only 1.67° (Schumann et al., 2010). It has also been applied to surface reconstruction, using freehand imaging and infrared tracking of probe position (Y. Zhang, Rohling, & Pai, 2002). However no accuracy figure was reported with the results. These methods employ 1D arrays that produce 2D images. The pixels of these images then become voxels by undergoing translation and transformation by using the positional data of the probe tracking system. Stereoscopic camera systems were employed in this study, which have varying levels of accuracy, with mean errors often approaching 0.5mm in accuracy tests (Elfring, Fuente, & Radermacher, 2009). However, such levels of accuracy may not be attained in practice, due to marker set complexity, tool movement and marker blockage. This is important to note, as the ultrasound images and extraction methods will also have inherent inaccuracies.

Other systems have achieved three dimensional imaging using sector scans, with one group creating a generic 3D imaging system allowing for full 3D surface reconstruction (Gee, Prager, Treece, & Berman, 2003). In addition to this, motorised systems have been employed in intravascular and transrectal ultrasound examinations (Klein et al., 1992).

Ultrasound imaging with robotically guided probe control has also been investigated using sector scans to provide three dimensional representations of soft tissue (Pua et al., 2006). Related more directly to this project is the work of (Torres & Goncalves, 2012; Torres, Sanches, Goncalves, & Martins, 2012), which sought to align markers with preoperative images intra-operatively using a robotically guided probe. Their system was tested on a bovine femur immersed in a water tank and also with human subjects, using a probe in contact with the skin. As a commercial biomedical ultrasound imaging system was employed, no data processing was possible, only image processing. This, of course, would limit the ability to improve image quality.

In addition, no accuracy results were provided.

The ability to image in real time is the reason B-scans generally remain so popular within the biomedical field. In NDT, however, the assessment of structures is usually not so urgent. However, B-scans are still often employed due to their real time capabilities and are considered a standard method of inspection (Holmes,

Drinkwater, & Wilcox, 2005). However, when being used in NDT, probes are often coupled to the piece by a wedge made of low loss material (Drinkwater & Wilcox, 2006). The wedge is kept at an angle, such that longitudinal waves will not travel through the test piece, due to the angles being greater than the critical angle. This allows inspection with shear waves only to take place. This mimics the same techniques that would be carried out manually with single element probes (Crowther, 2006), with mechanical movement being replaced by electronic control of the array.

Of the types of B-scan described herein, the focussed B-scan performs better than the others when imaging a perfect point reflector (Holmes et al., 2005). The evaluation of the ability to image an infinitely small perfect reflector, otherwise known as a Point Spread Function (PSF) (Dziewierz, Lardner, & Gachagan, 2013), is of relevance to NDT applications. This is because the goal in industrial situations is often to locate and assess very small defects that reflect diffusely, such as crack tips (Cochran, Earnshaw, Bargain, & Kirk, 2002). This is not the case in this project, where longer, largely specular surfaces are being imaged. It does, however, provide information on which imaging modality might perform best when looking at more difficult surfaces which might require the ability to image very small objects. These might arise due to symptoms of osteoarthritis, such as the development of osteophytes. In addition to this, the side lobe levels in a PSF indicate the level of material noise that might be encountered (Chiao & Thomas, 1994). It also provides a method of generally characterising array performance, with one study employing it to evaluate resolution, worst case imaging artefact and smallest detectable reflector (Dziewierz, Gachagan, Lord, & Mullholland, 2012), while others looked at image contrast and lateral resolution (Bannouf, Robert, Casula, & Prada, 2013) and Array Performance Indicator (API) (Holmes et al., 2005).

2.2.9 Full Matrix Capture

The data capture methods described above have been conceived so that as much information as possible can be gathered while maintaining real time image production. If, however, it is acceptable that the data can be post processed, more advanced data

capture methods become available. Within NDT research, the most prominent of these is FMC. This method captures every time domain signal for every transmit and receive pairing (Holmes et al., 2005; Sutcliffe, Weston, Dutton, Charlton, et al., 2012). In FMC, while one element fires, all elements receive. This is repeated until every element in the array has transmitted. It is usually carried out in sequence, from element 1 to N and is explained in Figure 2.10, which shows the firing sequence for the transmission, as well as the associated reception. Therefore, for each transmission, N A-scans are acquired, meaning that N^2 A-scans in total are captured for each FMC. The collection of all N^2 A-scans is known as the full matrix and is the maximum amount of data that can be captured in one acquisition (Grotenhuis & Hong, 2012).

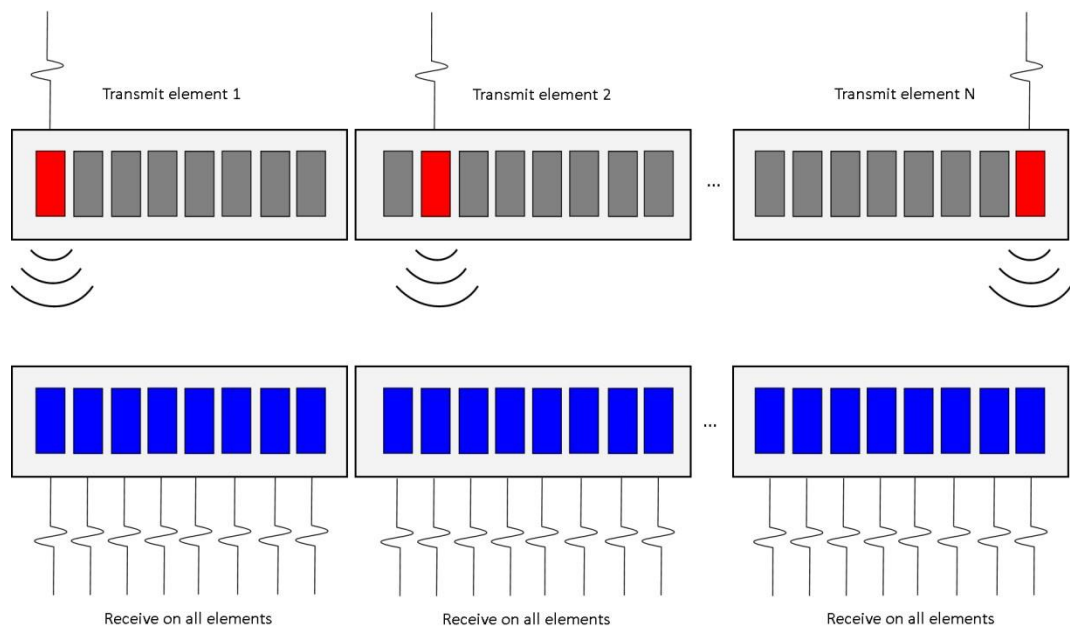


Figure 2.10: The FMC procedure showing transmission and reception cycles, with grey representing an inactive element, red an element being fired and blue a receiving element.

This high number of transmit and receive cycles can limit the ability of FMC to work in real time. The first bottleneck caused by this is a physical one and cannot be changed without altering the definition of FMC. When an element transmits, it takes time for the waves to travel through the medium and return to the receiving elements.

This time is multiplied by every transmitting element. In addition to this, the high element count of modern ultrasound probes poses a number of problems in parallel reception of high numbers of receiving elements, as well making data transfer difficult. These problems have been addressed by employing Field Programmable Gate Arrays (FPGA), which are large numbers of low level logic blocks which can be programmed at a hardware level. These have been employed to speed up FMC to achieve real time imaging (Fradkin et al., 2013; Njiki, Bouaziz, Elouardi, Casula, & Roy, 2013; Njiki, Elouardi, Bouaziz, Casula, & Roy, 2013), but have been held back in general uptake due to the necessity to program in Hardware Description Language (HDL).

Related to FMC is a method known as Full Raw Data (FRD), which is essentially a subset of the FMC data (Lines, Pettigrew, Kirk, Cochran, & Skamstad, 2006). The difference between the two methods being that in FRD, one element transmits and only a sub aperture receives, whereas in FMC, every element receives. This is a faster method of data capture than FMC, but does not command the same versatility in data processing, as the waveform data is not complete.

The roots of FMC are in Synthetic Aperture (SA) imaging (Sutcliffe, Weston, Dutton, Charlton, et al., 2012), which was, firstly, an innovation of radar engineering and was then taken up and modified for ultrasound by the NDT field (Jensen et al., 2006; Tapani & Helmut, 1994). This has been employed in biomedical imaging, in an attempt to achieve higher frame rates for purposes that required such an improvement, such as flow imaging or heart valve imaging (Jensen et al., 2006; Trots, Nowicki, Lewandowski, & Tasinkevych, 2010).

2.2.10 Total Focussing Method

Once an FMC data set has been acquired, there are a number of post processing algorithms that could be applied to produce an image. The Total Focussing Method (TFM) is arguably the most popular of these methods within NDT research at the moment. This technique allows for focussing to be performed for every point (or pixel) in the resulting image. It is considered by many to be the gold standard in ultrasound

inspection (Wilcox et al., 2007) and is often used as a benchmark for other imaging algorithms. For example, one study employed as a means of comparison with phase coherence imaging and spatial coherence imaging, finding TFM to be the most robust in terms of defect detection (Jie Zhang, Drinkwater, & Wilcox, 2013).

In order to produce a scalar image from FMC data via TFM, the area of interest must first be divided into a grid. Each of the points on this grid will represent both a focus point and a pixel in the final image. The intensity, I , of a particular pixel, p , is calculated by summation of the contribution from each delayed transmit-receive pairing. This is shown in equation 2.4 below, where i is the transmitter in question, j the receiver and S_{ij} is the A-scan relating to that particular pairing. T_{ip} and T_{jp} are the time delays, relative to the distance from the transmitter to the pixel and the pixel to the receiver, respectively. The term $S_{ij}(T_{ip} + T_{jp})$ represents the sample of the A-scan corresponding to the calculated delay.

$$2.4 \quad I(p) = \sum_{i,j=1}^N S_{i,j}(T_{ip} + T_{jp})$$

The delays, T_{ip} and T_{jp} , can be computed by calculating the time it would take the ultrasonic waves to travel from the transmit element to the receive element, via the point in question, in the given medium. This is shown below in equation 2.5, where $(x_i - x_p)$ is the distance from the transmitter to the pixel in the x direction, z_p is the distance from the transmitter to the pixel in the z direction and c is the speed of sound in the medium.

$$2.5 \quad T_{i,jp} = \frac{\sqrt{(x_{i,j} - x_p)^2 + z_p^2}}{c}$$

The variables described here can be seen in Figure 2.11, which shows the geometric configuration of the described situation. In this illustration, the transmitting element is shown in red, with a red line denoting the ray path of the waves radiating from it. This ray travels to the point of interest, (x_p, z_p) , where it is reflected and travels to the transmitting element, shown in blue. The origin is arbitrary, but must be invariable for all calculations within the TFM algorithm.

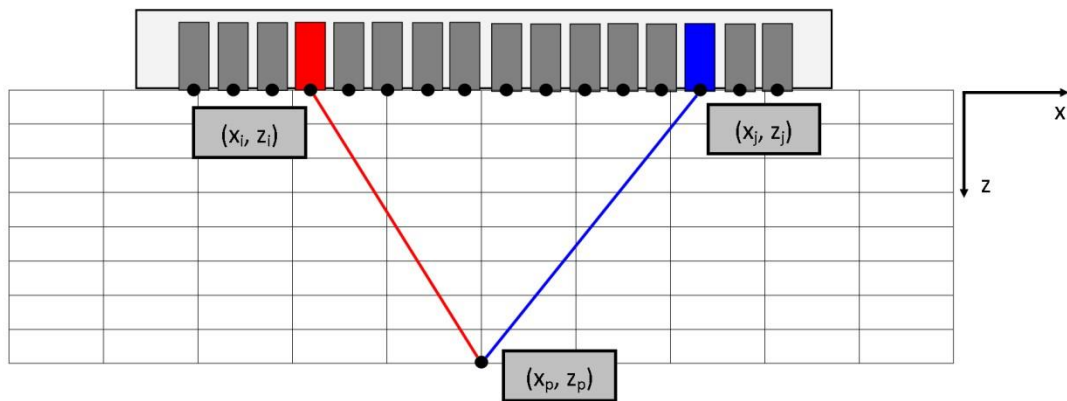


Figure 2.11: Geometric configuration for one iteration of the TFM algorithm. The red line represents the path of the ray from the transmitting element to the point under reconstruction (x_p, z_p) . The blue line is the path from that point to the receiving element.

When the delay for the particular pixel and transmit-receive pairing has been found, the amplitude at the time in the corresponding A-scan is added to the intensity of the pixel, \mathcal{P} . This is repeated for every transmit-receive pairing, yielding the final scalar value for the pixel in question. This is then repeated for every point in the inspection grid. The whole process is presented below in algorithm 2.1.

```

for k = 1 to number of pixels in x direction
  for l = 1 to number of pixels in z direction
    for i = 1 to number of elements
      calculate time of flight from transmitter(i) to pixel (k,l)
    for j = 1 to number of elements
      calculate time of flight from pixel (k,l) to receiver (j)
    calculate total time of flight in terms of sample number
    add value in A scan to element in
  
```

```
intensity          matrix, I(k,l)
end                end end
```

Algorithm 1: pseudocode describing TFM algorithm.

It should also be noted that while the grid representing the area of interest in Figure 2.11 begins at the face of the array, this is merely to emphasize the fact that this description is that of situations in which there is direct contact with the medium of propagation. In reality, the area of interest may be placed anywhere within the range allowed by the length of the reception period. Some applications where the area of interest might be some distance from the probe, such as weld inspections, make use of this fact. For example, it was employed in one inspect welds (Jobst & Connolly, 2010). In such an inspection scenario, it would be difficult to couple the probe to the weld surface, so it is placed at a more convenient place on the specimen and the imaging area is positioned such that the weld is inspected.

In addition, a Hilbert Transform might also be applied to the A-scans. Doing this changes the time domain signals into analytical signals and allows the signal envelope to be found (Sutcliffe, Weston, Dutton, Charlton, et al., 2012; Tweedie et al., 2007). This is a standard step in the TFM process (Fradkin et al., 2013; Holmes, Drinkwater, & Wilcox, 2008; Wilcox et al., 2007), the benefit of which is that the magnitude envelope of a returned signal is displayed in the returned TFM image. This is demonstrated in Figure 2.12, which shows in part (a) a TFM image constructed using only time domain signals while part (b) shows a TFM image produced with Hilbert transformed signals. The reason for this is that if the signals remain unaltered, reflections are returned in the form of a wave, an example of which can be seen in blue in Figure 2.12, part (c). As such, there will be parts of the reflected signal that cross the zero line. Employing the envelope eliminates these zero values, as can be seen in red in Figure 2.12 (c).

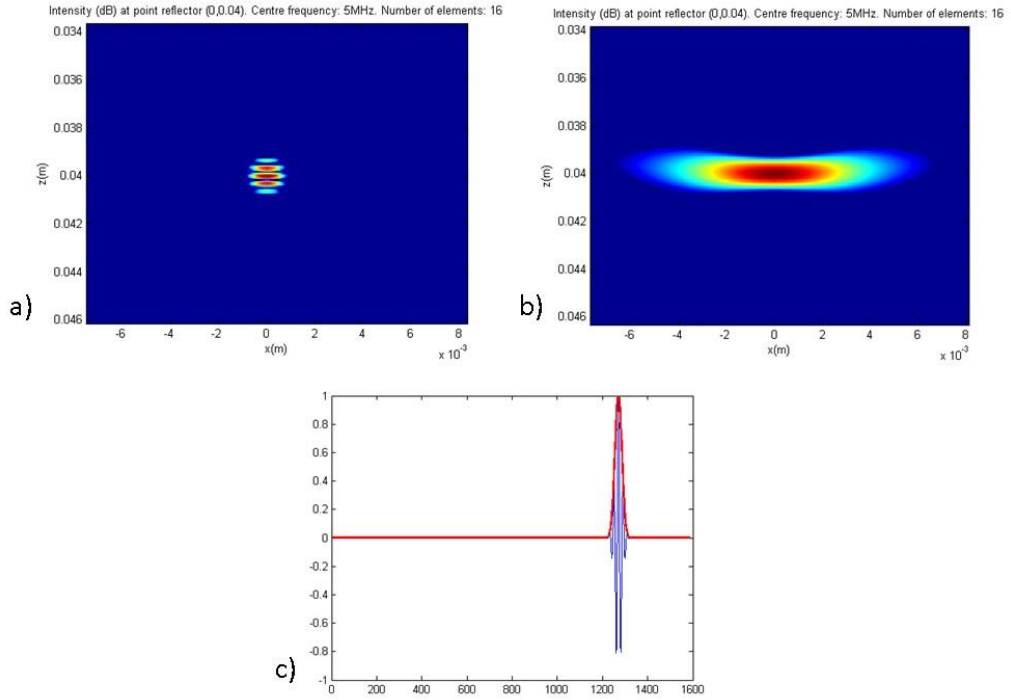


Figure 2.12: From a simulated FMC of a point reflector. The TFM image using only time domain signals (a), the TFM image using Hilbert-transformed signals (b), a typical unaltered A-scan (blue) with the envelope (red) as a consequence of the Hilbert Transformation (c).

One of the main disadvantages of TFM is that it is considered to be computationally expensive, which has denied it commercial acceptance (Dziewierz et al., 2013; Hunter, Drinkwater, & Wilcox, 2010). The most computationally expensive part of the TFM algorithm is the calculation of the delay associated with each transmitreceive pair. While this is a simple operation, it has to be performed (*number of pixels in image* $\times N^2$) times. This is a great challenge for traditional Central Processing Units (CPU) and has led to real time imaging being difficult. One way in which this demand can be alleviated is by exploiting symmetries in the FMC (Chiao & Thomas, 1994; Moreau, Drinkwater, & Wilcox, 2009). These arise in the FMC data due to the fact that the time of flight from transmission on element i , to point p , then to reception on element j is the same as the time of flight if the transmit and

receive elements were reversed. As such, there is no need to compute both. As a result, only $(N(N + 1)/2)$ time of flight calculations are needed per pixel.

Another way in which this problem has been addressed is with the use of General Purpose Graphics Processing Units (GPGPU). The delay calculations are independent from each other, in that the input variables required for each computation are the speed of sound and positional data. As such, these calculations are highly parallelisable. While CPU architecture has evolved to perform many different types of task in parallel, the Graphics Processing Unit (GPU) architecture has found its strength in performing the same mathematical operation on large data sets (Sutcliffe, Weston, Dutton, Charlton, et al., 2012). Numerous studies have implemented a parallelised TFM, with one achieving speedups over conventional CPUs of 762.1 times with a basic GPU implementation and a speedup of 2535.3 times with the inclusion of other optimisations (Lambert et al., 2012). Over the last couple of years, GPU implementation of the TFM algorithm has become commonplace, with numerous studies improving the performance (Lambert et al., n.d., 2012; Sutcliffe, Weston, Dutton, Charlton, et al., 2012), achieving image creation hundreds of times faster than would be possible with a traditional CPU implementation. Processing speedups have also been used to increase imaging capabilities, with a number of groups implementing a double medium refraction compensation method (Dziewierz et al., 2013; Sutcliffe, Weston, Dutton, & Cooper, 2012). This has allowed accurate imaging to be carried out in situations where direct access to the surface of the specimen is impractical or impossible. It has also been employed in simulation of PSF, so that different array parameters can be tested quickly (Dziewierz et al., 2012). This increases the speed of simulation and design of experimental arrays and can return simple metrics for the quality of performance.

Finally, GPU implementations of TFM have been used to test against other experimental imaging algorithms (T. Peng et al., 2014).

There are a number of advantages associated with TFM. Firstly, it has outperformed other traditional beamforming methods when analysing the PSF (Holmes et al., 2005). In this study, FMC data sets were simulated using a simple ray casting method. These were then processed with TFM and the images produced using a plane focussed and sector B-scan were simulated using the same FMC. The resulting PSFs were assessed

using API – an array performance indicator which is found by computing the area inside the PSF greater than -6dB and dividing it by the wavelength.

While such studies do not prove that the TFM algorithm is the best for every situation, they give an indication of the relatively impressive performance characteristics associated with it. In addition to this, the position of the grid representing pixels of the resultant image is not dependent on the position of the probe itself. Unlike traditional B-scans, TFM allows focussing past the edge of the probe, thus increasing the size of the possible viewing area. Additionally, the Signal to Noise Ratio (SNR) of the final image is decreased, due to averaging effects of the TFM algorithm. The decrease in SNR is proportional to the square root of the number of elements, with higher element counts producing a lower SNR (Holmes et al., 2005). It should be mentioned, however, in FMC, only one element fires at a time, meaning the level of random noise will be greater, due to decreased energy, relative to a steered or focussed beam.

While TFM has only been used in NDT research for around a decade, it has already found uses in many fields of research. Studies have employed it for weld inspection methods (Cunningham, Mulholland, Harvey, & Bird, 2012; J. Zhang, Hunter, Drinkwater, & Wilcox, 2012), exploiting the focussing power of TFM and the ability to position the imaging region arbitrarily. Additionally, weld imaging was investigated by a number of groups, one of which employed the technique to create a surface model of a welded pipe (Grotenhuis & Hong, 2012; Jobst & Connolly, 2010). Using this, dual media compensation was possible, which allowed for correct imaging of the weld and for accurate reconstruction of the inner surface. Another group employed TFM for weld imaging of gas pressure rail welds (C. Lu, Deng, & Li, 2012). An additional weld imaging study employed TFM and compared the ability to detect defects against a number of other popular imaging algorithms (Jobst & Connolly, 2010). TFM was found to be generally superior, with advantages including greater focussing power and less interference from shear waves.

Another study employed TFM within an automatic flaw detection system, which supported dual media refraction compensation (Fradkin et al., 2013). This group employed a library of defect characteristics that aided in the identification of defect type. It has also been used for defect imaging in composite materials (Xu, Li, & Zhou,

2013), with one group compensating for the anisotropy in the weld (Chuan Li, Pain, Wilcox, & Drinkwater, 2013). In this study, effect of angle on group velocity was measured so that this effect could be implemented into the TFM algorithm. This aided in compensating for the changes in velocity throughout the material. In doing so, the defects that were previously undetectable using standard TFM were found to be 15dB above the background noise. Another group investigated railway wheel inspection, testing TFM against a conventional sector scan (J. Peng et al., 2013). The group simulated a PSF for arrays with different element counts to determine a suitable array configuration. They then compared the results of a 32 element configuration with both TFM and sector scanning, with TFM providing a more focussed depiction of the defect, improving detection likelihood.

It should be said, however, that almost all of these methods aim to find and characterise small defects in relatively large pieces. This differs quite drastically from the imaging of large surface areas, as will be carried out here. In saying this, there are a number of dual media applications, some of which perform the TFM algorithm on the FMC data (Dziewierz & Gachagan, 2013; Humeida, Wilcox, Todd, & Drinkwater, 2014; Leary et al., 2007; Sutcliffe, Weston, Dutton, & Cooper, 2012). In practice, the TFM algorithm is applied twice. The first of these creates an image of the medium boundary and a distorted image of the structures within the second medium. From this, the boundary between the media can be modelled as a line, which allows for methods of refractive compensation to be applied. Future research in this area will likely create mutually helpful results, as it is essentially surface imaging.

2.2.11 Common Source Method

While TFM requires the full matrix, there are other algorithms that can make use of the FMC, but do not use every A-scan available. This applies to the common source method (CSM) which can be calculated from a subset of waveforms from the FMC. In CSM, every element is employed in reception, but only one element is needed in transmission (Chiao & Thomas, 1994; Hunter, Drinkwater, & Wilcox, 2008). This is explained in Figure 2.13, where part (a) shows how all the possible transmit-receive

pairs are necessary for TFM. Part (b), on the other hand, shows CSM, where only one transmit position is used, but all receive positions are employed. While FMC is made up of N^2 A-scans, CSM needs only N . This can be achieved either during inspection by only firing one transmission element and receiving on all. It could also be accomplished in post processing, by employing only one of the transmission cycles. The latter method provides the opportunity to test different sources. The same process as described for TFM is employed to generate an image, with delays being calculated for each pixel and each receiver, allowing for focussing at every point in the image.

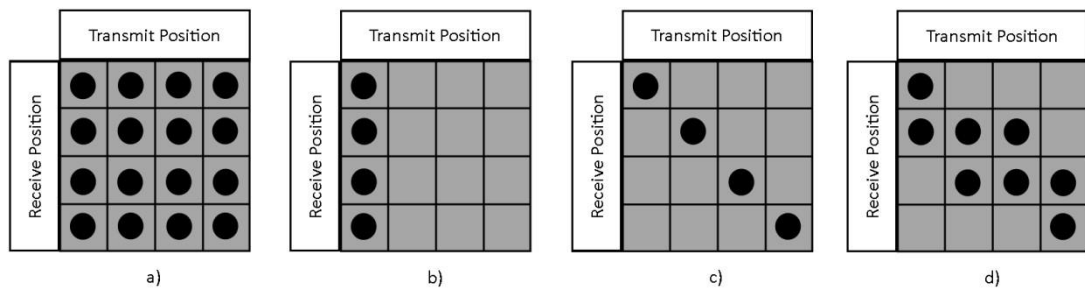


Figure 2.13: Transmit-receive matrices of a four element array for TFM (a), CSM (b), single element receive SAFT (c) and multiple element receive SAFT (d).

This method has not seen much uptake in NDT and is mentioned rather than researched. However, it has been compared against other advanced imaging algorithms by way of PSF (J. Davies, Simonetti, Lowe, & Cawley, 2006), where it was found to be inferior, with a greater simulated PSF width than found with TFM and in the ability to image holes in steel plates. It does, however, have the strength that processing time is low relative to TFM. In addition to this, if CSM is employed during data acquisition and only the required A-scans are captured, the acquisition is much faster.

2.2.12 Synthetic Aperture Focussing Method

As with CSM, SAFT can be performed as a subset of FMC. Again, the firing sequences can be altered so that only a subset of the A-scans are acquired. The most basic form

of this imaging method is a pulse-echo data capture method, where every element transmits and is received only on the same element. This can be seen in Figure 2.13 (c), which shows the receive element position following the transmit position. There are, however, variations to this method (Martín, Martínez, Octavio, & Godoy, 2008; Romero, Mart, & Mart, 2009). Part (d) of Figure 2.13, for example, shows the situation where for every one transmit element, a sub aperture of two elements is employed. This has been termed 2R-SAFT, based on the fact that two receive elements are employed. However, if the array driver is capable of providing FMC, the number of receive elements can be anything between 1 and N . As well as the number of elements in the receive aperture being variable, the method by which these elements are chosen is also flexible. It might, as in the given example, simply be the n nearest elements either side of the transmitting element. Another possibility is that the nearest n elements of the x coordinate of the pixel in question could be chosen. In addition to this, due to the large structure sizes, SAFT has, in certain cases, been employed using the mechanical movement of either single element probes or arrays, rather than the electrical scanning of a static array (Kaiys & Svilainis, 1995; Muller, Schmitz, & Schafer, 1986).

Once the signals have been acquired, the method is the same as TFM, whereby the area of interest is divided into a grid. However, unlike TFM, for every transmitter, only the contributions from a smaller number of receiving elements is considered. As such, this method provides the ability to focus at every point in the area of interest.

As with CSM, SAFT imaging has the potential to see increased processing speed, due to fewer receive elements being employed. Like TFM, it offers focussing at every point in an area of interest and has no limit on the viewing angle. When considering the point spread function and the ability to resolve two objects close to each other, single element SAFT has been shown to outperform both CSM and TFM (J. Davies et al., 2006). This was achieved by simulating two point scatterers near one another (ie. two PSFs). This simulation was repeated, increasing the distance between the two points with every iteration. When a dip appeared between the two points in the reconstructed image, in other words, when the two points were resolved, the distance was recorded. This gave a measure of resolution capabilities of the imaging methods. However, the same study also found that SAFT had a significantly higher side lobe level than TFM.

Unlike TFM, SAFT and the associated techniques have been used in NDT for decades. It has found employment in testing of concrete structures (C. Kohl et al., 2005; Krause et al., 1997; Ramm, Smith, & Pavy, 1991; M. Schickert, 2005; Martin Schickert, 2002), where the reduction in noise in the final image, as well as greater focussing power has allowed for more accurate detection of defects. In these studies, an array was not used, but instead a single transducer was moved along the surface of the specimen, recreating an array. Metal objects have also been imaged using SAFT (Stepinski & Lingvall, 2010). Mechanical scanning was also employed in this study, with the use of single element SAFT resulting in an increase in the resolution and decrease in background noise relative to traditional B-scan methods when imaging side drilled holes in a steel block.

Components in nuclear plants have been inspected using SAFT, allowing the group to image a wide variety of specimens due to the flexibility of the technique (Muller et al., 1986). Additionally, the algorithm was employed in the imaging of bridges, where the fusion of SAFT and radar data allowed for a 3D reconstruction of the inner parts of the structure (Christoph Kohl & Streicher, 2006). The abilities of SAFT in flaw sizing has also been investigated in austenitic welds (Holler & Schmitz, 1983).

In addition to practical use within industry, there are research efforts within NDT employing SAFT. An example of this is performance acceleration via GPU, where processing was implemented using CUDA, resulting impressive frame rates of 70fps (Romero et al., 2009). Reduction of grating lobes has also been investigated, with one group comparing PSFs created with different imaging methods (Martín et al., 2008). The study found that a two element receive aperture outperformed one element apertures, reducing side lobes and, therefore, improving contrast. Another group investigated array development, where a 2D annular array was compared to a rectangular array configuration using SAFT as the volumetric imaging method (Martinez, Akhnak, Ullate, & Espinosa, 2003). As well as this SAFT has been employed in defect characterisation in strongly attenuating materials (Spies & Rieder, 2010). In this study, the averaging effects of SAFT were exploited to improve the SNR, increasing the likelihood of defect detection relative to traditional beamforming methods.

These are just a sample of the uses of SAFT within industry and research. However, as with the employment of TFM, the current uses of SAFT within NDT have focussed on finding defects that are small in size relative to the piece of which they are a part. As such, the methods which have been investigated cannot be easily related to imaging large surfaces.

2.2.13 Biomedical Synthetic Aperture Techniques

Synthetic aperture techniques such as those described in the previous sections are also employed within the field of biomedical ultrasound imaging. These fall into five categories: SAFT, Multi-Element Synthetic Aperture Focussing (M-SAF), Synthetic Focussing method (SF), Synthetic Receive Aperture (SRA) and Synthetic Transmit Aperture (STA). Each of these techniques, like all previously described in this section, can be used to focus at every point in an area of interest through the described means. It should also be mentioned that naming conventions vary between fields and authors. As such, the names of techniques presented in this section may not conform with those employed by other authors.

The first of these is SAFT and is identical to the single element SAFT described above. Within biomedical ultrasound research, SAFT has not garnered as much interest as the other synthetic aperture techniques. One of the reasons for this is that it has a very low SNR because only one transmit element is used for every cycle. As biomedical applications often have to contend with high levels of absorption in soft tissues (Barnett, Rott, Ter Haar, Ziskin, & Maedaii, 1997), the power of the emitting transducer is vital to receiving discernible reflections (Synnevag, Austeng, & Holm, 2007). The technique does, however, boast an almost complete lack of the possibility of harm, seeing as only one transmitting element is employed at a time (Lockwood, Talman, & Brunke, 1998). Another of the disadvantages of SAFT is that once one element has fired, the next element can only be fired once reception has been completed. The time it takes to receive data for one A-scan can be found using equation 2.6, where T is the time it takes for reception of one transmission, D is the depth to which one wishes to image and c is the speed of sound in the medium.

$$T = \frac{2D}{c}$$

2.6

The length of time for one scan is *number of elements* $\times T$. So for arrays with high element counts, that can become a problem. The time to complete one scan might be so slow that motion artefacts appear (Karaman & O'Donnell, 1995).

The second method, M-SAF, transmits on a sub aperture of the full array and then receives on the elements of the same sub aperture (Karaman & O'Donnell, 1995). The purpose of this is to emulate the operation of SAFT, but to increase the power of the transmit beam. It is, like SAFT, not heavily employed in research. While it mimics the operation of SAFT, the amount of power in each transmission is greater because a sub aperture is used in place of a single element. The transmission sub apertures are delayed, such that natural focussing is eliminated and the emitted wave mimics a spherical wave as much as possible (Jensen et al., 2006). In addition to this, the number of sub apertures in a scan is likely to be much less than the number of elements in the array. As such, the acquisition time will be faster than in SAFT and the likelihood of motion artefacts will decrease (Karaman & O'Donnell, 1995). Additionally, the frame rates of M-SAF will be higher than those with SAFT. The maximum frame rate can be computed using equation 2.7, where F is the frame rate.

$$F = \frac{1}{\text{number of elements} \times T}$$

2.7

As T will be less with M-SAF than in SAFT, the frame rate will increase. This is of particular importance to medical imaging, where the ability to image in real time is often paramount. M-SAF also has the strength that, unlike SF which requires physically large hardware to complete the task of receiving on all elements at once, only a sub

aperture of the elements need be received at once. As such, it is more suitable for implementation in smaller systems, such as blood flow estimation (Bilge, Karaman, & O'Donnell, 1996). Studies into this method have found that it has a lower SNR than SAFT (Karaman & O'Donnell, 1995). While it should be said that this study found traditional phased array methods to perform better on both real and simulated data sets, only small receive apertures of 12 elements were employed. This reduced the complexity of the system, but did so at the expense of the reconstructive ability thereof. The third method, SF, is equivalent to TFM. Biomedical ultrasound imaging and ultrasonic NDT are separate fields of research and, as such, similar techniques are often named differently. Unlike the previously discussed biomedical synthetic aperture techniques, SF has found some level of research interest. The method of the operation of SF is almost identical to TFM. However, it does differ in one way. In order to retain real time imaging capability, recursive imaging is sometimes employed (Nikolov, Gammelmark, & Jensen, 1999). In SF, an intensity matrix is produced for every transmitting element, each of which is then summed to create a high resolution, focussed image. In TFM, this entire process is carried out for every frame created. In SF, however, a frame is displayed for every transmission made. As each matrix is produced, it is added to the total, I . When the $(N + 1)th$ transmission is made, the first intensity matrix is taken away from I and the $(N + 1)th$ is added. This enables much higher frame rates, as required in biomedical ultrasound imaging. This would, however, require N intensity matrices to be held in memory, which might become difficult, if not impossible, for systems with both a high element count and a high image resolution. Like TFM, SF also benefits from the ability to place the area of interest in any position (Jensen et al., 2006). This can be useful for imaging areas which have some kind of barrier with high acoustic impedance, such as the ribs. While imaging the heart, the probe can be placed some distance from the heart and still produce an image focussed at every pixel, by eliminating the need for the sound waves to navigate the ribs. In addition to this, while SA is computationally expensive when using every available element, it can be made to operate at very high speeds by employing a small subset of the full array (Jensen et al., 2006).

This method has been employed in flow imaging. The purpose of this is to track the velocity of fluids; usually through blood vessels. To measure blood flow, traditional beamforming techniques are capable only of finding the velocity along the direction of the focus of the beam (Jensen et al., 2006). Employing SF in this capacity allows for blood flow velocity estimation in any direction by considering the displacement of consequent images. A number of studies have employed this method, testing on both simulated data and from human inspections (Jensen & Nikolov, 2002, 2004). While the direction of flow was determined manually using B-mode images, the velocity calculations were carried out using SF, attaining a standard deviation between measurements of 0.65mm/s over the whole vessel profile. As well as being used in blood flow estimation, SF has also been employed in research into intraluminal imaging (O'Donnell & Thomas, 1992). An inter-luminal circular array configuration was employed in this study, producing cross-sectional images of arteries in the hope of detecting a thickening of the artery wall due to atherosclerosis. It was found that this method could produce high quality images and could reduce the amount of hardware needed within the artery relative to a traditional B-scan system with a high element count.

The reason SF and TFM might differ in operation is because the two methods have different demands set on them. SF has been used exclusively to image soft tissue through soft tissue, which is in complete opposition to its use in NDT. There is a large difference between the techniques required for bone imaging and soft tissue imaging. As such, the methods described for SA would not be applicable for the imaging of bony surfaces.

The penultimate imaging method is SRA, which derives its usefulness from being able to create a large effective receive aperture without the need of having a large number of receive channels in parallel (Nock & Trahey, 1992). Instead of receiving all the elements in the aperture at the same time, a number of sub apertures are employed which creates the large receive aperture. For every sub aperture in reception, there is one transmission which employs all the possible elements of the array. This is advantageous as having a high number of transmit elements is not costly, while having a high number of receive elements is (Nock & Trahey, 1992). Additionally, as all the

transmitters are fired to create the transmit aperture, the power will be higher, such that the signals in reception will be much stronger. It is also advantageous in that the data processing is less computationally expensive than SF, for example. As only one effective transmit aperture is in use, the delay law calculations are comparatively minimal. This method, however, may suffer from motion artefacts, as the image will only be constructed when all receive sub apertures have been used.

The final imaging method, STA, transmits only on a sub aperture, while receiving on all elements (Nikolov et al., 1999). It has the advantage that while it maintains the ability to transmit a beam from many locations on the array, it does not need the as much time to complete the data capture as SF, for example. The problem with this method relative to SRA is that it has a lower SNR, due to the reduction in transmit beam power, but this has, to an extent, been addressed (Chiao, Thomas, & Silverstein, 1997). In this study, the transmit apertures were spatially encoded, meaning that instead of firing a group of connected elements, the subset of elements were dispersed throughout the array, resulting in an increase in SNR.

As with B-scan imaging, biomedical synthetic aperture imaging offers much to the areas of research for which it was designed. As such, the strengths of these methods have been tailored towards soft tissue imaging, rather than that of hard, bony surfaces. Their implementation in such an application would likely not bear as much fruit as certain methods developed in NDT, which would have more in common with the imaging of bone.

2.2.14 Bone Surface Contour Extraction

While this study has not investigated many of the available bone surface extraction methods, there is an entire field of study dedicated to this purpose. Most of this research has been focussed on registering intraoperative ultrasound data with patientspecific CT data. While the difference in the purpose of this research relative to this body of research can cause the methods to be unsuitable, many of the problems encountered are common. As such, some of the solutions are applicable to both fields.

There have been a number of studies which have employed intensity-based bone surface extraction methods, which were designed for the registration of CT-derived, patient specific models. One method employed by others was to compute a probability map of the image based on the intensity gradient between pixels (Penney et al., 2006). This was achieved by interrogating the image column by column, moving from the probe face towards the bottom of the image. When a pixel with an intensity greater than a given threshold was reached, a marker was set. All pixels in that column between the marker and the probe face were eliminated from further processing. Doing this for all columns created an exclusion map. Another group employed a backward scan line tracing method, which made use of the shadowing effects of cortical bone in sectorial scan images (Yan et al., 2011). A further study simulated ultrasound within the CT slices, in order to make the CT and ultrasound data more similar by applying weights generated from the ultrasound simulation, thereby aiding in registration (Gill et al., 2012). Some groups have employed pointbased registration methods, depending on manual selection on a number of points representing the bone surface in each image (Barratt et al., 2006a). Given the high volume of data captures and the fact that intra- and inter-operator variability is to be avoided, this method would not be suitable.

More complex statistical methods have also been applied, such as the use of Bayesian probability to determine whether or not a pixel represents a surface (Jain & Taylor, 2004). This method employed knowledge of the appearance of bony surfaces in ultrasound B-mode images and more standard feature recognition techniques to determine the separate probabilities contributing to the Bayesian probability of the pixel representing a bony surface. A number of groups have also employed dynamic programming to extract surface contours, using a cost function to ensure continuity and smoothness (Foroughi, Boctor, Swartz, Taylor, & Fichtinger, 2007). Finally, level set methods have been employed before (Hanisch, Hsu, Noorman, & Radermacher, 2015a, 2015b; Chunming Li, Xu, Gui, & Fox, 2005) and could be of use in this body of work. These methods evolve an estimated initial curve so that it is optimised such that the bone surface is identified.

In addition to employing complex methods of contour extraction, another way in which the bone surface can be readily identified is to improve the images from which the

contours are being extracted. There are a number of methods that can achieve this. For example, depth weighted thresholding has been employed (Kowal, Amstutz, Langlotz, Talib, & Ballester, 2006), which favours those pixels furthest from the transducer face. Another method that has been used is the implementation of a shadow map, which makes use of the lack of reflection beneath the surface of the bone (Foroughi et al., 2007). In doing this, the bone surfaces are further enhanced, making surface identification simpler.

2.2.15 Bone Surface Reconstruction

Once the ultrasound imaging method has been decided and the necessary data extracted from the images, a number of means by which the final bone surface mesh can be reconstructed are available. The most common method by which this is achieved in knee joint reconstruction has been with deformable models. These are a group of methods which employ *a priori* combined with patient-specific, ultrasound-derived data to reconstruct bony surfaces. In the field of knee joint imaging, the *a priori* data, or training set, is usually a number of 3D models of either knee joints or full legs retrieved from CT data (Hanisch et al., 2015b). Including a large number of samples in the training set will help to compensate for natural variability of the shape of the bones in knee joints (Heimann & Meinzer, 2009). Using these models, landmarks are determined which represent the most important regions of variation within the set, termed modes (Fleute & Lavall, 1998). These regions may not correspond to actual anatomical landmarks and are, therefore, sometimes referred to as Point Distribution Models (PDM) (Heimann & Meinzer, 2009). Once this has been achieved, a mean shape is computed for the training set. The combination of the mean shape and the modes of variation are known collectively as Statistical Shape Models (SSM). This mean shape is then deformed in the regions of variation, with the linear deformations being weighted by the landmarks (in the form of 3D point coordinates) provided by the ultrasound data (Talib et al., 2005).

This method has been employed for bone surface reconstruction using ultrasound by a number of groups (Barratt et al., 2008b; C. S. K. Chan, Barratt, Edwards, & Penney,

2004; Khallaghi, 2010; Schumann et al., 2010; Talib et al., 2005), with one even achieving sub millimetre accuracy for the distal femur (Hanisch et al., 2015b). In addition to this, it is employed as an imageless surface reconstruction method by the Blue Belt (Blue Belt Technologies Inc., PA, USA) Navio Surgical System (Jaramaz, Hafez, & Digioia, 2006). The use of this product amongst surgeons for partial knee arthroplasties suggests an acceptance of both this method and the associated performance. As such, it could be said that the use of these SSMs would be justified through acceptance within the robotic surgery community.

However, it should be said that the accuracy of these methods is questionable, as a small sample size in the training set can cause there to be constricted deformation (Heimann & Meinzer, 2009). While this would be true for reconstructing healthy knees, the effects would be even more pronounced in the knees of those suffering from OA. Deformations in the general shape of the bone and the development of osteophytes would imply further variation in the bone shapes. As such, a lack of variation in the training set would cause even more pronounced inaccuracy in the surface reconstruction. Therefore, it would be better to avoid their use and employ only real, patient-specific data in surface reconstruction. However, if this was not possible and the use of SSMs became a necessity, the training set should be large (in the hundreds) and should include CT data relating to subjects with late stage OA. In doing this, the types of variation seen in such subjects would be included in the calculation of the modes of variation and would, therefore, allow for more accurate reconstruction of such knee joints.

2.3 Conclusion

A number of data capture and processing methods have been discussed. Details of the operation of each of these techniques have been presented, along with the associated strengths and weaknesses. In addition to this, examples of practical use and employment within research have been detailed, giving more insight into the strengths

of each method by providing context. Some of the methods were shown to be common to both biomedical ultrasound and to NDT, but a number of higher level techniques unique to each field were also presented. Despite the depth and variety of uses, it was shown that none were directly related to the imaging of bony surfaces. However, a number of the synthetic aperture techniques displayed the potential to produce an accurate image of a bony surface.

The information and results presented in this chapter are of importance, in that they aid in deciding the imaging method best suited to the task at hand. In addition to this, it has shown that no one method is entirely suited to imaging highly specular surfaces. In doing so, it has been determined that physical investigation into these methods and their performance is required to decide which method will perform to the highest standards in this task.

Chapter 3

System Component Design Consideration and Testing

3.1 Introduction

This chapter investigates a number of components that would be necessary in a robotically controlled, ultrasound knee imaging system. This begins with a short discussion of possible coupling media. After this, probe mount design, tool calibration and path planning were discussed in order to manipulate the probe automatically. Finally, to assess the accuracy of the calibration parts, the robotdefined tool position was compared to that found using a six camera Vicon motion capture system. In doing so, the ability of the robot to accurately report on the position of the probe was analysed.

3.2 Coupling Media

Ultrasonic measurements, be they industrial or biomedical in nature, require a coupling medium to link the face of the probe and the surface of the specimen. While considerable research efforts have focussed on employing air-coupled transducers in NDT (Grandia & Fortunko, 1995), non-gaseous materials remain the most popular means of coupling. This is due to the large differences in acoustic impedance between gasses and most common materials under inspection. For example, while air has an acoustic impedance of only 447Rayl, that of aluminium is 17MRayl (Halmshaw, 1991). Most ultrasonic inspection scenarios in both NDT and biomedicine employ contact probes, which are in direct contact with the surface of the specimen. In this situation, any air gaps are usually filled by a coupling medium. In biomedical applications, where the examination is external, this coupling medium is usually a water-based coupling gel (Szabo, 2014), while in NDT a number of different materials are commonly employed, including water, glycerine, silicon-oil, grease, water-soluble paste and petroleum jelly (Halmshaw, 1991). The choice of couplant is dependent on the acoustic impedance of the material, its toxicity, whether or not it is corrosive and, importantly, how expensive it is.

In this application, contact measurements are not practical, as industrial robots will be employed as probe manipulators. As such, immersion inspections are employed. The choice of coupling medium, then, is of great importance, not only in terms of ultrasonic propagation, but also in terms of cost and patient friendliness. With this in mind, the couplants listed above would not be suitable. The employment of gels would likely result in air bubbles, given the volume of couplant needed. Additionally, the cost would be prohibitive, given the need to replace gel after every patient. The oils employed as coupling media would also be unsuitable, in that such substances would be unpleasing, both practically and aesthetically. Such couplants would be difficult to clean from patients and might make some uncomfortable.

Further, as with gels, the cost of large volumes of oils could be prohibitive.

With this in mind, the most suitable coupling medium would be water. The acoustic impedance of water is very similar to soft tissue and would, therefore, satisfy the

ultrasonic properties required. Further, there would be little chance of air bubbles, as would be the case with different gels. Additionally, the cost would be minimal, if anything at all. This would allow for large volumes to be employed and replaced after every subject. Finally, most patients would be comfortable being partially submerged water. As such, it satisfies the requirement for patient-friendliness.

3.3 Probe Positioning and Tracking

Having concluded that water should be used as the coupling medium, other aspects of the setup of the system were considered. These included the type of housing for the samples and the coupling medium, the means by which the probe should be positioned, probe characteristics, tracking of the probe position and the path taken by the probe during inspection. Decisions relating to the nature of each of these areas were of importance, as the success of the system could be hindered by a weakness in any one. As such, the advantages and disadvantages of each possibility were carefully considered in an attempt to find a balance between practicality, cost, availability and fitness for purpose.

3.3.1 Probe Manipulation

The first area to be considered was the method by which the probe should be manipulated. As discussed in Section 1.3, there are a number of possible techniques available for this task, found in both biomedical ultrasound imaging and in NDT. Within biomedicine, real time, untracked imaging is most prevalent. As such, a hand held device using manual manipulation and contact scanning is the most common type of configuration. This, however, increases intra- and inter-operator variability (Vieyres et al., 2006), with the best results only possible with an experienced operator. In NDT, manual positioning is still common (Gordon Dobie et al., 2011). As in biomedical imaging, quality of inspection is highly dependent on the operator, with tiredness and

motivation thought to affect performance (Farley, 2004). With this in mind, an automated probe manipulation system was employed in the form of a 6 axis robot. Offering the versatility needed for system development and the repeatability required for implementation, these manipulators were considered to be well suited for the task of probe control.

3.3.1.1 Probe Manipulation Methods

Located within the Facility of Innovation and Research in Structural Testing (FIRST) laboratory, CUE has a pair of Kuka KR5 Arc HW 6 axis robots, as can be seen in Figure 3.1. Funded by grants from the ESPRC and the Department of Electric and Electronic Engineering (University of Strathclyde), the robots were acquired along with a number of pieces of equipment intended to help bridge the gap between research and industry in the field of NDT. The robots are capable of working either independently or as a cooperative master/slave combination. As can be seen in Table 3.1, the robots offer a high level of repeatability, but no information on accuracy is quoted. Designed specifically for arc welding, these robots are intended for applications in which the robot performs the same movements many times over a long period of time. Such tasks require high levels of repeatability, but not necessarily high levels of accuracy (Greenway, 2012; Liu & Wang, 1994). Quotes of accuracy are difficult to provide generally, but can be estimated for specific applications, with accuracy thought to be dependent on a number of factors including velocity, payload and pose.

Table 3.1: Specifications of the Kuka KR5 Arc HW 6 axis robots in CUE (C Mineo et al., 2012).

Specification	Value
<i>Maximum Reach</i>	1423mm
<i>Rated Payload</i>	5kg
<i>Maximum Total Payload</i>	37kg
<i>Positional Repeatability</i>	±0.04mm
<i>Number of Axis</i>	6
<i>Precision</i>	0.017mm




Figure 3.1: A pair of Kuka KR5 Arc HW robots within the robot cell in the FIRST laboratory in CUE.

In addition to the availability of these 6 axis robots, the employment of this family of manipulators can be further justified in that they offer a high level of versatility. While many scanning systems in both NDT and biomedicine are well suited for the prescribed task, it is this quality that makes them unsuitable for development of a system. Usually created in the form of gantries, these systems do not provide the ability to investigate different scanning paths, as would be necessary during development.

With the research themes in CUE generally being driven by industrial NDT demands, robotic manipulation is generally employed to increase both the speed and quality of inspections. The fixed robots are, therefore, often employed to manoeuvre different sensor types around complex geometries, maintaining either contact or close proximity. This has manifested in ultrasonic inspection of fan blades (R. H. Brown et al., 2015) and carbon fibre aeroplane wing covers, eddy current inspection of complex components found in the nuclear industry and for photogrammetric surface inspection

of wing covers. Using fixed 6 axis robots, therefore, allows for parts of various shapes and sizes to be inspected using complex tool paths.

Operation of one of the robots requires three main components: the robotic arm itself, the robot controller and the teach pendant. The teach pendant is the means by which the operator controls the robot and also provides the capability to program the motions of the robot. The pendant allows for four modes of operation: manual reduced velocity (T1), manual high velocity (T2), automatic (AUT) and automatic external (AUT EXT). Both manual modes require that the operator have a dead man switch on the pendant depressed at all times. If this switch is released, the brakes are enabled and the robot comes to an immediate halt. As a safety precaution, all work using the robots in CUE is carried out in T1, unless another mode is necessary. In this mode, the maximum velocities are limited, reducing the risk of collision. Paths that have been programmed offline are initiated using the teach pendant.

The robot configuration employs a Kuka KR C2 robot controller for each of the robots in the cell. The controller unit consists of two main features: power unit and the controlling computer. The purpose of the power unit is to amplify the signals sent to it by the controlling computer via a servo driving unit. These units control the six axes of the manipulator and are powered by an external power source. The controlling computer interprets Kuka Robot Language (KRL) input commands and sends an output to the servo driving units. In addition to this, the computer can act as a server, communicating positional information to an external computer.

In order for the controlling computer to translate the KRL input code into useable signals, the individual Cartesian coordinates of each point must be translated into axis angles for the servos. This is carried out using inverse kinematics, in which a tool point is known and the individual joint positions are calculated from this. To do this, the robot itself must first be described mathematically, for which DenavitHartenberg parameters are usually employed. These are four parameters that describe each link: three of which are static, while the fourth, being the angle between adjacent joints, is variable. Using these, a transformation matrix can be created which allows for the input of Cartesian coordinates and outputs joint angles (Greenway, 2012).

3.3.1.2 Probe Mount Design and TCP Calibration

While these robots have been designed with arc welding in mind, they have been repurposed as general sensor manipulators. No probe mount is common to all sensor types, so custom designed and manufactured mounts are a necessity to enable the manipulation of different probe types. These mounts must not only allow the probe to be attached rigidly to the end effector of the robot, but the design must take into consideration the manipulation requirements, as well as the tool calibration.

The first consideration to be made before the probe mount can be designed is the determination of which probe to employ. A number of different probe types are available within CUE, but not all were suitable for this task. It was decided that TFM and SAFT imaging were the best candidates for imaging the bony surfaces of the knee based on a review of the literature. Single element transducers were, therefore, redundant. As such, linear and 2D arrays were the only candidates. While CUE has a number of experimental 2D arrays, no commercial probes were available at the time. However, a number of linear arrays were available. All of the available linear arrays were manufactured by Vermon (Tours, France), were submersible, employed Hypertronics connectors (Smith Connectors, Hypertac, MA, USA), were flat and, as such, met the minimum requirements for consideration. The desired clinical accuracy for the system is 1mm, implying that the resolution of any probe should be less than this. While each probe met the criteria, this was merely a theoretical limit, computed as half the wavelength of the centre frequency. In practice, the probes would likely have resolutions at least twice that calculated (Angelsena, Tarp, Holmb, Kristoffersen, & Whittingham, 1995; Mautz & Tilch, 2011; Wilhjelm, Pedersen, & Jacobsen, 2001). This ruled out the 1MHz probe and brought the 2.25MHz close to the acceptable limit. While the 10MHz probe could provide a better resolution than the 5MHz, it has a number of relative weaknesses. Firstly, the full aperture of the probe was significantly smaller than that of the others, limiting the possible area of coverage. Additionally, the high centre frequency would produce a higher level of frequency-dependent attenuation. This effect can be characterised by a power law relation (Oralkan et al.,

2002), shown in equation 3.1, where α is the attenuation coefficient in $dBcm^{-1}$, f is the centre frequency and both α_0 and b are material dependent constants. It is clear from this law, then, that an increase in the frequency will also increase the attenuation.

$$\alpha = \alpha_0 f^b \quad 3.1$$

Table 3.2: Properties of the four available probe types.

Centre Frequency (MHz)	1	2.25	5	10
Number of Elements	128	128	128	128
Pitch (mm)	0.75	0.75	0.7	0.5
Elevation (mm)	15	12	10	7
Full Aperture Width (mm)	96	96	89.6	64
Minimum Theoretical Resolution (mm)	0.74	0.33	0.15	0.074
Rough Bandwidth (MHz)	0.5	1.125	2.5	5

With the reasons presented above, the 5MHz appears to be the best suited choice of probe, due to the balance it has between resolution, aperture width and attenuation. More practically, the 5MHz transducer, unlike the others, had four lips with holes. This allowed for a simple attaching mechanism that would ensure reliable and repeatable screw points. With these reasons in mind, the Vermon 5MHz linear array was chosen.

It was known that the samples to be inspected would be housed in a water bath, being completely submerged. With water being the ultrasound couplant, the probe would also have to be completely submerged. A constraint to such a setup was the fact that it was not permissible to allow any part of the robot to become submerged. To overcome this,

a probe mount was designed with a significant gap of 10cm between the end of the sixth axis of the robot and the probe. This can be viewed in Figure 3.2, which shows a separation between the base of the mount and the probe holder.




Figure 3.2: Exploded view of probe mount and the probe. The unlabelled white lines illustrate how the three parts come together.

The base was designed so that it would not connect directly to the robot itself, but would rather attach to an ATI MC-16 robot adaptor (ATI Industrial Automation, NC, USA). This is a manual tool changer that comes in two parts: the robot side and tool side. The tool (in this case, the probe mount) can be attached to the tool side of the changer, while the robot side is attached to the robot. This allows for simple and quick removal and application of the tool to the robot, while maintaining a rigid and repeatable tool position relative to the robot. The third part of the mount was the probe holder, which was designed to hug the probe tightly, preventing movement, as can be seen in Figure 3.2. Additionally, four screws were inserted, ensuring both rigidity and

accurate probe placement relative to the robot. As well as this, a second screw hole was added to the probe holder after the design stage. This was to provide more stability to the probe holder position by preventing it swivelling on the main screw. The design of all the parts was completed in Pro/ENGINEER Wildfire 5.0 (PTC, MA, USA) and the manufacture of the probe mount was carried out by Alec Ward (technician, Department of Electric and Electronic Engineering, University of Strathclyde).

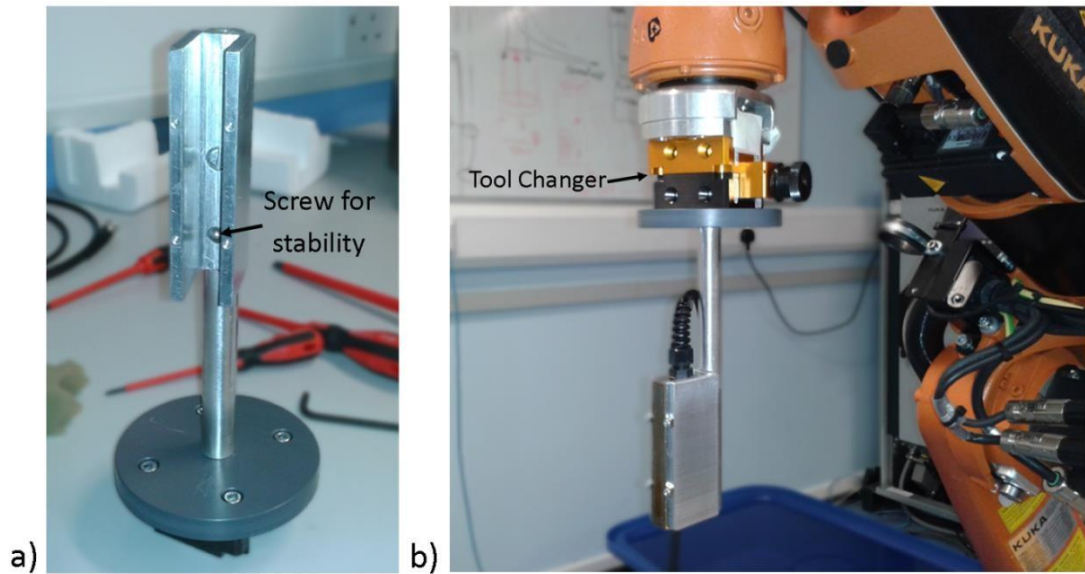


Figure 3.3: The probe mount showing the extra screw added for stability (a) and the probe and probe mount attached to the robot via the tool changer (b).

Given the shape of the mount, certain poses would be difficult, if not impossible, to implement, whilst ensuring full immersion of the probe and a complete lack of contact between the water and the robot. This is exemplified graphically in Figure 3.4, which shows a possible scenario in which the probe is moved such that the top side of the sample can be accessed. In doing so, the robotic arm collides with the side of the tank. Therefore, a second probe mount was required, which would increase the versatility of the system in terms of path possibilities.




Figure 3.4: Possible collision scenario with first probe mount design, showing the robot colliding with the tank. With the dotted line representing the water line, collisions could be possible when attempting to achieve a top-down perspective of the sample, while keeping the probe face fully submerged.

While the first probe mount placed the probe parallel to the sixth axis of the robot, the second mount was designed such that the face of the probe was perpendicular. For clarity, the first described probe will henceforth be referred to as Tool 1, while the second will be known as Tool 2. Doing this would solve the problem shown in Figure 3.4 and would also allow for the execution of paths that might have caused the robot to reach axis limits using the parallel probe. The perpendicular probe mount employed the same base, separator and probe holder as the parallel mount. However, the probe holder was attached to the end of the separator, rather than along its body, as can be seen in Figure 3.5. In addition to this, a stabiliser component was designed and manufactured that ensured the probe holder did not rotate around the main attachment screw. The fully completed Tool 2 can be seen in Figure 3.6.




Figure 3.5: Exploded view of the second probe mount for Tool 2. The stabiliser, which ensures the probe maintains the appropriate placement relative to the robot, is marked.




Figure 3.6: The probe mount (a) and the probe in place on the robot (b). The base was remade in plastic to minimise electrical interference from the robot.

3.3.1.3 Path Planning

Before path implementation or planning could take place, TCP calibration had to take place. This allowed the robot to relay the position of the probe face and the orientation of the image plane to the controlling computer accurately. The procedure describing this can be found in section A1. With both probes fully calibrated, it was then possible to both design and implement the path which the probe would follow around a sample. It was thought following the surface of a sphere would give the most complete coverage of any sample placed at the centre of the sphere. However, given that specimens would be fully immersed and would require some kind of frame to maintain balance, it was not possible to implement a full sphere, as the probe would collide with the sample or that which propped it up. As such, a hemisphere was seen as a good alternative. In this arrangement, the face of the probe would follow curved paths on the surface of a virtual hemisphere in the form of latitudes, as can be seen in Figure 3.7. At all times, the TCP (ie. the probe face) would be tangential to the surface of the probe, implying that the positive X axis of the probe would be pointing towards the centre point of the hemisphere. This would, in turn, maintain the field of view on the object under investigation at all times. Additionally, the hemisphere was augmented, so that the underside of such a sample could be accessed, as can be seen in Figure 3.7, where the path extends beyond the bottom of the hemisphere.

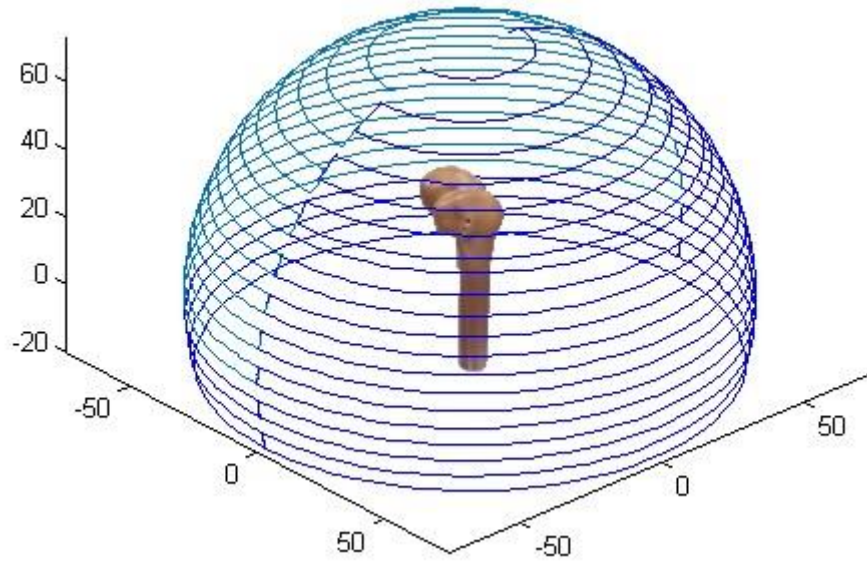


Figure 3.7: A representation of the augmented hemispherical path followed by the probe, with an example of a sample at the centre. The probe face would be tangential to the surface of the hemisphere, with the TCP pointing at the centre of the hemisphere at all times.

There are many of methods by which this could be implemented, including a number of third party robotic path planning software packages and the ability to hard code directly into the robot controller. KUKA robot controllers are capable of translating KRL - abstract, human-readable commands which define TCP position – into low level rotational joint motion (R. H. Brown et al., 2015). While KRL is employed directly for hard coding, it is also the language into which higher level path planning tools must convert commands.

There are four methods of motion which can be defined in KRL: linear (command name LIN), point to point (command name PTP), circular (command name CIRC) and spline (command name SPL). Through these motion commands, every possible movement is defined. Linear motion is the simplest of these and is defined using a start and end point. From the start point, the shortest path by Euclidean distance is calculated to the end point. This is always a straight line, which the TCP will follow, as is shown in Figure 3.8 (a). Point to point motion is similar to linear, in that it is defined using a start

and end point. However, instead of taking the shortest path by distance, the shortest path in terms of time will be taken. Intuitively, it might be thought that the shortest travel time would equate to the shortest distance between two points, but this is often not the case. A more curved path that deviates from the line between the start and end points may be quicker, as it may require less individual joint movement. Circular motion is quite different to these methods, in that it requires not only a start and end point positional definition, but also an auxiliary point. The auxiliary point is the position through which the TCP must move through whilst following an arc from the start to the end point, as is illustrated in Figure 3.8 (c). The final motion type is spline, which differs from the other motion types, in that only an end point is defined. The robot controller decides an appropriate path from which to move the TCP to the end point. This type of motion is usually reserved only for particularly complex paths, defined by many consecutive end points.




Figure 3.8: Graphical representation of three motion types defined in KRL: linear (a), point to point (b) and circular (c) (KUKA, 2010).

Coding the path directly using KRL (ie. hard coding) using these motion types is generally only employed in two cases. The first such scenario is where the path to be followed is very simple and, as such, can be defined by a relatively small number of lines and arcs. This method is advantageous in situations where the surface geometry of the sample under investigation is unknown, meaning a generic tool path is necessary to avoid collisions. The second case in which hard coding would be used is when a complex path is being followed. In this situation, the robot would be jogged to individual positions by an operator. Those positions would be recorded and employed in the simple motion commands described above, producing a complex path. This is often employed in industry (R. H. Brown et al., 2015; Mohr et al., 2014) for part inspection and for manufacturing processes. While this process can be very time consuming, this is often considered acceptable, as the motions will be repeated consistently over long periods of time, with little or no need to re-teach the positions. However, in an application such as joint scanning, the samples under inspection will change regularly. As such, this method would not be suitable.

Third party path planning software packages have been designed specifically to allow for complex paths to be defined based on accurate models of the object under inspection. These models are either provided by the manufacturer of the sample in the form of CAD models or the models can be reverse engineered using a Coordinate Measuring Machine (CMM). There are two such software packages available in CUE that allow for such complex path planning. The first of these is a combination of two packages: Mastercam X6 (Mastercam, CT, USA) and Robotmaster (Robotmaster, QC, Canada). The role of Mastercam is to receive CAD models of the part and create a wireframe of the surface. With this, tool paths are created as a dense series of points along the wireframe. In addition to this, normals to the surface are calculated at each of the points making up the path, allowing for probe orientation to be defined. This software does not, however, produce results that include actual tool and base data (Carmelo Mineo, Morozov, Pierce, Nicholson, & Cooper, 2013) that can be employed by particular robots. For this, Robotmaster is required, as it has the ability to take the path data produced in Mastercam and apply realistic parameters, such as tool and base frames, tool rotation and workspace boundaries, which allow for collision checking.

Additionally, Robotmaster allows for realistic visual simulation using a library of KUKA robot models. Finally, it produces KRL code representative of the generated path, which can be exported to the robot controller. As the path is broken up into a dense line of discrete coordinates, the KRL file is simply a long list of PTP or LIN commands which appear to be a continuous curve when executed.

The second package of this nature available within CUE bears similarities to this in terms of operation. This package employs Rhinoceros (McNeel & Associates, WA, USA) and a number of add-ons. To begin the path planning process, the mesh representing the object to be examined is imported into Rhinoceros. Using the Rhinoceros plugin Grasshopper 3D (McNeel & Associates, WA, USA) and KUKA|prc (Robots in Architecture, Vienna, Austria) – a plugin for Grasshopper – tool paths can be created in the form of discrete coordinates. Tool and base frames, tool rotations, robot definition and code output is then carried out by KUKA|prc.

While these methods are not only helpful, but in some cases necessary for tool path generation, the workflow is problematic for this particular task. The major prerequisite for these methods is an accurate model of the sample under inspection. The work presented herein seeks to create an accurate model of a sample of unknown geometry. As such, the work flow associated with path planning software is incompatible with this task. As well as the requirement that the tool path follow the surface of a hemisphere, it was also necessary that the dimensions of the hemisphere and the number of lines along the height of the surface be easily modified. Therefore, the method of teaching positions was eliminated, as it would be both time-consuming and inaccurate. Therefore, hard coding using KRL was employed as the means of tool path creation. While this method was constrictive, in that only simple motions were possible, this was suitable for the proposed hemispherical path.

In order to implement the hemispherical path using KRL, the start, end and auxiliary points had to be defined. All that was necessary for this was the radius of the hemisphere, R , and the difference in height between the latitudes making up the scan path, dH . The X coordinates for the first CIRC command were obvious, with the start coordinate being R , the auxiliary being 0 and the end being $-R$. While the X coordinate for the auxiliary point would always be 0 , as denoted by A in Figure 3.9, that of the

start and end points would have to decrease in magnitude at every iteration. As is illustrated in Figure 3.9, this could be simply calculated using the knowledge that the tool would always have to be in contact with the surface of the hemisphere and, therefore, the distance of the point to the centre of the sphere would remain constant, as R . For the start and end points, the X coordinate was found using Pythagoras. While the X coordinate of the start and end points would have the same magnitude, one would be negative and the other positive, as they would be at opposite side of the hemisphere, as indicated by the S and E in Figure 3.9. The Y coordinate of the auxiliary point was equal to the magnitude of the X coordinate, as it was on the same latitude. This procedure is summarised in the pseudo code of algorithm 3.1, where the coordinates of the start, end and auxiliary points are defined as S , E and A . The Z coordinate for all points was found straightforwardly by multiplication of dH by the latitude number, $(i + 1)$, where the first CIRC command was $i = -1$.




Figure 3.9: formulation of the coordinates of start, end and auxiliary points for the hemispherical path, with a sphere shown for reference.

The change in orientation is defined by the change in the angle α between each iteration. This calculation of this change in orientation is shown in Algorithm 3.1, where it is defined as Θ .

```

i = -1 while
α < 90      i
= i + 1
      α = sin-1(((i+1)*dH)/R)
          if
i>0        Θ
= 0        else
          Θ = α - α_old
end if      α_old = α

          if i is even
            Si(x) = (R2 - (i*dH)^2)^1/2
          else if
            Si(x) = -(R2 - (i*dH)^2)^1/2
          end if
          Ei(x) = -Si(x)
          Ai(y) = ((Si(x))^2)^1/2
          Ei(z) = (i+1)*dH
          Ai(z) = (i+1)*dH
          Si(z) = (i+1)*dH

          Si = (Si(x), 0, Si(z))
          Ai = (0, Ai(y), Ai(z))
          Ei = (Ei(x), 0, Ei(z)) end
while

```

Algorithm 3.1: the formulation of the coordinates of the start (*S*), end (*E*) and auxiliary (*A*) points at latitude number *i*. α is the angle between the horizontal as shown in Figure 3.9, α_old is the same value from the previous iteration and Θ is the change in orientation needed to ensure the probe face was normal to the centre of the hemisphere.

The above algorithm was implemented in KRL and, in combination, with CIRC commands, allowed for a hemispherical path to be followed. During testing of this code, the hemisphere was split into regions, as shown in Figure 3.10. The reasoning behind this was, firstly, for development of the code. In addition to this, Tool 1 was employed for this testing phase, but it was found that there was a particular region that the tool was unable to enter, due to axis limits being reached. With the radius of the hemisphere, *R*, set to 75mm and *dH* set to 5mm, the angle at which the probe in combination with Tool 1 could no longer continue was found to be 55°. Therefore, when this angle was reached, that part of the scan (shown in blue in Figure 3.10) was

stopped and was continued using Tool 2, which allowed for access in this region (black path).




Figure 3.10: The different components of the hemispherical scan represented with different colours, with the part of the scan completed by Tool 2 in black.

3.3.2 Path Accuracy

Calibration of the TCP and associated axes are of high importance to the operation of the imaging system. Any inaccuracy in the tool calibration would lead to inaccuracy in data placement during reconstruction. Therefore, calibration accuracy verification was required to eliminate the returned TCP coordinates as a major source of error. Within CUE, there are two systems routinely employed to carry out such accuracy experiments: laser tracking and infra-red motion capture. High levels of positional accuracy are possible using the Leica Absolute Tracker A2901-B (Leica Geosystems, Heerbrugg, Switzerland), with errors as low as $2\mu\text{m}$ (Summan, 2012). Despite the impressive accuracy, laser tracking systems such as this have limitations in terms of line of sight, due to the point and shoot method employed. Given the type of motion

involved, the line of sight between the single tracked marker and the laser source would likely be broken. Additionally, laser trackers provide no rotational data.

Vicon motion tracking systems are not capable of such high accuracies, with the best results being an order of magnitude greater than those achieved with the Leica tracking system (Windolf, Götzen, & Morlock, 2008). In addition to this, the accuracies achieved with Vicon tracking systems are highly dependent on a number of factors including camera position (Windolf et al., 2008), camera generation, position of marker set within measurement volume (Summan, 2012) and marker size (Windolf et al., 2008). In addition to causing possible reductions in accuracy, the issue of camera positions puts further restrictions on use by requiring a relatively large space in which to position the cameras. This point is important, as there was limited space in the FIRST laboratory at that time.

However, despite the large variability in the accuracy of Vicon systems and the problems relating to space, it was chosen over the Leica system. In order to properly characterise the accuracy of the calibration of the tools, it was necessary to measure the entire path, which would not be possible with the laser tracker. Therefore, a six camera Vicon system was used to track the position and orientation of a marker set simulating the TCP of Tool 1. For comparison, the TCP position was provided by the robot controller simultaneously, allowing for direct comparison.

3.3.2.1 Methods and Materials

A system consisting of six Vicon T160 cameras (Vicon Motion Systems, Oxford, UK) was employed, an example of which can be seen in Figure 3.11. In addition to the cameras, an MX Giganet (Vicon Motion Systems, Oxford, UK) was required to synchronise all the recordings from the cameras and to stream this data to the host computer. An explanation of the principles of operation of Vicon optical tracking systems can be found in (Summan, 2012).



Figure 3.11: The camera setup relative to the robot position, with all cameras on one side of the robot, due to spatial constrictions in the robot cell. Also detailed is the face of one of the cameras, showing the LED strobe ring and the lens in the centre.

While both the robot controller and the Vicon systems could provide positional data, it was necessary that these data acquisitions be synchronised. This was achieved using a custom program written in C# by Rahul Summan (Research Fellow, Department of Electronic and Electrical Engineering, University of Strathclyde). This program recorded the position of the TCP and the associated rotations provided by the robot controller on the host computer via the Robot Sensor Interface (RSI). This allows the robot controller to act as a server, transferring data over a network created between it and the host computer. Similar functionality is available with the Vicon software package Tracker. As such, data is streamed to the host computer as it becomes available, with data packages being sent from the robot controller every 12ms (KUKA, 2009). The Vicon system provided data at a higher and less stable rate of approximately every 8.3ms (120Hz sampling). As such, each sample comprised of a single robot-derived position and orientation and multiple Vicon readings.

It should be mentioned that the program only allowed for measurements to be made when the TCP was at rest. Therefore, the KRL code could not be employed, as much

of the movements were continuous. To remedy this, the path was executed and only the robot positions were recorded. These positions were then used to create a KRL script with a large number of small PTP motions. This, in effect, reproduced the path, with the change that instead of a continuous motion, there were many decelerations and accelerations at each PTP end and start point.

As mentioned earlier, there are numerous factors which influence the accuracy of Vicon positional measurements, with camera position being one of these. It is beneficial to have a large variation in the positions of the individual cameras, which allows for a greater field of view and minimises the chance of blind spots. However, the extent to which this could be implemented was limited by the spatial restrictions of the robot cell. As such, the cameras were organised such that all were placed on the same side of the robot, as can be seen in Figure 3.11. While it was not possible to position any cameras on the other side, it was hoped that spreading the cameras out across 180° from the centre of the robot would provide enough variation in position. Before data recording took place, the path was executed and the positions of the cameras and the marker set were visualised using Tracker, an example of which can be seen in Figure 3.12. Doing this ensured that the TCP could be reconstructed in every position.




Figure 3.12: The positions of the six cameras relative to each other within Vicon Tracker software. Also shown is the marker set and the TCP (white box).

Two calibrations were necessary before measurements could be made: camera calibration and marker set calibration. The first of these to be carried out was the camera calibration, the purpose of which was to define the relative positions of the cameras and, in turn, allow for accurate reconstruction of marker positions. A standard calibration method was employed, whereby a known marker set was moved by hand throughout the volume under inspection in a figure of eight motion. Each camera captured a user-defined number of frames. From this, point cloud data was created and employed by the Vicon software to determine the relative positions of the cameras. The marker set used to do this was a calibration wand, which can be seen in Figure 3.13 (b). This is a standard tool of known dimensions and marker positions.




Figure 3.13: The marker set employed, with the sphere representing the TCP attached to the gold coloured screw (a) and the calibration wand (b). The silver screw in part (a) was employed to adjust the position of the TCP sphere.

Once this procedure had been carried out, the system was capable of determining the position of any markers within the field of view, allowing for any marker set to be defined. The set consisted of five 12mm markers, four of which were attached to the base of the probe mount using double sided tape. The fifth was screwed on the end of

a rod which was attached to the probe separator of the mount. As can be seen in Figure 3.13 (a), there is another screw on the opposite side of the separator. This acted as a stabiliser for the gold rod and allowed the position of the marker representing the TCP to be accurately placed. Using the same procedure described in Section the appendices, the distance from the top of the sphere to the flat surface of the separator was determined using the Faro touch probe. The distance was lessened by 6mm to account for the radius of the sphere, as the point which is tracked is the centre of the sphere. The position of the gold rod was moved and measurements taken incrementally to accurately place the sphere. Once in the field of view, the marker set was defined as having a centre point at the centre of the sphere representing the TCP, as can be seen in Figure 3.12.

Once the data had been gathered, to match one Vicon measurement with one robot measurement, the mean of the multiple associated Vicon measurements was taken. However, direct comparison was not possible, as the robot and Vicon measurements were relative to different reference frames. As such, the resultant point clouds were matched using Iterative Closest Point (ICP). This algorithm consists of three main phases: matching each point in cloud 1 to the nearest neighbour in cloud 2, minimisation of the sum of the square of the distance between point pairings and finally, transformation of cloud 1 using the minimisation (H. M. Kjer & Wilm, 2010). The second and third phases are repeated until either a certain number of iterations have been implemented or a specified limit on minimisation has been reached. An example of the results of such a procedure can be seen in Figure 3.14 where robot data has been matched to Vicon data. Formulated and built upon by a number of authors (Besl & McKay, 1992; Y Chen & Medioni, 1991), the method is a standard in point cloud matching. A custom Matlab function 'icp' (M. Kjer & Wilm, 2012) was employed to carry out the matching process using ICP. However, the first phase of ICP is to match each point in cloud 1 to the nearest neighbour in cloud 2. This was unnecessary, as it was known that point i in the Vicon data matched point i in the robot data. As such, the function was altered to account for this. When this had been completed the absolute Euclidean distance between each point pairing was calculated.

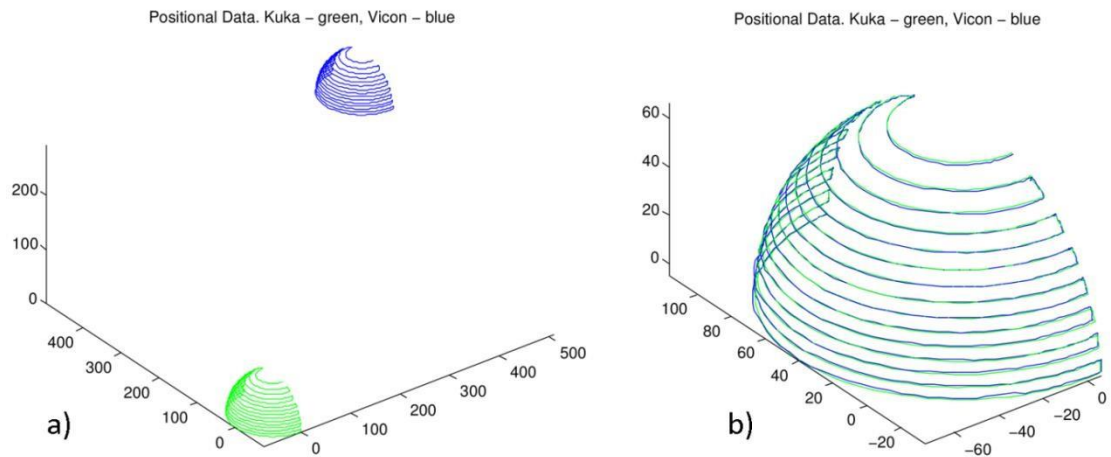


Figure 3.14: An example of the results of ICP, taking unmatched data (a) and transforming one set to minimise error between the sets (b).

There is, however, a weakness in this method of comparison, in that it was highly dependent on the point cloud matching process. In an attempt to eliminate this dependence, a second method of comparison was also employed, in which the distance between each consecutive point in each of the data sets was calculated. The difference between associated pairs was then calculated. In other words, the Euclidean distance between point i and $(i + 1)$ in the Vicon data was calculated. Alongside this, the distance between point i and $(i + 1)$ in the robot data was calculated. The difference between these two values was then found. While this did not give information about absolute position, it provided a comparison related to the change in position between each measurement point.

3.3.2.2 Results

The absolute error between paired Vicon and KUKA robot measurement points had a low mean error and standard deviation, as can be seen in Table 3.3. However, a large maximum error was returned, which can be seen concentrated in a small area in Figure 3.15 (b). In addition to this small incidence of large error, there are regions with errors around 1mm. This can be seen in Figure 3.15 (a), where two curves are almost completely yellow, indicating errors in excess of 1mm. Also, the side of the hemisphere

shown in Figure 3.15 (a) appears to have a lower level of error generally than the other side.

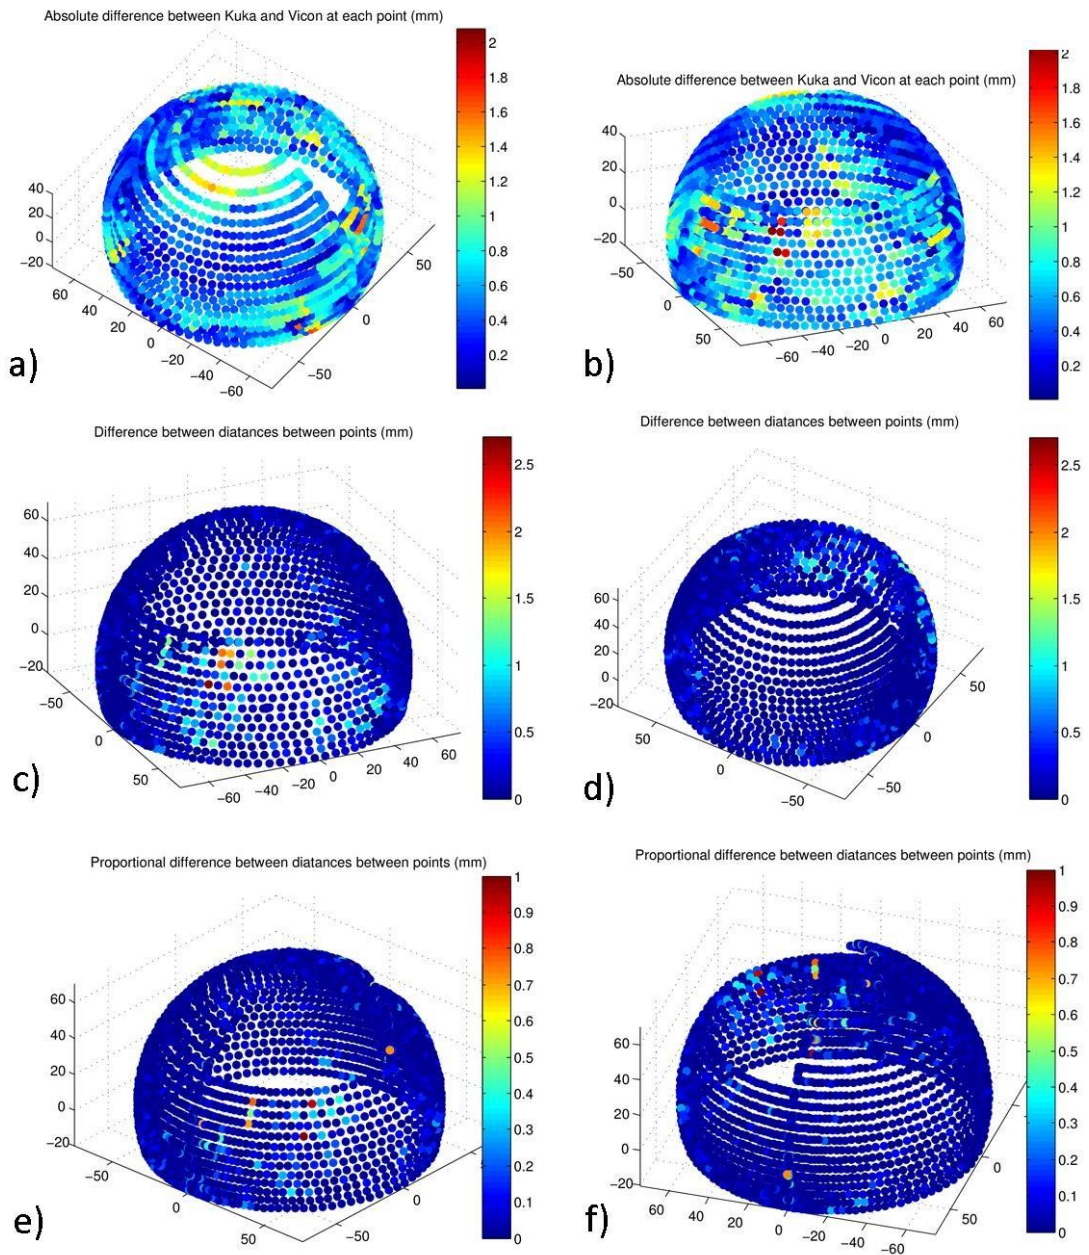


Figure 3.15: Absolute error between paired points of Vicon and Kuka robot data (a, b), difference between distance of consecutive points in each data set (c, d) and the previous value normalised by the corresponding distance between consecutive points provided by Vicon (e, f).

Table 3.3: summary of the results of the comparison between the Vicon positional data and the robot positional data.

Comparison Method	Mean Error	Maximum Error	Standard Deviation
Absolute Euclidean Distance	0.5 mm	2.1 mm	0.3 mm
Difference of Distance Between Consecutive Points	0.1 mm	2.7 mm	0.17 mm
Normalised Difference of Distance Between Consecutive Points	0.1	1	0.2

The second error measurement method had a very low mean error and standard deviation relative to that found in the first method. However, relative to the mean error, the standard deviation was found to be high. The maximum error was, once again, very high and occurred in the same small pockets that were seen in the first method. As with the first method, the side presented in Figure 3.15 (d) generally showed a lower error than that illustrated in Figure 3.15 (c).

The final method of error measurement yielded a mean error of 9.3% and a standard deviation of 19.13%. The mean absolute distance between each recording point provided by Vicon was 2.36mm. Using the mean distance between measurement points, the mean error can be becomes 0.22mm. As such, the error is quite low, when considered as a measurement of distance, rather than a proportion. While the maximum proportion was 100%, this may have been a very low distance, as the associated distance between the Vicon measurement points could have been very small in terms of Euclidean distance.

3.3.2.3 Discussion

The results show an acceptable level of mean error, but suffer from regions of high levels of error. However, the fact that the highest errors were found in pockets suggest

there may be an underlying cause that does not represent true error. Focussing on the highest errors found in the absolute error results, these are found only in very small pockets of a few points, as seen in Figure 3.15 (b). It is believed that these were, in fact, caused by wobble of the tool mount. When the path is being executed, a dead man switch is held on the teach pendant. If this switch is released, the robot brakes are enabled immediately, causing a sudden stop in movement. The abruptness of this can cause tools to visibly wobble, with this effect being augmented with longer tool mounts. It is believed that a mistake on the part of the operator caused the depression of the dead man switch to be interrupted and, as such, the marker set wobbled. This effect was measured accurately by the Vicon system, but would not be registered by the encoders in the robot. This almost certainly accounts for the error mentioned above and may also be responsible for other small pockets of high error. While the absolute error has been referenced, this effect is also present in the second error measurement method, where high error occurs in this same place, as is shown in Figure 3.15 (c).

In addition to these small pockets of a few points, there were larger regions with high levels of absolute error. As before, it is believed this was not a true error, but rather a badly reconstructed Vicon TCP position. Figure 3.16 displays the robot position on one of the two curves with high levels of error, as seen in Figure 3.15 (a). Most of the cameras are positioned to the left of the marker set. It is clear that the robot is at least partially blocking the line of sight between the cameras and the markers. The Vicon system does not require every marker to be visible to every camera to reconstruct the position of the TCP, instead using as many markers and cameras as possible. However, the fewer images used to reconstruct the TCP position, the less accurate it will be. It is believed, therefore, that the regions of high error identified were, in fact, points where the Vicon system had failed to accurately reconstruct the TCP position. It is, however, difficult to verify this and identifying the regions in which this effect was responsible for inaccuracy is merely educated estimation.




Figure 3.16: The position of the robot at a point on one of the two curves identified as having high error. Most of the cameras are positioned to the left of the robot, meaning the markers are, for the most part, blocked by the robot itself. With fewer rays intersecting the spheres because of line of sight blockages, the accuracy of the reconstruction degraded.

All three of the error measurement methods provided a reasonably low mean error, while maximum error in each method was quite high. While they had these characteristics in common, the value these errors represented were vastly different in each case. The absolute error measurements gave the simplest and most useful representation of error, being a direct comparison between the two systems. It is believed that the ICP matching algorithm worked well and, as such, the results are reliable.

With the other two methods being a change in position, this eliminated the dependence on the matching algorithm. The third method put the values found using the second into context, as the distance between points changed depending on what part of the path was being traversed. Around the centre of the arc, the measurement points were sparse, while they were denser at the start and end points of the CIRC motion, which can be

seen in Figure 3.15. The reason for this is that all motions defined in KRL cause the tool to accelerate from the start point and then begin to decelerate as they pass the half-way point in the motion. Doing this causes less abrupt braking and acceleration, thereby reducing wear and increasing longevity. The results found with these two methods yielded a low mean error, high maximum error and relatively high standard deviation. Despite the difference in standard deviations, the three methods did return similar results. As can be seen in Figure 3.16 (a), (c) and (e), maximum errors occurred in similar regions. Additionally, all three yielded low mean errors, implying a level of concurrence.

While this entire section relates to the accuracy of the calibration parts, it should be made clear that all 6-axis manipulators will be inherently lacking in accuracy to some level (Young & Pickin, 2000). This intrinsic error is defined as the difference between the actual pose of the robot and that calculated by the robot controller and is known as pose error. There are a number of factors affecting pose error, the first of which is that the link lengths employed by the robot controller to calculate the TCP position are generic and, as such, are inaccurate. Each unit has minor variations in link length and while the lengths for individual units can be calculated, this has not been carried out for the two robots in the cell (Greenway, 2012). Secondly, these robots are designed specifically to have high levels of repeatability, but not to have high levels of accuracy (R. H. Brown et al., 2015; Carmelo Mineo et al., 2013). It is near impossible to quantify the contribution of this inherent error, as the level of error appears to be related to a number of factors, one of which is velocity.

While the velocity of the TCP is an indirect factor, the rate of change in angle at individual joints is where the true impact occurs. The faster the velocity is, the less likely it will be that the encoders will be capable of accurately estimating the angle at each axis. If the angle measured at each of the axes is inaccurate, this error will be multiplied in the reconstruction of the TCP position. Additionally, the pose of the robot causes variation in the level of accuracy. When the robot is calibrated, a static reference point is employed. As such the variation in angles is limited, as the TCP must be in contact with the reference point at every measurement. Therefore, when the TCP moves away from the region which was inhabited by the reference point, the angles at each

axis become less like the angles found when calibrating the robot. This can be minimised by employing as large a variation as possible in pose at each of the four calibration points, but cannot be eliminated.

As well as this, both the weight and the length of the probe mount can cause wobble, which cannot be accounted for by the robot controller. A further problem is that the resolution of the encoders at each joint is a limiting factor, as this limits the level of accuracy possible at each joint (Greenway, 2012). Additionally, thermal effects caused environmentally or by the operation of the robot itself can cause errors. Many of the components that make up a robot are susceptible to expansion, affecting many facets of the operation of the robot (Greenway, 2012; Weill & Shani, 1991).

3.3.2.4 Conclusion

While three methods of error measurement have been presented, the most simple of these – the mean absolute error – was sub-millimetre. Given the contribution of error from the braking effect, line of sight issues and the high complexity of the path, the errors found herein can be thought of as a worst case scenario. The standard knee scans carried out using CT - ‘gold standard’ method - employs a resolution of 1mm (Citak et al., 2013; Conditt & Roche, 2009; B. Davies et al., 2006; Roche et al., 2009). As the worst-case scenario mean absolute error was significantly lower than the resolution of the ‘gold standard’ method against which this system is to be compared, the robotic positing system was judged to be trustworthy and capable of returning positional data of an acceptable accuracy.

3.4 Conclusion

A number of the fundamental, non-imaging parts of the system were discussed, investigated and evaluated. Taking into account a number of factors, the most suitable coupling medium was found to be water. Additionally, the possible path planning methods were considered, with KRL deemed the most appropriate method for imaging samples with unknown geometries. Further to this, the design of the probe mounts and

calibration parts that allowed for probe manipulation was presented. Finally, the accuracy of the TCP calibration and, more generally, the robot accuracy was tested using a six camera Vicon system, finding the positioning system to be acceptably accurate.

The work presented in this chapter confirmed some of the important aspects of the system, essentially preparing it for imaging tasks. With accurate robot positioning now possible, it was necessary to test the system. The ability to accurately reconstruct surfaces, however, relies not only on accurate probe position reporting, but on the effect of the imaging method employed. It was, therefore, necessary to test using physical samples in order to fully understand the capabilities of the system.

Chapter 4 Imaging Known Geometries

Imaging complex surfaces, such as those found in the bony surfaces of the knee, requires consideration of both data capture and data processing methods. The techniques in both areas discussed thus far have been largely limited to NDT; research in which does not generally focus on acquiring surface profiles. Therefore, while many of the possible issues associated with this kind of imaging can be predicted and at least partially accounted for, the performance of the previously discussed methods is, for the most part, unknown. It was thought, therefore, that it would be prudent to test the imaging capabilities of the system with more simple geometric shapes before moving on to more complex surfaces. Doing so would allow the strengths and weaknesses of the system to be identified, so that attempts could be made to account for them.

4.1 Materials and Methods

As simple geometric shapes were to be tested, it was thought that three surface characteristics should be focussed on: curves, flat surfaces and sharp edges. The presented surface types would provide different varieties of reflection, which would test the ability of the imaging processing methods to account for the differences. Additionally, it was hoped that the simple flat surfaces would test the level of accuracy that could be achieved. Therefore, three shapes were chosen: a sphere, a cuboid and a partly flattened cylinder; all of which can be seen in Figure 4.1.



Figure 4.1: The three test pieces used: a sphere, cuboid and flattened off cylinder. Also shown is the mount, to which each was attached to allow full access.

The sphere provided, in all regions, highly curved surfaces, while the cuboid offered both simple, flat surfaces and sharp corners. The cylinder had one side flattened off, which was done so that the sample could provide curved, flat and sharp surfaces in close proximity. The cylinder and cuboid were made from aluminium, while the sphere was made from brass (Dejay Distribution Ltd., Cornwall, UK). The sphere had a diameter of 25mm and the cuboid had a height of 25.12mm, a length of 35.4mm and a width of 35.2mm. The cylinder had a height of 62.5mm and a diameter of 50mm, with the edge being flattened 15mm from the central axis. As such, the maximum dimension of all the samples were smaller than the full aperture width of the probe (89.6mm). In addition to the tests pieces, a mount was also manufactured, allowing the samples to be placed in a water bath such that they were not in contact with the bottom surface. This avoided collisions and ensured the accessibility of more of the surfaces. All work to create and alter the test pieces was carried out by Alec Ward (technician, Department of Electronic and Electrical Engineering).

Each of the samples was placed in a water bath in close proximity to the robot, as can be seen in Figure 4.2. Two paths were employed with each of the samples, the first of which was the hemispherical path described in the previous chapter. The centre of the hemisphere was aligned approximately with the centre of mass of each of the samples by eye and the radius was defined such that there existed a clearance of a few centimetres between the probe and the samples at all times. The second path was a simplified cuboidal path, as can be seen in Figure 4.3. An additional two passes were made over the top, so that coverage could be more complete. This path was executed simple by jogging the robot (ie. moved using simple axis-based movements using the pendant), while keeping linear movements on the world axes. Unlike with the hemispherical path, the probe did not rotate during these motions, so the probe face was not pointing towards any centre point. In addition to the cuboidal and hemispherical probe paths, three more scan paths were employed with the cylindrical sample. As is detailed in Figure 4.4, the paths were intended to probe the curved part of the cylinder.




Figure 4.2: The probe being manoeuvred around the sphere sample, ensuring immersion of the probe and a lack of contact between robot and water at all times.




Figure 4.3: The cuboidal path employed in the cylindrical sample, showing two passes over the top to extend coverage. Also shown are the probe orientations relative to the surface.




Figure 4.4: The three additional scan paths employed for inspection of the cylindrical sample.

A 5MHz, 128 element linear array (Vermon) was employed as the ultrasonic transceiver. This was driven by a Flawinspecta phased array controller (Diagnostic Sonar Ltd., Livingston, UK). This is a modular system, allowing for implementation of further components that increase parallelism and, as a consequence, increase performance. At the time of use, the system had two digitisers in the form of FlexRIO FPGA cards (National Instruments Corporation, TX, USA), each providing 32 parallel channels. This configuration in conjunction with multiplexers was capable of 128 element FMC. As such, each data capture took this form.

Diagnostic Sonar provided controlling software, FIToolbox, which could be modified and built upon using Labview (National Instruments Corporation, TX, USA). In order for the acquired data to be properly associated with the correct probe position, the output FMC files needed to be synchronised with the closest corresponding KUKA-derived probe position. The ability to do this was integrated into the FIToolbox by Timothy Lardner (Research Assistant, Department of Electronic and Electrical Engineering) and Joseph Jackson (Lecturer, Department of Electronic and Electrical Engineering). In addition to this, a stripped-down, custom version of the CUE implementation of the FIToolbox was written in conjunction with Joseph Jackson, so that there were no limitations in the number of FMC files that could be captured, with the exception of the amount of free space of the hard drive. As well as this, the custom implementation enabled the FMC files to be saved in a single precision binary format, rather than the standard double precision format. This was necessary for two reasons. Firstly, in order to accurately represent a 3D sample with 2D images, large numbers of images and, therefore, FMCs would be required. Reducing the precision brought that data set sizes down to an acceptable size. Additionally, the measurements were taken in a depth of 8cm in water. Ultrasound travels slowly in water relative to solid materials and, as such, requires more time to capture data from the same depth of material. As the sample rate would not change, this implies that the number of samples in an individual A-scan would be large in water, because of the increased recording time per A-scan. Due to this, individual FMC file sizes were large and it was necessary to reduce the precision to make the data set sizes more manageable.

The positional increment size at which the each FMC was captured was dependant on a number of factors, the first of which was the TCP velocity. For the cuboidal paths, the robot was jogged at 1% velocity in T1 mode, providing a near constant velocity, with very little acceleration or deceleration. However, velocity of the TCP during the execution of the hemispherical path was variable, accelerating from the start of the CIRC motion and decelerating towards the end. Therefore, while the CIRC commands were carried out at 10% velocity in T1 mode, the measurement points were sparse in the middle of the arcs of the CIRC motions and dense at the beginning and ends. The second factor in increment size was the FMC sample rate. This was dependent on the number of elements used and the A-scan sample size. The latter was determined by the A-scan sample rate, the ultrasound material velocity and the depth to which the piece was to be inspected. Given an A-scan sample rate of 40MHz and an imaging depth of 8cm, the FMC sample rate, while highly variable, was roughly 0.3Hz. This was uncharacteristically slow, but expected, given the high element count, the slow ultrasound velocity within the material and large imaging depth.

To retrieve useable images from each of the FMC files, a synthetic aperture data processing technique was required. Despite the improvements in imaging capability offered by such methods, computational expense is still a major issue. This expense results in a long processing time, which is particularly troublesome for this task. The major influences in processing time include the number of elements used, the number of pixels in the resultant image (ie. the resolution of the image) and the number of images. The number of 2D images needed for 3D reconstruction is high, increasing processing time for each data set. In order to provide acceptable levels of sample coverage and to give the best possible image reconstruction, all possible 128 elements of the probe were employed. This is particularly time consuming in that the processing time increases like N^2 . Finally, given the imaging depth and the need for high resolution images, the number of pixels in each image was in excess of 10^6 . Processing a single FMC using an algorithm lacking optimisation and a normal CPU would take in excess of 30 minutes. Therefore, processing hundreds of FMCs per data set was not viable, particularly if multiple processing methods were to be used with each FMC.

To remedy this, CUE created CueART: a software platform designed to exploit the parallelism of synthetic aperture imaging methods in order to implement massive reductions in computation time per image (Dziewierz et al., 2012; Mcgilp, Dziewierz, Lardner, Gachagan, & Bird, 2014). The platform has been developed in Compute Unified Direct Architecture (CUDA), which is a parallel computing platform designed to allow exploitation of NVIDIA GPUs for general purpose tasks. While CueART has been incorporated into the FIToolbox to move towards real time TFM imaging, it can also be employed in post processing on FMC data sets. The platform can be called in Matlab by way of a Matlab Executable (MEX) file, which allows Matlab-defined variables to be employed in CueART functions, with the results returned to Matlab. Following installation of the CueART platform, the FMC files were processed on a Dell Alienware X51 R2, with an Intel® Core™ i7-4790 processor and an NVIDIA GeForce GTX 760Ti (1344 CUDA cores).

The resultant images had a resolution of 0.1mm, which was chosen as it is below the theoretical resolving limit of the 5MHz probe of 0.15mm. Each image consisted of 1,204,224 pixels in a 1344×896 configuration. The maximum depth was 8.95cm, while the total width was 13.42cm. This corresponded to 1.5 times the length of the probe, which allowed for surface coverage past the limits of the probe edges.

Production of a single image took approximately 30 seconds.

As well as TFM, CueART has the functionality to provide SAFT imaging. This variation of SAFT defined the width of the synthetic aperture based on the X coordinate of the pixel under reconstruction. This is illustrated in Figure 4.5, where the width of the SAFT aperture is defined by both the position of the pixel in question and an isosceles triangle defined by Θ . An increase in the size of Θ would result, therefore, in an increase in the aperture width, with the number of elements between α and β increasing also. A pixel can, therefore, be reconstructed using equation 4.1, with the same definitions as given to equation 2.4. A 20° aperture SAFT was chosen, as this is quite small relative to the full aperture used by TFM. It was hoped this would provide a profound dissimilarity in resulting images, so that the relative strengths and weaknesses could be identified.

$$4.1 \quad I(p) = \sum_{i=1}^N \sum_{j=\alpha}^{\beta} S_{i,j}(T_{i(p)} + T_{j(p)})$$




Figure 4.5: Graphical representation of the TFM and SAFT synthetic aperture definitions, with the elements constituting the SAFT aperture shown in a lighter shade. Defining θ as a relatively small angle (eg. $<40^\circ$) will ensure that $\beta - \alpha$ is less than N , so that the reconstructing aperture will always be smaller than with TFM.

To provide a representation of the surface profile in each image, the pixels representing it had to be identified and isolated. Typical images showed not only the reflections from the surface, but often displayed a strong reflection caused by the waves having a double reflection: from the probe to the surface, to the probe to the surface and back again. Additionally, a number of backwall reflections were present, close together because speed of sound in aluminium is significantly greater than that in water. This is illustrated in Figure 4.6, which is a typical SAFT image from the cuboidal path inspection of the cuboid sample. The first step in identifying the appropriate pixels was to apply global thresholding, eliminating the pixels below a certain decibel (dB) limit. However, as can be seen in Figure 4.7 (a), this did not only identify the initial surface

reflections, but the other reflections also. In an effort to rid the contour extraction of these erroneous reflections, the coordinates of every point above the threshold was examined. If any of the pixels contained the same X coordinate, they were discarded, with the exception of the pixel with the smallest Z coordinate. The smallest Z value corresponded to the pixel that would be closest to the probe surface and would, therefore, be the most likely to represent the initial surface reflection. The result of this change can be viewed in Figure 4.7 (b), which shows a relatively accurate surface profile extraction and an almost complete elimination of unwanted reflections.




Figure 4.6: A typical SAFT image from the cuboid test piece, showing the element positions as red circles. The surface profile is seen at the top of the image, while the return signal reflection can be seen as a strong reflection towards the bottom of the image. Also present is an obvious backwall reflection, shown as the second surface representation.

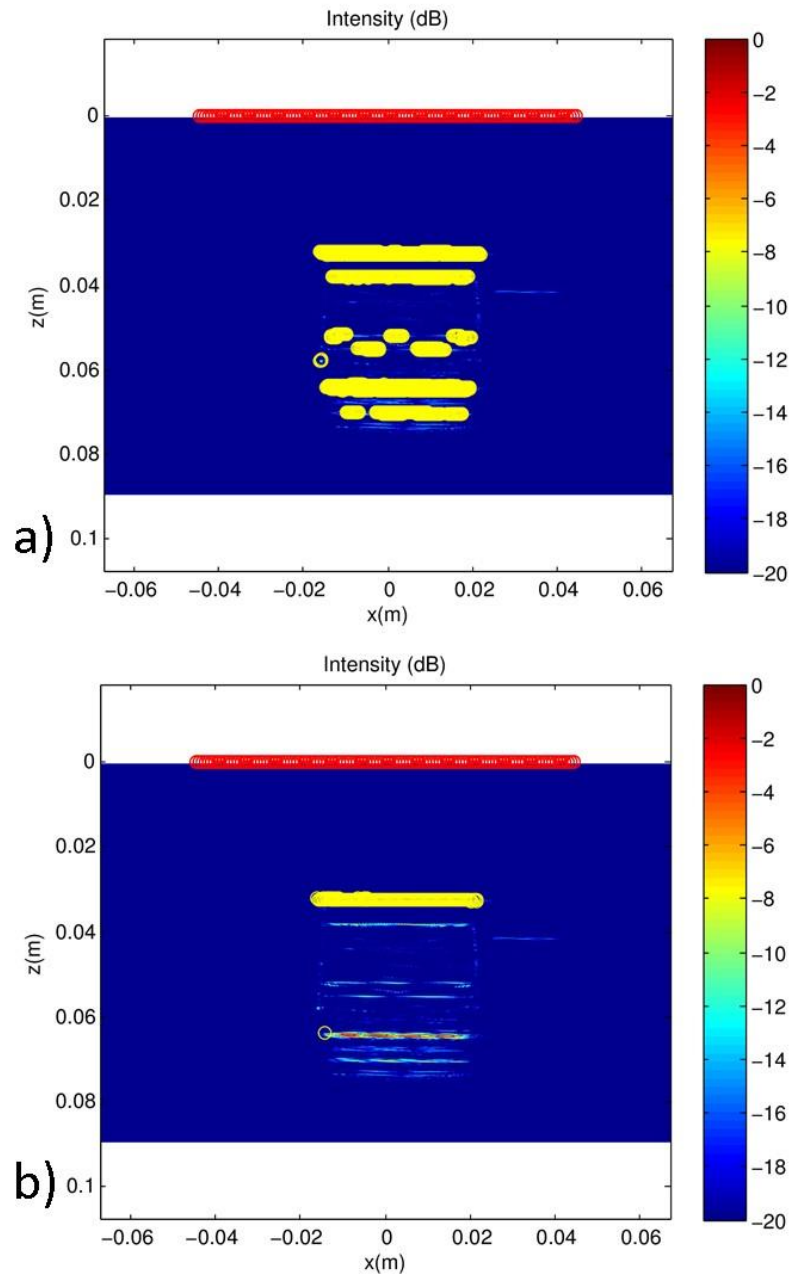


Figure 4.7: The result of coordinate extraction using only global thresholding, with extracted points represented as yellow circles (a). Also shown is the final contour extraction method (b), where all the unwanted reflections have been eliminated from the extraction.

The surface contours for each image were then saved in the form of 2D points, with coordinates relative to the probe face. In order to position these coordinates in 3D

space, the coordinates, firstly, were assigned a dummy Y coordinate of 0. These coordinates, now in a 2D plane in 3D space were rotated using the KUKA-derived orientations describing the rotations of the probe at the moment the FMC data was captured. These rotations were labelled A, B and C, denoting orientations around the Z, Y and X axes of the World coordinate system, respectively. To rotate each pixel coordinate to the correct 3D position, three rotation matrices were applied. These are displayed in equations 4.2, 4.3 and 4.4, describing the rotations around the X, Y and Z axes, respectively. It was then necessary to translate each of the newly rotated coordinates, so that they were positioned correctly relative to the probe position during capture. This was achieved by simply adding the X, Y and Z coordinates of the associated TCP. This process is summarised in equation 4.5, where the subscripts f , o and T represent the final 3D coordinates, the original coordinates and the TCP coordinates needed for translation, respectively. The result of this process is a 3D point cloud, as can be seen in Figure 4.8.

$$4.2 \quad R_x = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(C) & -\sin(C) \\ 0 & \sin(C) & \cos(C) \end{pmatrix}$$

$$4.3 \quad R_y = \begin{pmatrix} \cos(B) & 0 & \sin(B) \\ 0 & 1 & 0 \\ -\sin(B) & 0 & \cos(B) \end{pmatrix}$$

$$4.4 \quad R_z = \begin{pmatrix} \cos(A) & -\sin(A) & 0 \\ \sin(A) & \cos(A) & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

$$4.5 \quad \begin{pmatrix} X_f \\ Y_f \\ Z_f \end{pmatrix} = R_x R_y R_z \begin{pmatrix} X_o \\ Y_o \\ Z_o \end{pmatrix} + \begin{pmatrix} X_T \\ Y_T \\ Z_T \end{pmatrix}$$




Figure 4.8: the resultant point cloud from the scan, of which the image in Figure 4.6 was a part. The coordinates are relative to the KUKA World axes.

It should also be mentioned that a number of pixels from each image were eliminated from the extraction process completely. This was achieved by setting the pixel intensity values to zero before the image was set to a decibel scale, rendering them undetectable through global thresholding. The first region to be set to zero was that directly in front of the probe face. As can be seen in Figure 4.9, in an image in which no strong reflections are present, there is a thin but intense group of pixels at the top of the image, which is caused by ‘ring down’. In biomedical ultrasound imaging, this refers to an artefact caused by resonances in air bubbles (Hindi et al., 2013). However, in transducer design and in this case, it refers to an effect found in piezoelectric transducers, in which the material continues to vibrate after the electrical pulse has been

applied to it. This is presented in A-scans as a small collection of peaks and troughs of large magnitude, found within the first few samples. While this effect can be reduced by a number of methods, including adhering an appropriate backing material to the transducer (Ramadas, Leary, & Gachagan, 2009), it cannot be eliminated completely. Therefore, eliminating this artefact from each image by setting the first 20 rows of pixels to zero ensured that it is not confused with a true reflection. The other region of pixels set to zero were those with Z coordinates greater than 8cm. The reason for this was that the actual imaging depth was 8cm, but CueART returns images with depths greater than that recorded in the A-scans. While this would likely have caused little problem, as it was assumed any value beyond the maximum sample in the A-scans was set to 0, the pixels were eliminated for the sake of accuracy.




Figure 4.9: A TFM image produced from a data capture when the probe was operating into water, with no reflective materials present. Shown in detail is the result of ring down.

Finally, to provide a reference model for each of the samples a Faro Quantum laser line probe was employed. A plugin for Geomagic Studio (Geomagic Inc., Morrisville, NC,

USA) allowed for communication with the Faro arm. The measurements from laser line probe were converted into points in the form of 3D coordinates, which were then converted into .stl files by Geomagic. Surface capture using the laser line probe provided CAD models of the surfaces, which allowed for simple comparison. The reference models can be seen in Figure 4.10. The comparison procedure was carried out in CloudCompare (EDF/Telecom ParisTech, Paris, France) – an open source software designed for comparing 3D point clouds. The comparison process began by matching the ultrasound-derived point cloud to the reference mesh. This was achieved in two steps: firstly the point cloud was manually manoeuvred such that an approximate match was achieved. The second step employed ICP to finely match the point cloud to the mesh. The comparison procedure in CloudCompare involves finding the shortest Euclidean distance between the point in question and the surface of the mesh. Absolute errors were employed, making no distinction between errors above and below the target surfaces. The reason these were employed was that the main measure of error was mean error. If signed errors were employed, this would likely results in a lower mean error, which would not be representative of the level of error.




Figure 4.10: The Faro-derived reference models of the sphere (a), the cylinder (b) and the cuboid (c), each of which shows the support onto which they were each screwed.

The matching step described above was necessary due to the different coordinate systems employed by the robot and the CMM. In this situation, it would have been advantageous to employ absolute error measurements, which would have eliminated the dependence on the matching algorithm and the error of 0.1mm associated with the Faro arm (Rajna, Herold, & Baylard, 2009), which was employed to generate the reference models. This could be achieved by accurately positioning the samples, such that the true location of the sample was known relative to each packet of KUKA positional data. However, when imaging bone in the ultimate application, no reference models would be available. Therefore, in an attempt to emulate the conditions for which the system was intended, the samples were positioned manually.

4.2 Results

With the different geometries of the samples being capable of eliciting very different responses both at a fundamental wave physics level and in data processing, this section has been separated into three parts. The results from the cylinder, sphere and cuboid occupy each of these parts, with a description of both the resulting 2D images and the subsequent 3D point clouds.

4.2.1 Cuboid

The difference in results between TFM and SAFT images during the cuboidal path inspection of the cuboid was pronounced. The same FMC data has been used to create the TFM image found in Figure 4.11 as was to produce the SAFT image in Figure 4.6. While both images clearly show the initial surface reflection, numerous backwall reflections and the reflections of a return journey, the TFM image displays far higher levels of noise. This was caused generally by side lobe formation. The practical result of this was that when using the same contour extraction technique on both images, the SAFT image yielded generally accurate surface profiles, with a threshold of -12dB.

The TFM image, on the other hand, returned almost no representation of the initial surface. Even with a reduction of the threshold to -10dB, the vast majority of the points extracted were from side lobes, rather than the surface itself, as can be seen in part (b) of Figure 4.11. When the threshold was reduced further, parts of the initial surface were missed by the extraction algorithm in favour of the intense reflections caused by the return journey of the waves. It was, therefore, almost impossible to choose a threshold limit which would yield an accurate representation of the surface.

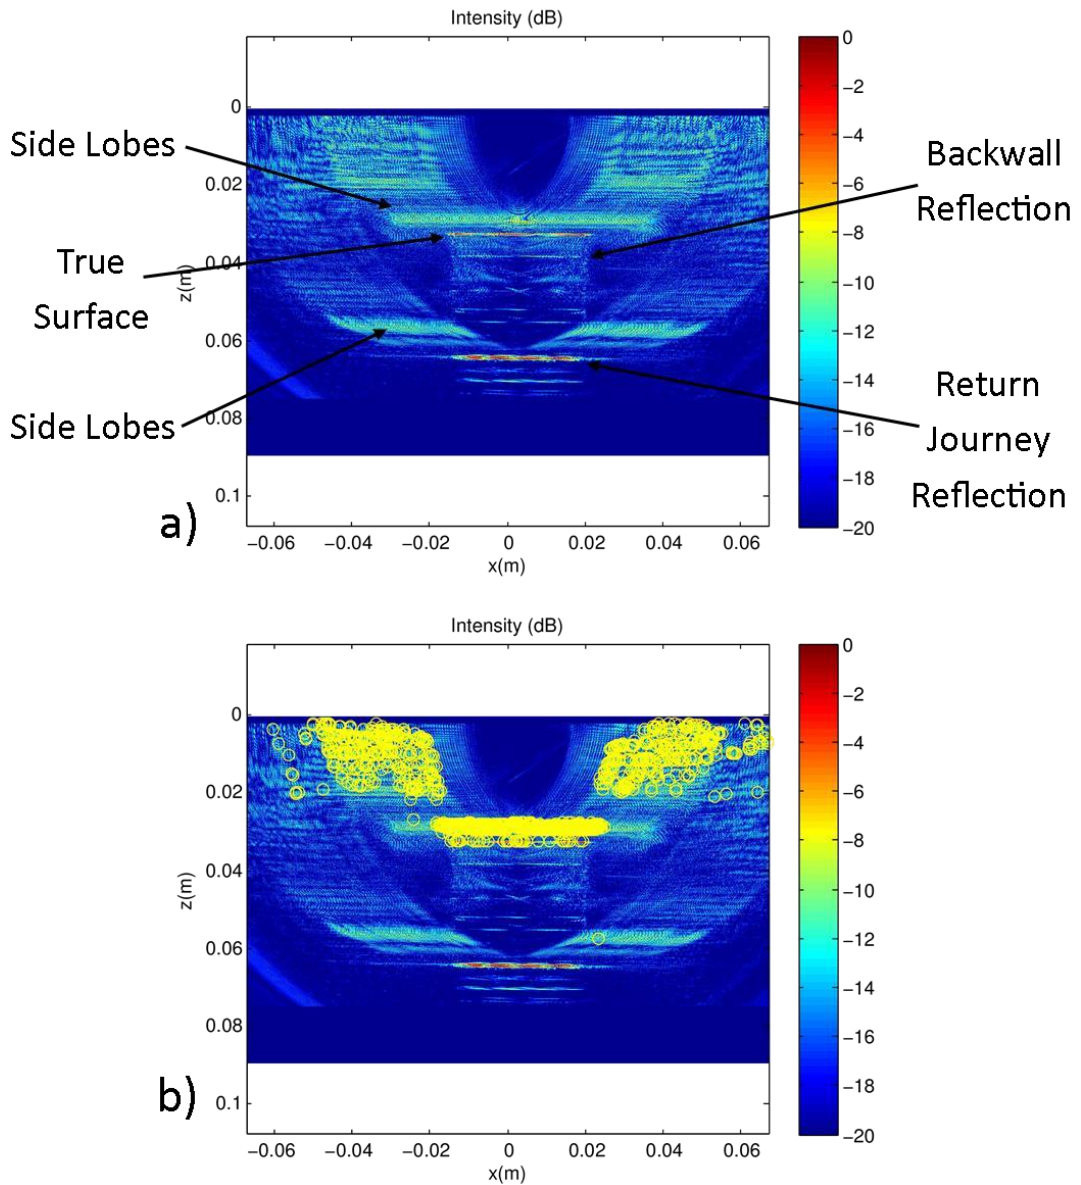


Figure 4.11: A typical TFM image from the cuboidal path inspection of the cuboid test piece (a). Also displayed is the contour extraction for the image (b), where unwanted regions have been extracted.

The point cloud derived from the SAFT images can be seen in Figure 4.12. There is an obvious representation of the four outer faces of the cuboid. However, each face extends beyond where it should. Additionally, instead of being cylindrical, the representation of the probe mount has been squared off, being more cuboidal in shape. Finally, there are a relatively small, but still significant number of outliers spread across

the point cloud. The resulting point cloud from the TFM images can be viewed in Figure 4.13. These results show absolutely no resemblance to the test piece, with the side lobe contribution in each image dominating.




Figure 4.12: The resultant point cloud from the cuboidal path inspection of the cuboid test piece using SAFT. A side on view (a) and a top down view (b) are displayed. While the cuboid shape is obvious, the sides have extended past the edge of the sample.




Figure 4.13: The resultant point cloud from the cuboidal path inspection of the cuboid test piece using TFM. A side on view (a) and a top down view (b) are displayed. The side lobes have dominated the surface extraction, resulting in a noisy surface representation.

As both point clouds suffered from serious inaccuracies, neither was matched to the associated reference model for comparison. In the case of the TFM-derived point cloud, the reason for this is obvious, as there was no clear representation of the cuboid. While the SAFT-derived point cloud displayed an apparent cuboidal shape, comparison was still not possible. Before a direct comparison can take place, the reference model and the point cloud had to be matched, which is achieved using ICP. However, this method matched every point in the cloud, meaning that the translations and rotations were based on the reduction of error for not only the points representing the cuboid, but erroneous points also. As such, no accurate matching was possible and, as a consequence, no comparison could be completed.

In an effort to improve the results of the TFM-derived point cloud, the threshold limit was decreased from -10dB to -5dB. It was hoped that doing this would alleviate the negative impact of the side lobes on the resultant point clouds. The results of this endeavour can be seen in Figure 4.14. As with the SAFT-derived point cloud, the representations of each face extend beyond the physical location. In addition to this, the easing of the threshold limit allowed the erroneous reflections to dominate in certain images, meaning two cuboids appear, rather than one.




Figure 4.14: The resultant point cloud from the cuboidal path inspection of the cuboid test piece using TFM, with a -5dB threshold limit. Showing a side on view (a) and a top down view (b), it can be seen that multiple surface representations exist.

The images achieved during the hemispherical scan were of a similar standard as those found with the cuboidal path. However, there were fewer contributions from inner reflections, as can be seen in Figure 4.15. As with the previous results, the TFM images showed considerable levels of side lobes. The SAFT images showed a reduction in noise, but lacked the resolution of the TFM images. This is exemplified in Figure 4.15 (b), in which the bottom right hand corner of the surface profile has thickened significantly. In contrast, this has not happened to the same extent in the TFM image, as shown in Figure 4.15 (a).

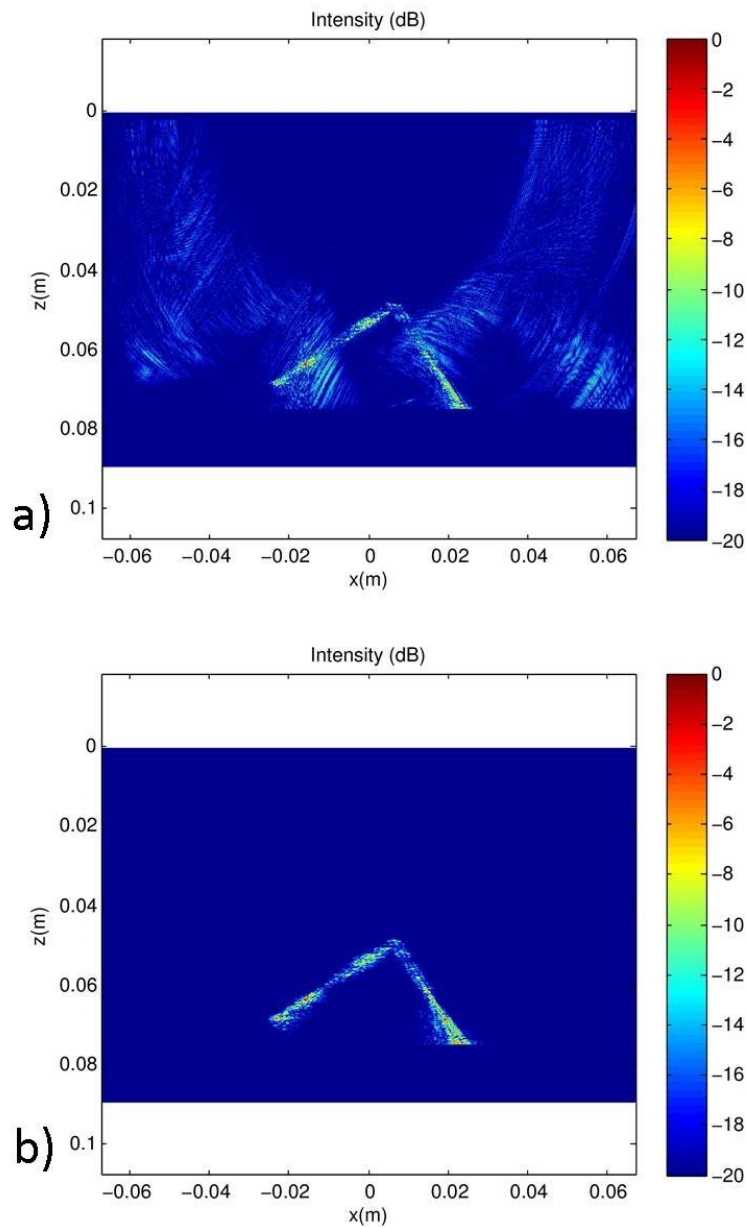


Figure 4.15: Typical TFM (a) and SAFT (b) images from the hemispherical scan of the cuboidal sample. Sidelobes are present in the TFM image, while none can be seen in the SAFT image.

As with the TFM-derived point cloud found using the cuboidal path, that resulting from the hemispherical probe path was unrecognisable, as is evidenced in Figure 4.16 (a). The dominance of the side lobes eclipsed the true signals, creating an unclear image of the sample. However, when the threshold level was, again, switched from -10dB to -

5dB, the results shown in Figure 4.17 were achieved. A large increase in clarity was achieved, with the cuboid and mount obviously visible. However, large outliers were present, causing the maximum error of 72mm. The mean error was 2.06mm, while the standard deviation was 4.65mm. As well as all this, the points representing the faces of the side of the sample were very sparse.




Figure 4.16: The TFM (a) and SAFT (b) derived point clouds, resulting from the hemispherical scan of the cuboidal sample. Side lobes have resulted in the noisy TFM-based point cloud in (a), with the poor surface representation in (b) resulting from the inaccuracy associated with the robot path.




Figure 4.17: The result of the full hemispherical scan of the cuboid (processed with TFM) with a threshold of -5dB rather than -10dB (a), resulting in a cleaner point cloud. Also shown is the comparison with the reference model. The threshold in the colour scale has been set at 5mm to make comparison of the most accurate points more simple.

The SAFT point cloud, while more clear than the original TFM point cloud (derived from images with a -10dB threshold), yielded very inaccurate results, as can be seen in Figure 4.16 (b). By splitting the assessment of this error into point clouds representing individual parts of the full hemispherical scan, the problems manifested in Figure 4.16 (b) can more clearly be examined. Figure 4.18 (a) shows the part of the scan associated with the green path in Figure 3.10, while part (b) shows the results of the path in yellow in the same figure. Both of these point clouds appear to show some level of depiction of the cuboidal shape. However, as well as the expected common level of noise, there seems to be a number of very inaccurately placed points from particular images. This effect was far less pronounced in the paths shown in blue and red in Figure 3.10. The combined point clouds of these two paths can be seen in Figure 4.18 (c), where the cuboid and the mount are quite visible. This point cloud was matched using ICP and compared to the reference model, returning a mean error of 1.5mm, a maximum error of 27.4mm and a standard deviation of 2.04mm. The largest errors are in a number of lines of points which appear to have had the associated coordinates flipped around one or more axes.

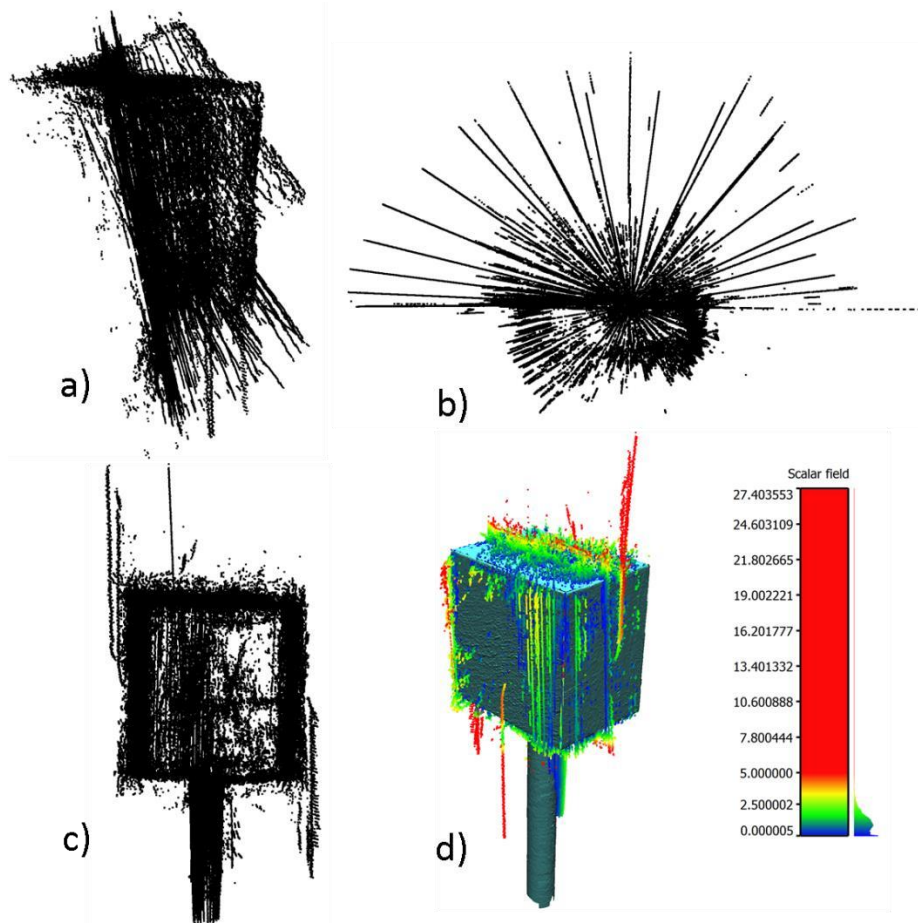


Figure 4.18: Results of SAFT processing from individual parts of the hemispherical scan that show positional errors (a), (b) associated with specific arts of the probe path. Also shown are two parts of the scan combined which show more accurate point placement (c) and that combined cloud when compared to the reference model (d).

4.2.2 Sphere

Typical TFM and SAFT images from the cuboidal scan of the spherical sample can be seen in Figure 4.19. Significant side lobes were present in the TFM images, increasing the general level of noise and causing substantial obstruction of the surface profile. There was a particularly strong reflection at the part of the sphere closest to the probe

face, with particularly strong side lobes surrounding it. Far less noise and side lobes were present in the SAFT images, as can be seen in Figure 4.19 (b). However, the profile of the sphere and holder are of low resolution, with a relatively thick line representing the surface. Additionally, there were a number of reflections resulting from multiple reflections from the backwall. However, these had little effect on the profile extraction, as strong reflections directly above were chosen in favour of them. In addition to this, when the probe was positioned such that the image slice was far from the centre of the sphere, images such as that illustrated in Figure 4.20 were produced. In these cases, the surface profile was smeared and the arced profile was not continuous, resulting in a poor surface profile extraction.

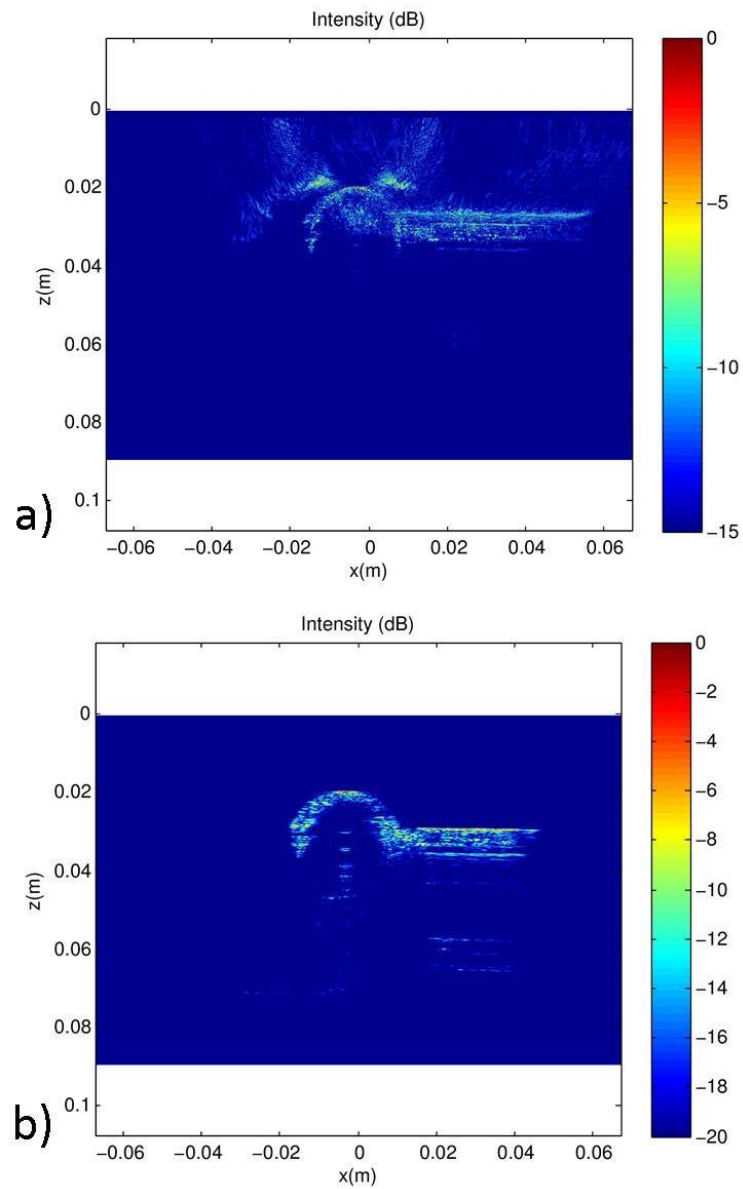


Figure 4.19: Typical TFM (a) and SAFT (b) images from the cuboidal scan of the spherical sample. While side lobes are present in (a), a thicker surface representation has been reconstructed in (b).




Figure 4.20: An example of a poor surface profile from SAFT processing found with the sphere sample, showing a smearing effect with the surface.

As can be seen in Figure 4.21 (a), the point cloud resulting from the TFM images showed little resemblance to the spherical sample. The significant side lobe contribution became dominant, serving to completely obstruct any representation of the sphere. The SAFT-derived point cloud showed a much higher level of likeness to the sample. However, as before, outliers were present and, more importantly, the main body of the representation of the sphere was highly inaccurate. This inaccuracy is displayed in more detail in Figure 4.22, which shows point clouds from individual passes. It would be expected that the contour locations would change according to the curvature of the sphere. This, however, did not happen, with little or no change in curvature along the scan direction.




Figure 4.21: The TFM (a) and SAFT (b) derived point clouds, resulting from the cuboidal scan of the spherical sample.

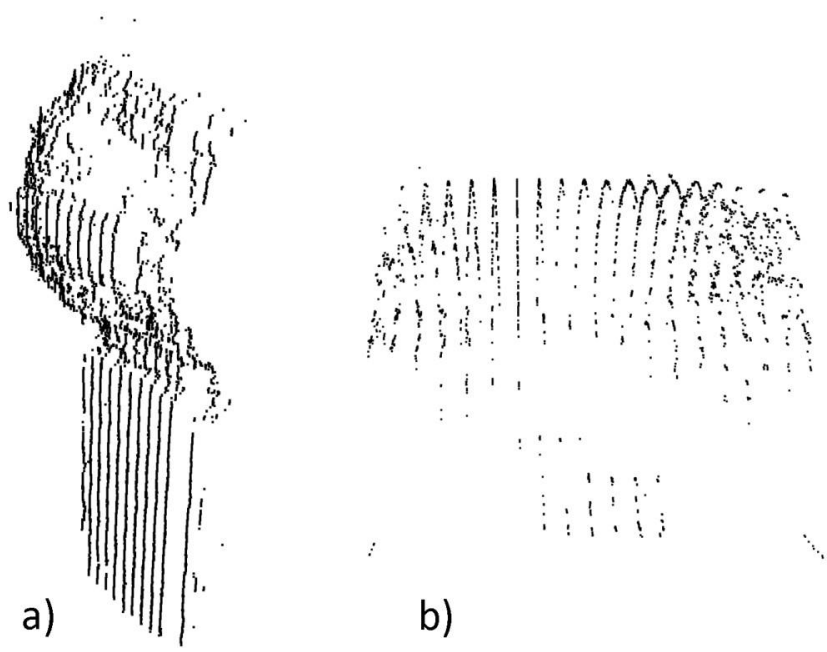


Figure 4.22: Resultant SAFT point clouds from a single pass along one side of the sample (a) and a top side pass (b), both showing a lack of curvature expected from the arc of the sphere.

While no accurate matching was possible using ICP, manual matching was possible. This was difficult, given the complexity of the point clouds shown in Figure 4.21. However, using only the topside passes over the sphere, the comparison shown in Figure 4.23 was possible. This yielded a mean error of 1.58mm, a maximum error of 25.21mm and a standard deviation of 1.84mm. The largest errors appeared in the form of outliers and the most accurate points were found at the top of the sphere, where the normal to the face of the sphere would be perpendicular to the probe face. Additionally, there are, as with the cuboidal sample, extensions past where would be expected, from which the large standard deviation was derived.




Figure 4.23: The comparison between the topside scan of the sphere processed with SAFT and the reference model, matched manually. As the probe moves from the centre of the sphere, the error increases.

When the same process was followed using higher threshold levels, the results shown in Figure 4.24 were achieved. The SAFT images were processed using a -6dB threshold, while that of the TFM images was -5dB. The TFM-derived point cloud

showed a large reduction in side lobe contribution, with a reasonably accurate representation of the top of the sphere. The mean error achieved was 1.44mm, the maximum was 22.53 and the standard deviation was 1.35mm. The SAFT-derived point cloud also showed improvement, yielding a mean error of 1.06mm, a maximum of 4.73mm and a standard deviation of 0.87mm. However, the problems displayed in the previous topside scan were also present in the SAFT-derived point cloud using a higher threshold, but to a lesser degree. This can be seen in Figure 4.24 (a), where the points have extended past where the sphere physically ends.

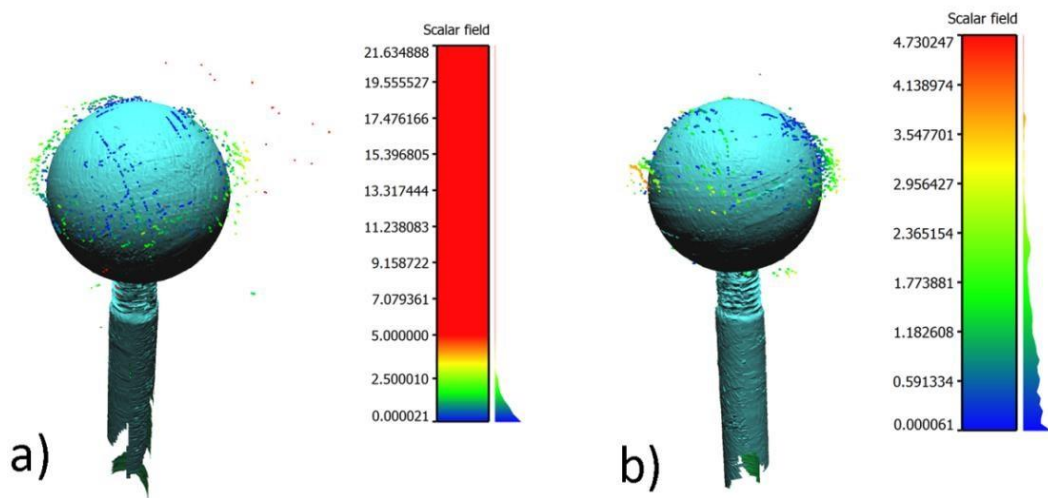


Figure 4.24: The comparison between both the topside scan of the sphere processed with SAFT, with a -6dB threshold (a) and that with TFM, with a -5dB threshold (b) and the reference model. Less side lobe contribution is obvious in (b), due to a lower threshold level.

The point clouds resulting from the hemispherical scans can be seen in Figure 4.25. The TFM-derived point cloud showed an obvious spherical representation, but, in addition, displayed high levels of side lobes. Therefore, the levels of inaccuracy were very high. The SAFT point cloud, on the other hand, yielded an obvious representation of a sphere, albeit with a lack of clarity and a significant number of outliers. Due to this improvement over the cuboidal scan path data, it was possible to match the SAFT-derived point cloud using ICP. The result of this can be seen in Figure 4.26, where the

highest error levels can be seen when the probe passed one of the axes. The mean error was 1.08mm, while the maximum error was 39.44mm and the standard deviation was found to be 1.08mm. Reducing the threshold limit to -5dB on the TFM data produced the results shown in Figure 4.27. The representation of the sphere was very similar to that found with the original threshold, but there was a massive reduction in the side lobe contribution. In return for this, however, the spherical representation lost density. The point cloud was compared to the reference mesh, yielding a mean error of 1.44mm, a maximum of 22.53 and a standard deviation of 1.35mm.




Figure 4.25: The TFM (a) and SAFT (b) derived point clouds, resulting from the hemispherical scan of the spherical sample.




Figure 4.26: The comparison between the full hemispherical scan of the sphere, processed with SAFT and the reference model, matched using ICP. The top view of the point cloud is detailed, showing the high levels of error found at axis crossings.




Figure 4.27: The resulting point cloud from the full hemispherical scan of the sphere (processed with TFM) (a) and the comparison with the reference model, matched using ICP.

4.2.3 Cylinder

Comparing the TFM and SAFT point clouds created with the data associated with the cylindrical sample, it was clear that The SAFT-derived point cloud was superior. As can be seen in Figure 4.28 (a), the TFM point cloud bore no resemblance to the shape of the cylinder, with side lobes obstructing any true data. The SAFT point cloud was much clearer in appearance, with an apparent cylindrical shape being present. However, different parts of the scan seem to have been positioned incorrectly. For example, in Figure 4.28 (b), the points representing the flat face of the sample appear to be too far forward, residing within the rest of the cloud. When compared to the reference model, a mean error of 2.06mm, maximum error of 24.18 and standard deviation of 2.11mm were found. As can be seen in Figure 4.29 (a), the point cloud continues on past where it should stop on the flat face of the cylinder, as was seen in previous samples. Additionally, the lowest errors appear to be found on the parts of the piece where the surface was parallel to the face of the probe. This can be seen in the visible curved face in Figure 4.29 (a), where the errors are low close to the flat face. As the measurements are made further along, where the parallelism between the probe face and the sample surface diminishes, the errors increase. When the TFM images were processed using a -5dB threshold, the point cloud became more clear. However, as can be seen in Figure 4.29 (b), the side lobes were still made large contributions to the point cloud. This resulted in a mean error of 3.37mm, a maximum error of 50.76mm and a standard deviation of 4.56mm.




Figure 4.28: The TFM (a) and SAFT (b) derived point clouds, resulting from the cuboidal scan of the cylindrical sample. Part (a) shows high levels of side lobe extraction, while (b) displays a cleaner representation, but with inaccurate placement.

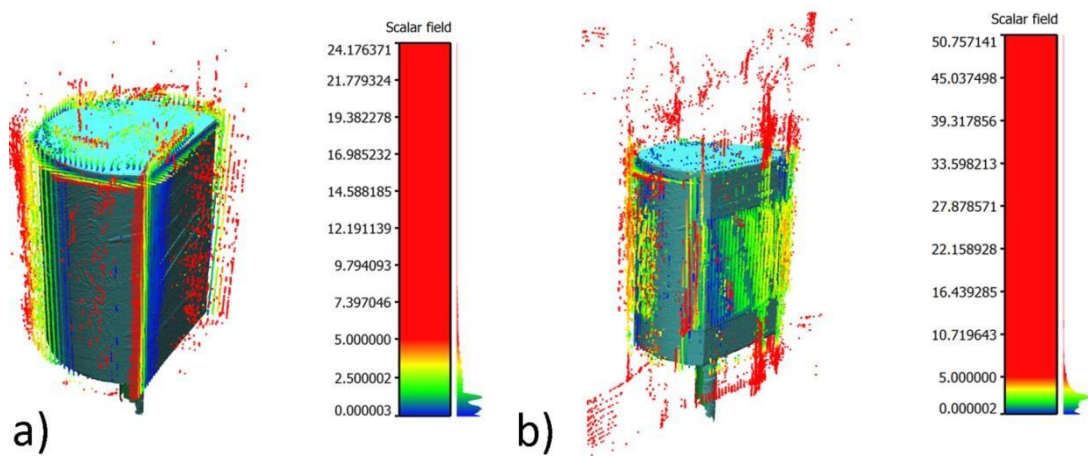


Figure 4.29: The comparison between both cuboidal scan of the cylinder processed with SAFT, with a -12dB threshold (a) and that with TFM, with a -5dB threshold (b) and the reference model. Side extension is present in part (a), while (b) shows a high side lobe contribution.

The results of the hemispherical scan of the cylindrical sample can be viewed in Figure 4.30. As expected, the TFM-derived sample bore little resemblance to the geometry of the sample, with high side lobe contribution. The point cloud resultant from the SAFT images had a clear similarity to the test piece. However, as can be seen in Figure 4.30 (b), there were particularly significant inaccuracies in the form of lines of points that appeared to have been positioned incorrectly. These large inaccuracies rendered it difficult to properly compare the point cloud to the reference data. The point cloud resulting from a -5dB threshold limit in the TFM data set can be seen in Figure 4.31. The sample surface profile was much more apparent than was found with the -12dB threshold. However, the profile still bore little resemblance to the physical surface shape of the sample.

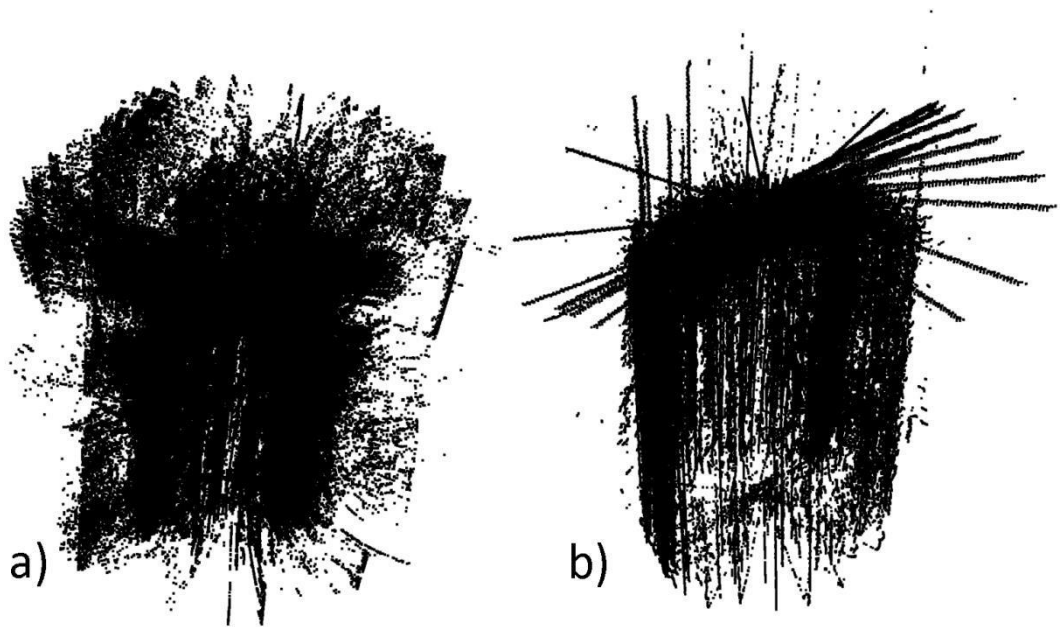


Figure 4.30: The TFM (a) and SAFT (b) derived point clouds, resulting from the hemispherical scan of the cylindrical sample. Side lobes have resulted in the noisy TFM-based point cloud in (a), with the poor surface representation in (b) resulting from the inaccuracy associated with the robot path.



Figure 4.31: The TFM-derived point cloud resulting from the hemispherical scan. The threshold limit employed was -5dB. The increase in threshold has resulted in lower side lobe extraction, but has also resulted in true surface extraction also.

The results presented in Figure 4.32 are from pass 1 of the additional scan paths undertaken on the cylindrical sample, a depiction of which can be seen in Figure 4.4 (a). The TFM-derived point cloud was very accurate, as evidenced in Table 4.1. As can be seen in Figure 4.32 (a), there were small areas of sparseness at the centre of the point cloud and in the regions between the centre and the edges. The SAFT point cloud was both sparse and inaccurate, as can be seen in Figure 4.32 (b). A number of inner reflections were present, as were side lobes, making the surface virtually impossible to distinguish.

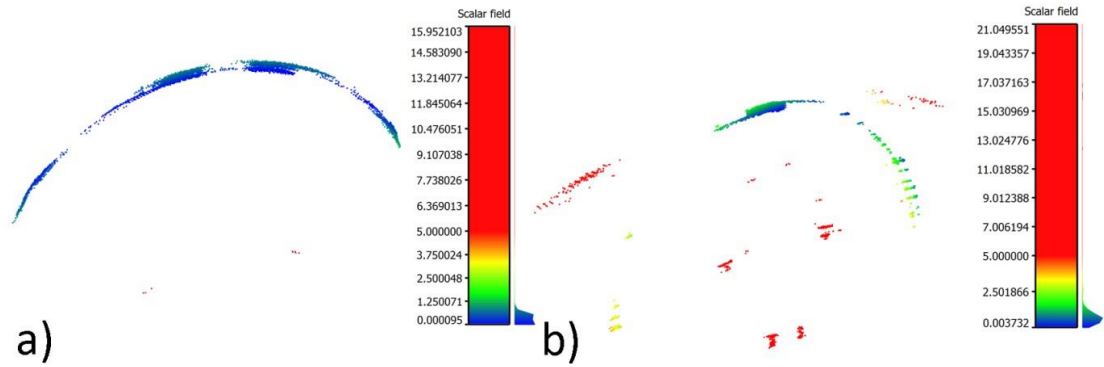


Figure 4.32: The comparison between both the TFM (a) and SAFT (b) derived point clouds and the reference model for Pass 1. TFM reconstruction resulted in an accurate surface representation, while SAFT displayed side lobe extraction and backwall reflections.

Table 4.1: The error results from the comparison between the TFM-derived point clouds and the reference model.

Path	Mean Error (mm)	Maximum Error (mm)	Standard Deviation (mm)
Pass 1	0.5	16.0	0.8
Pass 2	0.2	3.7	0.4
Pass 3	1.8	20.3	2.5

Table 4.2: The error results from the comparison between the SAFT-derived point clouds and the reference model.

Path	Mean Error (mm)	Maximum Error (mm)	Standard Deviation (mm)
Pass 1	2.8	21.1	4.9
Pass 2	2.5	17.2	4.8
Pass 3	0.8	21.3	0.8

Presented in Figure 4.33 are the results for the second pass of the additional paths for the cylindrical sample, which relate to part (b) of Figure 4.4. The TFM results provided the highest accuracy of any of the scans in the chapter, as is evidenced in Table 4.1.

The largest errors occurred at the edges of the point cloud. The SAFT results showed a large reduction in accuracy relative to the TFM error levels. Only a small band of accurate points was achieved across the centre of the sample and there was a large contribution of points from inner reflections. Combined, this led to relatively poor accuracy.

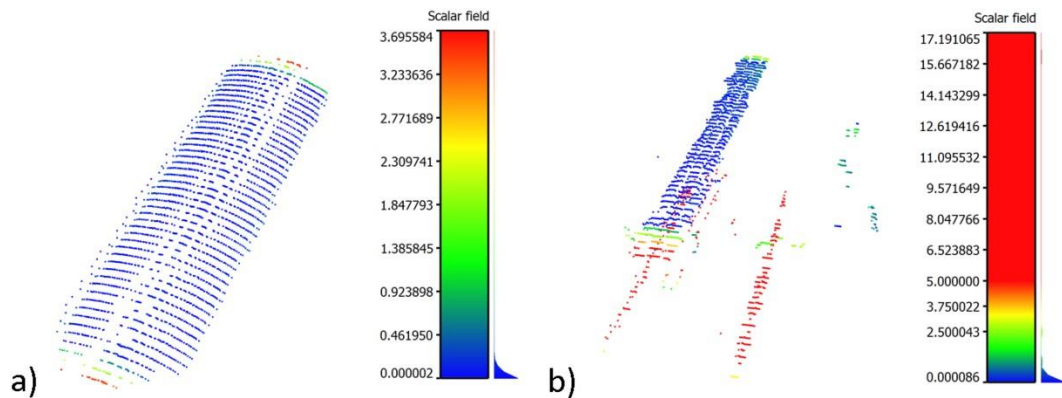


Figure 4.33: The comparison between both the TFM (a) and SAFT (b) derived point clouds and the reference model for Pass 2. TFM produced an accurate surface representation, while SAFT resulted in backwall extraction.

The third pass of the cylindrical sample provided a change in results, with the SAFT point cloud achieving a higher level of accuracy than the TFM. The TFM-derived point cloud suffered from erroneous reflections and was generally inaccurate, as can be seen in Figure 4.34. However, there were few contributions from side lobes. The SAFT point cloud fared better, with a good level of accuracy across the centre of the sample. However, the patch in the centre seemed to lack the correct level of curvature, appearing flat. Moving away from the centre, the level of error increased.

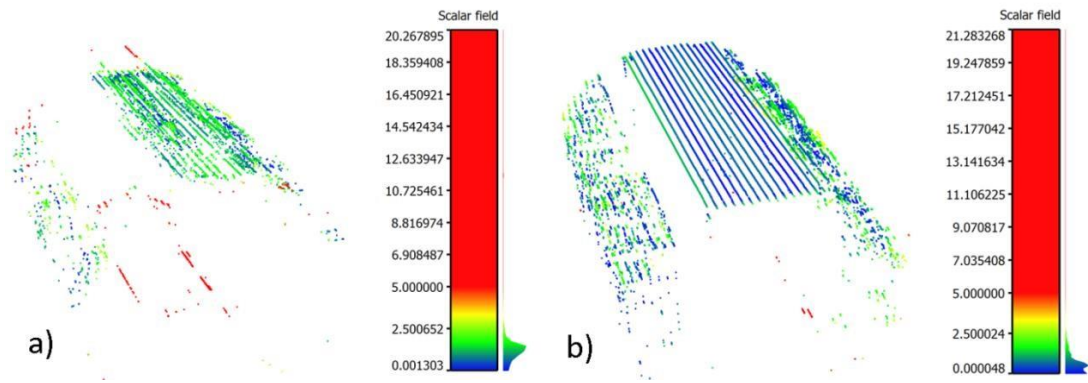


Figure 4.34: The comparison between both the TFM (a) and SAFT (b) derived point clouds and the reference model for Pass 3. Both point clouds show a flattening effect towards the centre of the arc.

The errors for all the samples and different parameter combinations are presented in Table 4.3. In this table, the point clouds which could not be accurately matched using ICP were simply matched manually. As can be seen, the mean error was reduced for every parameter combination when SAFT was compared to the equivalent TFM-derived point cloud (with the exception of two cases). However, on certain curved surfaces, TFM was found to outperform SAFT, with SAFT becoming susceptible to inner reflections and side lobes.

Table 4.3: Summary of the comparisons for the different samples and parameters.

Sample	Path	Processing Method	Threshold Limit (dB)	Mean Error (mm)	Maximum Error (mm)	Standard Deviation (mm)
Sphere	Cuboidal	TFM	5	3.0	25.2	2.6
Sphere	Cuboidal	SAFT	12	2.9	42.9	4.4
Sphere	Hemispherical	TFM	5	1.4	22.0	1.3
Sphere	Hemispherical	SAFT	12	1.1	39.3	1.1
Cuboid	Cuboidal	TFM	5	3.7	44.1	3.3
Cuboid	Cuboidal	SAFT	12	1.3	32.1	1.8
Cuboid	Hemispherical	TFM	5	2.1	72.0	4.6
Cuboid	Hemispherical	SAFT	12	2.5	41.0	3.1
Cylinder	Cuboidal	TFM	5	3.1	52.2	4.5
Cylinder	Cuboidal	SAFT	12	2.0	26.5	2.1
Cylinder	Hemispherical	TFM	5	1.9	65.1	3.6
Cylinder	Hemispherical	SAFT	12	2.2	32.3	3.0

4.3 Discussion

While a number of the inspection scenarios presented above achieved sub-millimetre mean accuracy, many weaknesses in the system were apparent. While some of these were due to the design of particular components of the system, many were due to the inherent nature of ultrasound physics and the data processing techniques employed.

One such problem was that found in the cuboidal scan of the cuboidal test piece. As can be seen in Figure 4.12, the points representing each of the faces extended beyond the edges of the test piece. The cause of this was not positional inaccuracy on the part of the robot controller or in the point cloud construction procedure, but rather the physical nature of the ultrasound beam itself. As can be seen in Figure 4.35, the beam shape of an individual element of the array has a thickness associated with it, not only in the width along the direction of the full aperture, but also associated with the height of each element.

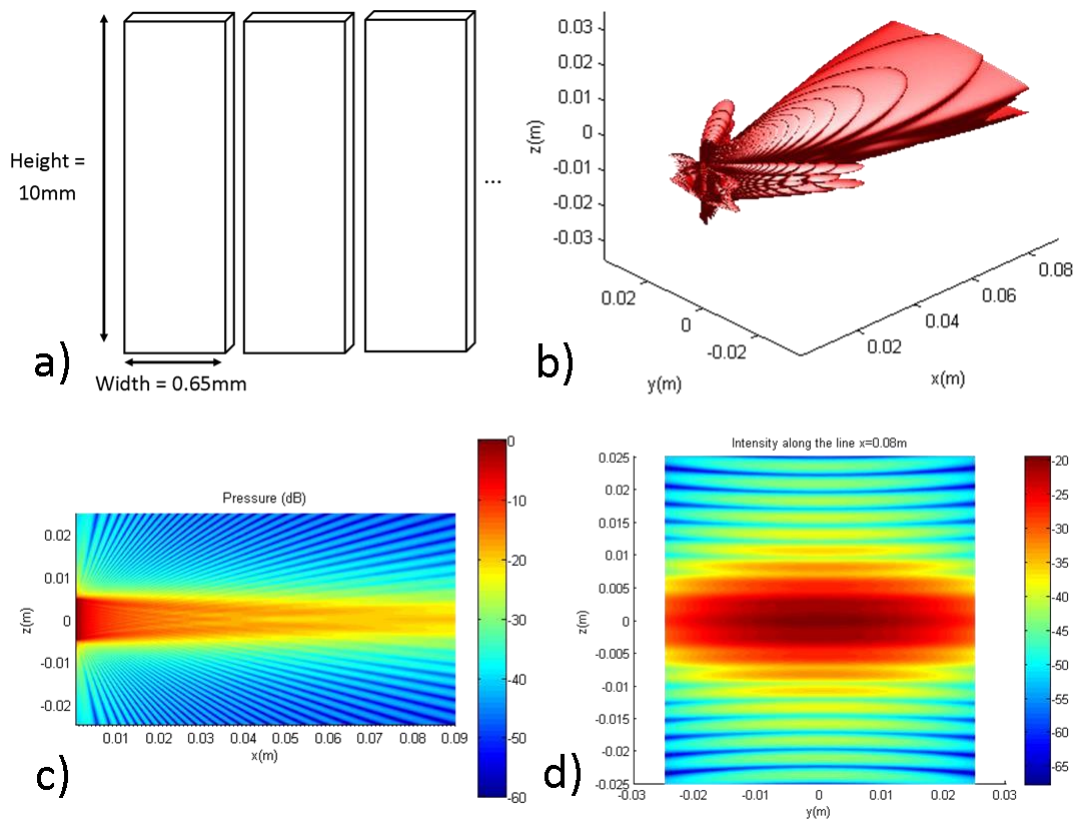


Figure 4.35: The dimensions of the individual elements from the 5MHz linear array (a), the 3D beam shape from a single element in water (b), a two dimensional cut of the beam intensity at $y = 0$ and a further two dimensional cut of the beam intensity at $x = 0.08$ (d). The X dimension relates to the distance from the probe face, while the Z dimension relates to the height of the individual elements, meaning the Y dimension is where the full probe aperture is located.

Considering, then, a single element in transmission over a planar surface, the scenario shown in Figure 4.36 is found. Due to the relatively flat, polished nature of the surfaces under inspection, specular reflection dominates. As such, the centre part of the beam is reflected perpendicularly back to the transducer. The outer parts of the beam, if thought of as rays, encounter the plane at an angle, reflecting away from the transducer, as illustrated in Figure 4.36. The result of this is that the signals registered by the transducer are mainly from the reflections located in the image plane, which is shown in blue in Figure 4.36. This effect is compounded by the fact that the most intense part

of the beam is located at the centre, as is shown in Figure 4.35 (d). Therefore, any small diffuse reflections off the image axis would be overpowered by the on-axis reflections.




Figure 4.36: A single ultrasonic transducer (grey) emits a wide beam shape, the outer parts of which do not return to the transducer, due to specular reflection on a planar surface. The centre of the beam is reflected back to the transducer. The beam is depicted in the height (ie. elevation) dimension of the transducer and the blue represents where the resultant image plane would be located.

When the transducer passed the edge of the flat surface, the imaging plane did not come into contact with the reflective surface. However, an off-axis part of the main lobe would still have come into contact with the edge of the flat surface, as shown in Figure 4.37 (a). The elevation of the transducers was only 10mm, implying that the situation depicted in part (a) would only be valid for 5mm beyond the edge. However, in SAFT-derived point clouds, points were recorded in excess of 9mm beyond the edge of the sample. This occurred because on the relatively polished, planar surface, specular reflection would have dominated. At the sharp edge, on the other hand, diffuse reflection would begin to dominate. The rays representing the main lobe would not,

then, simply be reflected away from the transducer, but would instead scatter in all directions, as illustrated in Figure 4.37 (b). The result of this scenario would be that reflections from the edge would be recorded as being in the image axis. The extracted points from each such image would, as a result, be inaccurately placed.




Figure 4.37: The effect of diffuse reflection from a sharpened corner, with rays travelling back from outside the image plane due to the height of the element (a) and the diffuse reflections caused by the edge of the sample.

This was responsible for producing points that did not represent the surface of the samples. This effect can be seen in the cuboidal scan of the cuboid, where points have continued to be produced past the edge of the sample. This was present in both the SAFT and TFM derived point clouds, as can be seen in Figure 4.12 and Figure 4.14. Similar results were found with the cuboidal scan of the cylinder. The effect is

particularly obvious in Figure 4.29 (a), where the points representing the flat surface extend beyond from the expected positions. Extensions can also be seen from the top, flat surface. In both described parts, the rest of the points from the specular surface were found to be reasonably accurate. However, when the probe moved past the sharp edges, the described effects caused significant inaccuracy. Finally, the effect can be seen in pass 2 of the additional paths for the cylindrical sample. Unlike the other examples, the probe did not move from the inspection of a flat surface to a sharp edge, but rather from a curved surface to a sharp edge. The effects still occurred, as can be seen in Figure 4.33, where the only noticeable inaccurate points were found in the regions extending beyond the edges.

It should, however, be noted that such scenarios would be less likely when imaging the bony surfaces of the knee, where sharp edges and highly specular surfaces are unlikely. Reflections dominated by diffuse reflection might occur in osteoarthritic patients due to the increased risk of osteophytes (Caine & Golightly, 2011), but these formations would tend to be small. As such, the risk of the effects described above would be significantly less than found in the presented results.

Another of the problems caused by the nature of ultrasound itself can be seen in Figure 4.20, where instead of a clear surface, the profile is smeared. This can be explained by considering the scenario illustrated in Figure 4.38, where a single element is located above a curved surface. The surface is relatively polished, implying that specular reflection would be dominant. Due to the position of the element, the rays emanating from the centre of the transducer would be reflected away from the transceiver. However, off-axis rays from the outer parts of the main lobe would interact with the transducer. The result of this would be that reflections from a part of the curve outside the image plane would be registered. As mentioned before, all reflection is a combination of specular and diffuse reflection. Therefore, while specular reflection dominates, rays would arrive from a number of locations on the curve. This accounts for the smeared appearance of the profile shown in Figure 4.20.




Figure 4.38: The effect of specular reflection from a curved surface, with rays travelling back from outside the image plane due to the angle of reflection from the surface.

This effect can be seen, firstly, in the cuboidal scan of the spherical sample, with the most obvious problem illustrated in Figure 4.22. When near the centre of the sphere, the effects described above are lessened, as the rays from the transducer come close to being normal to the surface of the sphere. Moving away from the centre, the steepness of the curve increases relative to the transducer, producing the effect described above. The result of this was a smeared surface profile, with the part closest to the probe face being from towards the centre of the sphere, while the part furthest away from the probe face related to the part of the sphere nearest the image plane. As the surface profile extraction algorithm sought to find the point closest to the probe face, the points in the final point cloud mirrored the surface closer to the centre of the sphere. As can be seen in Figure 4.22, the surface profile at the centre appears to extend beyond where it should, as would be expected, given the above explanation. The effects of this could be greatly reduced by altering the extraction algorithm, such that the points furthest from the probe face are chosen, rather than those closest to it. This effect might also be reduced by decreasing the threshold limit, which would make the size of the smear smaller. This would result in extracted points being further from the probe face and, therefore, closer to the correct position.

The same effect was present in the cuboidal path inspection of the cylindrical sample. As can be seen in Figure 4.29 (a), as the probe moved to regions where the curve became steeper, the point cloud representation became less accurate. This effect was more clearly demonstrated in pass 3 of the additional cylindrical inspections. At the centre of the sample, the point cloud is accurate. As the probe moved away from the centre, the accuracy remained relatively high, but it is clear from Figure 4.34 (b), instead of following the curve, the point cloud representation remained flat. This mirrored the effects seen in the sphere inspection. As the probe moved further from the centre and the steepness of the curve increased, the location of the extracted points became more erratic and less accurate.

The surfaces of the bones of the knee are considered long and smooth relative to the wavelength of the ultrasound in use, implying that specular reflection would dominate (V. Chan & Perlas, 2011). While this is most likely the case, all reflection is, as stated earlier, a combination of specular and diffuse reflection. The surfaces inspected in this section were, relative to the surface of bone, highly polished. The proportion of specular reflection would, therefore, be greater than expected with the surface of bone. As such, the problems encountered here due to the dominance of specular reflections would not be expected to be so significant when imaging bony surfaces. Additionally, the surfaces that cause these effects would simply not be encountered when imaging bony surfaces, as they are rudimentary forms of surface types, which would amplify the effects described thus far.

Another of the causes of inaccuracy was reflection from within the sample. While the profile extraction algorithm sought to minimise this, it did not always achieve this goal. As can be seen in Figure 4.33 (b), while the correct points were captured in some areas, the surface could not be identified in others and inner reflections dominated. There are two reasons for this, the first of which is a weakness in the profile extraction algorithm. As the algorithm only discards any points with the same X coordinate, if there is any small break in the surface profile, the reflections from the backwall or from return reflections could be registered.

In addition to this, the effect can be caused by the nature of the data processing algorithms. Referencing Figure 4.33, it can be seen that the TFM point cloud does not

suffer from the effect of inner reflection as much as that derived from the SAFT images. Considering a pixel located at the centre of a surface of the curved section of a shape similar the cylindrical sample, it can be seen in Figure 4.39 that the specular reflection is in the centre of the receiving aperture. As such, the reflection is registered by the aperture. Considering point 2, on the other hand, it can be seen that the angle at which the rays are reflected is such that the receiving aperture does not register any reflections. Finally, point 3 has the same effect as can be seen with point 1, with the only difference being that the reflection is from the backwall, rather than the outer surface. It should be noted that refractive effects have been discounted in this example to simplify the explanation. The result of this is that highly curved surfaces would likely not be registered in SAFT and but flat surfaces directly below could be registered instead. The likelihood of this would be greatly reduced in TFM processing, as the receiving aperture is not limited as it is in the example below. There would be, therefore, a higher chance of the full aperture registering the signals from curved surfaces.




Figure 4.39: Reconstruction scenarios from a number of pixels on a shape similar to the cylindrical sample. The blue elements represent the associated receiving apertures, the arrows show the main direction of the specular reflection, the green points denote the pixel positions and the triangles show how the receiving aperture is determined.

The effects of signals generated from reflections from within the structures under inspection would be greatly reduced if bone were being imaged. The reason for this is that the bone of the knee joint has two material constituents: cortical and trabecular bone. While the dense cortical tissue is thick on the long parts of the bones, such as the shaft of the femur, this cortical layer is thin at the knee joint. Therefore, when the ultrasound waves interact with the thin cortical layer of the bony surface, the vast majority of the energy will be lost at the water/cortical layer and cortical /trabecular bone interfaces. This was not the case with the test pieces employed herein, as the second interface encountered by the waves would have been the backwall. In addition to this, the reflecting surfaces associated with the test samples were often flat and, therefore, reflected well. This would not be the case when imaging the knee, as surfaces would be, by comparison, complex.

A further cause of the inaccuracies observed in the results was positional error from the robot controller. This can be observed in Figure 4.26, where the level of inaccuracy increased significantly when the probe was crossing the axis. The effects are even more noticeable in Figure 4.30 (b), where lines of points representing the flat face of the cylindrical sample appear to have been inverted. This not only created a level of inaccuracy, but the effects were so significant that it brought the entire shape of the point cloud into question. The most likely and obvious cause of this was inaccuracy in the calibration procedure. While the returned TCP position might be accurate, the associated orientations could be incorrect. The likelihood of this is high, seeing as the variability of pose was limited during the calibration procedure.

Generally speaking, the coverage offered by the TFM images surpassed that of the SAFT images. This can be seen in Figure 4.33, where the correct surface profile of the TFM point cloud was, simply put, longer than that of the SAFT point cloud. This has been further depicted in Figure 4.40 where the result of the curved part of the cylindrical sample being imaged in pass 1 is shown. As can be seen in part (a), the cylinder is located just beyond the edge of the probe. Despite this, when processed with TFM, the shape of the cylinder was clearly visible. The SAFT image, on the other hand, in no way resembled the shape of the cylinder. Instead side lobes and noise due to electrical interference from power sources dominated. The reason for this effect is

the same as described for specular reflection on curved surfaces, with the small apertures unable to register the signals from outside the full probe aperture.

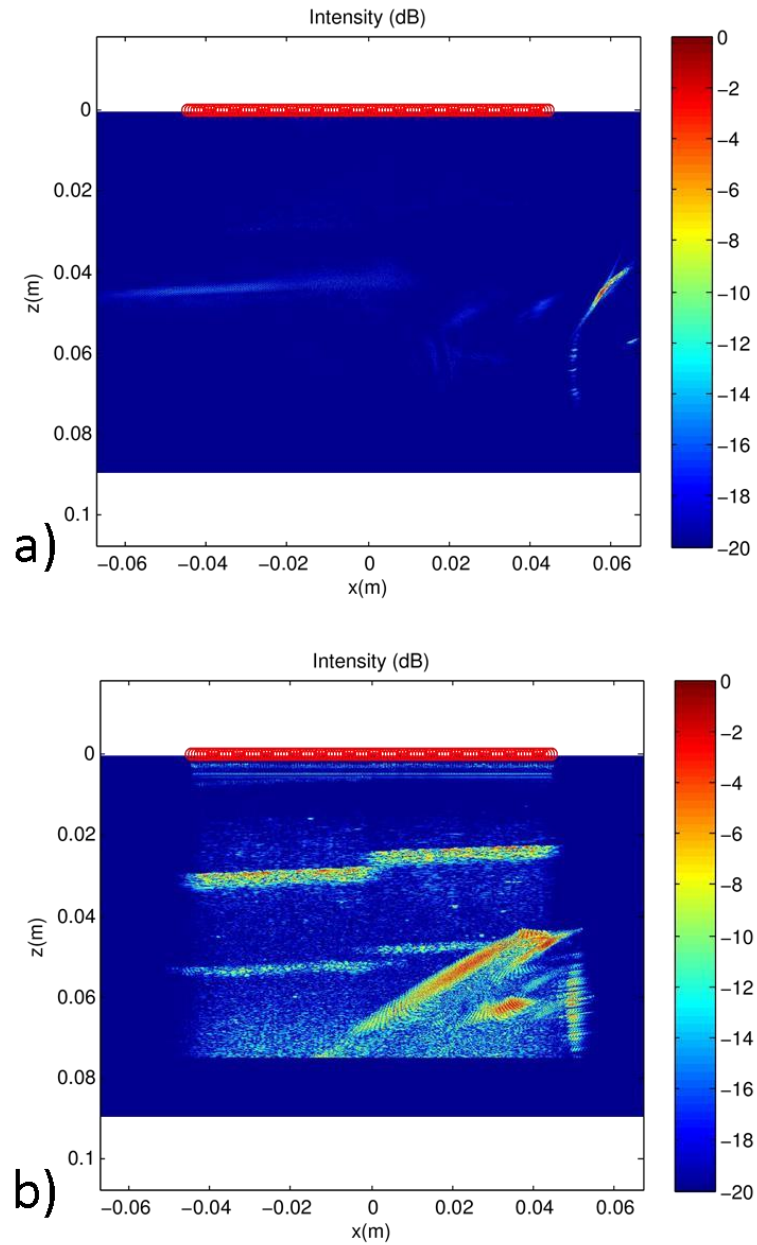


Figure 4.40: TFM (a) and SAFT (b) images of the cylindrical sample when located beyond the edge of the probe. TFM reconstructed the surface, while SAFT was incapable of doing so.

The final major problem encountered during the tests was the presence of significant levels of side lobes. These were most evident in the TFM images, an example of which can be seen in Figure 4.19. In fact, in every TFM point cloud created, the side lobes dominated, with the true representation of the samples only becoming clear when the threshold level was lowered considerably. The problem in doing that was that the coverage was effectively decreased, as can be seen in the comparison between the TFM and SAFT derived point clouds in Figure 4.17. This also had a further effect, in that parts of the initial surface could be lost with too low a threshold level, which sometimes gave way to backwall reflections. This effect is illustrated in Figure 4.14.

The reasons for the discrepancy in performance between TFM and SAFT in terms of side lobes can be explained by considering a simulated point reflector, reconstructed using TFM. As can be seen in Figure 4.41, the main point reflector is obvious at point (a,b). However, in addition to this, there are very obvious side lobes either side of the true reflector which appear in a manner very similar to that seen in Figure 4.19 (a). The reason that regions of intensity occur outside the true reflector is that the algorithm does not discriminate based on position. To explain further, for transmitting element i and receiving element j , the time of flight to pixel locations (a, b) – the position of the point reflector – and (c, d) is the same. In other words, with reference to equation 2.4, $S_{i,j}(T_{i(a,b)} + T_{j(a,b)}) = S_{i,j}(T_{i(c,d)} + T_{j(c,d)})$. As such, the contribution to the intensity of the pixel, for that particular transmit-receive pair will be the same at both locations.




Figure 4.41: TFM reconstruction of a simulated FMC of a point reflector. Element positions are displayed as red circles, while the paths of travel for two pixels at (a, b) and (c, d) are shown for the same transmission and reception. It is clear that the path to the centre of the side lobe is the same distance as the path to the point reflector.

This principle is repeated across the whole image, resulting in significant regions of non-zero intensity which contain no real reflector. Because reflections are not represented in A-scans as infinitely thin peaks, but instead have a width, the pixels which receive a non-zero contribution need not possess the exact time of travel as the position of the true reflector. Instead, the time of travel must only be within a range close to the true reflector. This is illustrated in Figure 4.42, which shows the exact time of travel of the point reflector, D , along with the region in which the contribution will be non-zero, P . This leads to numerous non-zero contributions throughout the image, confined by the geometrical combinations of pixel position, transmit-receive pairings and reflector position. There are particular regions where more of the transmit-receive pairs and pixel positions meet the criteria for non-zero contribution. This leads to regions where the side lobes have a particularly high intensity, as can be seen in Figure 4.41.




Figure 4.42: An A-scan from the point reflector FMC, showing the time of travel for the point reflector, D , and the range in which pixel contributions will be non-zero, P .

While this example employed a simple point reflector for simplicity, the principle presented is true for any reflector. As such, the flat surface shown in Figure 4.11 can be thought of as a line of strong reflections, each creating side lobes. These have merged to form a thick line above the true surface. Thinking of the true surface as a line of reflectors, the side lobe effects seen and explained in Figure 4.41 would be extended. In particular, the combination of the regions in which the side lobes are of particularly high intensity can, in fact, lead to ‘ghost’ surfaces, where the entire surface can be mirrored either side of the true surface. An example of this can be seen in Figure 4.43 (a), where the low resolution mirrored regions of intensity are marked by red lines. This effect is amplified for strong surface reflections, as is exemplified in part (b) of Figure 4.43.

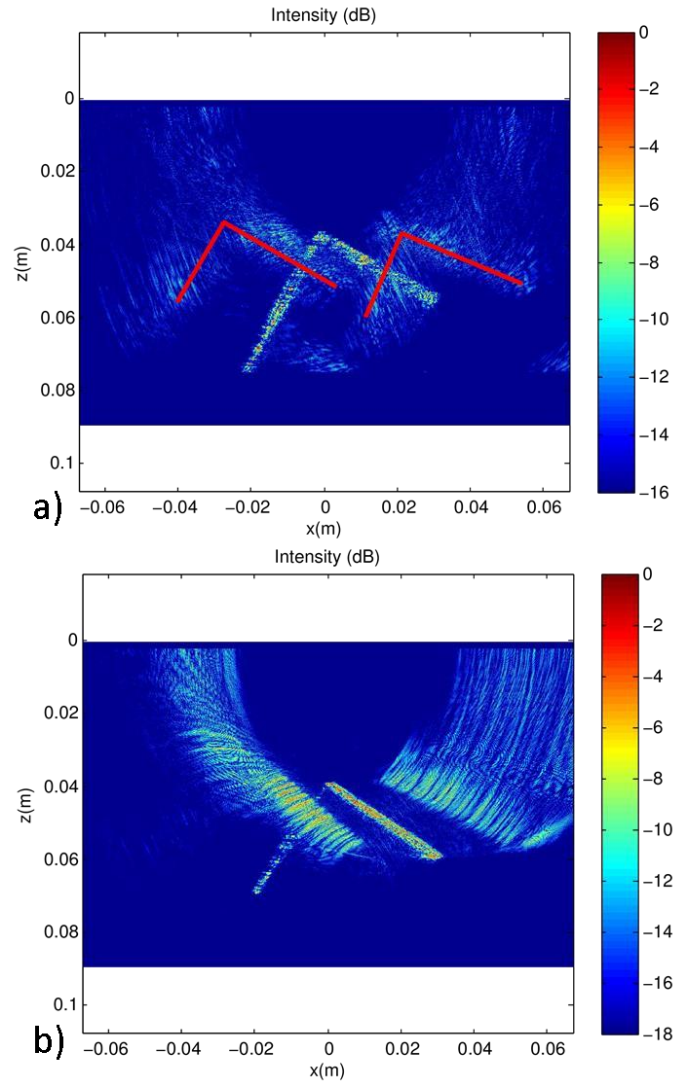


Figure 4.43: Examples of ‘ghost’ surfaces where the whole surface has been mirrored either side of the true surface (a) and one particular strong surface has been mirrored, while the weaker has not (b).

This effect is significantly lessened in the SAFT images, as the number of possible transmit-receive pairs is limited by the number of elements considered in reception is significantly less than that in TFM. This geometrically restricts the regions in which side lobes can be formed. Given that a narrow sub aperture has been employed for reception in this study, the likelihood of side lobes is low. If, however, the sub aperture size was increased, the possibility of side lobes would increase.

4.4 Conclusion

A number of simple, polished, geometrically-shaped test pieces were imaged using the system, providing 3D point clouds representative of the samples themselves. The differences in performance between TFM and SAFT were exposed, with the conclusion being that SAFT provides images with fewer artefacts and, therefore fewer errors. The fact is further supported by the errors displayed in Table 4.3. In this, the point clouds which could not be accurately matched using ICP were simply matched manually. As can be seen, the mean error was reduced for every parameter combination when SAFT was compared to the equivalent TFM-derived point cloud (with the exception of two cases). However, on certain curved surfaces, TFM was found to outperform SAFT, with SAFT becoming susceptible to inner reflections and side lobes. Both data processing techniques, then, have strengths in different scenarios. However, the result of the presented work was that SAFT would be more likely to allow for accurate reconstruction of complex, bony surfaces.

Both curved surfaces and edges caused errors, which were amplified due to the rudimentary nature of the shapes employed. However, such effects would be considerably lessened when imaging bony surfaces, due to the increase in the complexity of the shapes and the decrease in the level of surface specularities. Given this, the likelihood of the system being capable of accurately reconstructing complex, bony surfaces was high. As such, the next stage of investigation was to image shapes and surfaces representative of the bony surfaces of the knee joint.

Chapter 5 Composite Human Distal Femur Surface Reconstruction

5.1 Introduction

The reflective response of ultrasonic waves varies considerably between different surface materials and geometries, affecting the quality and accuracy of resulting images. The errors displayed in the surface point clouds found in the previous chapter were a result of the type of reflection associated with the shapes and material of the test pieces. As such, imaging bone would illicit a different response and, therefore, a different quality of image.

While the ultimate goal is to image an intact human knee joint *in situ*, it is prudent to provide a proof of concept at a lower level of complexity. In particular, it is important that the system be capable of accurately reconstructing a 3D surface with characteristics similar to that found in human knees. This chapter, therefore, describes the process and results of imaging an epoxy fibre model, based on the CT data of a human distal femur. First, the general setup of the experiment is explained, which leads on to a description of the data processing methods. Surface mesh construction is then described and the results are compared to reference models.

5.2 Methods and Materials

While it would have been ideal to use a real human distal femur, there were a number of restrictions preventing this. Firstly and most importantly, the use of human tissue is tightly regulated, with use only possible in accredited buildings with a licensed anatomist. While these are available in the Department of Biomedical Engineering, the KUKA robot used in the experiments is not mobile and could not have been moved to a suitable building. As such, use of human tissue in the FIRST lab would be illegal. Additionally, as the sample would be fully submerged, a real bone sample might become susceptible to degradation through not only prolonged submersion, but through continual drying and soaking that would occur. Therefore, a fibrestrengthened epoxy human distal femur phatom was employed instead. This was a Sawbones® composite bone model, intended for mechanical testing (Pacific Research Laboratories, Inc., Vashon Island, WA, USA). This sample simulated the mechanical,

but not the acoustic properties of bone. Despite this, the shape of the sample was typical of human distal femurs, offering a good substitute for real bone.

The sample was acquired from Daniel Gilroy (EngD candidate, Department of Biomedical Engineering, University of Strathclyde) as a leftover from mechanical testing. As such, the sample had a drilled hole through the centre for mounting purposes, as can be seen in Figure 5.1(a). Additionally, small holes were present on each side. A cylindrical base was fashioned from aluminium, with a rod coming from the centre of the base and up through the drilled hole. This was manufactured by Alec Ward (technician, Department of Electric and Electronic Engineering, University of Strathclyde) and served to stabilise the sample when submersed.




Figure 5.1: The composite human distal femur model, showing drilled hole through the centre and metallic base for stability (a). Four screws have secured the base to the table to eliminate movement and the sample has been spray painted to account for laser absorption (b).

In order to provide an accurate reference model for the sample, it was laser scanned using a Faro Quantum (Faro Technologies, FL, USA), in the same way as described in

Section 4.1. To achieve this, the base was attached to an optical table using four screws, maintaining rigidity during surface measurement. Additionally, a thin layer of spray paint was applied to the surface of the sample. This was necessary due to the optical characteristics of the epoxy, which absorbed much of the laser light. As such, it was difficult to achieve any kind of reading, particularly given the requirement of large amounts of diffuse surface reflections. The resulting reference mesh can be seen in Figure 5.2 (a).

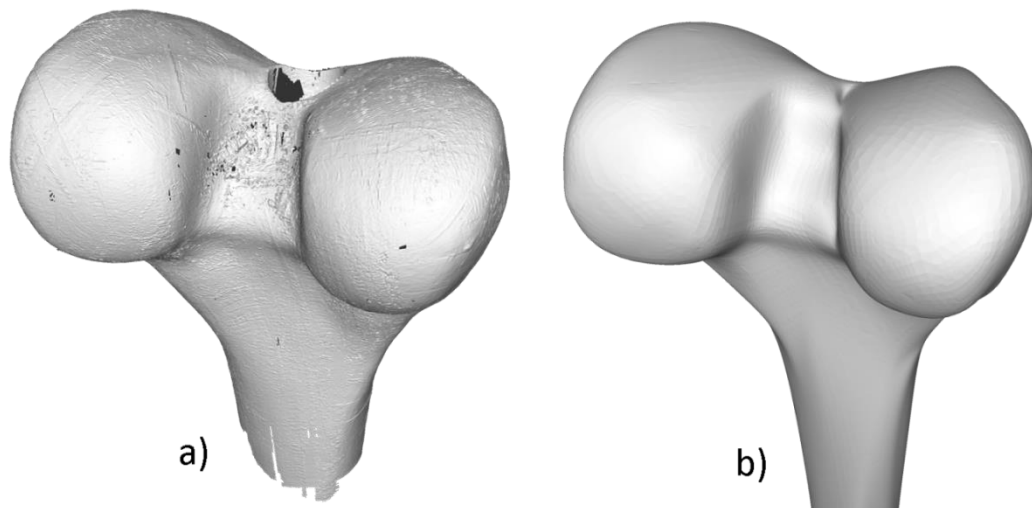


Figure 5.2: Reference model captured with Faro laser scanner (a) and the representative mesh supplied by Sawbones, with no drilled hole (b).

In addition to the laser scanned model, Sawbones provide a CAD model of the entire femur. The sample had two physical layers, the outer part of which was made from fibre strengthened epoxy, serving to simulate the cortical layer. The second, inner part of the model was made of a foam substance, which replicated the cancellous part of the bone. This was recreated in the CAD model in the form of two surfaces. As only the outer surface was being imaged, the two CAD layers were separated in Blender (Blender Association, Amsterdam, Netherlands), leaving only the outer layer.

Two paths were employed for imaging, the first of which was the hemispherical path, an illustration of which can be seen in Figure 3.10. For this, the sample was manually positioned so that the centre of the hemisphere was located approximately in the centre of the distal femur. The radius of the hemisphere was defined as 70mm, which allowed

for a clearance between the probe face and the sample surface of at least a centimetre at all times during acquisition. The second path was cuboidal, as shown in Figure 5.3. As in Chapter 4, the probe was jogged manually, maintaining a clearance of at least one centimetre at all times. The path was split into six segments between which data collection was stopped, producing six data sets for one full scan. The segments were defined by simple straight lines, as is represented by different colours in Figure 5.3.

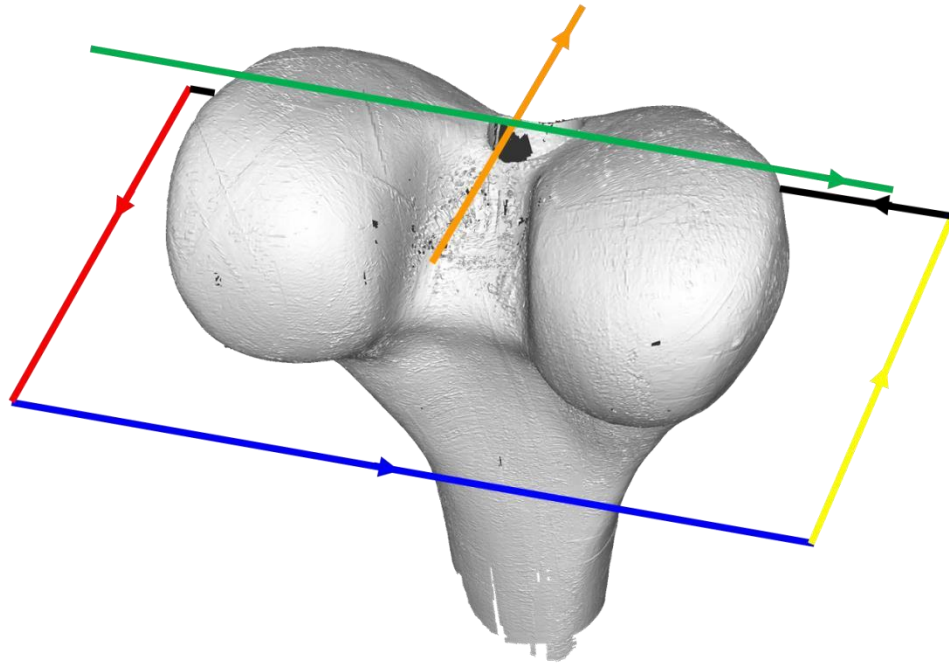


Figure 5.3: The second path used for the inspection of the composite human distal femur, with separate segments denoted by different colours.

The composite distal femur was placed in a water bath, such that it was fully submerged, as can be seen in Figure 5.4. While being manoeuvred around the composite bone using the paths described above, the face of the probe remained beneath the water at all times, so as to maintain the ultrasound coupling. For the hemispherical path, a 10% T1 velocity was employed. Due to the accelerations and decelerations associated with CIRC commands, the data captures would be bunched up at the beginning and end of the commands. In the middle of the arc, on the other hand, the data captures would be sparse. In an attempt to compensate for this, the

velocity was increased to 30% manually when at the start and end of CIRC commands and as the TCP was moving to succeeding latitudes. For the cuboidal path, a TCP velocity of 1% in T1 was employed. This resulted in approximately 1mm spacing between each data capture position. As the tool was jogged manually, no acceleration or decelerations were present, resulting in roughly equal spacing between captures. The distance between each capture was determined by the time between data captures, which varied slightly.



Figure 5.4: The physical setup of the experiment, showing the robot with probe attached, manoeuvring around the composite bone. The bone was completely submerged in tap water, as was the face of the probe.

The Labview script described in Section 4.1 was used to capture the FMC files, allowing for single precision storage to be employed. All 128 elements were utilised and an imaging depth of 8cm was defined. This resulted in a variable sample rate of approximately 0.3Hz. Each FMC file was stored alongside the corresponding KUKA-derived TCP position and orientation.

Four methods of data processing were carried out on each FMC file: TFM, SAFT with a 20° aperture, SAFT with a 30° aperture and SAFT with a 40° aperture. The resulting images for each processing method had the same dimensions, with a width of 1344 pixels and a depth of 896 pixels. The width and depth of each pixel was 0.1mm, which corresponded to a depth of 89.6mm and a width of 13.44mm – 1.5 times the length of the probe. This ensured that all signals within the imaging depth of 80mm were accounted for and that it was possible to image objects outside the width of the probe. All data processing was performed on a Dell Alienware X51 R2, with an Intel® Core™ i7-4790 processor and an NVIDIA GeForce GTX 760Ti (1344 CUDA cores), using the cueART platform to significantly speed up the TFM and SAFT algorithms.

In order to capture the surface contours of the sample, a method similar to that employed in Section 4.1 was employed. For each image, thresholding was applied with a threshold limit of -8dB for SAFT images and -4dB for TFM. These threshold limits were found through a method of trial and error. The X coordinates of the pixels that passed the thresholding procedure were compared. Any pixels which had the same X coordinate were discarded, with the exception of that with the smallest Z coordinate (ie. the pixel closest to the probe face). This was an attempt to counter the errors caused by inner and double reflections. This method, however, was responsible for jagged, unrealistic surface contour. This is exemplified in Figure 5.5, which shows a typical SAFT-derived image and the extracted surface, found using the described method. As can be seen in the detail, the contour is jagged – a feature caused by the fact that the image would not have a solid line representing the surface, but broken clusters of high intensity. Therefore, when the point closest to the probe face was calculated, the Z value could vary over a few tenths of a millimetre.

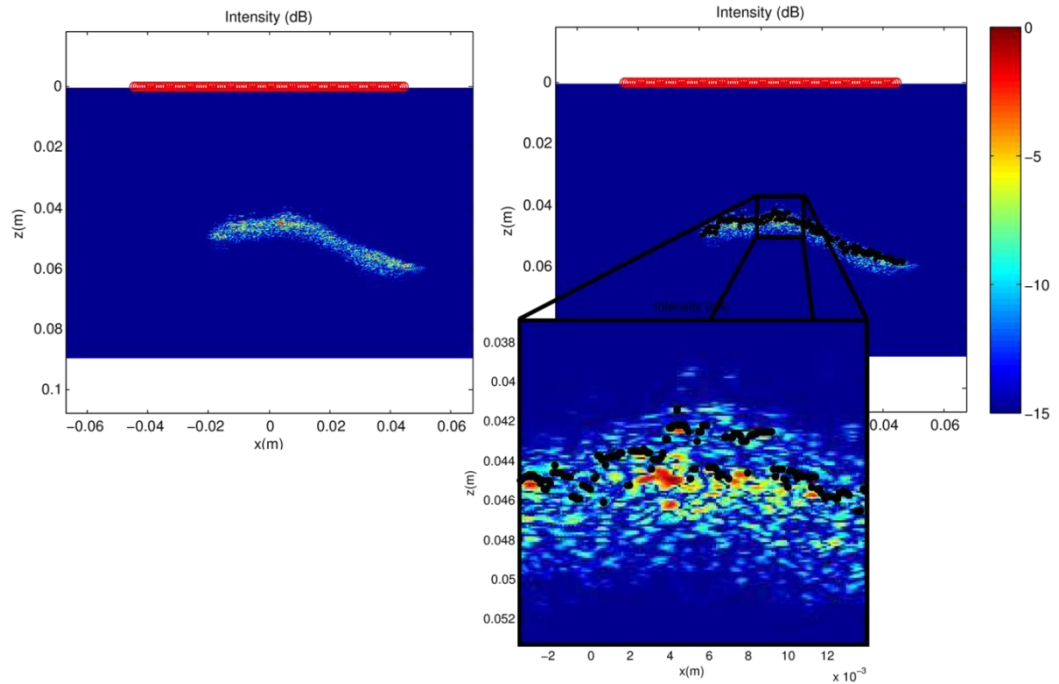


Figure 5.5: A typical 20° SAFT-derived image (left), with the extracted surface (right), with extracted points displayed as black dots. The detail of the extracted points shows the jagged nature of the extracted contour.

In order to provide the ability to create a smooth, continuous line, a smoothing function was applied to each of the 2D contours. The ‘smooth’ function in Matlab was employed for this purpose, using a moving average and a span of 0.3. The effects of this can be seen in Figure 5.6, with a smooth contour produced.




Figure 5.6: The extracted contour (in black) after the application of a smoothing function, eliminating the jagged appearance seen in Figure 5.5.

The 2D coordinates resulting from the above processes were translated and rotated using the associated KUKA TCP positional coordinates and orientation, as described in Section 4.1. In an attempt to rid the resultant surface point cloud of erroneous points, points from each segment of the cuboidal path were removed based on those from the other segments. For example, if points from one of the segments positioned at the side of the sample had a Z coordinate greater than the maximum Z value for the top side segments, those points would be eliminated. Equally, referring to Figure 5.3, if any points from the black segment had a Y coordinate less than the minimum Y coordinate of the yellow segment, the points would, again, be eliminated. This example is illustrated in Figure 5.7, which shows the eliminated points in red and the boundary as a transparent blue plane. The virtual boundary created by the purging of points below a minimum Y value provided by the yellow segment. In doing this, any points resulting from either side lobes or the reflective problems described in Section 4.3 which caused points to be created outside these bounding boxes could be discarded. This is shown in Figure 5.7, where all the eliminated points are towards the outer limits of the different axes. Such erroneous points could be caused by the effect seen in Section 4.2.1, where the imaging plane is no longer over a surface, but the 3D ultrasound beam still interacts with the surfaces and signals are returned.




Figure 5.7: The point cloud resulting from a cuboidal path scan, with no smoothing applied. The blue points represent the final point cloud after point elimination, while the red points represent those points that have been eliminated. The blue panel represents one of the boundaries past which points were eliminated.

In addition to this, a second point elimination method was introduced which employed a nearest neighbour algorithm to rid the cloud of sparse points. For every point in the cloud, the Euclidean distance between that point and every other point was calculated. In order to reduce the number of calculations and, therefore, speed up the algorithm, the distance calculations were limited to those points within 15mm in the Z direction. If a particular calculated distance was less than a given limit (5mm in this case), a counter for the point in question was increased by one. This was carried out for every point in the cloud, resulting in each point having a number associated with it, representing the number of points within a surrounding sphere of radius 5mm. The points were then reordered based on this value and the lowest 1% was discarded. An example of the results of this can be seen in Figure 5.8, where the eliminated points

are shown in red. As can be seen, the outliers towards the back of the point cloud have been purged, as have a few others. However, genuine points were also lost, as can be seen towards the front, centre of the image. Both the radius of influence and the proportion of elimination were defined by way of trial and error, in an attempt to rid the cloud of as many outliers as possible, while still maintaining valid points.




Figure 5.8: The result of the nearest neighbour point elimination algorithm on a cuboidal scan, processed using SAFT, with no smoothing applied. The final points are shown in blue, while the eliminated points are displayed in red. While outliers have been eliminated, so have genuine surface points.

Once this procedure had been performed for a number of images from different data processing methods, the resulting point clouds were imported into CloudCompare. Using the method described in Section 4.1, the mean error, maximum error and standard deviation were calculated for each point cloud, relative to the laser scanned

reference model. Only the laser scan model was employed as a ground truth, as it was the most accurate representation of the physical object.

Each point cloud was imported into Geomagic Wrap (Geomagic Inc., Morrisville, NC, USA), after which the mesh wrapping algorithm was applied. This is a proprietary algorithm owned by Geomagic Inc. (Fletcher, Gloth, Edelsbrumer, & Fu, 2006) that reconstructs surfaces from 3D point cloud data. However, the algorithm is sensitive to point density inconsistencies. As can be seen in Figure 5.8, not all of the cloud is equally dense, so sampling was employed to remedy this. Two sampling methods were employed, the first of which randomly eliminated points, such that a defined number of points were returned. Using the 'rand' function in Matlab, a pseudo-random value between 0 and 1 was assigned to each point in the cloud. The number of desired points was converted into a proportion of the total number of points. If the random number assigned to the point was below the desired proportion, it was retained and if it was greater, it was discarded. In this case approximately 10,000 points were retained.

The second sampling method ensured that the minimum distance between points was greater than a specified value. This was achieved using CloudCompare, which has a spatial sampling function which ensures local densities are constant (GirardeauMontaut, 2013). In this sampling method, 1mm and 2mm minimum distances were employed. The sampling methods were always applied in isolation.

When the point clouds had been properly prepared, the wrapping function was applied. As can be seen in Figure 5.9 (b), a number of holes might have existed in the surface mesh. In order to create a continuous surface, the hole filling function in Geomagic was employed. While the option to fill individual holes was available, all holes were filled at once. This approach was employed in an effort to increase the level of automation, so that operator variance could be minimised. The surface was then smoothed and the polygons relaxed using the respective functions, as can be seen in Figure 5.9(d). This eliminated many of the peaks and created a more realistic surface. However, due to this process, the shaft was closed, as it was considered as a hole. To remedy this, the planar cutting tool was employed to exclude the bottom of the mesh, as illustrated in Figure 5.9(e). This positioning of the trimming tool was applied

manually, with attempts being made to retain as much of the original shape, whilst eliminating the filled hole.




Figure 5.9: The mesh reconstruction process carried out in Geomagic Wrap. The original point cloud (a) was wrapped, resulting in a jagged surface mesh with a number of holes (b). The holes were then filled (c), after which smoothing and polygon relaxation were applied (d). This resulted in a continuous surface mesh, which was representative neither of the physical bone or the original point cloud. As such, a planar trimming tool was employed (e) to remove the unwanted part of the mesh (f).

While CloudCompare has the ability to evaluate the difference between two meshes, this method is weak. During mesh comparison, the first mesh is considered to be the reference, while the second is that which is to be compared. This comparison process entails matching each vertex of the second mesh to the closest vertex of the first.

However, this is not representative of the surface itself, merely the points that define it. The vertices that make up the meshes representing the surface could be sparse, meaning that the distance between the centre of the faces and the associated vertices could be quite far apart. This problem was less likely to occur in the laser scanner-derived reference meshes, in that the vertices were based on highly dense point clouds. Therefore, the centre of a face from the data could be very close to the nearest vertex in the reference mesh, but if the vertices associated with the face were far from the nearest reference vertices, the error assessment would show only high errors. This problem is illustrated in Figure 5.10 (a), where if a face is defined by two vertices, the distance between those two vertices and the closest point on the reference surface is not a fair representation of the distance between the two surfaces generally. If, however, the compared surface was densely sampled, as shown in Figure 5.10 (b), the distances between the discrete points and the reference surface would be more representative of the distance between the surfaces.




Figure 5.10: Idealised example illustrating the weakness in mesh to mesh comparisons. The red lines represent the surface of the reference mesh, while the blue lines represent that of the compared mesh. The red points signify the vertices of the reference mesh, while the black points in (a) represent the vertices defining that part of the surface and the green points in (b) denote the points sampled across the whole surface.

Therefore, during comparison, each surface mesh was densely discretised using the mesh subsampling function in CloudCompare. Random sampling was employed, with a target of 10^7 points. The results of this process are exemplified in Figure 5.11, which shows the original surface mesh and, in part (b), the representative point cloud. It

should be noted that sampling of the reference mesh was not carried out, as the vertices describing it were very dense and, as such, it was unnecessary. During comparison of the ultrasound-derived meshes and the reference meshes, both the data mesh itself and the dense point cloud derived from it were analysed, providing mean error, maximum error and standard deviation.




Figure 5.11: Final surface mesh of an unsmoothed SAFT data set (a) and the densely sampled point cloud (b). Also shown is the detail of the point cloud, illustrating the transparency and, therefore, discontinuous nature of the model.

5.3 Results

Some of the error reduction methods described in the previous section were designed specifically to improve the outcome for point cloud reconstruction. However, an improvement in the accuracy of the point cloud would not necessarily improve the accuracy of the resulting mesh. The results of the two phases are, therefore, presented separately.

5.3.1 Point Cloud Comparison

As can be seen in Table 5.1, the point clouds resulting from SAFT data processing with 30° and 40° apertures, with the application of a smoothing function achieved a sub-

millimetre mean error. While the 40° aperture achieved the lowest mean error, the 30° results showed the lowest standard deviation and maximum error. The TFM results had the worst results for mean error and standard deviation, but had a lower maximum error than found with the 20° SAFT data. As well as this, there was a significant difference in the number of points in the SAFT clouds and the TFM-derived clouds. Additionally, the application of a smoothing function appeared to improve the results in all the SAFT data sets. However, the opposite effect was produced in the TFM results.

Table 5.1: Results of the original point cloud data from the cuboidal scan, before the application of surface meshing. The bracketed angle next to all SAFT results denotes the angular aperture employed.

Processing Method	Smoothing	Mean Error (mm)	Standard Deviation (mm)	Maximum Error (mm)	Number of Points
SAFT (20°)	yes	1.0	0.9	13.2	77106
SAFT (20°)	no	1.1	0.9	13.1	77121
SAFT (30°)	yes	1.0	0.8	8.0	68765
SAFT (30°)	no	1.1	1.1	32.1	68859
SAFT (40°)	yes	1.0	0.8	10.2	66589
SAFT (40°)	no	1.0	0.9	10.4	66664
TFM	yes	1.1	1.1	10.0	12284
TFM	no	1.1	1.1	12.0	12310

The general shape of the non-smoothed SAFT point clouds was very similar to that which was expected, being formed by criss-crossing lines of points. This was as expected, given the method of probe manipulation and data capture. However, as can be seen in Figure 5.12, there were still a noticeable number of outliers, which were responsible for the high maximum error. These appeared in two forms, the first of which was created by side lobes. The effect of this can be seen in the top right of the bottom left viewing angle. The second form was in specular effects, which caused points to be produced when the probe was no longer physically in line with the sample

surface. This is made obvious in the top right of the top left viewing angle in Figure 5.12.

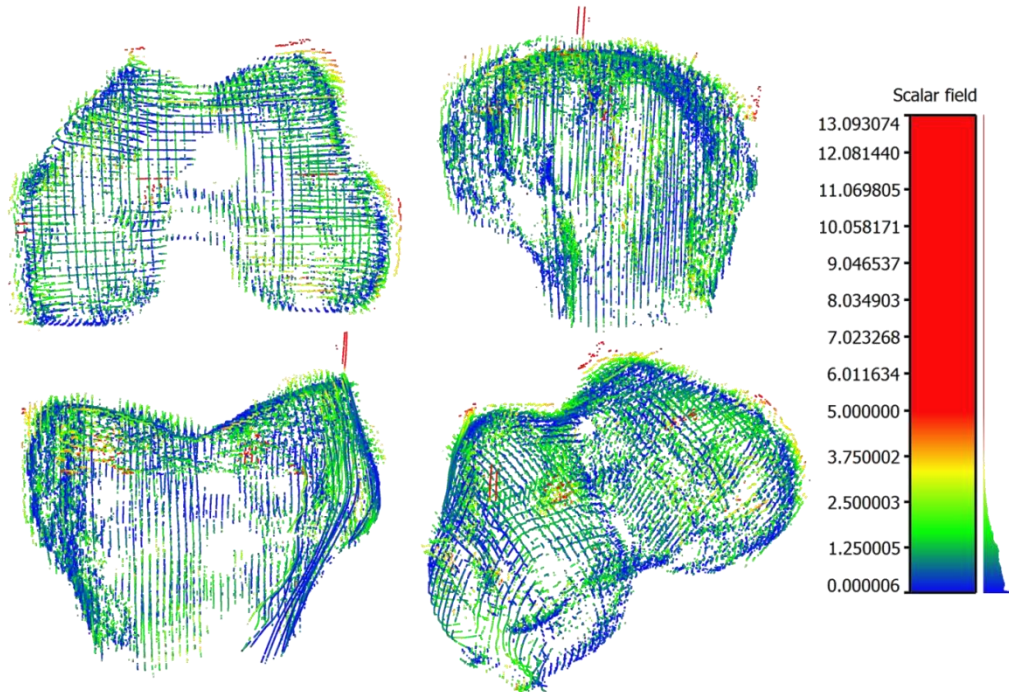


Figure 5.12: Four viewing angles of the compared point cloud produced using 20° aperture SAFT processing, with no smoothing function.

In addition to this, the points did not form an infinitely thin surface. Instead, there was a group of points shown in blue that displayed high levels of accuracy, while a number of points existed above and beneath this surface, as shown in green and yellow. The drilled hole shown in the physical sample in Figure 5.1 was apparent, as can be seen in the top left of Figure 5.12. There were also other regions where the point density decreased, such as the area between the condyles, which can be seen just below the hole in the top left of Figure 5.12. This lack of density was, however, minimised in the larger SAFT apertures. While the different apertures in the SAFT data displayed varying results, the shape and features were similar.

The smoothing function produced a very similar shape to that seen in Figure 5.12. In addition to this, the outliers caused by specular effects were produced in the same

regions as the unsmoothed, as can be seen in the bottom left image in Figure 5.13. The smoothed and unsmoothed point clouds were, however, different in that the areas of low error were filled by continuous lines of points, rather than by detached individual points. This is more obviously illustrated in the differences between the extracted contours shown in Figure 5.5 and Figure 5.6.

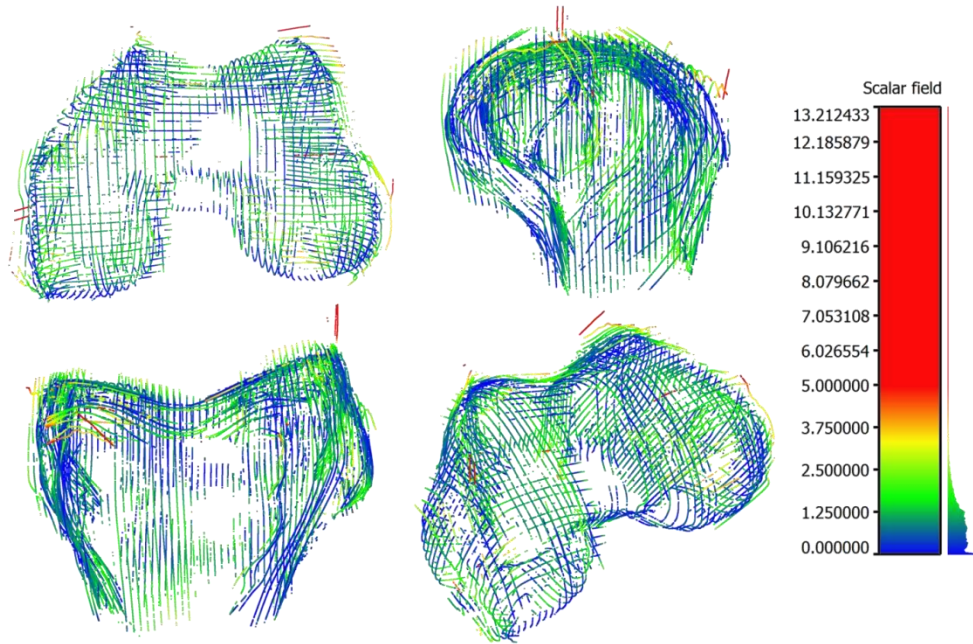


Figure 5.13: Four viewing angles of the compared point cloud produced using 20° aperture SAFT processing, with smoothing function. While the individual lines are cleaner than in Figure 5.12, a number lines have been located inside the sample.

While all the SAFT-derived point clouds were very similar, the TFM point clouds were significantly different. Firstly, the number of outliers appeared to increase. This can be seen in Figure 5.14 as groups of red points and was caused primarily by side lobes. Additionally, comparing Figure 5.14 to Figure 5.12, it can be seen that there is a large difference in density. As shown in Table 5.1, the number of points in the TFM clouds was massively lower than that in the SAFT clouds. This was, once again, caused by side lobes. As demonstrated in Figure 5.15, side lobes were not present in the SAFT image, while there was a significant side lobe contribution in the TFM image. In order

to avoid extracting the side lobes instead of the true surface, the threshold limit was made higher. In doing this, however, much of the true surface was lost, resulting in sparser point clouds.

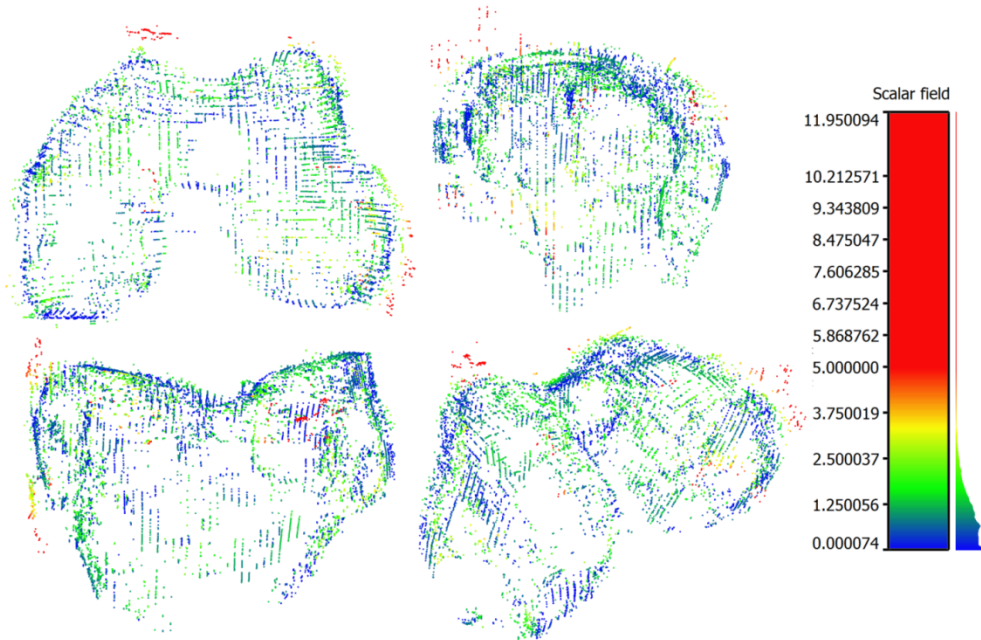


Figure 5.14: Four viewing angles of the compared point cloud produced using TFM processing, with smoothing function. The point clouds are significantly more sparse than with SAFT reconstruction.




Figure 5.15: A typical FMC processed using 30° aperture SAFT (a) and TFM (b), with part (b) displaying side lobes.

When the hemispherical scan path was employed, the results became significantly poorer. As can be seen in Table 5.2, none of the processing techniques were capable of delivering a sub-millimetre mean error. The standard deviation and maximum errors found were also significantly higher than those produced with the cuboidal path, with no sub-millimetre standard deviations achieved and all the maximum errors greater than 10mm. Additionally, the number of points in each cloud was significantly higher than those found with the cuboidal path. This was the result of increased image numbers, with the cuboidal path producing 349 images and the hemispherical path producing 1042.

Table 5.2: Results of the original point cloud data from the hemispherical scan, before the application of surface meshing. The bracketed angle next to all SAFT results denotes the angular aperture employed.

Processing Method	Smoothing	Mean Error (mm)	Standard Deviation (mm)	Maximum Error (mm)	Number of Points
SAFT (20°)	yes	1.3	1.1	14.0	232,950
SAFT (20°)	no	1.3	1.1	12.0	232,950
SAFT (30°)	yes	1.4	1.3	15.9	210,123
SAFT (30°)	no	1.4	1.3	15.6	210,123
SAFT (40°)	yes	1.5	1.9	31.3	199,994
SAFT (40°)	no	1.4	1.9	31.2	199,994
TFM	yes	1.4	2.4	40.3	31,194
TFM	no	1.5	2.4	39.8	31,194

In both the SAFT and TFM derived point clouds, it can be seen in Table 5.2, it can be seen that the effect of smoothing was negative, but bordering on insignificant. Further, it is clear that an increased synthetic aperture size was associated with a general decline of results.

As is evident in Figure 5.16, while the shape of the sample was represented, the point clouds did not depict the shape of the bone as well as was achieved with the cuboidal path results. The unsmoothed results showed a number of regions of particularly high

error, including the back part of the cloud, as can be seen in the top right view in Figure 5.16, and across the midline, as can be seen in the bottom left view. As expected with unsmoothed clouds, the points did not appear to form continuous lines, but rather clusters of points, with some above and others below the target surface. Also, as can be seen in the top left view of Figure 5.16, there was a decrease in the density of the points towards the front of the sample. This sparsity led to an almost complete lack of representation for that region of the point cloud. Additionally, the density of the points appeared to increase across the midline of the cloud. This can be seen in the bottom left view of Figure 5.16 and was in contrast to the relatively equal dispersion of points found with the cuboidal scan path. Also shown in the bottom left view are a significant number of points which appear to be embedded quite deeply within the point cloud, rather than on the surface.

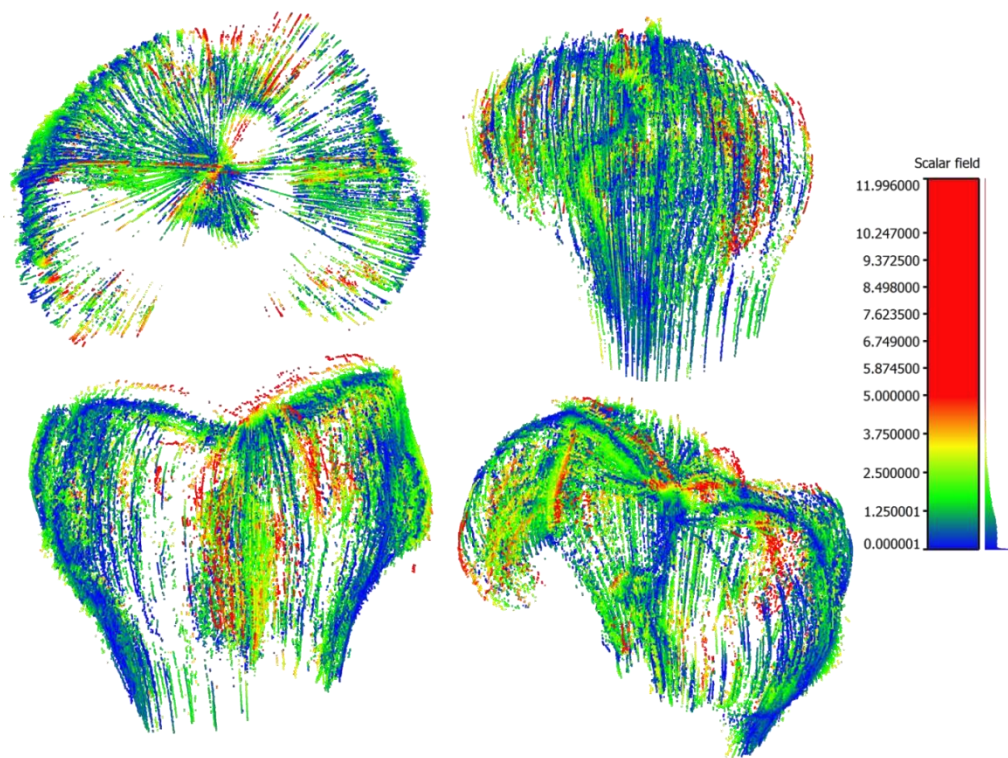


Figure 5.16: Four viewing angles of the compared point cloud produced using 20° aperture SAFT processing and the hemispherical scan path, with no smoothing

function. The intercondylar notch has not been reconstructed as well as was the case with the cuboidal scan path.

Comparing the above results to those found with smoothing applied, it can be seen in Figure 5.17 that instead of clusters of points, definite lines of points were formed. Additionally, while the unsmoothed point cloud maintained a representation of the drilled hole in the sample, the smoothed cloud did not. This can be seen in the top left view of Figure 5.17, where an area of high error can be seen with a circular shape. The regions found to have large errors corresponded with those found in the unsmoothed point cloud. However, the points found beneath the surface were more obvious, as can be seen in the bottom left hand view of Figure 5.17. Additionally, the sparsely populated regions in the unsmoothed point clouds were found also with the smoothed results.

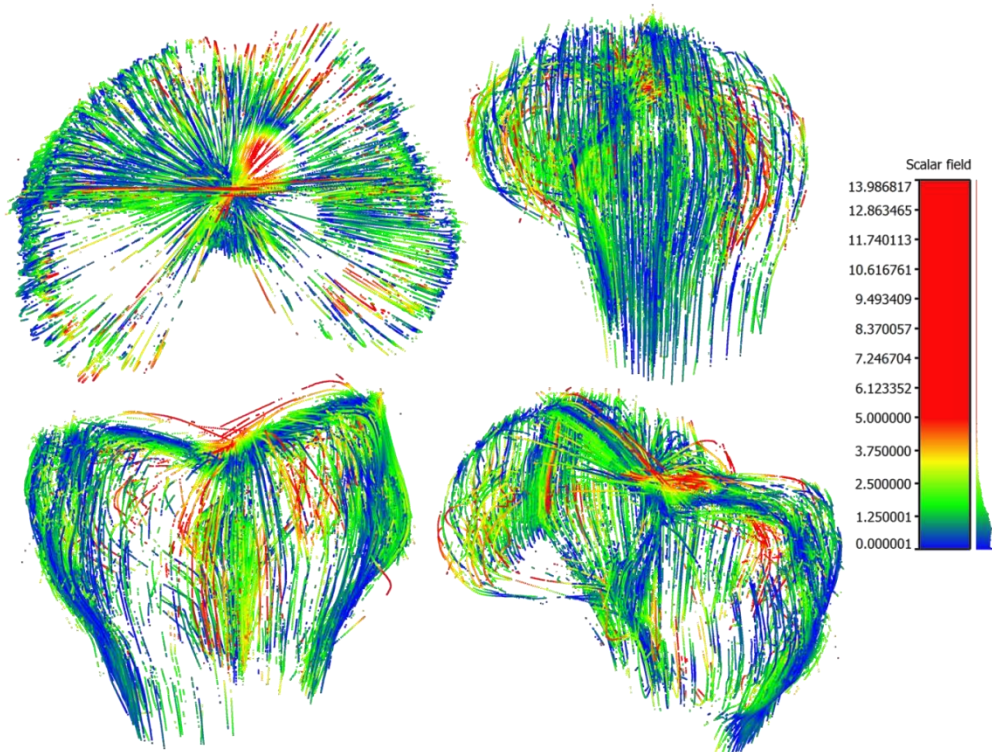


Figure 5.17: Four viewing angles of the compared point cloud produced using 20° aperture SAFT processing and the hemispherical scan path, with smoothing function.

The comments and comparisons on the point clouds generated using a 20° aperture SAFT processing method could be applied to the other two aperture widths. While error sizes were magnified as the aperture increased, the features remained the same. However, it should be mentioned that the number of outliers and the associated error magnitude increased drastically as the aperture width increased. As can be seen in Figure 5.18, significant numbers of points on the far right and left of the image are far from the surface. These were caused by high intensity side lobes that were not thresholded out, as the intensity was equal to that of the surface reflections.




Figure 5.18: The point cloud produced using 40° aperture SAFT processing and the hemispherical scan path, with no smoothing function. Large errors caused by side lobes are present on the far left and right of the image, as well as at the top.

While there was much agreement between the different SAFT aperture sizes in terms of performance, the TFM-derived point cloud differed significantly. The first obvious difference was the point density, which is exemplified in Figure 5.19. The reason for the drop in density was because the threshold level was set higher for the TFM images, at -4dB. While this helped eliminate many of the side lobes, it entailed that much of

the true surface was lost. Additionally, the errors were, generally speaking, higher, as evidenced by both Figure 5.19 and Table 5.2.

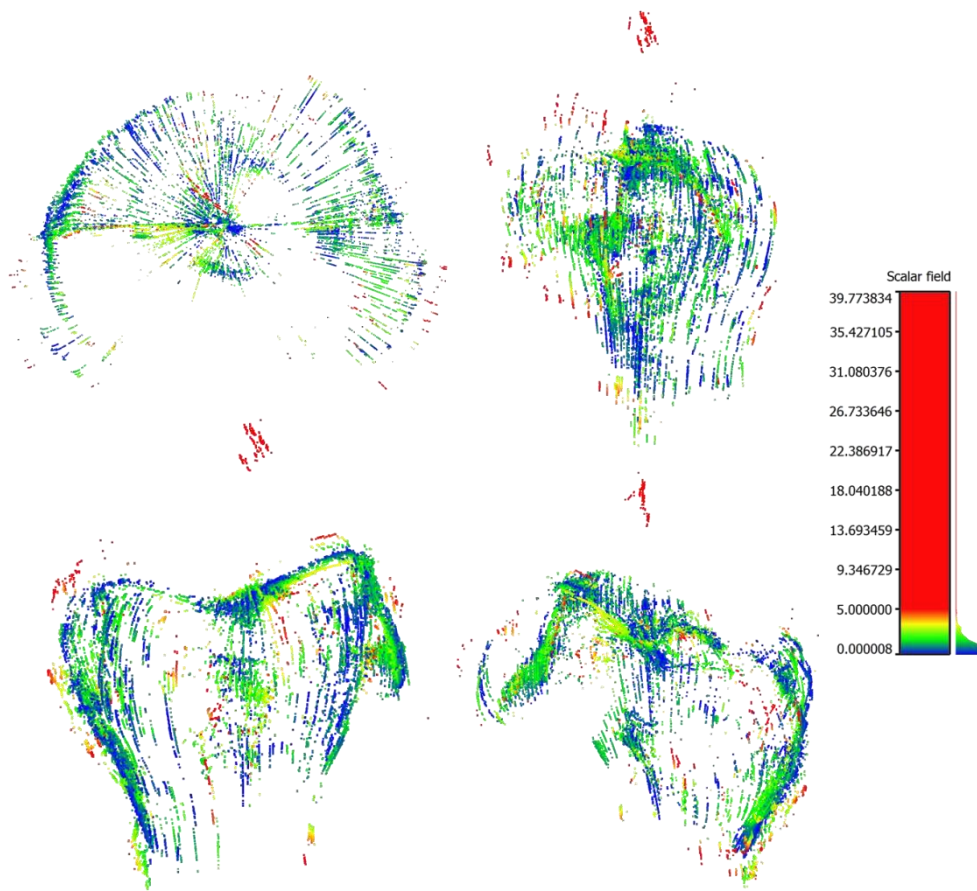


Figure 5.19: Four viewing angles of the compared point cloud produced using TFM processing and the hemispherical scan path, with no smoothing function. The higher thresholding associated with TFM processing resulted in a sparser point cloud.

As the synthetic aperture width of the SAFT methods was increased, the magnitude of the errors increased, as did the number of highly erroneous points. This can be seen by comparing Figure 5.16 and Figure 5.18. As TFM is essentially SAFT with as large an aperture as possible, the progression of error magnitude continued from the large aperture SAFT methods to the wider aperture TFM. This is further demonstrated in Figure 5.20, which shows how the mean error increases with the aperture for the

hemispherical scan path. The conclusion of this, then, was that large synthetic apertures were not suitable for this scan path.




Figure 5.20: Mean errors for the point clouds associated with the two scan paths, with smoothed and unsmoothed examples. The most important factor in mean error was the probe path, with the cuboidal path producing the smallest errors.

As well as the problems listed above, the TFM-derived point clouds also suffered during the error reduction process. While the SAFT point clouds lost a proportionally small number of true surface points, the TFM clouds lost a large number. This left large parts of the surface unaccounted for, with little or no representation. In addition to this, the process did not eliminate any significant part of the outliers, as evidenced by Figure 5.21. Therefore, the error reduction process only served to increase the mean error of the TFM point clouds.




Figure 5.21: Result of error reduction procedure on unsmoothed, TFM point cloud. The procedure has not eliminated the outliers, instead removing the true surface.

5.3.2 Mesh Comparison

Focussing, firstly, on the results of the comparison of meshes generated from the cuboidal robot path, it can be seen in Figure 5.22 that the mean error was below the desired limit of 1mm in a number of cases. It is clear from the graph that the TFM-derived meshes were poorer than all three aperture widths employed using SAFT. In the TFM-derived meshes, the mean error achieved for every parameter combination was greater than 1mm. This can be seen in Figure 5.23, which shows two point clouds with colours defined by the error levels, exemplifying the difference between the TFM and SAFT results. While the largest errors appear in the same region, the general magnitude of the errors appears to be much higher for the TFM results, which would be detrimental to the mean error. In addition to this, it can be seen from the left hand images that the profile of the surface representations appear different, with the TFM-derived cloud showing less definition. While this figure only shows two of the possible

parameter combinations, the results were typical of the comparison between the different SAFT processing methods and TFM errors in both the point clouds and meshes.




Figure 5.22: The mean error achieved for data compared with the laser-scanned reference mesh for the cuboidal scan path. ‘Smoothing’ refers to the application of a smoothing function to the 2D points forming contours in each image, ‘nearest neighbour’ denotes the application of the outlier elimination algorithm and ‘sampling’ indicates whether or not the mesh itself or a densely sampled point cloud representing the mesh was employed




Figure 5.23: Two viewing angles of errors resulting from matching TFM-derived (a) and 20° SAFT-derived (b) meshes with no smoothing and no outlier elimination applied that have been densely sampled and compared with the laser scan data.

The application of a smoothing function at the contour extraction stage had a significant effect on the mean error produced with the associated mesh. The results in Figure 5.22 for the 20° SAFT meshes show a clear distinction, with smoothing causing higher mean errors. This is reversed, however, for the 30° SAFT results, with the unsmoothed errors being larger than the smoothed. The effects of this leveled out to an extent with the meshes derived from 40° SAFT processing, with the smoothed errors decreasing slightly and the unsmoothed increasing marginally. Finally, the application of a smoothing function appears to have had little effect on the TFM-derived meshes. On the whole, smoothing appeared to be of benefit to meshes derived from 30° SAFT processing, while the opposite was true for 20° SAFT processing.

While the application of the nearest neighbour outlier elimination algorithm always resulted in a lower mean error for the point clouds, this was not the case for the final

meshes. Due to this, the point clouds which did not have the algorithm applied were also surface wrapped and compared to the reference mesh. The results of this followed no obvious pattern, with the exception of the 20° SAFT meshes, in which use of the algorithm always resulted in a higher mean error.

Comparisons made directly between the ultrasound derived meshes and the laser scan-derived reference mesh always resulted in a lower mean errors than those made between the reference mesh and the densely sampled point clouds. The differences between the two comparison methods are further highlighted in Figure 5.24, where the directly compared mesh shows far less detail than the sampled mesh.




Figure 5.24: Two viewing angles of errors resulting from matching 20° SAFT-derived meshes with no smoothing and no outlier elimination applied. One has been compared directly with the laser scan data (a), while the second (b) has been densely sampled then compared. Part (b) offers greater detail, as (a) has a low vertex density.

The resulting standard deviations (Appendix B) followed much the same patterns as observed with the mean error. As before, TFM-derived meshes showed the worst results, while the best performing SAFT aperture width was dependant on other parameters. The maximum errors (Appendix B) showed a striking contrast between the values associated with sampled mesh and the direct mesh to mesh comparison, with sampled cases being much higher and more consistent. The sampled results show little difference due to changes in other parameters, all being approximately 7.6mm. The non-sampled results, on the other hand, show a much higher dependence on the parameter permutations.

As can be seen in Figure 5.23, there exists a region of particularly high error corresponding to the drilled hole present in the physical sample, as shown in Figure 5.1. However, the non-sampled results showed not only lower levels of maximum error, but a higher level of variation. This can be visualised in Figure 5.24, which shows a very clear circle of error, corresponding to the hole. However, in the direct mesh comparison, while there is an obvious area of higher error, the circular shape is not obvious. The vertices making up that region of error were relatively sparse in distribution and the location varied from mesh to mesh. As the maximum error was based on the location of these vertices with variable location, the vertex furthest from the reference model also varied.

The mean error results from the comparison with the CAD model from Sawbones differed from those from the laser-scan-derived reference model. The results (Appendix B) show that, unlike the other comparison method, the application of the smoothing function resulted in a lower mean error for the 20° SAFT meshes. For most parameter combinations in the 20°, 30° and 40° SAFT results, use of the outlier elimination resulted in poorer mean error, with the effect being particularly noticeable with the 20° and 40° SAFT meshes. These effects were, however, reversed in the TFM results. As seen previously, the sampled comparisons resulted in a lower mean error for all parameter permutations.

Assessing the visualisation of the error for both reference models, the most striking difference between the two is the lack of the circular region of error, as exemplified in Figure 5.25. The reason for this dissimilarity was the lack of the hole in the CAD

model. This affected the mean error, with the elimination of an area with high and sustained error. The two comparison methods were, however, similar, in that the areas in which low and high errors appeared to be replicated between methods. For example, a region of high error can be seen in the centre of the medial condyle in both examples.




Figure 5.25: Two viewing angles of errors resulting from matching 20° SAFT-derived meshes with no smoothing and no outlier elimination applied. One set has been sampled and compared with the laser scan data (a), while the second (b) was compared with the CAD reference model.

The standard deviations found for the CAD model comparisons differed from those found from the laser scan, with smoothing yielding lower values. This effect was reproduced for all processing methods (Appendix B).

The most striking difference in numerical results between the laser and CAD reference model comparisons was in the maximum error. The values were smaller and there is a

general rise in error as the synthetic aperture width increases. This major difference is obvious also in Figure 5.25, which shows the area of the highest error very obviously in the laser scan compared data. However, as this hole was filled in both the final meshes and the CAD reference model, this area of error does not exist and the highest error is found at some other point. Additionally, the clear distinction between the sampled and non-sampled was not found in the CAD reference results

Focussing on the results of the hemispherical scan, was found (Appendix B) that the levels of mean error were much higher than found with the cuboidal scan path, evidenced by the fact that the lowest mean error was 1.34mm and the highest almost 3mm. Additionally, the TFM mean errors were not as poor as before, with the 30° SAFT values reaching the same levels of inaccuracy. These results are shown in Figure 5.26, which compares the results from the hemispherical and cuboidal paths. In the hemispherical results, the number of areas with very high (>5mm) levels of error was larger than the single area found with the SAFT results. This was caused by the regions of the mesh which have bulged out, creating a shape that was very unlike that which was desired. As such, the mean error would have been increased by the existence of a number of regions of high error. The lack of distinction in performance between the different processing methods in terms of mean error could be explained by the inaccurate meshes exemplified in Figure 5.26 and replicated across all the different parameter combinations.




Figure 5.26: Two viewing angles of errors resulting from matching 30° SAFT-derived sampled meshes with no smoothing and no outlier elimination applied. One set has been produced using the hemispherical scan path (a), while the second (b) was generated using the cuboidal scan path. The inaccurate surface point placement caused by use of the hemispherical path caused the construction of erroneous surfaces.

The effects of smoothing were most noticeable with the 30° SAFT and TFM results. While the other processing methods saw fairly small changes due to the application of the smoothing function, the 30° SAFT and TFM errors were much higher for the unsmoothed meshes. Unlike the results for the cuboidal path, the mean error for those meshes compared with the laser scanned reference was lower. With a couple of exceptions in the 40° SAFT results, the application of sampling, as before, improved the mean error.

The standard deviations achieved with the hemispherical path showed no obvious pattern, with the poor reconstruction being ubiquitous with all parameter combinations. The same could be said of the maximum errors, where changes in parameters caused differences in result, but followed no pattern. However, those maximum errors tended

to come from two distinct features. The first of these was a surface mesh separate to the main mesh, which can be seen in Figure 5.26 (a), in which another mesh has formed within the main mesh. This effect can also be seen in Figure 5.27 (a), but the small surface mesh was instead formed outside the main mesh. The second feature causing the high maximum errors for the hemispherical scan results was a complete loss of shape and definition in the reconstructed meshes, as illustrated in Figure 5.27 (b). In this example, the gap between the condyles was completely lost, along with many of the other defining features. Due to the simplified shape of these meshes, the errors were very high.

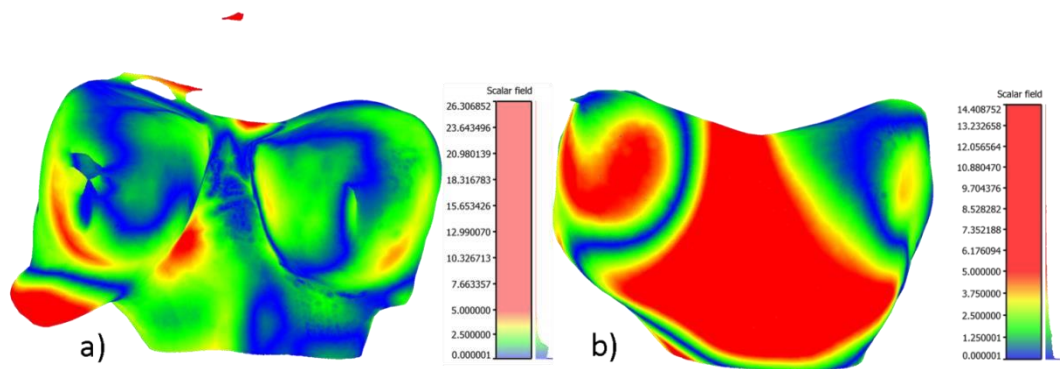


Figure 5.27: Errors resulting from matching 40° SAFT-derived (a) and TFM (b) sampled meshes with no smoothing and no outlier elimination applied with a hemispherical path. The meshes produced with the hemispherical path showed no pattern in the placement of inaccurate surfaces.

Speaking generally, the meshes derived from the parameter permutations which achieved the best results possessed similar attributes. In particular, the regions in which the error was highest was replicated across the described meshes. In addition to the obvious circular area of high error in the region of the drilled hole, there were areas both on the left sided corner of the lateral condyle and on the front of the medial condyle, both of which can be seen in Figure 5.28 (a). In addition to this, there was a region on the underside of the right side of the right condyle, as is shown in Figure 5.28 (b). Further to this, there was a large region of high error on the left side of the

back of the epiphyses, which is exemplified in Figure 5.28 (c). Finally, as can be seen in Figure 5.28 (d), there was a small region of error on the right hand side, at the metaphysis.




Figure 5.28: Errors resulting from matching 20° SAFT-derived sampled meshes with no smoothing and no outlier elimination applied. The colour limit has been set to 3mm to better expose the regions containing the highest error. Regions with a high gradient tend to have higher levels of error.

5.4 Discussion

A number of the point clouds achieved the desired sub-millimetre mean error, but only certain parameter combinations allowed for this. This implied a strong dependence on parameter selection. Most of the different parameters were defined before the point cloud was created, which demonstrated the importance of discovering the most effective parameter combination, as the earlier in the process the parameter definition is, the more costly it would be to change it.

One such parameter was the choice to apply a smoothing function to each of the extracted contours. For the cuboidal scan path, the use of a smoothing function improved all three of the presented result values. For the hemispherical scan, while smoothing worsened both the mean error and the maximum error, the standard deviation was still enhanced. Improvements in the standard deviation were expected, given that the contours were smoothed and, therefore, more likely to be less erratic. This principle is illustrated in Figure 5.29, which shows a scenario in which a smoothed contour can have a higher mean error than an unsmoothed contour, but still maintain a lower standard deviation, due to the lack of variability from one point to the next. Generally speaking, smoothing appeared to work well with images in which the surface was very obvious, while it was detrimental to those that had a poor representation. Therefore, if the quality of the images could not be guaranteed, it would be prudent to omit the application of the smoothing function.




Figure 5.29: Representation of a possible contour extraction scenario, displaying the difference between points with smoothing applied (a) and those without (b), showing how smoothing can increase mean error.

Another parameter which greatly affected the error levels of the final point cloud was the data processing method. The results for the cuboidal scan path showed little difference between the SAFT aperture width definitions, but showed a decrease in performance when TFM was employed. For the hemispherical path, the increase in aperture width was always accompanied with an increase in error levels. The most likely reason for this was the rise in both side lobe area and intensity. The slight improvement in results seen for the 40° SAFT point cloud from the cuboidal scan could be explained by the improvement in resolution offered by the increased aperture width

relative to the 20° SAFT. This would be balanced by a decrease in aperture width compared with TFM, implying smaller and less intense side lobes.

Comparing the TFM results against those from the SAFT methods generally, there were far fewer points and, therefore, far less coverage. The reason for this was the threshold limit for the TFM-derived images was set to -4dB, while that for the SAFT images was -8dB. While this helped eliminate much of the influence of the side lobes associated with TFM, the extracted contours were smaller, giving a less comprehensive representation of the surface.

A further factor affecting the performance of the point clouds was the choice of probe path. A number of differences were apparent between the point clouds resulting from the different paths, one of which was the lack of coverage in certain regions relative to the cuboidal path. The reason for this was the probe face was always oriented towards the centre of the sample and, as a consequence, did not have a direct line of sight to certain surfaces. This is illustrated in Figure 5.30, which shows a number of probe positions which do not allow for access to the surfaces of the intercondylar notch. The result of this was a sparseness of the points in that area, as can be seen in Figure 5.16.

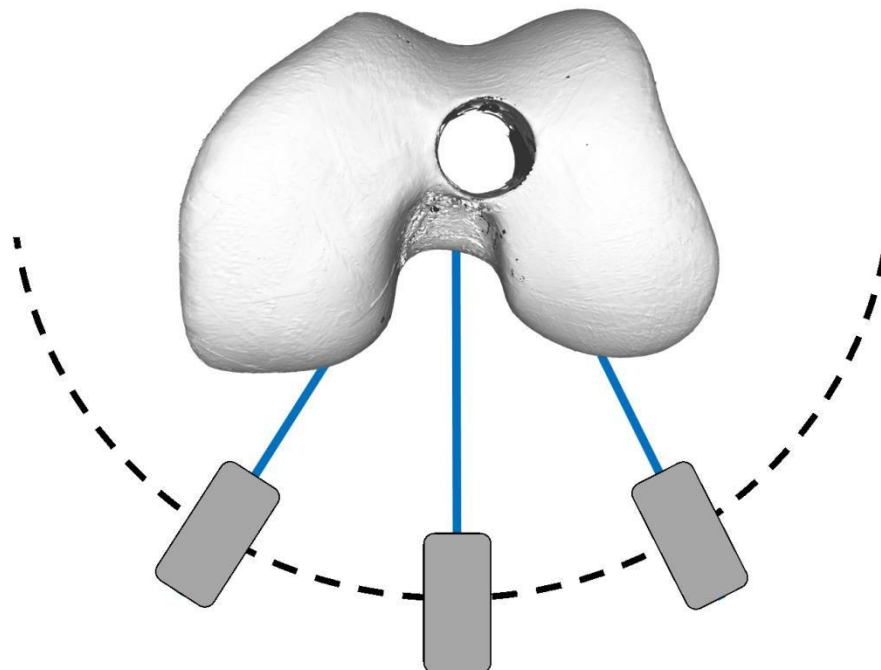


Figure 5.30: The line of sight (blue line) of the probe relative to the sample at a number of positions along the circular path, illustrating how the hemispherical path allowed no access to the intercondylar surfaces.

A second implication of the hemispherical path was the change in velocity as a consequence of the accelerations and decelerations of the CIRC commands. While the spacing of the data captures and, therefore, the contours was equal for the cuboidal path, that of the hemispherical path was sparse at the midway point of the CIRC commands and dense at the beginning and end. Based on these two disadvantages, the cuboidal path was superior, in that it offered more comprehensive and uniform coverage of the surfaces of the knee.

The point clouds associated with the hemispherical path also displayed a large number of points inside the main surface representation, which was not seen in the clouds resulting from the cuboidal paths. It is difficult to identify exactly what caused this, but one possible explanation is that the complexity of the hemispherical path was much greater than that of the cuboidal path. In particular, the changes of orientation were much greater, which resulted in large angular changes at each individual joint of the robot. If the orientation calibration of the tool was not optimal, this could result in orientation errors and, as a consequence, misplacement of points. As the orientation calibration was carried out with relatively small changes in joint angles, the likelihood of errors in orientation reconstruction would increase as the joint angles moved further from the joint angle spectrum used during calibration.

Another related cause could be due to phase discontinuities during KUKA angle reporting. More specifically, there might be occasions when the angle passed a threshold and ‘flipped’ (ie. the angle would be recorded 180° lower than it should have been). Such a problem should not happen, but could be helped by applying phase unwrapping to the angular outputs.

Two processes were applied after the point cloud had been established: the nearest neighbour outlier elimination algorithm and the segmentation based point elimination. The former had a positive effect for most of the point clouds, in that it improved the

associated results. However, as demonstrated in Figure 5.21, this removed many points which were accurate. The reason for this was that the parameters governing the smoothing function were chosen by trialling, resulting in the best results for the majority of the point clouds. As the TFM point clouds were strikingly different in composition to the SAFT clouds, the effect was different. In future, an effort could be made to improve the effectiveness of this method for all cloud types, the parameters used could be dynamic, being based on features associated with each cloud. For example, a measure of sparseness could be found for each point cloud, which would be based on the mean distance of the closest 10 points for every point. Those point clouds with a higher level of sparseness would then employ a larger catchment radius, as described in Section 5.2.

The mean errors achieved when comparing the meshes found with the cuboidal scan path achieved the desired sub-millimetre status for almost all of the SAFT parameter combinations. With a small number of exceptions, the parameter changes only resulted in relatively small variations in results. The application of a smoothing function to the extracted contours, for example, while affecting the performance of the meshes, had a major effect on the 20° SAFT, but little on both the 30° and 40° SAFT results. Similarly with the TFM results, the effect of these parameters was minor.

While the meshes were derived from the point clouds, this did not imply that the performances of the point clouds would be replicated once meshed. For example, while the smoothing function appeared to have a detrimental effect on all the point clouds, it might actually be beneficial during the wrapping procedure. The reason this might happen is that the point clouds were compared based on every point in the cloud, whereas the point cloud was spatially sampled before the wrapping procedure was applied. This fundamentally changed the cloud and would have an influence on the wrapping procedure and the surface it produced.

The same idea could be applied to the nearest neighbour outlier elimination, in that the differences in the results based on the parameter combinations did not always translate from the point clouds to the meshes. In the case of outlier elimination, while it might have helped eliminate points of high error in the point cloud, it might also have resulted in a loss of accurate points. However, this might not have had much of an impact on

the results for the point cloud comparison. Such a loss of accurate points would, however, become a problem during the wrapping procedure, as fewer points accurately representing the surface would be available and the function would have to approximate the rest of the surface.

Comparing the performance of the different aperture widths, the TFM-derived meshes proved to provide the poorest results. This was expected, given that the TFM point cloud results were inferior to the SAFT results. However, the point cloud results suggested that an increase in aperture width led to an increase in the error, but this was not the case with the mesh results. The reason for this was that while the TFM point clouds were significantly worse than the SAFT results, the SAFT meshes were not considerably different. During the wrapping procedure, these small differences in the SAFT point clouds could be eliminated.

As with the point cloud results, the meshes derived from the hemispherical scan produced significantly poorer results than the cuboidal scan. This was contributed to both by the inaccurate points in the associated point clouds and the sparsity of points in certain areas. The former contributed to the formation of surface meshes separate from the main mesh, as is displayed in Figure 5.27 (a). The latter produced meshes which had large areas which were incomplete, as demonstrated in Figure 5.31. In such cases, the completion of the mesh was dependant on the hole filling function, which simply used the closest surfaces to reconstruct the entire mesh. In such situations, the results were similar to that seen in Figure 5.27 (b). As such, the cuboidal path was found to be superior in mesh reconstruction, made possible by the more uniform coverage of the entire surface and the lack of inaccurate points.




Figure 5.31: The result of the wrapping procedure on the 20° SAFT-derived point cloud with no smoothing and no outlier elimination applied showing the incomplete nature of some of the surface meshes produced.

A further parameter which affected the results of the mesh comparison was the use of mesh sampling. For the meshes created via the cuboidal scan paths (which has been shown to be the most reliable of the two scan paths), the direct mesh-to-mesh comparisons produced lower mean errors, maximum errors and standard deviations. This reason for this was that the regions with the highest level of error were also likely to be the areas which had either only a small number of points, if any, representing the surface or had small numbers of erratically positioned points. In the latter case, the wrapping function would estimate the positions of the vertices making up the surface, which would likely result in a simplified representation, made up of a less dense group of vertices than seen in a highly populated area. In the former case, the surface would not be reconstructed during the wrapping procedure, but would, instead, be reconstructed using the hole filling function. This, as before, would result in a highly estimated surface, consisting of a relatively sparse distribution of vertices.

Therefore, in the regions with poor surface representation in the associated point clouds, the consequence in the resulting mesh was that these areas would have a relatively small number of vertices. During the comparison procedure, this implied that

the areas likely to be accurate would have a greater contribution to error measurement than those likely to be inaccurate. This problem was eliminated by dense sampling of the meshes, allowing for equal representation of all parts of the mesh, be they accurate or otherwise. This was most obviously demonstrated in the maximum error results, where the values for the sampled meshes were very similar, as they corresponded to the radius of the drilled hole. However, this value was highly variable for the mesh-to-mesh comparison, due to the low number of vertices representing the area. The conclusion, therefore, was that sampling the meshes before comparison provided a more accurate result.

The purpose of employing the CAD model rather than the more accurate laser scanned reference model was to assess the effect of removing the drilled hole from the reference model, as such a feature would be unlikely to occur in real knees. The effect of the use of the CAD reference was minimal for both the mean error and standard deviation. However, the maximum error was much smaller for SAFTderived meshes. This was as expected, given the elimination of the area of highest error. It should be mentioned, however, that the CAD reference was not an accurate method of comparison, but merely an attempt to test the influence of the drilled hole.

The positions of the regions in which error was high were replicated for all the meshes which performed well. The area depicted at the bottom of Figure 5.28 (d) can be explained by examination of the surface mesh reconstruction method. After the hole filling function had been executed, the surface was smoothed, which caused the part of the surface representing the shank to pull towards the centre of the mesh. This can be seen in Figure 5.9 (d) and was caused by the fact that there were no points representing the closed shank, implying that the surface was defined by very few vertices. In regions with few vertices, the smoothing function was capable of more influence and, as such, distorted the shape considerably. The error could, therefore, be explained by this effect.

The other regions illustrated in Figure 5.28 all had in common the fact that they were located in areas with steep angles relative to the position of the probe face. This effect is demonstrated in Figure 5.32 (a), which shows two specular ray paths at a part of the sample which is almost perpendicular to the probe face. These reflections would be

easily registered by the transducers. The bottom part, however, depicts a ray path in the area of high error shown in Figure 5.28 (b). In such a region, the specular rays were reflected away from the probe face and the only signal returned from that region would be from small amounts of diffuse reflection. Given the strength of the specular reflections above, the bottom surface would have a very weak intensity relative to the top part. This is demonstrated in Figure 5.33, which shows the same two 20° SAFT images, with one employing an image threshold of 18dB and the other -8dB. In the first image, all parts of the surface are clear. However, when examining the second image, which shows the points that would be employed during contour extraction, the bottom part of the surface disappears due to the relative lack of intensity. While this example only describes in detail one occurrence of this, the principle of the explanation in all the described areas of high error remains valid.




Figure 5.32: The incident (blue) and specularly reflected (green) rays, showing how certain regions can be rendered inaccessible due to specular reflection and the probe angle.




Figure 5.33: 20° SAFT images from the region described in Figure 5.32, showing the image threshold limits of -18dB (a) and -8dB (b), so that the angled surface was below the extraction limit.

The result of such underrepresentation of the surface in each image resulted in regions which had very few points, if any, to define the surface. While some of the influence of this sparseness of points would be diminished by the application of spatial sampling, the effects would still be noticeable. A possible solution that might aid in alleviating these problems would be to alter the cuboidal path, such that two similar paths would be executed. The first of these would be the same as that employed herein, while the second would be lowered in height and the probe would be angled, as shown in Figure 5.32 (b). In doing this, some of the regions that were reflecting the specular rays away from the probe face would become accessible and would, therefore, be represented in the surface point cloud.

5.5 Conclusion

A human, epoxy resin distal femur was imaged using two different scan paths. This provided surface point clouds which were compared to a reference model and then employed to create surface meshes. Numerous meshes were produced using various combinations of parameter definitions, which were then compared to two reference

models. This allowed for analysis of the parameters which contributed to the accuracy of the final surface mesh.

The result of this analysis was that the system was fundamentally capable of reconstructing the surface of the kind of complex shapes found in human knee joints and was able to do so to the desired sub-millimetre accuracy. Achieving a mean absolute error of 0.8mm using the laser-derived and CAD model as references, use of the 20° SAFT processing method allowed for surface reconstruction well within the desired mean error of 1mm. However, not every parameter permutation provided this level of accuracy, as certain variables were damaging to the accuracy of the final mesh. Such detrimental variables included the use of the hemispherical scan path and TFM data processing. For the meshes which achieved the desired level of accuracy, there were differences in performance and these were heavily dependent on the parameter combination employed, with no permutation performing significantly better than any other.

Therefore, while it was established that the choice of parameters was highly influential in the success of the system, it was unclear which should be employed generally, as only one sample geometry was tested. In order to establish a parameter combination which would be robust enough to reconstruct different shapes to the same sub-millimetre accuracy, imaging of a second sample would be necessary. In order to test the system appropriately, it would be sensible to employ a sample shape that, while preserving similar features to that which would be found in real human knee joint, was sufficiently different from the sample presented herein.

Additionally, while the experiment was a success, it should be mentioned that the sample was made from an epoxy resin and, as such, did not possess the acoustic properties of real bone. While the epoxy was unknown, epoxies often possess an acoustic impedance in the region of 2MRayl (Grewe & Gururaja, 1989), while that of bone is around 7.75MRayl (H. Azhari, 2010). Therefore, it would be prudent to employ a sample that was not only sufficiently different in shape, but that was made of real bone. In doing this, not only would it be possible to determine the most suitable set of parameters, but the robustness of the system would be put to the test.

Chapter 6

Bovine Distal Femur Surface Reconstruction

6.1 Introduction

Despite having similar surface shapes, different samples of different materials may elicit different reflective responses due to associated acoustic properties. As such, the success found in the previous chapter might not be reproduced whilst employing real bone. Therefore, in an effort to further the proof of concept of the imaging system, a bovine distal femur was imaged in a manner similar to that seen in the previous chapter. This sample would provide similar acoustic properties to real human bone, but would present complications in the form of changes in both scale and shape. However, these changes offered the opportunity to test the robustness of the system.

This chapter begins with a description of the methods employed for preparation of the bovine sample. The experimental setup employed for scanning of the sample is then described, after which data processing techniques are discussed. The surface point cloud and mesh reconstruction processes are then explained and the results are compared to ground truth reference models.

6.2 Bovine Sample Preparation

In order to recreate the experiment described in Chapter 5 with a bovine sample, the bone had to undergo a complete removal of the surrounding soft tissue. Defleshing such that soft tissue is removed and only bare bone remains is often employed by forensic pathologists (Rennick, Fenton, & Foran, 2005) and by bone preservationists, but is a rarely the subject of research. As such, the body of literature describing possible

techniques is small. However, three methods have been described by a number of authors (Bandini et al., 2013; Rennick et al., 2005): maceration, degradation by cooking and chemical removal. However, the third method was eliminated from use, due to the possibility of the destruction of the cortical layer (Rennick et al., 2005).

The first method employed was degradation by maceration. A stripped bovine leg was acquired from a local butcher. The leg was disjuncted and the distal femur was sawn away from the rest of the leg approximately 20cm from the condyles. Any large pieces of soft tissue remaining on the sample were removed using a scalpel. The scheme followed that suggested by Bandini et al. (2013). However, after two months, the soft tissue was only partially degraded, suggesting soft tissue destruction would require an impractical amount of time. As such, this method was abandoned in favour of the cooking method.

A second bovine leg was acquired and was prepared in the same way as described above. Once the distal femur had been defleshed with a scalpel, it was placed in a large pot and immersed in tap water. The pot was placed on an electric heating plate and was heated until boiling point. The water was changed once and the loosened soft tissue being removed at the same time. After approximately six hours, almost all of the soft tissue had been removed. While a small amount of soft tissue remained at the tendon attachment sites (Figure 6.1), small holes began to appear in the surface of the sample, possibly caused by pressure increases within the bone, which were, in turn, caused by temperature increases. Additionally, the bottom of the sample, which was in contact with the surface of the pot, had started to denature. In order to stop any further degradation of the osseous tissue, the sample was removed and dried overnight at room temperature.

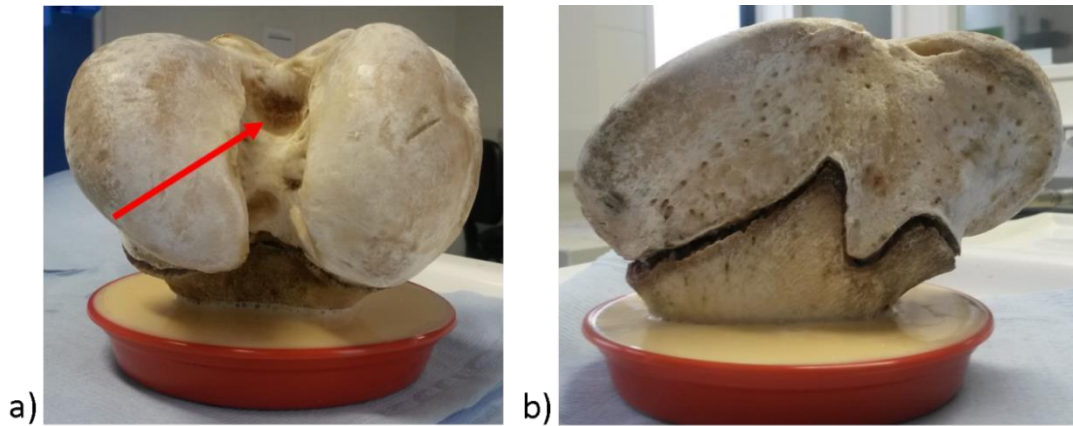


Figure 6.1: Front view of the bovine sample (a), showing a tendon attachment site with non-osseous tissue still intact. Also shown is a side view of the sample (b), showing both small holes and the separation of the shaft and the epiphysis.

After drying, the sample was set in an epoxy resin, so as to provide a base for stability. A fast cast polyurethane casting resin (Easy Composites, Stoke-on-Trent, UK) was employed for this purpose, providing a hard resin capable of holding the weight of the sample. It should also be mentioned that after drying, the distal epiphysis became detached from the diaphysis, as can be seen in Figure 6.2. In humans, before adulthood, the epiphysis and femoral shaft are joined by connective cartilage, which then ossifies once adulthood has been reached. While the age of the sacrificed cow was unknown, the observed effect implied that it must have been in young adulthood, meaning the two parts would yet to have developed into solid bone. To compensate for this, the two parts were joined using the same epoxy resin used for the base. While this created a single piece, there was a large crack running around the sample, as can be seen in Figure 6.1(b).



Figure 6.2: The separation of the femoral shaft and the distal epiphysis, with the remains of the connective cartilage visible on both portions.

6.3 Bovine Distal Femur Surface Imaging and Reconstruction

The purpose of the sample preparation process was to modify the characteristics of the bone, such that the experiment carried out in Chapter 5 could be recreated as closely as possible with a real bone sample. In doing this, the performance of the imaging system could be tested with more realistic and arguably more challenging sample features using the same experimental procedure as employed in a successful scan.

6.3.1 Materials and Methods

In order to provide a reliable ground truth against which to judge any results, the bovine sample was laser scanned using Faro Quantum (Faro Technologies, FL, USA) in the same way as described in Section 4.1. To maintain rigidity during surface capture, the base of the sample was secured to an optical table using four screws. The topology of the bone allowed for diffuse optical reflection – a feature necessary for operation of

the laser scanner. As such, it was not necessary to apply any spray paint to the surface, as was the case in the previously.

Only the cuboidal scan path was employed during this experiment, as shown in Figure 6.3. The hemispherical path was abandoned for two reasons: firstly, the practical limitations of the experimental setup made it very difficult to implement the hemispherical path. The bovine sample was of a much larger scale, which would have required changes in the controlling KRL code. In addition to this, the sample sat much lower in the water bath than was the case with the epoxy resin sample. Therefore, as the sides of the tank were set at a gradient, the distance between the sample and the sides was smaller, reducing the room in which the probe could be manoeuvred. Secondly, the results for the hemispherical scan were poor in the previous chapter, suggesting the same could be true with the bovine sample. As such, the cuboidal scan alone was seen as adequate.

However, as the sample was larger than the width of the probe, additional paths were included in the topside passes. As shown in Figure 6.3(b), the two extra passes were added, which allowed for full coverage of the top of the sample. While this increased coverage, it also increased scan time, doubling the length of time for the top side scan.




Figure 6.3: The path employed for inspection of the bovine distal femur, showing the six divisions of code execution.

While the same path shape was employed as in the previous chapter, instead of the robot being jogged, the path was implemented using KRL. This was beneficial in that

it increased the level of automation and, therefore, decreased the intra- and interoperator variability. The code was executed in six parts, as shown in Figure 6.3, using simple LIN commands. Four input parameters were required for the KRL script to operate correctly: rough estimates of the starting height, length, width and depth of the regions of inspection. The starting height could be found either by jogging the probe to the desired start position and noting the coordinates or by measuring the height of the start point relative to the bottom of the sample. The other three parameters were found simply by measuring the largest dimensions of the sample with a ruler. The code took these parameters and added a clearance of 15mm to ensure the probe did not make contact with the sample. The data captures associated with each of the path executions were stored separately, allowing for the entire scan path to be executed without exhausting the hard drive on the controlling computer.

The sample was placed in a water bath and fully submerged under tap water, which acted as the coupling medium between the sample and the probe face. Two full scans were executed, with one executing the KRL code at 1% in T1 mode and the other at 3%. The purpose of this was to determine whether accuracy could be maintained with a lower image count, which would, in turn, decrease the necessary scan time.

The same Labview script as was used previously was employed, allowing for FMC files to be captured and written directly to disk, with single precision storage. All 128 elements of the 5MHz linear probe were employed and an imaging depth of 11cm was defined along with an image width of 13.44cm (1.5 times the length of the probe). This resulted in a depth of 1152 pixels and a width of 1344 pixels. Each FMC capture was stored with the associated positional data. The same processing procedures and hardware as applied in Chapter 5 were employed.

As well as the surface contour extraction algorithm described in Chapter 5, another algorithm was additionally applied to each image. In a fashion similar to the previously described method, thresholding was applied to each image, eliminating any pixels with an intensity below -8dB for both the SAFT and -4dB for TFM images. In the previous chapter, any pixels with the same X value were discarded, with the exception of that with the smallest Z value (ie. the pixel closest to the probe face). However, in this case, the pixel with the largest Z value was retained, while the rest were discarded. The

reason for this change was that it was observed that the resulting images showed no signs of reflections from beneath the surface of the bone, as was often the case in the previous chapter. In addition to this, retaining only the pixel furthest away from the probe was thought to be beneficial, in that any side lobes were likely to appear above the surface contour. Therefore, the likelihood of extracted points being from side lobes rather than the true surface was lessened. This is exemplified in Figure 6.4, which shows a typical TFM image from the data set and the two different extraction methods. Part (b) illustrates the result of the method used in Chapter 5, with parts of the side lobes being incorrectly extracted as surface points. This problem was remedied to an extent by applying the new method, as shown in Figure 6.4 (c), where only one point has been extracted due to the side lobes.

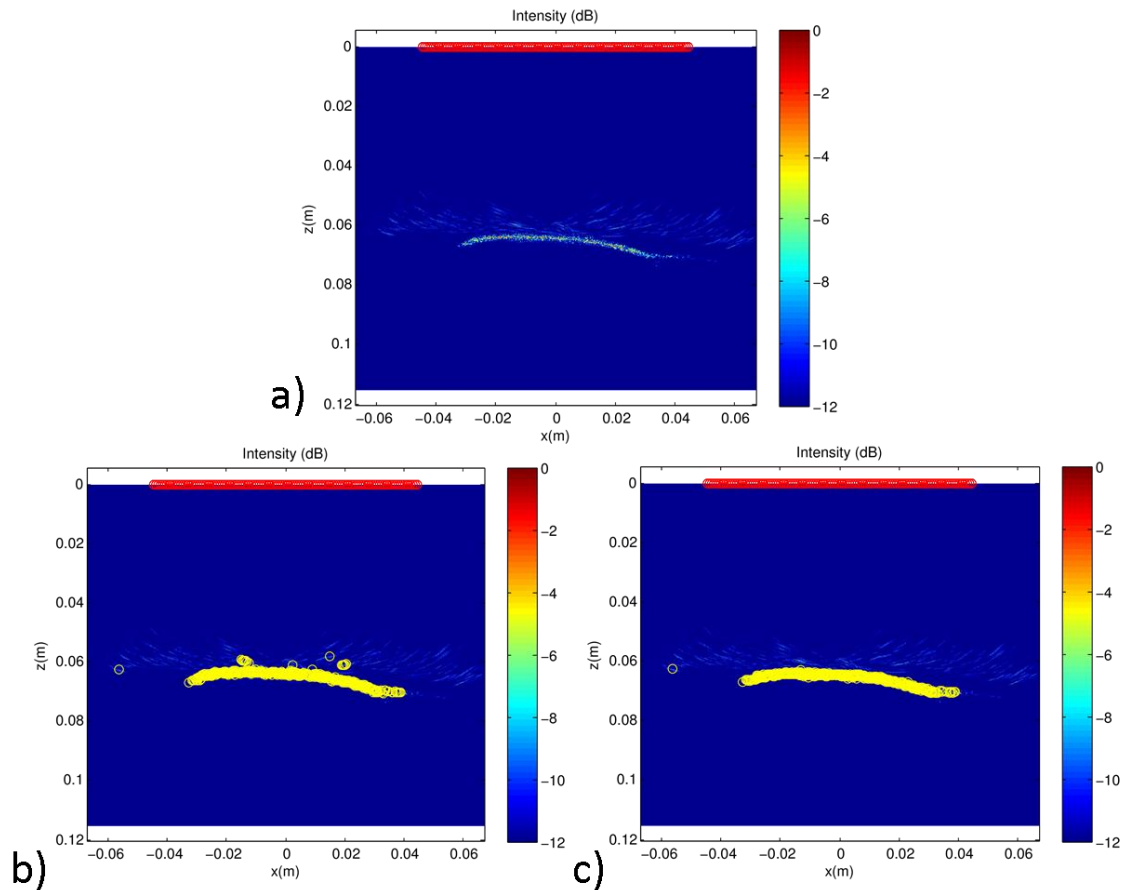


Figure 6.4: A typical TFM image from the data set (a), and the result of surface extraction using both the method described in Chapter 5 (b) - the ‘Top’ method - and that herein (c) – the ‘Bottom’ method.

For clarity, from here on, the contour extraction method described in Chapter 5 will be referred to as the ‘Top’ method. This is due to the extracted pixels being closer to the top of the image, as can be seen in part (b) of Figure 6.4. The new method described in this chapter will, on the other hand, be referred to as the ‘Bottom’ method, due to the extracted pixels being closer to the bottom of the image.

In addition to the application of the method described above, smoothing was also applied to the extracted surface contours. As in Chapter 5, the ‘smooth’ function in Matlab was used, with a moving average and a span of 0.3. Furthermore, the point elimination techniques described in Section 5.2 were applied to each resultant point cloud. Due to the size of the point clouds, random spatial down sampling was

employed before the outlier elimination algorithm was applied. For those point clouds created using a 1% velocity robot path, 90% of the points were removed using down sampling, while for 3% velocity, 50% were excluded. These values were chosen based on both the amount of processing time required to carry out the outlier elimination algorithm on the entire point cloud and on the number of points in the original cloud. At this point, each point cloud was compared to the reference laser surface scan. Using the methods described in Chapter 5, the point clouds were then converted into surface meshes and were compared to the reference surface mesh.

6.3.2 Results

As in the previous chapter, many of the techniques employed herein were formulated with surface point cloud construction, rather than surface mesh, in mind. As such, the effects could be different for the reconstruction method and the two are presented separately.

6.3.2.1 Point Cloud Comparison

Two contour extraction methods were employed to produce the surface point clouds, resulting in significant differences in performance. As shown in Figure 6.5, the ‘Top’ extraction method provided considerably lower mean errors for all SAFT-derived point clouds. This trend, however, was reversed for TFM-derived clouds. In addition to this, the use of a single top side pass resulted in a small but consistent increase in mean error, relative to the double top side pass alternative. While Figure 6.5 does not exhibit the results for all parameter combinations, the results presented therein are representative of those for other permutations.




Figure 6.5: Comparison of the mean error attained using different contour extraction methods. ‘Top’ refers to the extraction method described in Chapter 5, due to the extracted points being closer to the top of the image. ‘Bottom’ refers to the method described in Chapter 6 for the same reason. The compared point clouds employed no smoothing function and had the nearest neighbour point elimination algorithm applied.

A further parameter which was highly influential on the quality of the resulting point clouds was the velocity of the robot. For all SAFT-derived point clouds, the use of 3% velocity brought about significant increases in mean error relative to those in which 1% velocity was employed (Appendix C). The TFM-derived clouds, however, saw little change due to the velocity employed. The application of the smoothing function only showed itself to be highly influential in certain cases (Appendix C). In particular, it imposed an increase in mean error for the ‘Top’ extraction method, but only for small synthetic aperture widths.

In order to provide a more generalised view of the effects of the processing method, the mean value for the mean error, maximum error and standard deviation for every parameter combination was calculated. While the effects varied for different parameter permutations, this provided the opportunity to quickly and simply observe the influence of the processing method. As shown in Figure 6.6, the processing method had little effect on the mean error. However, this result was not likely caused by a lack

of influence of the processing method on the mean error. Instead, it was likely an averaging effect caused by the contour extraction methods, which behaved differently for different aperture widths.




Figure 6.6: Comparison of mean error, standard deviation and maximum error for the four processing methods. The values were calculated as a mean of every parameter permutation, with a general deterioration visible as aperture increases.

Based on the mean of the all the parameter permutations, it is clear that the standard deviation increased as the synthetic aperture width increased. While, Figure 6.6 presents merely a general assessment of the effect of the processing method on the standard deviation this effect was replicated for the majority of the individual parameter combinations. Much like the standard deviation, the mean maximum errors shown in Figure 6.6 displayed a dependence on the processing method, with higher aperture width being associated with a higher maximum error. In addition to the processing method, the maximum error was also influenced by the use of a single top side pass and the outlier elimination algorithm. As expected, the use of the nearest neighbour outlier elimination method resulted in a lower maximum error in all cases. This effect became more potent as the synthetic aperture width increased. The use of the single top side pass had little effect for the SAFT-derived point clouds. However,

for the TFM point clouds, the use of the single top side pass resulted in a higher maximum error.

Assessing the physical appearance of the point clouds, it can be seen in Figure 6.7 that the cloud has a number of groups of points which are of high error. These are located at the edges of a number of structures, with more accurate surface points located below. Additionally, while the point density is generally high, there are regions of particular sparsity, particularly on the surfaces of the intercondylar region.




Figure 6.7: Two viewing angles of the compared point cloud produced using 30° aperture SAFT processing, with no smoothing function at 1% velocity.

The point cloud produced using the same parameters as those in Figure 6.7, but with a 3% TCP velocity is displayed in Figure 6.8. The most obvious difference between the point cloud shown above and that below is density of points. While the cloud created using a 1% velocity has a fairly uniform distribution of points, the 3% counterpart shows very clear lines of points. Despite this difference, the regions in which the surface has little or no representation are the same in both point clouds. Additionally, the areas of high error are shown to be in the same positions. However, there appears to be a proportionally higher amount of points of high error (>5mm) in the 3% cloud

than the 1%. However, both clouds show a clear representation of the surface of the sample.




Figure 6.8: Two viewing angles of the compared point cloud produced using 30° aperture SAFT processing, with no smoothing function at 3% velocity. The density of the point cloud was decreased significantly relative to that in Figure 6.7.

The TFM-derived counterpart to the point cloud shown in Figure 6.7 is displayed in Figure 6.9. The most striking difference between the two point clouds is the point density. Due to the higher threshold limit, less of the surface contour was extracted from each image. This resulted in a significantly more sparse representation of the surface. Additionally, a group of points were produced which were very far away from the true surface. This accounted for the large maximum error of over 6cm. The other areas of high error mimicked those seen in the other point clouds. While the first two point clouds demonstrated clear and obvious representations of the surface of the sample, the TFM-derived point clouds were so sparse in certain regions that the surface was completely lost.




Figure 6.9: Two viewing angles of the compared point cloud produced using TFM processing, with no smoothing function at 1% velocity. This shows an even greater decrease in density than that seen in Figure 6.8

6.3.2.2 Mesh Comparison

Focussing on the comparisons of the fully reconstructed surface meshes, it can be seen in Figure 6.10 that for the representative group of parameter combinations, the ‘Top’ contour extraction method resulted in lower mean error for smaller synthetic aperture widths. However, as the aperture width increased, the ‘Bottom’ method provided the lowest levels of mean error. In addition to this, Figure 6.10 also displays the influence of the use of the single top side pass. For narrower aperture widths, the single pass meshes yielded smaller mean errors, while the double top side pass meshes performed better for large apertures. Further to this, the use of TFM processing proved to produce significantly higher mean errors than the SAFT-derived meshes. Finally, five of the presented parameter combinations achieved submillimetre mean errors, with, in particular, the 20° SAFT returning 0.895mm for the double top side pass and 0.877mm for the single pass.




Figure 6.10: Comparison of the mean error attained using different contour extraction methods. The compared meshes employed no smoothing function, did not have the nearest neighbour point elimination algorithm applied and employed a 1% TCP velocity.

The TCP velocity at which the data was captured also influenced the resulting surface meshes (Appendix C). For smaller synthetic aperture widths, the use of 3% velocity resulted in higher mean error, while the opposite was true for larger apertures. The influence of the application of a smoothing function to the extracted contours is also demonstrated in Appendix C. For small synthetic aperture widths, smoothing provided poorer mean errors, while for larger apertures, the opposite was true. This was the case for both contour extraction methods.

Shown in Figure 6.11 is a comparison of the three main methods of numerical assessment based on the mean values for all parameter permutations for different synthetic aperture widths. The purpose of this was to provide a general view of the effects of the aperture width across all parameters. Assessing mean error, it is clear that TFM produced the worst results. While the mean errors achieved using SAFT were all significantly better than that found with TFM, 30° SAFT yielded the best

results. However, it should be noted that the results presented in Figure 6.11 are means of all the possible parameter combinations. As such, while the 30° SAFT produced lower mean errors on average, the 20° SAFT yielded the lowest mean errors, albeit only with specific parameter combinations.




Figure 6.11: Comparison of mean error, standard deviation and maximum error for the four processing methods. The values were calculated as a mean of every parameter permutation.

The standard deviations shown in Figure 6.11 follow the same pattern as the mean errors. The only obvious difference was a greater dependence on aperture width, with a greater difference between the standard deviation at the 30° SAFT and TFM. The pattern formed by both the mean error and standard deviation was mirrored in the maximum error, as demonstrated in Figure 6.11. As further evidenced in Appendix C, for small synthetic aperture widths, there was little difference between the parameter combinations in maximum error. However, the error increased significantly with TFM. Assessing the reconstructed meshes visually, it can be seen in Figure 6.12 that the shape of the bony surface was recreated relatively accurately. The areas in which the highest level of error can be seen in the example below were common to all the meshes. These regions were found on the intercondylar surfaces, as well as underneath the

condyles. In addition to this, there are a number of artefacts located on the outer sides of the condyles. Similar artefacts were found on most of the meshes recreated using the ‘Top’ contour extraction method. However, in the best performing meshes reconstructed using the ‘Bottom’ method described in this chapter, such artefacts were not present. The crack that was present in the sample between the epiphysis and the diaphysis could also be seen in the errors, with the edges of the crack yielding errors of approximately 2mm. Finally, a line of low error is present along the front of both condyles, as can be seen in the two bottom images of Figure 6.12.




Figure 6.12: Four viewing angles of absolute errors resulting from matching 20° SAFT-derived mesh with no smoothing, no outlier elimination and using a single top side pass.

Comparing the performance of the meshes created using 3% and 1% TCP velocities, it can be seen in Figure 6.13 that the 3% mesh lacks the shape found in the 1% results. In particular, the intercondylar notch was completely lost. While this did not happen

in all 3% results, it was common. The 1% results, on the other hand, almost always preserved the intercondylar notch. Despite this very obvious difference in performance, the meshes resulting from the different TCP velocities had in common other areas of high error. For example, both meshes show regions of high error at the top of the intercondylar notch and share the same artefact on the outer surface of the left condyle.

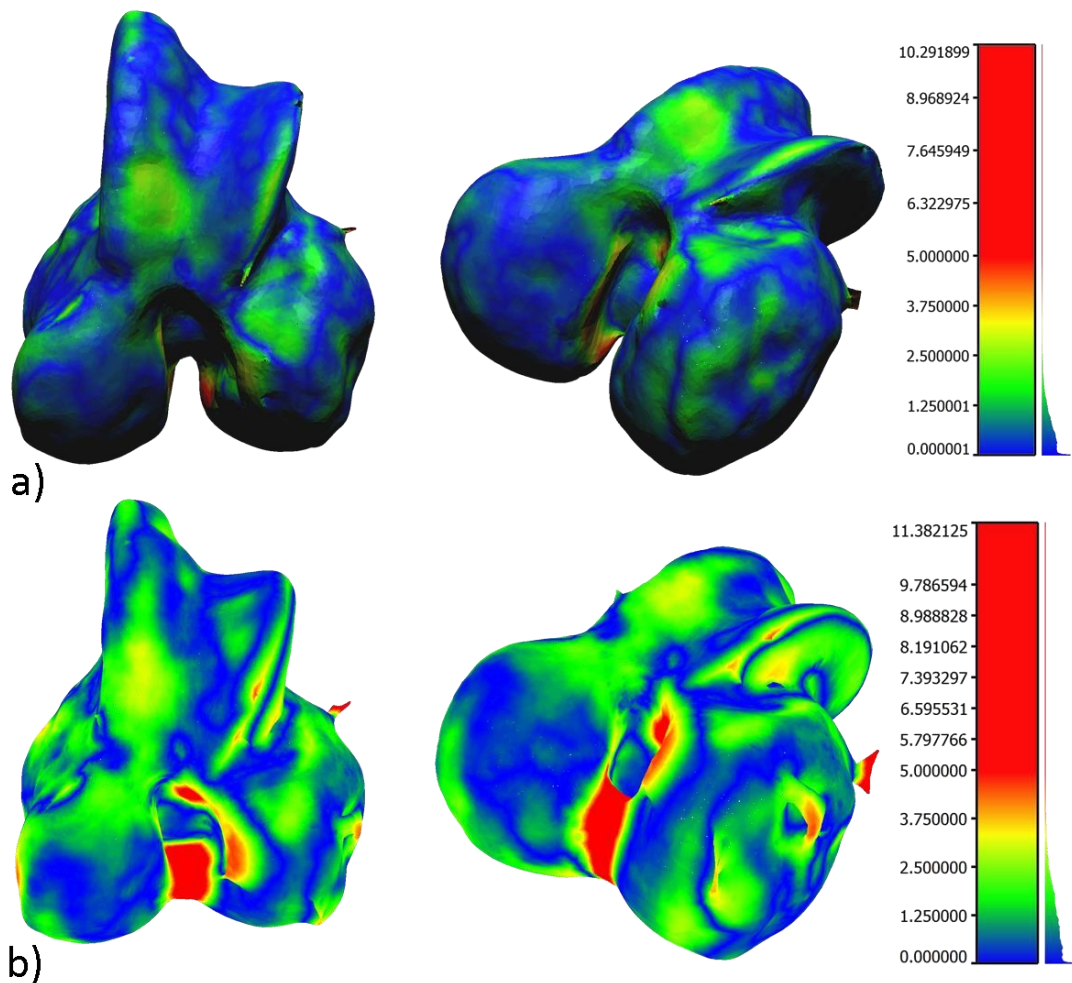


Figure 6.13: Two viewing angles of errors resulting from matching 30° SAFT-derived sampled meshes with no smoothing, no outlier elimination applied and a single top side pass. One mesh has been produced using a 1% TCP velocity (a), while the other has been created using a 3% velocity (b), showing how the decrease in point cloud density can be detrimental to the mesh reconstruction.

In order to assess the differences between the meshes resulting from the ‘Top’ and ‘Bottom’ contour extraction methods, two meshes of the same parameter definition (with the exception of the extraction method) were compared. Employing the mesh derived from the ‘Top’ extraction method as the reference and the ‘Bottom’ mesh as the comparison, Figure 6.14 demonstrates the errors found between the two. In this scenario, negative errors represent the reference mesh being outside the compared mesh, while positive numbers signify the points being inside the mesh. Therefore, it is clear that most of the surface of the ‘Top’ method mesh was outside that of the ‘Bottom’ method. In other words, the mesh resulting from the ‘Top’ method was larger than that of ‘Bottom’. This assessment was confirmed by the mean signed error of -0.813mm, implying that most of the compared surface was smaller than the reference mesh. There was only one area located in the intercondylar notch in which the error was significantly positive. This region corresponded roughly to the area of high error shown in the bottom left of Figure 6.12.

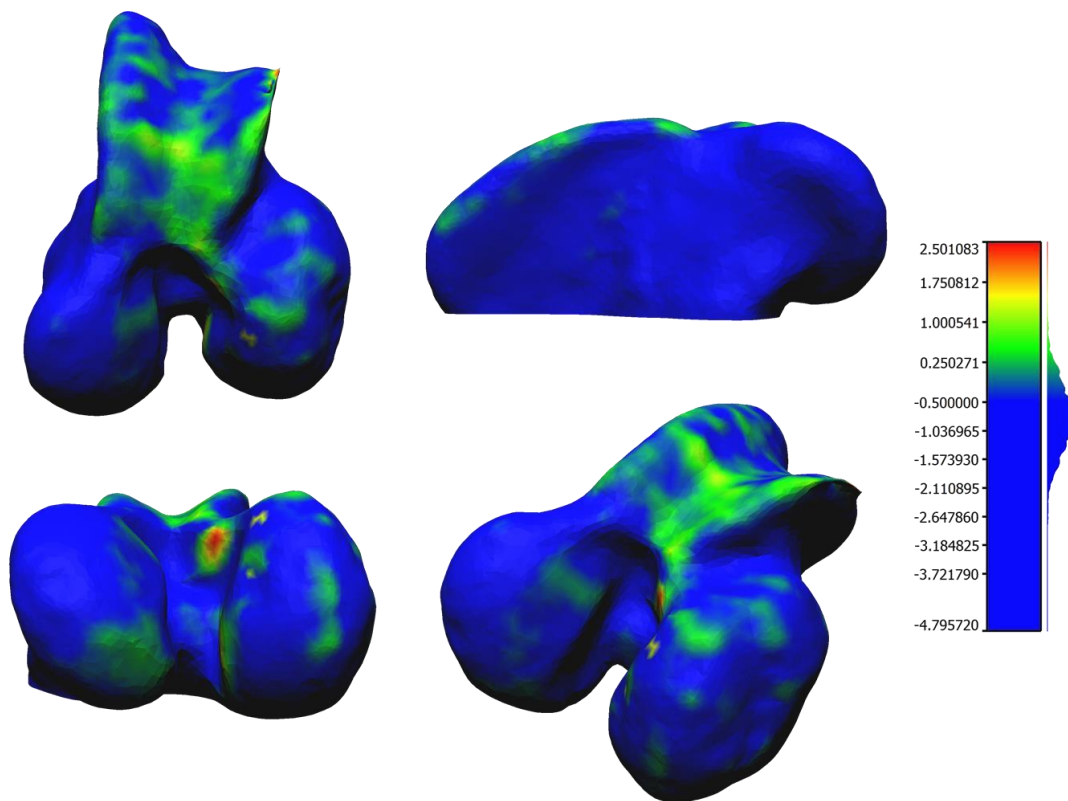


Figure 6.14: Four viewing angles of the 20° SAFT-derived mesh with no smoothing, no outlier elimination, using a single top side pass and using the 'Bottom' contour extraction method. The errors displayed are from comparison with the mesh of the same parameters, but created using the 'Top' extraction method. Negative errors denote that the reference mesh (the 'Top' mesh) was outside the surface, while positive numbers imply that the reference mesh was inside the compared mesh.

A further parameter which affected the accuracy of the final reconstructed mesh was the synthetic aperture width. Therefore, in an attempt to assess the effect of the width of the aperture, the errors of two meshes with the same parameter combination (with the exception of the aperture width) are shown in Figure 6.15. While the TFM mesh appears to have lost some of the clarity shape seen in the SAFT mesh, the shape is not, visibly speaking, particularly poor. In spite of this, the errors across the majority of the surface are much higher than those seen in the SAFT comparison. The TFM mesh, on the other hand did not suffer from the artefact on the condyle seen on all the SAFT meshes.

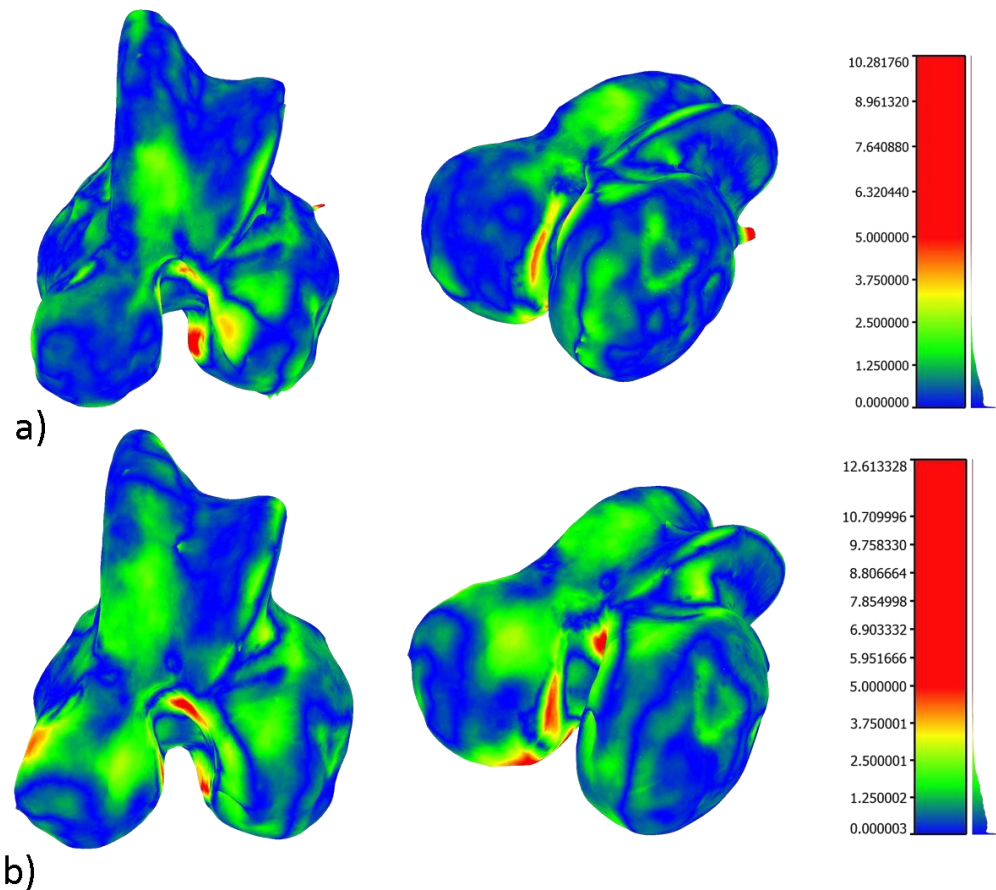


Figure 6.15: Two viewing angles of errors resulting from matching sampled meshes with no smoothing, with no outlier elimination applied and a single top side pass, created using a 1% TCP velocity. One mesh has been produced using 40° SAFT (a), while the other has been created using TFM (b).

6.3.3 Discussion

There were a greater number of parameters trialled during the reconstruction of the surface of the bovine sample than with the human. This was performed in an effort to determine the combination of parameters that would provide an accurate reconstruction of the surface, not only for one particular sample, but for significantly different samples.

The first of these new parameters was the different contour extraction methods. The ‘Top’ method provided the most accurate results, but was highly dependent on other

parameters. As such, it also produced the worst results, with the ‘Bottom’ method showing less variation over different parameter permutations. A possible reason why the ‘Top’ method produced significantly poorer results in some circumstances was that it was much more likely to allow the extraction of side lobes. This effect is visualised in Figure 6.4, which shows how the ‘Bottom’ method minimised the influence of side lobes. The effect of this in the final meshes came in the form of the small artefacts seen in Figure 6.13 and Figure 6.12.

In addition to these points, a very obvious effect was the general difference in volume. As demonstrated in Figure 6.14, the ‘Top’ method resulted in larger meshes, stemming from the fact that the contour extracted from each image would be closer to the probe face. This reason this occurred was due to certain contours having a thickness associated with them. Such features usually appeared when the probe face was not parallel to the surface. An example of this situation is displayed in Figure 6.16, which shows the extracted points from the two methods. While both sets of points are extracted from within the surface representation, the ‘Top’ points (shown in yellow) are all closer to the probe surface. As such, the resulting mesh would be of greater volume than that created using the ‘Bottom’ method (shown as red points). While the ‘Top’ method showed a high dependence on the other parameters, it allowed for the best results in terms of mean error. As such, it is considered to be the superior method.




Figure 6.16: A typical 20° SAFT image from the bovine scan. The red points show the points extracted using the ‘Bottom’ method, while the yellow show that for the ‘Top’ method. As the ‘Top’ points are higher, this can result in a larger overall surface reconstruction.

Another of the parameters that was not varied during the surface capture for the human sample was the TCP velocity. The meshes resulting from 1% velocity performed significantly better than those using a 3% velocity. The reason for this was that the point clouds resulting from the 1% TCP velocity had a much more uniform distribution of points, as can be seen when comparing Figure 6.7 and Figure 6.8. For a number of meshes, the lack of representation resulted in gross deformation. An example of this can be seen in Figure 6.13, in which the intercondylar notch was completely lost due to lack of representation in the point cloud. With these points in mind, the use of the 3% TCP velocity could not be justified, due to the severe decrease in accuracy, despite the considerable reduction in both data capture and processing time.

The final parameter new to Chapter 6 was whether or not the probe took a single pass each way across the top of the sample, or increased coverage and took two passes. Use of the single pass performed slightly better, providing lower mean errors. The reason the results were so similar was that the point clouds produced with the different capture methods were very similar in terms of both point density and general coverage, as can be seen in Figure 6.17. The slight difference in performance can be attributed to the fact that the clouds created using a double pass were denser in the regions which lacked perpendicularity relative to the probe face. This can be seen in Figure 6.17 (b), in which the point density increases at the edges of the cloud, where the gradient increases dramatically. It is in these areas, where the surface is far from parallel to the probe that erroneous point extraction is likely to happen.




Figure 6.17: Two point clouds resulting from 20° SAFT processing, 1% TCP velocity, no smoothing, with one using a single top side pass (a) and the other using a double (b), showing an increase in density in regions of high error.

A further parameter which influenced the final results was the application of a smoothing function to the extracted contours. The general outcome of the use of the smoothing function was an increase in mean error for small synthetic aperture widths and a decrease for large apertures. Therefore, if a 20° SAFT image was being used, smoothing should not be applied, but the opposite would be true for a TFM image. However, this is irrelevant, as the TFM-derived meshes were not capable of attaining sub-millimetre mean error, meaning only smaller aperture widths would be employed. As such, a smoothing function should not be applied.

The final parameter which was tested for the bovine data set was the outlier elimination algorithm, which employed a nearest neighbour method. When employed with the point clouds, the algorithm resulted in improved results in all cases, with particular success in reduction of the maximum error. However, this success was not translated into the mesh results, with it resulting in deterioration in results in almost all cases. The reason for this is unclear, but the outlying points may have helped to offset other points which were erroneous, but more accurate than the eliminated points. In any case, while the nearest neighbour elimination algorithm resulted in improvements in the

point clouds, it resulted in increases in error for the associated meshes and, as such, should not be employed.

Many of the parameters described above were designed specifically to improve the accuracy of the point clouds. However, it is the fully reconstructed surface meshes which are of importance, as they would be employed in the planning procedure. Therefore, while the application of certain parameters resulted in point cloud improvements, the same parameters were of detriment to the associated mesh. As such, it was important to fully test the effect of the parameters on both the point clouds and the final meshes. This allowed for the identification of the most suitable parameter combinations.

In addition to the parameters stated above, there were other causes of error on the final reconstructed surfaces in the form of an incomplete reference model. As can be seen in Figure 6.18, there were a number of areas of high error typically found in the compared sampled meshes that corresponded directly to the regions which were incomplete in the reference model. There were a number of reasons for this lack of completeness, the first of which was the inability to access certain areas due to the position of the sample on the optical table during surface capture. As can be seen in Figure 6.19 (b), the underside of the condyles were close to the surface of the table and, as such, were impossible to access with a direct line of sight using the laser. The result of this can be seen in Figure 6.18 (b) and (a), which show the incomplete parts of the reference mesh and the resulting inaccuracy in the compared mesh, respectively. Line of sight access was also a problem in the sides of the intercondylar region, as can be seen in Figure 6.18 (c-f). Again, the laser was unable to access these areas, resulting in high levels of error in the final compared meshes.

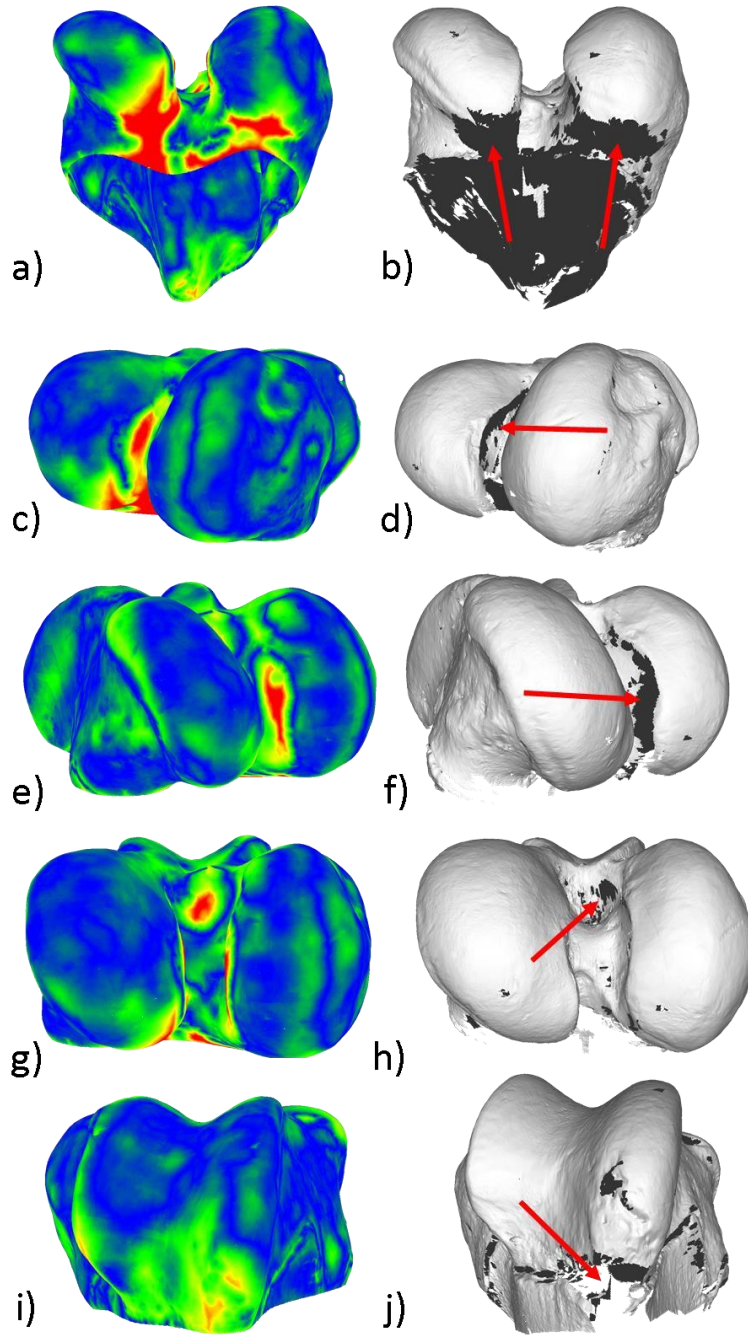


Figure 6.18: Typical regions of high error displayed in red (a, c, e, g) and the accompanying incomplete area in the reference laser scan data.

There were, additionally, problems with the surfaces themselves which made interaction with the laser more difficult. The problems presented in Figure 6.18 (g) and (h) were caused by the presence of some soft tissue. This area in particular was the

attachment site of a tendon, the remnants of which were not fully removed by the boiling process and were difficult to remove by scalpel. Given that the surface was not highly reflective like bone, the laser would likely have been absorbed, leading to a lack of representation of that area. As well as this, the bone itself had undergone degradation due to the boiling process. The most obvious effect of this was the large crack which ran round the entire sample. This, however, did not cause many problems. More problems were, however, caused by the degradation of an area of bone at the back of the sample, as can be seen in Figure 6.19 (a). In this area, the bone had softened and did not provide a continuous surface on which the laser light could easily reflect. The effects of this on both the reference model and the errors in the final mesh comparison can be seen in Figure 6.18 (j) and (i), respectively.

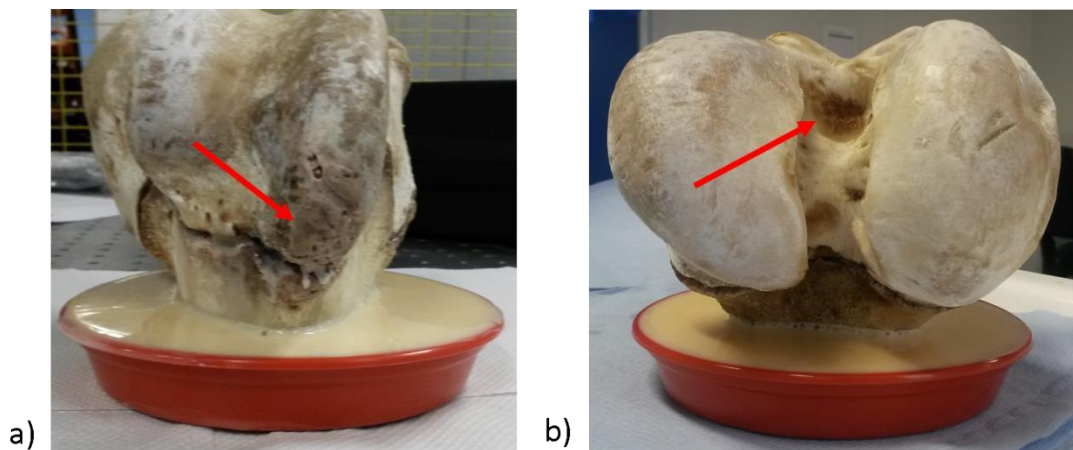


Figure 6.19: Causes of incompleteness in the laser scan data denoted with a red arrow. One section shows general degradation of the osseous tissue caused by boiling (a), while the other shows remnants of soft tissue (b).

It is a reasonable assumption that if the reference model had been complete, particularly in the areas in which line of sight access was impossible, the results would have improved significantly. It is clear from the results presented above that the mean errors, standard deviations and maximum errors were of a poorer standard than was achieved with the artificial human sample. While part of this may be related to the difference in shape, it is believed that much of the difference in performance arose

from areas in which the reference model was incomplete, which led to large regions of high error.

There were a number of purposes of this study, the first of which was to ensure that the acoustic properties of bone were such that synthetic aperture methods would allow for a faithful reconstruction of the surfaces. The images produced using data from the bovine scan appeared, in fact, to be superior to the artificial human distal femur studied in Chapter 5. The main improvement in the individual images came in the form of a lack of reflections from beneath the surface of the sample. This was possible due to the large impedance mismatch between water and cortical bone. While the acoustic properties of the composite were unknown, the acoustic impedance was likely not as high as that of cortical bone, implying a better chance of ultrasonic energy penetrating the outer surface of the composite sample.

The second purpose of the study was to determine the most robust set of parameters, so that significantly different shapes could be reconstructed to a similar accuracy. The parameter combinations with the five lowest mean errors for both the composite and bovine studies are displayed below in Table 6.1 and Table 6.2. The most noticeable common factor is the use of SAFT processing. From these results, it is clear that TFM would not be capable of achieving errors anywhere near the desired range. In saying this, in the studies were carried out in very specific conditions, with only one probe type being employed and only bare surfaces being examined. In practical application, different probe types might be employed and, more importantly, the bony surfaces would be covered in a number of soft tissues – factors which might influence the performance of different data processing methods.

Table 6.1: The parameters and results from the five smallest mean errors for the composite sample.

Processing Method	Nearest Neighbour?	Smoothing?	Mean Error (mm)	Standard Deviation (mm)	Maximum Error (mm)
20° SAFT	no	no	0.8	0.6	7.6
30° SAFT	no	yes	0.8	0.7	7.6
20° SAFT	yes	no	0.8	0.6	7.6

40° SAFT	no	yes	0.9	0.7	7.7
40° SAFT	no	no	0.9	0.7	7.7

Table 6.2: The parameters and results from the five smallest mean errors for the bovine sample. All values were achieved using a 1% TCP velocity and the ‘Top’ contour extraction method.

Processing Method	Single Pass?	Nearest Neighbour?	Smoothing?	Mean Error (mm)	Standard Deviation (mm)	Maximum Error (mm)
20° SAFT	yes	no	no	0.9	1.0	10.5
20° SAFT	no	no	no	0.9	1.0	10.5
30° SAFT	yes	no	no	0.9	1.0	10.2
40° SAFT	yes	no	no	0.9	1.0	10.3
40° SAFT	no	yes	no	0.9	1.0	10.4

The bovine results had three extra parameters which could influence the results: the TCP velocity, the number of top side passes and the contour extraction method. It would be advantageous to use the single pass and 3% TCP velocity, given the reduction in data capture and processing time offered by these choices. However, it was clear that the vast reduction in coverage given by the use of the 3% velocity would not allow for the desired level of accuracy in the surface reconstruction. Therefore, with the probe employed herein, a 1% TCP velocity would be suggested. The use of a single pass, on the other hand, provided a more ambiguous outcome. It was expected that the reduction in coverage associated with the use of the single pass would reduce the in accuracy. With this in mind, the purpose of the use of the single pass was to determine whether or not the reduction in accuracy implied with the use of a single top side pass was of an acceptably low level. However, the results showed that the single pass offered the most accurate results. The outcome of this was that an increase in the

coverage of areas which include steep angles (through a second pass), which result in a lack of parallelism relative to the probe face, will likely lead to inaccuracy. Finally, the ‘Top’ contour extraction method was superior to the ‘Bottom’, as it proved to be capable of significantly lower mean errors. As such, it was shown that the ‘Top’ contour extraction method should be employed over the ‘Bottom’.

The parameters under which the lowest mean errors were achieved were the same for both the bovine and human shaped samples. While these permutations did not achieve the lowest standard deviation and maximum error in both cases, these values were not significantly higher than the lowest achieved. As such, it was concluded that the most robust parameter combination was found to be a 20° SAFT with no smoothing and no outlier elimination. Additionally, uniform coverage was found to be important with the most accurate results being found while using a 1% TCP velocity in both the human and bovine cases.

6.4 Conclusion

The bare surface of a real bovine distal femur was imaged using two different scan paths. Employing the captured data, surface point clouds were produced, which were then compared to a reference model. These point clouds were then used to create representative surface meshes. Using a number of parameter combinations, the meshes were compared to a reference model derived from a surface laser scan. This allowed for analysis of the parameters which contributed to the accuracy of the final surface mesh. Using the results from the scans of both real bovine and composite, human-shaped distal femurs, the parameters which allowed for the most robust surface reconstruction were found.

While it was known the system was capable of reconstructing the surface of the human-shaped sample to the desired accuracy, it was not yet known if the same could be achieved for real bone or for significantly different shapes. This study found that this was, in fact, possible, albeit with slightly lower levels of accuracy. While the system could comfortably reconstruct the surface of the composite sample to submillimetre

accuracy using a variety of parameter combinations, this could only be achieved for the bovine sample with a limited number of parameter combinations. The fact that the system was capable of handling different shapes is important, in that knees in the late stages of OA tend to have undergone bone remodelling (Pritzker et al., 2006). As such, the surface shapes found in the populations likely to benefit from knee arthroplasty would likely differ significantly from patient to patient. Therefore, it would be a requirement of an accurate imaging system used for pre-operative imaging for robotic knee arthroplasty that it be capable of accurately reconstructing significantly different shapes.

The study was, therefore, considered to be a success, with the mean errors being comfortably below the desired limit and the system being proven to be capable of accurately reconstructing both bony and highly different shaped surfaces. However, the conditions under which these results were achieved remained unrealistic. The knee joint *in situ* would have the added complications of both soft tissue and line of sight problems. As such, the next stage in this proof of concept should be to include soft tissue with the bone. In doing this, the consequences of the inclusion of soft tissue could be assessed and any detrimental effects could be tackled and alleviated.

Chapter 7 Bone Surface Imaging with Soft Tissue – a Modelling Approach

7.1 Introduction

As of yet, only bare surfaces have been imaged, both experimentally and through simulation. While this has provided a fundamental proof of concept, there remains a significant lack of realism. The goal of this body of research is to image knee joints *in*

situ, which implies two major associated difficulties: line of sight issues and soft tissue penetration problems. Soft tissues pose significant problems within biomedical bone surface imaging (Hacihaliloglu, Rasoulian, & Rohling, 2014; Vincent, Jocelyne, & Tonetti, 2004), but these have yet to be tackled using synthetic aperture imaging methods.

With this in mind, this chapter presents an investigation through which the effects of soft tissue were studied. To achieve this, FMC simulations were carried out using PZFlex, providing understanding of the ability to image intact, human knee joints.

7.2 Finite Element Modelling of Soft Tissue Interfaces

Given the success of the previous experiments, the required improvement in methods was a higher level of realism. In particular, this would include multiple soft tissue layers, bones of the correct shape and scale and intact joints. The best way to achieve this would be to carry out full scans on cadaveric specimens. This, however, could not be achieved, as in order to employ the KUKA robot as the positioning system, the experiment would have to be carried out the FIRST laboratory. As this area was not under licence to employ human specimens, it would have been illegal to carry out such work. To circumvent this limitation, simulation could be employed, which would allow for a number of probe parameters and positions to be tested on a single sample *in silico*. To this end, PZFlex (Weidlinger Associates Inc., CA, USA), a Finite Element (FE) modelling platform, was employed to simulate different biological materials and their properties, as well as FMC using representative probe parameters. This was carried out for two samples, with different probe parameters also being modelled, providing an estimation of the performance of FMC and synthetic aperture methods in fully intact knee joints.

7.3 Materials and Methods

PZFlex is a commercial FE analysis program created specifically for ultrasonic applications. Originally employed solely as a means of modelling ultrasound probes, the software has grown to incorporate a range of facilities, including wave propagation (Gachagan, McNab, & Reynolds, 2004) and simulation of thermal effects (C. Wang, Gachagan, & Leary, 2011; G Wojcik et al., 1995). All of these have since been employed in a variety of applications. In particular, the platform has been employed for wave propagation in biological tissue (G Wojcik, Jr, & Carcione, 1999) , using tissue maps which allow material properties to be assigned to each pixel in the image. For the purposes of probe design, geometric models are often implemented using scripting. This method allows for the creation of simple geometric shapes which are approximated on a defined Cartesian grid, with each element on the grid being assigned a set of material properties. The same can be said of wave propagation simulations involving simple samples which only require rudimentary shapes for definition. When defining more complex models, such as those involving biological materials, scripting would not be possible. In answer to this, PZFlex offers an image importation tool, which allows for the assignment of material properties to each pixel based on the associated colour or intensity.

In order to do this, images representative of a human knee were required. Cross sectional images acquired using either CT or MRI are such that different materials can be in the same intensity range. It can, therefore, be difficult for a non-expert to identify different structures and materials. To overcome this, MRI images were provided by Christian Paulus (Medical Media Images), which have been coloured manually by a physician medical illustrator. This allowed for simple identification of different material types, whilst retaining an accurate shape. These images can be seen in Figure 7.1, with one image showing the coronal plane and the other the sagittal. In part (a), fat is shown in white, muscle in dark red, trabecular bone in weak yellow, meniscus in blue, articular cartilage in green, tendon in bright yellow and cortical bone in a very dark red. The same colour assignment was employed in the second image, with the exception of trabecular bone, which was shown as orange. These images were chosen

in particular, firstly, because of their availability and, secondly, because they offered two different but representative views of the knee joint.




Figure 7.1: Custom coloured MRI images showing a coronal (a) and sagittal (b) cross section.

The image importation tool could not directly segment the images in the form displayed in Figure 7.1, as the colours representing different materials were not consistent throughout the image. In the bottom right of part (b) of Figure 7.1, for example, it can be seen that fatty tissue is represented by bright white. However, towards the top of the image the colours become darker to the extent that the colour matches that of muscle in the brighter part of the image. Additionally, boundaries between different materials were not abrupt, but instead showed colour gradients containing a range of colours. As such, numerous materials would be assigned to the small number of pixels between each obvious colour. For the same reasons, simple segmentation tools in Matlab were unable to accurately distinguish between different materials.

While it would likely have been possible to develop an automatic segmentation tool in Matlab which would have performed reasonably well, it was decided, given the shortness of time and the small number of samples, that the image should be manually segmented. This was performed in Microsoft Office Publisher (Microsoft Corporation,

WA, USA) by tracing the presumed interfaces between different materials. Initially, the thickness of the cortical layer was such that the lines representing it crossed – essentially eliminating the cortical layer from the image. To remedy this, the cortical layer was artificially thickened, ensuring it was fully encompassed the trabecular bone. Additionally, the location of the skin layer was not obvious in the images. To ensure the inclusion thereof, a thin layer was added to the outer materials based on a colour gradient between the outer surface and the background black. The results of this process are displayed in Figure 7.2.




Figure 7.2: Segmented images of both samples, with materials separated by lines.

The resulting images were then opened in Microsoft Paint (Microsoft Corporation, WA, USA), where the closed objects created by the lines were filled with different colours. The images resulting from this contained black lines between each coloured material. This was corrected in Matlab by identifying every pixel assigned black, finding, on an individual basis, the nearest non-black pixel and assigning the same colour to the pixel in question. Based on the associated colour, each pixel was then assigned a simple number, in order to make material assignment simpler, the result of which can be seen in parts (a) and (b) of Figure 7.3. These were then imported to the image importation tool in PZFlex, where each pixel was assigned material properties based on the associated intensity. The resulting tissue maps can be seen in parts (c) and

(d) of Figure 7.3, with water in dark green, skin in dark blue, fat in orange, muscle in dark purple, cortical bone in green, meniscus in yellow, articular cartilage in light purple, tendon in light blue and trabecular bone in red.

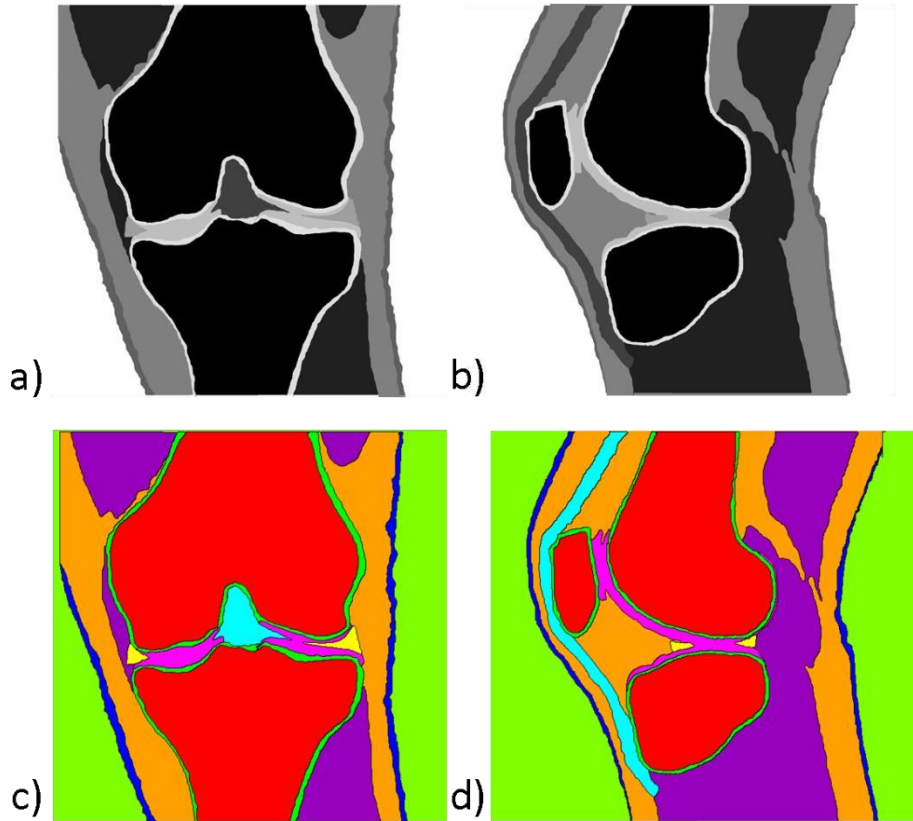


Figure 7.3: Final segmented images before input to PZFlex (a, b) and after material assignment in the image importation tool (c, d).

To define a material to be employed within PZFlex, the minimum requirement for known material properties is the longitudinal velocity, shear velocity and density. While a number of other characteristics can be applied, including permittivity and damping properties, these were excluded due to lack of available known quantities. Liquids do not have a shear velocity component (Halmshaw, 1991), so that of water was set to zero. Additionally, while soft tissues do have a shear component, they are often as low as 1% of the longitudinal velocity and, as such, are usually omitted in FE analysis, essentially treating soft tissue as a fluid (Reynolds, 2005). Only cortical bone, therefore, included a shear wave velocity, defined as 2800ms^{-1} (Haim Azhari, 2010).

The other values and the associated sources are listed below in Table 7.1. Meniscus and articular cartilage were assigned the same material properties as they are both cartilaginous tissues. The term ‘PZFlex’ in the sources column indicates that the values were taken from the materials list supplied with PZFlex. The sources for these values are not stated specifically, but the materials file does list a number of generic sources, which would suggest the properties are, at least, in the correct range.

Table 7.1: Material properties and the associated sources as employed in the PZFlex model.

Material	Longitudinal Velocity (ms⁻¹)	Source	Density (kgm⁻³)	Source
Trabecular Bone	1827	(Trebacz & Natali, 1999)	1120	(Trebacz & Natali, 1999)
Cortical Bone	4080	(Haim Azhari, 2010)	1900	(Haim Azhari, 2010)
Muscle	1571	PZFlex	1041	PZFlex
Tendon	1650	(Garcia, Hornof, & Insana, 2003)	1600	(Hashemi, Chandrashekar, & Slauterbeck, 2005)
Skin	1605	(Rigal et al., 1989)	1020	(Liang & Boppert, 2010)
Fat	1427	PZFlex	928	PZFlex
Meniscus	1600	(Kohles et al., 2012)	1330	(Kohles et al., 2012)
Articular Cartilage	1600	(Kohles et al., 2012)	1330	(Kohles et al., 2012)
Water	1496	PZFlex	1000	PZFlex

It should be noted that only the longitudinal velocity of trabecular bone has been listed. The use of one longitudinal velocity throughout all the trabecular bone would suggest that it is an isotropic material. While the compact cellular structure of cortical bone allows it to be thought of as a homogeneous material relative to the wavelengths employed in medical ultrasound imaging (Clarke, Evans, Truscott, Milner, & Smith, 1994; Talmant et al., 2011), trabecular bone is considered a heterogeneous material

with a complex, porous structure (Culjat, Goldenberg, Tewari, & Singh, 2010). It, therefore, contains numerous materials, each with different acoustic properties. In order to model trabecular bone, therefore, the image itself would have to be modified to include the porous structure of trabecular bone. Given time limitations and the fact that it was expected that almost all of the wave energy would be reflected by the cortical layer, it was decided to simply employ a single velocity for trabecular bone.

The image importation tool outputs a text file which has two elements: firstly, a list of coordinates, representing the position of each pixel. Secondly, a list of numbers, each representing a particular material and each associated with a particular coordinate. The tool automatically generates the pixel coordinates, whilst maintaining the proportional dimensions defined by the resolution of the image. As the coordinates were not in the correct scale, a new coordinate list was produced using average bone sizes. Two dimensions were assumed in order to do this: firstly, the maximum width of the femoral condyles from the lateral to medial side and, secondly, the maximum dimension of the medial condyle. These are shown in parts (a) and (b) of Figure 7.4, respectively. The mean reported values for these dimensions for adult males (Seedhom, Longton, Wright, & Dowson, 1972) were employed to calculate the width of the entire image, based on the number of pixels within the described dimensions. Using this, the coordinates were altered, such that the scale was correct.

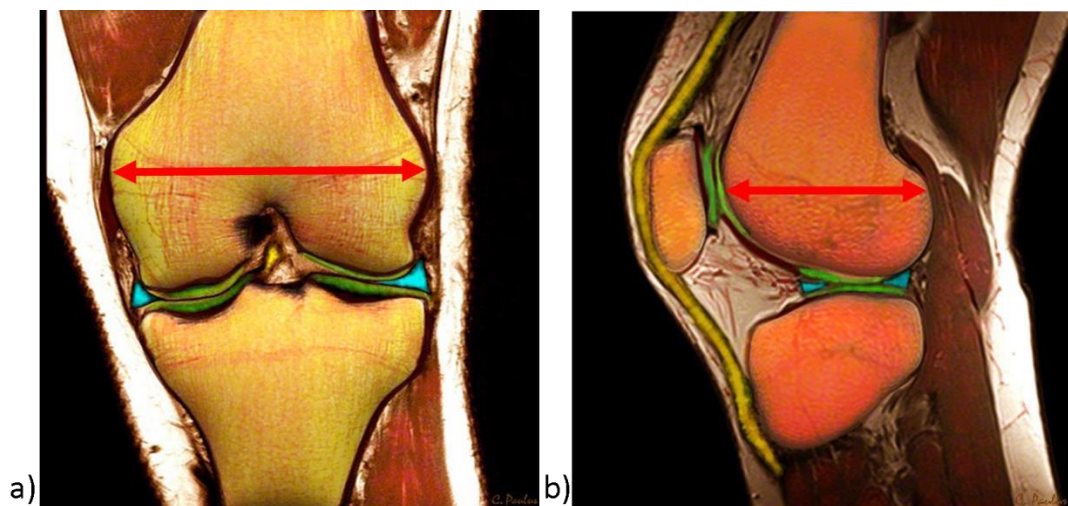


Figure 7.4: The dimensions used to scale the pixel coordinates in the materials table; the maximum width of the femoral condyle (a) and that of the medial condyle (b).

Once the tissue map was in place, the probe was modelled. Rather than simulating all the individual components of the probe and applying a drive function to each transmitting element, the probe elements were represented by a line of finite elements, to which a pressure load was applied. This decreased the amount of time required to model the probe and decreased the complexity of the whole model, speeding up the solution time. The pressure load could have been measured and imported via text file, but in order to simplify the process, a wavelet was employed.

This was achieved using the 'wavelet' function, which requires only the centre frequency of the transducer. The receiving elements were defined in the same way, as a line of elements defined by the start and end coordinates, which were, in turn, based on real probe properties. An average pressure was calculated for each receiving probe elements by taking the mean of the associated FE elements, with the representative time-amplitude signals of each element being stored. Each transmission was carried out in the same way it would be performed in reality, with one element firing and every element receiving. In order to fully recreate FMC, the batch controller found in PZFlex was employed which allowed the position of the transmitting element to be changed. Each batch run resulted in a Matlab data file containing a 2D matrix, which contained the time-amplitude signals of every receiving element. These were then combined, so that a full, useable FMC file was produced.

Multiple simulations were carried out on each image with different probe positions, so that the joints could be imaged fully. For the image shown in Figure 7.1 (a), only two probe positions were employed, as this was believed to fully represent the joint. These positions can be seen in parts (a) and (b) of Figure 7.5, where the probe is shown as a grey bar. As can be seen in part (a), a third of the image was removed. This was carried out so that the number of finite elements in the model would be reduced and, therefore, the solution would be provided more quickly. This was also carried out with the second image, with only one example employing the full image. It can also be seen that the thickness of the probe changed between images. This was performed simply to provide more space between the probe face and the skin surface and made no difference to the data capture method. Finally, it can be seen in Figure 7.5 that additional probe positions

were employed in the second image. The reason for this was that the second image represented larger physical objects than the first and required more data captures to fully image the bony surface.




Figure 7.5: Every probe position (grey rectangle) and tissue map employed for the 2.25MHz probe simulations. Parts (a) and (b) are not of the same scale as (c-h).

In all the examples shown in Figure 7.5, a 2.25MHz Vermon array was modelled. While a 5MHz array has been employed almost exclusively until now, the amount of time taken for the model to arrive at a solution was heavily influenced by the centre frequency of the employed probe. The reason for this was that it is recommended that the elemental density should be such that there are at least 12 elements per wavelength to maintain accuracy (Weidlinger-Associates, 2012). In order to ensure accuracy, 15 elements per wavelength were employed. For a 5MHz transducer operating in water, the wavelength is 0.3mm, while that for a 2.25MHz transducer is 0.66mm. Therefore, for a 100mm×100mm image, a 2.25MHz transducer would require approximately 6 million elements, while the 5MHz transducer would require 24 million elements. This imposed significantly longer solutions times. In addition to this, the 2.25MHz probe was justified, in that the purpose of the study was to investigate the effect of soft tissue layers in bone imaging and was not, therefore, limited to the experimental specifications employed herein. Further to this, the 2.25MHz probe still provided an acceptable theoretical resolving power of 0.33mm.

However, in order to investigate whether changes in the centre frequency affected the principle results significantly, both a 5MHz and 1MHz Vermon probe were modelled in addition to the 2.25MHz probe. This was carried out in the image shown in Figure 7.6, as the image was small, with a relatively low element count and, therefore, realistic solution times.




Figure 7.6: Probe position and tissue map employed for the 5MHz and 1MHz probe simulations.

In order to provide a contrast of synthetic aperture widths, TFM and a 20° SAFT were employed to process the resulting FMC files. Once the images had been produced, the contours were extracted. In previous chapters, the most successful surface profile extraction method consisted of finding the pixels above the threshold which were, for each column of pixels, closest to the probe face. Given that a number of tissue types would precede the bone surface, this method was likely to extract some of the soft tissue interfaces. The contour extraction method presented in the previous section did not provide an acceptable solution to this. Therefore, in order to compensate for strong soft tissue interfaces; a new contour extraction method was developed.

As was demonstrated in previous chapters, if the surface profile can be clearly identified through thresholding, it is advantageous to seek the pixels closest to the probe face. However, as is demonstrated in Figure 7.7 (a), the inclusion of soft tissue can cause there to be two or more clear and intense surface contours in the same column of pixels. As such, when the contours are extracted, the soft tissue interface could be extracted rather than the bony surface. This is demonstrated in Figure 7.7 (b), which shows the extraction of tendon surface in red, while the underlying bone surface is lost.




Figure 7.7: Part (a) shows the TFM image from part (e) of Figure 7.5 and part (b) shows the two contour extraction methods, with black representing the new extraction method.

To combat this, it was necessary to identify the individual surface contours in each column of pixels, from which the bone surface contour could be isolated, whereby the pixel closest to the probe face in that particular surface could be found and recorded. To achieve this, the algorithm cycled through each column of pixels, in which the pixels above the threshold were identified. The points were held in a matrix, in which they were rearranged so they were in order of the associated Z value. From this 1D array of Z values, the distance between pixel i and $(i + 1)$ was calculated. If the distance was less than 2mm, the points were added to a further array representing a particular interface. If however, the distance was greater than 2mm, the previous array (if it existed) was discarded and a new array was created.

This process was repeated until the final pixel was compared.

At this point, at most one array would be available, which would represent the bone surface. This principle relied on the fact that the interface with the greatest Z values would represent the bone, as no surfaces would be present beyond the bone, due to the high acoustic impedance mismatch between soft tissue and the cortical layer. From this final array of Z values, the smallest was chosen, the associated X coordinate was identified and the pixel position was recorded. The result of this process can be seen in Figure 7.7 (b), in which the red points represent the ‘Top’ contour extraction method and the black represents the ‘New’ method, presented in this chapter. The ‘Top’ method extracted the upper interface, which represented the tendon, while missing out the bony surface below. However, the ‘New’ extraction method disregarded the tendon interface and identified the bony surface.

Once the 2D coordinates representing the surface contours had been extracted from each reconstructed image, the coordinates were altered to reflect the probe position from which they were captured. As the probe positions were known relative to one another, the extracted contours were combined in each image. Therefore, as well as

visual inspection of the resulting synthetic aperture images, it was possible to numerically assess the accuracy of the extracted surface contours.

In order to provide a reference against which to compare the extracted data, the traced images shown in Figure 7.2 were employed. All line segments were deleted, with the exception of those representing the outer surface of the cortical layer. This image was then saved and imported into Matlab. Each black pixel was identified and assigned a position using the same method as was employed to scale the coordinates of the output file provided by the image importation tool. Additionally, in order to ensure that the line representing the surface was as close to continuous as possible, it was upsampled, as demonstrated in Figure 7.8. It should also be mentioned that that the images from which the scales were derived were different from those employed to make the tissue maps. There would, therefore, have been slight discrepancies between the scale of the data and the reference.




Figure 7.8: Upsampling of the reference model, with the original points in blue and the new, upsampled points in red.

The combined coordinates and reference data were imported into CloudCompare, in which they were matched and compared, as has been performed previously. This provided mean error, maximum error and standard deviation. The probe frequency comparison data was compared in the same way, but with a single contour from a single reconstructed image.

7.4 Results

It can be seen in Figure 7.9 that SAFT and TFM processing resulted in images with different characteristics. Firstly, TFM suffered from weak but obvious side lobes, whereas the SAFT equivalent showed no side lobes. In addition to this, the TFM image displayed very weak representations of the soft tissue preceding the strong reflections of the cortical bone surface. The SAFT image, on the other hand, only presented the bone surface. As well as this, the intensity of the bone surface profile was quite uniform in TFM, while it was patchy in the SAFT image, with regions where the surface was perpendicular to the probe face receiving disproportionately higher intensities. Further to this, the bone surface profile in the TFM image was slightly thinner than that seen in the SAFT image. Finally, the TFM algorithm allowed for construction of surfaces past the edge of the probe surface, meaning more of the surface was represented in the final image than was the case in the SAFT image.




Figure 7.9: SAFT (a) and TFM (b) images from the FMC resulting from the image shown in Figure 7.5 (b), showing a more accurate surface representation with TFM.

Termining the image shown in Figure 7.1 (a) as knee 1 and that in part (b) as knee 2, it can be seen in Figure 7.10 that the outline of the cortical surfaces is very clear in both cases. In the composite image representing knee 1, it can be seen that the soft tissue is, in both SAFT and TFM, very weak compared to the bone surfaces. In the images

representing knee 2, however, while most of the soft tissue layers were absent, the tendon produced high intensity reflections. This not only produced an obvious visual representation, but pixels of intensity greater than the threshold limit.

Common to both knees, the inner surfaces of the joint could not be accessed.




Figure 7.10: Composite images of TFM images of knee 1 (a), SAFT images of knee 1 (b), TFM images of knee 2 (c) and SAFT images of knee 2 (d).

Comparing the results numerically, it can be seen in Table 7.2 that the results for knee 2 are worse than those for knee 1. Assessing the errors visually, it can be seen that the knee 2 results show erroneous points, due to the extraction of the tendon surface. As tendon was not present in the knee 1 image, the highest possible errors were eliminated. It can also be seen in Table 7.2 that TFM yielded superior results in both images, with both extraction methods. The performance of the extraction method was dependant on the data processing method employed, with the ‘New’ method providing the best results for TFM, while the opposite was true for the ‘Top’ extraction method.

Table 7.2: Results of the comparison between the resultant 2D point clouds and the reference model.

Knee Image	Processing Method	Contour Extraction Method	Mean Error (mm)	Maximum Error (mm)	Standard Deviation (mm)
1	TFM	‘New’	0.2	3.3	0.2
1	TFM	‘Top’	0.3	3.2	0.3
1	SAFT	‘New’	0.8	17.8	2.2
1	SAFT	‘Top’	0.4	2.1	0.3
2	TFM	‘New’	0.4	7.7	0.7
2	TFM	‘Top’	0.7	14.3	1.3
2	SAFT	‘New’	0.5	8.8	0.9
2	SAFT	‘Top’	0.5	14.5	0.9

It can also be seen in Figure 7.11 and Figure 7.12 that the TFM-derived point clouds provided greater coverage of the surface than the SAFT-derived clouds. In particular, those regions of the surface where there was a lack of perpendicularity to the probe face suffered from a lack of representation. Additionally, it can be seen in parts (c) and (d) in Figure 7.11 that the difference between the old and new extraction methods are minimal. This is reflected in the numerical assessment, as presented in Table 7.2. This, however, is not the case for parts (c) and (d) of Figure 7.12, in which much of the high errors representing the tendon surface have been eliminated.

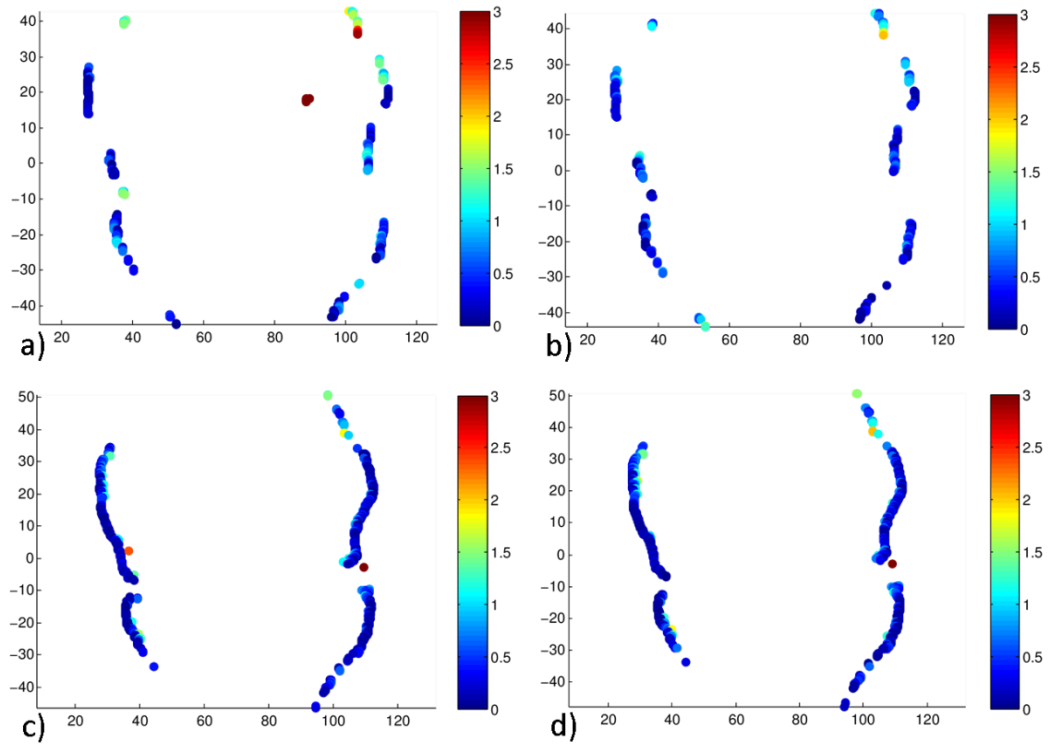


Figure 7.11: Errors resulting from the comparison of the data and the reference model form knee 1, showing the SAFT-derived points obtained using the ‘New’ (a) and ‘Top’ (b) contour extraction methods. Also shown are the TFM-derived points, found using the ‘New’ (c) and ‘Top’ (d) extraction methods.

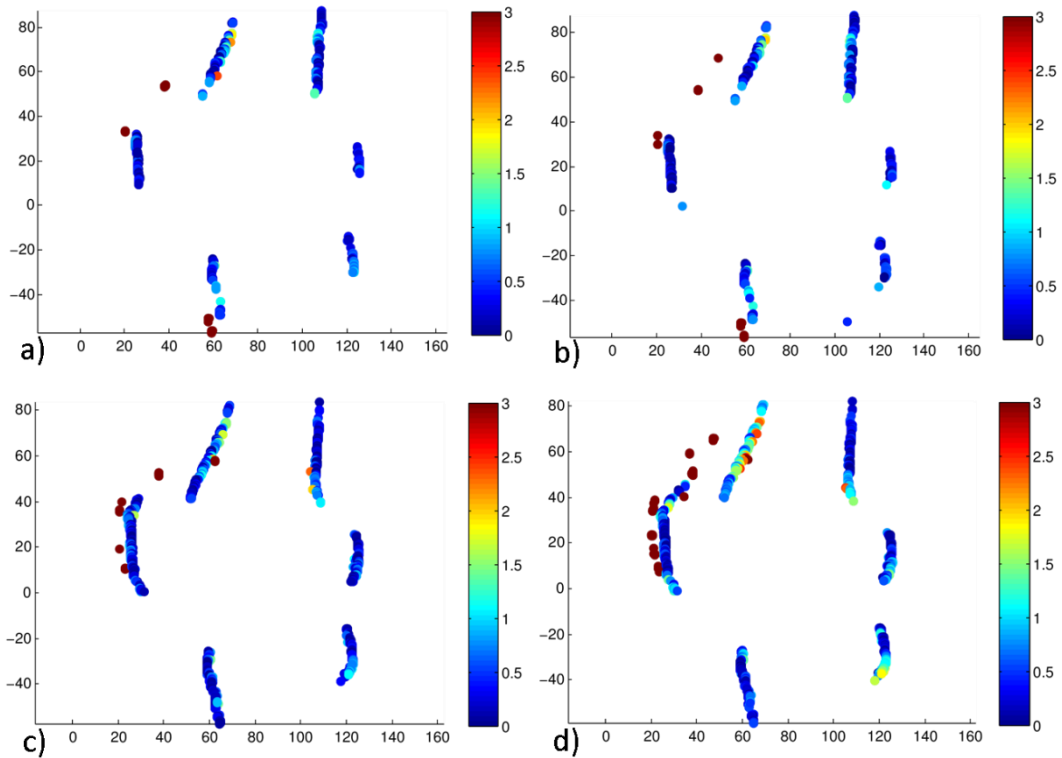


Figure 7.12: Errors resulting from the comparison of the data and the reference model from knee 2, showing the SAFT-derived points obtained using the ‘New’ (a) and ‘Top’ (b) contour extraction methods. Also shown are the TFM-derived points, found using the ‘New’ (c) and ‘Top’ (d) extraction methods.

Comparing the performance of probes with different centre frequencies, it can be seen in Figure 7.13 that as the frequency decreased, the thickness of the surface profile increased. In other words, higher frequencies allowed for higher resolutions. In addition to this, lower frequencies seemed to incur higher levels of non-side lobe noise. For example, comparing Figure 7.13 (a) and (e), it can be seen that no side lobes are present in either, but while part (e) has little or no noise, part (a) has a substantial amount. As with previous examples, no side lobes were obvious in any of the SAFT results. In the TFM images, however, higher frequencies incurred more side lobe contributions.

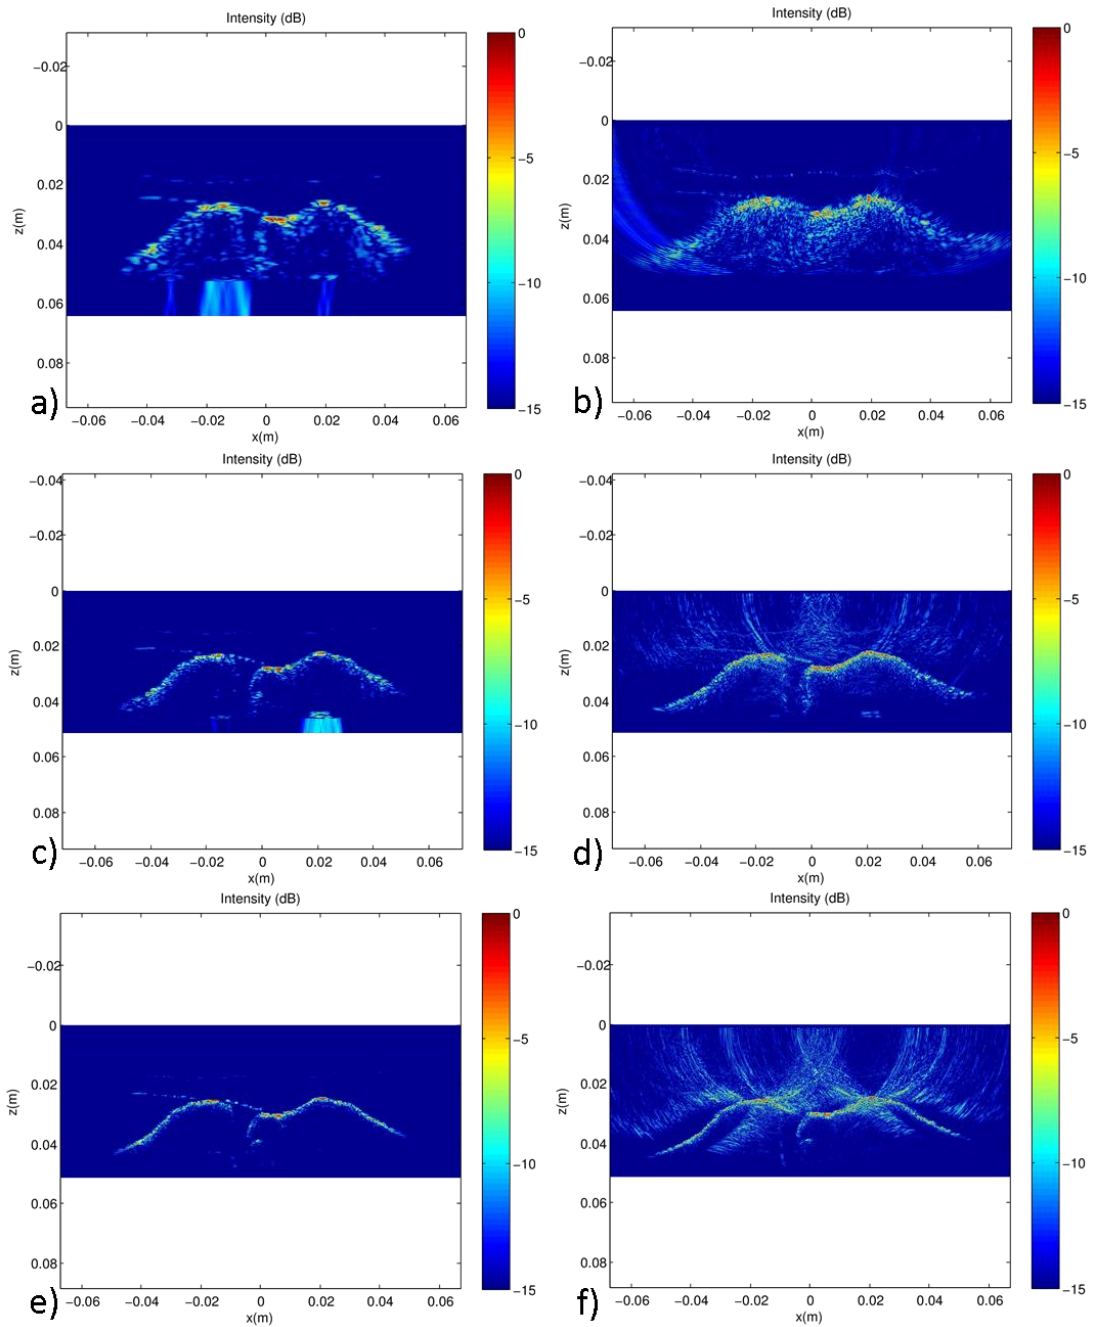


Figure 7.13: SAFT images created using 1MHz (a), 2.25MHz (b) and 5MHz (c) probes. Also shown are TFM images produced using 1MHz (b), 2.25MHz (d) and 5MHz (f) probes.

Assessing the extracted surfaces, the main outcome was that if the frequency was too low, the performance would also decrease. This is demonstrated in Table 7.3, where,

with a number of exceptions, the mean errors were higher for the 1MHz derived clouds than the other two centre frequencies. Unlike the previous results, the choice of processing method and surface contour extraction method made little difference, with a small number of exceptions.

Table 7.3: Results of the comparison between the resultant 2D point clouds from different frequencies and the reference model.

Probe Centre Frequency (MHz)	Processing Method	Contour Extraction Method	Mean Error (mm)	Maximum Error (mm)	Standard Deviation (mm)
1	SAFT	'Top'	0.6	2.6	0.4
1	SAFT	'New'	0.6	2.4	0.5
1	TFM	'Top'	0.5	2.2	0.5
1	TFM	'New'	0.6	3.1	0.6
2.25	SAFT	'Top'	0.1	3.0	0.2
2.25	SAFT	'New'	0.2	3.0	0.2
2.25	TFM	'Top'	1.2	11.9	1.9
2.25	TFM	'New'	0.8	8.3	1.2
5	SAFT	'Top'	0.2	1.1	0.2
5	SAFT	'New'	0.9	19.1	3.2
5	TFM	'Top'	0.2	3.3	0.2
5	TFM	'New'	0.2	3.4	0.2

Evaluating the surface contour extraction comparisons visually, it is clear from Figure 7.14 that lower frequencies resulted in less coverage of the surface and sparser point clouds. Employment of SAFT also resulted in a reduction in surface coverage, relative to TFM. Considering the difference between parts (a) and (g) of Figure 7.14, for example, it is clear that the surface definition is almost completely lost when employing SAFT. However, the TFM extractions were much more susceptible to side lobes, but only at high frequencies, as illustrated in parts (i) and (l).

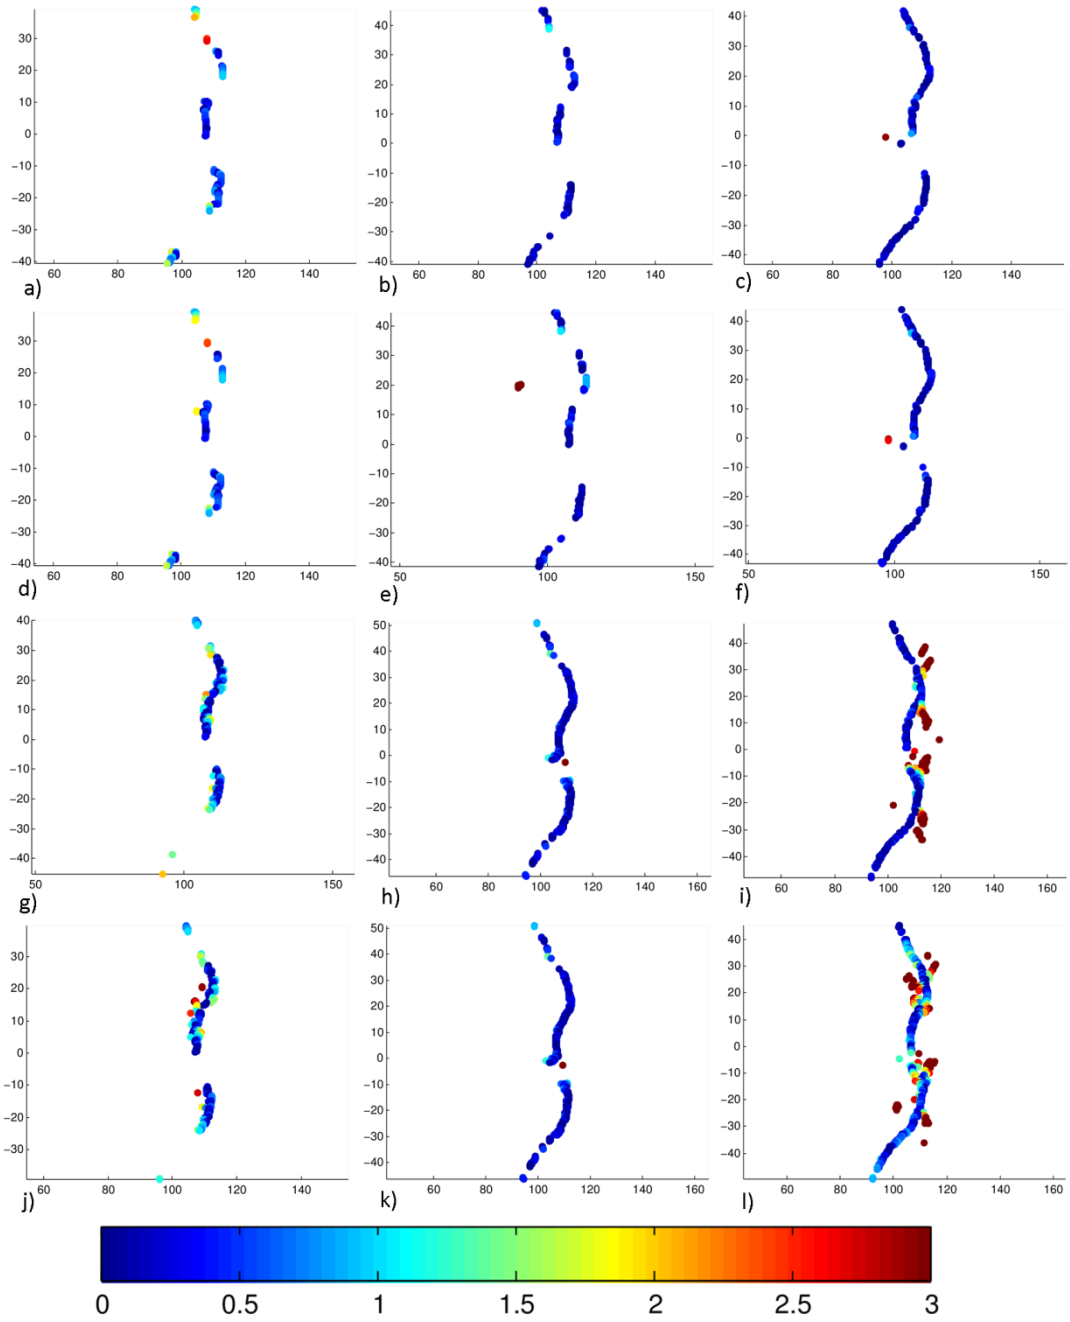


Figure 7.14: Errors resulting from the comparison of the data and the reference model from knee 1, with the columns showing 1MHz, 2.25Mhz and 5MHz probes, from left to right. The rows display SAFT processing with the ‘Top’ extraction method, SAFT processing with the ‘New’ extraction method, TFM processing with the ‘Top’ extraction method and the TFM processing method with the ‘New’ extraction method, from top to bottom.

7.5 Discussion

The purpose of this study was to assess the influence of soft tissue inclusion on imaging the bony surfaces of the knee joint when using FMC and synthetic aperture methods. The main outcome was that the synthetic aperture methods employed were capable of accurately representing the bony surfaces. All the interfaces with different soft tissue types, while visible, were of such low intensity that thresholding easily discarded them. The exception to this was tendon, which resulted in strong reflections with magnitude comparable to that of bone. The reason for this was that while the rest of the soft tissues employed in this study fell within a similar range in terms of acoustic properties. Tendon was, on the other hand, much closer to cortical bone in density, as shown in Table 7.1. This is due to tendon being dense and highly organised compared to other soft tissue types.

Despite this success, there were a number of reasons this study may not have accurately represented the behaviour of FMC-based imaging in the knee joint. For example, the segmented images employed herein were not truly representative of the original images. As can be seen in Figure 7.1, there are streaks of fat within regions that were defined solely as muscle. The same can be said of other regions which were defined as homogeneous when, in fact, other tissue types were present in small amounts. The inclusion of these features may have contributed to noise and decreased the accuracy of the image. Further to this, the internal structure of the materials was not taken into account, which may have produced noise. For example, if the striations of the muscle had spacing of the order of the ultrasound wavelength, reflections could have been produced which would cause noise (Hill, 1976). It is, however, unclear how much of difference this would have imposed, as other authors have stated that only interfaces produce any prominent reflections (Friedman, Finlay, & Jurriaans, 2001).

In a similar vein, a further reason the simulations may have generally been inaccurate was that the material definitions were rudimentary in nature. In order to define a material at a basic level, PZFlex requires only a longitudinal velocity and a density. While these values were accurate, the platform also allows for the definition of

damping properties, which were not forthcoming in the literature. Additionally, the materials were defined as isotropic, which was, as discussed above, incorrect.

Further to this, trabecular bone is a heterogeneous material, with two main constituent material types - each with highly different acoustic properties. In order to simplify this, isotropic material properties were assigned to those regions. As mentioned previously, ultrasonic waves do not travel through trabecular bone as would be expected in an isotropic material and, as such, trabecular bone should be considered anechoic. In this model, however, waves travelled normally. This is illustrated in Figure 7.15 which shows that while there is a major reflection at the cortical surface of the patella, the pressure waves penetrated the inner parts. This was not only an inaccurate representation of the behaviour of the tissue, but caused there to be reflections inside the patella that were then transmitted through the cortical surface and registered by the receivers. As such, there were a number of lines of high intensity beneath the outer surface of the patella in the resulting images. Therefore, a portion of the affected images were set to zero to eliminate these effects, as can be seen in Figure 7.7.




Figure 7.15: Summation of the pressures at all time steps for the 2.25MHz probe on knee 2, showing how ultrasonic energy inaccurately penetrated to the back of the patella.

In order to correct these problems, an exhaustive literature search would be required to acquire the necessary material properties to accurately model the materials. Additionally, the images employed should be high resolution, photographic images to allow for accurate segmentation of the materials. This would include segmentation of trabecular tissue, which would allow for solid and fluid parts, resulting in a more realistic acoustic interaction.

The surface extraction method presented in this section performed well. In reconstructing the surface of knee 1, the difference in results produced using the 'Top' and 'New' methods were minimal when using TFM. Employment of the 'New' method in knee 2, however, saw a significant improvement. The reason for this was that knee 2 included tendinous tissue, while knee 1 did not. The 'New' method was designed specifically to exclude such surfaces from extraction and was, therefore, successful in that task. The weakness of the method is that it relies on the fact that if there are any pixels above the threshold in a column, the pixel furthest from the probe face will always be representative of bone. This, however, may not always be the case. For example, in Figure 7.1 there is a region beneath the patella in which the tendon is present, but there is no bony surface beneath it. As such, the tendon surface could be extracted. In saying this, such situations would be in the minority and the other presented surface extraction methods would suffer from the same problems.

In the previous experiments, most, if not all of the surfaces under inspection were within a line of sight of the transducer. In this case, however, many of the surfaces were inaccessible. These included the region between the articulating surfaces and in knee 2, the area directly behind the patella. While this is a limit imposed by the physical characteristics of ultrasound, it could be solved, firstly, by defining a tool path that would allow line of sight access to the surfaces covered by the patella. Additionally, in order to gain access to the articulating surfaces, the knee could be scanned in extension and then flexion, with each scan revealing different surfaces.

In terms of data processing methods, TFM produced superior images for the 2.25MHz probe results. Relative to the SAFT results, the TFM images showed a thinner surface profiles, with a more uniform intensity distribution across the surface. This was made

possible by the increased synthetic aperture width. As well as this, however, the increased aperture imposed side lobes. These were not, however, any side lobes of significant intensity and most could be eliminated through thresholding. Even if the side lobes had intensities above the threshold limit, they were mainly located above the true surface image and, as such, would be eliminated by the surface extraction method described in this section.

Focussing on the dependence of the performance on the centre frequency of the probe, the higher frequencies achieved higher resolutions than the lower frequencies. For the 1MHz probe, this not only led to the surface becoming slightly thicker, but to there being a large amount of noise beneath the bony surface, as is shown in parts (a) and (b) of Figure 7.13. The reason for this was the increased pulse width associated with a decrease in centre frequency. This is represented in the A-scans shown in parts (c) and (d) of Figure 7.16, where although the shape is maintained, the width of the pulse is much greater for the 1MHz example. The effect this has on the physical waves produced is that those of lower frequency are of a significantly higher wavelength. Further to this, the width of the wave packet increases, as can be seen in parts (a) and (b) of Figure 7.16. Therefore, the difference in time between the start of the wave packet and the end of the wave packet interacting with a surface will be greater for lower frequencies. This will be mirrored in the A-scans, which will, in turn, cause the thickness of surfaces to increase. This effect was the most prominent contributor to the poor performance of the 1MHz probe. It would, therefore, be prudent to use a transducer with centre frequency greater than 2.25MHz.

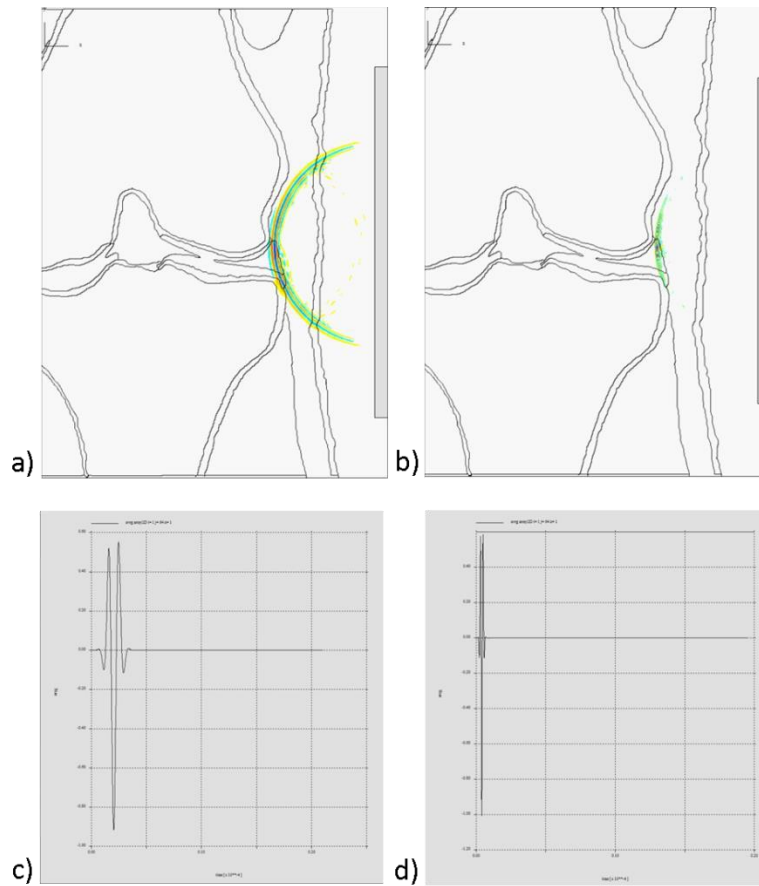


Figure 7.16: Snapshots of the travelling pressure wave and the associated A-scans for the 1MHz (a, c) and 5MHz (b, d) probes, firing and receiving on element 64. There is a clear decrease in the width of the wave packet with the 5MHz transducer.

Changes in centre frequency also affected the performance of the imaging algorithms. For example, increases in the frequency resulted in significant increases in side lobe levels. The reason for this effect is due to the width of the pulse. With reference to Section 4.3, when the frequency is low, the width of P is increased. As such, more of the image will be non-zero. For higher frequencies, P will be narrower, meaning fewer pixels will receive a non-zero contribution. However, due to more pixels being non-zero with lower frequencies the side lobes will, on a decibel scale, be less intense. With higher frequencies, on the other hand, the intensity of the side lobes will be relatively higher, as fewer regions contain non-zero pixels. This is demonstrated in Figure 7.17, which shows how the side lobes created using a 2.25MHz probe have more coverage

than those simulated with a 5MHz probe. However, the intensity of the side lobes shown in part (b) is significantly higher. In addition to this effect, the 5MHz results herein provided a more uniform representation of the surface. There were, therefore, more pixels with a high intensity attached to them. As side lobes are formed by high intensity pixels, the side lobe levels will also increase. The result of these two effects was that for the 2.25MHz probe, TFM resulted in the best results, as evidenced in both parts of the study. However, for the 5MHz probe, TFM resulted in high levels of extracted side lobes, as can be seen in parts (i) and (l) of Figure 7.14. It is, therefore, necessary to select the imaging method relative to the probe centre frequency.

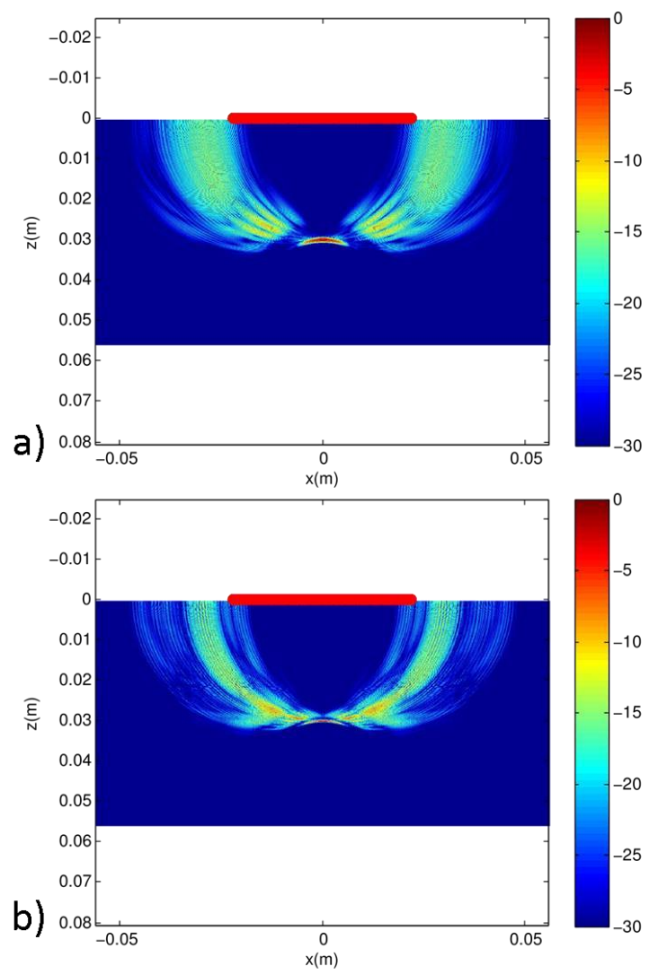


Figure 7.17: Simulated surface reflections from a sphere using 2.25MHz (a) and 5MHz (b) 64 element probes, showing more concentrated and intense side lobes in part (b).

Although not necessarily represented in the numerical results, it was clear from the visual representations in Figure 7.13 and Figure 7.14 that the performance of SAFT improved as the centre frequency increased. This was due to the increased performance provided by the reduced wave packet width offered by higher transducer frequencies. A common problem in all the SAFT images, however, was the lack of uniformity in the intensity of the surface, with the images from each of the three probe frequencies showing high intensities in the same areas. Therefore, while the choice of whether or not to employ TFM was governed by the centre frequency of the probe, the performance of SAFT did not show the same level of dependence.

7.6 Conclusion

An FE modelling platform was used to simulate FMC with intact knee joint models, allowing for different transducer properties to be tested. The prime purpose of the study was to determine whether the presence of a number of soft tissue layers and types would prevent the accurate extraction of surface profiles. It was found that soft tissues imposed no strong reflections and could, therefore be easily eliminated. The only exception to this was tendon, which has a density comparable to bone. However, a new surface extraction method was developed which significantly decreased the number of points extracted from the tendon surface.

As such, the study was a success in that soft tissue layers provided no real obstacle in retrieving the bony surfaces. It was, however, recognised that the model was simplified in many areas, suggesting that the results might not be as accurate as could possibly be achieved with the software.

Additionally, it was found that the centre frequency of the probe was influential in the ability to extract the bony surface contours. In particular, it was found that the 1MHz probe did not offer - under any circumstances - image quality comparable to the other two probes. As such, the probes employed in imaging bony surfaces should be equal to or greater than 2.25MHz. Further, the imaging method employed should be based on the centre frequency of the probe, with higher frequencies achieving better results

with SAFT. Lower frequencies, on the other hand, achieve better results with TFM, due to the decrease in side lobe intensity and coverage.

To conclude, the problems caused by the soft tissue were minimal, but the accuracy of the results was questionable. Therefore, in order to ensure that the soft tissue effects are as minimal as predicted, further experimental data were required. In particular, experimental studies would eliminate much of the simulation uncertainty and possible inaccuracy and provide a more realistic result.

Chapter 8 Bone Surface Imaging with Soft Tissue – an Experimental Approach

8.1 Introduction

Based on the results of the previous study, it was concluded that more experimental studies were required in order to further evaluate the impact of the inclusion of soft tissue. While the FIRST laboratory was unavailable due to a building move, the FlawInspecta phased array controller was accessible. As such, experiments were possible, but were required to be simple, as no positioning capability was available. Given these constraints, imaging an intact knee joint, be it human or animal, would not have been prudent. An alternative to employing real soft tissue was to employ tissue mimicking materials. This would allow for controlled evaluation of the sample before and after the application of the tissue mimic.

Therefore, a single bovine condyle was inspected in water. A known tissue mimic material was then manufactured and used to cover the sample, which was then inspected in the same way. The results were compared to determine if the inclusion of the soft tissue mimic resulted in problems in surface extraction relative to the best case scenario of a bare bone.

8.2 Materials and Methods

Materials mimicking the acoustic properties of biological tissues are used widely within biomedical ultrasound imaging (Culjat et al., 2010). There are many examples of mimics with a simple manufacture process, which use readily available materials and require little experience (Bude & Adler, 1995; Kendall & Faragher, 2007; Osmer, 2008). The purpose of these, however, is not to accurately imitate soft tissue, but rather to provide quick, simple and replaceable phantoms which reproduce the sonographic appearance of soft tissue (Bude & Adler, 1995). There are a great many alternatives to these simple methods, employing a host of materials, but many require specialist techniques, materials, and equipment. Therefore, the author of a recent publication describing a number of manufacture methods (A. J. Weir, Sayer, Wang, & Parks, 2015) was consulted, so that the exact detail of the method could be obtained. It was advised that of the base materials employed in the study, Polyvinyl Alcohol (PVA) provided the best results.

However, before the mimic material was manufactured, the bone sample was prepared. A bovine condyle was employed, as it provided a complex surface similar to that likely to be found in the ultimate application. Additionally, it provided the acoustic properties of bone that would be found in human knee joints. The sample was placed in a plastic container, in which it was secured to the bottom by a hard set epoxy resin, which allowed for the position to be maintained during the setting of the tissue mimic material.

In order to attain reference data, the sample was immersed in tap water and taped to the bottom of the surface of the container to maintain a standoff between the probe and the sample. A 5MHz, 128 element Vermon array was employed, the height of which was fixed using a clamp stand, as can be seen in Figure 8.1. The probe was positioned manually, such that between 10 to 15 FMCs were captured along one axis. Care was taken to attempt to position the probe such that the data capture positions were representative of the whole shape of the sample. This was then repeated along the other axis. This technique did not allow for 3D reconstruction, as no positional data were available. For the same reason, direct comparison was not possible between data and

reference data, but visual comparison could be carried out with similar data sets, showing similar areas of the bony surface.




Figure 8.1: The apparatus used to retrieve reference FMC data from the bare condylar surface.

In order to cover the sample with the soft tissue mimic, approximately 1L of the material was required. Therefore, based on the methods presented in the article, 1L of distilled water was mixed with 100g of 99% hydrolysed PVA (Sigma-Aldrich Corporation, MO, USA), gently heated using a hot plate and stirred until thick. While application of an antibacterial agent was advised, this was not carried out. The reason for this omission was that the phantoms described in the paper were designed for long term use, which would have made them susceptible to microbial infection. In this situation, on the other hand, the phantom was to be used only once, soon after setting had occurred. It was, therefore, considered to be unnecessary for this study and omitted.

Once thick, the solution was allowed to cool to room temperature and 10g of calcined aluminium oxide, with a $3\mu\text{m}$ grain size (Logitech Ltd., Glasgow, UK) was added to the solution, which acted as a scattering agent. This was then mixed and degassed using a Thinky ARE-310 planetary centrifugal mixer (Thinky Corporation, Tokyo, Japan).

The solution was then poured into the plastic container, so that it not only covered the top of the bony surface, but allowed for a gap between the phantom surface and the bony surface of at least 1cm. The solution then underwent two freeze-thaw cycles, which consisted of freezing the solution at approximately -20°C for 12 hours, then thawing at approximately $+20^{\circ}\text{C}$ for another 12 hours.

The same method of inspection as described above was employed. While the exact location of the sample was unknown due to the opacity of the mimic material, it was assumed to be central. Therefore, the probe was positioned accordingly, as can be seen in Figure 8.2. In addition to mirroring the reference data capture method, data was also captured while the probe was in contact with the phantom surface. This was achieved by pushing the probe down on to the surface, with the probe being positioned and FMCs being captured in the same way as before. In doing this, the method by which most biomedical ultrasound imaging is performed was emulated.



Figure 8.2: The apparatus used to retrieve FMC data from the condylar surface whilst situated in soft tissue mimicking material.

Once the FMC data had been acquired, it was processed using both TFM and a 20° SAFT. Both the ‘New’ and ‘Top’ contour extraction methods were employed on the

resulting images. However, a further extraction method was employed which exploited the nature of the contours displayed in the resulting images. This process began by defining a centre point unique to each image by finding the mean X coordinate of all the pixels above the threshold limit. The Z coordinate, meanwhile, was simply defined as Z value of the non-zero pixel with the largest Z coordinate. This centre point formed the apex of 90 triangles, the first of which was defined by drawing one line the length of the image from the apex to the same Z coordinate as the apex. The second line was created with the same length, but with an inclination of 2° relative to the first. The ends of the two lines were then connected to create a polygon in the shape of a triangle. The first line of the second triangle was simply the second line of the first triangle. The second line of the second triangle was defined by adding 2° on to the first. This process is illustrated in Figure 8.3, which shows the formation of the first two triangles. This process was repeated over 180° , so that 90 triangles were created.

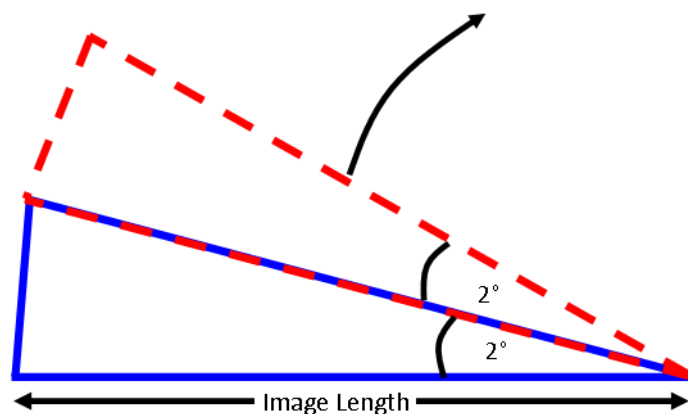


Figure 8.3: Formation of the first two triangles during the contour extraction, with the first being solid blue and the second being dashed red.

For each triangle, all the pixels above the threshold limit and within the polygon created by the triangle were identified. The Euclidean distance between the triangle apex and each of the points within the above criteria was then calculated, with the pixel achieving the shortest distance being retained, while the rest were discarded.

An example of the process of pixel elimination for one triangle can be seen in Figure 8.4. The process was repeated for every triangle created in each image, resulting in a representation of the inner surface. From this point on, this method will be termed the ‘Radial’ contour extraction method.

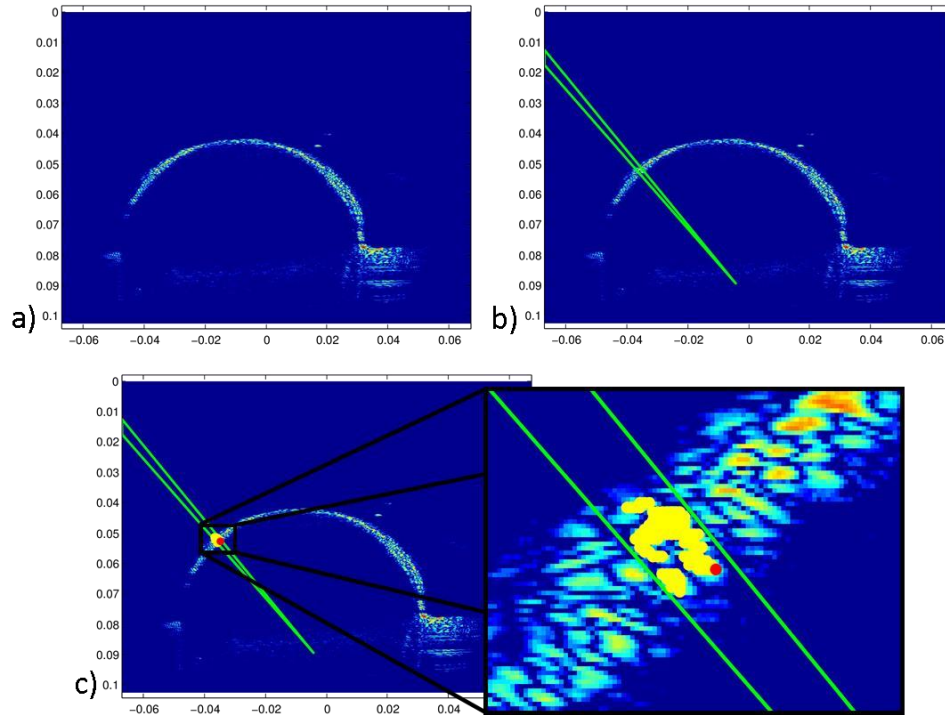


Figure 8.4: Graphical representation of the ‘Radial’ extraction method. From the original image (a), a centre point is created. From this, a triangle is produced (b) and any pixels above the threshold (yellow points) within the triangle are identified (c). The distance of each of these pixels to the centre point is calculated and all the points are discarded, with the exception of that with the smallest distance (shown in red). This process is repeated over 180° , creating the final contour.

As no positional data was acquired and the reference and phantom data had to be acquired separately, there was no method of direct, numerical comparison. In place of that, each of the FMCs was processed using both TFM and a 20° SAFT. A number of representative images from the phantom data set were chosen and matched with similar images from the reference data set. From each of these, the bony surfaces were

extracted using three contour extraction methods: the 'New' method, the 'Top' method and the 'Radial' method. The comparison of the performance of these methods was then performed visually.

No acoustic characterisation was performed on the materials, meaning it was unknown whether or not it complied with standards relating to soft tissue phantoms. However, it was possible to check for more obvious and potentially problematic characteristics. This was achieved by cutting away part of the material away from the sample after the inspection had taken place. In doing so, the point of contact between the bone and the mimic could be visually assessed to check for any obvious air pockets. Additionally, it would allow for inspection of the material itself, to ascertain whether or not the composition of the material was consistent throughout.

8.3 Results

Two image examples were chosen for visual inspection, each of which were approximately positioned at the central axes of the sample. Essentially, one example was picked from the midpoint of one scan line, while another was picked from the midpoint of the other scan line. As can be seen from Figure 8.5, use of different synthetic aperture methods provided varying results. As would be expected based on previous results, the thickness of the line representing the surface was greater for the SAFT images. Additionally, the TFM images contained side lobes, while the SAFT images did not.

The differences between the bare bone images and those resulting from the inclusion of the soft tissue mimic were slight. Comparing parts (b) and (d) of Figure 8.5, it can be seen that the same shape was maintained between the two examples, as was a similar line resolution. The most obvious difference is the inclusion of mimic interface representation above the bony surface. However, it is of relatively low intensity. In addition to this, there is a small area of noise directly above the centre of the surface of the bone. Further to this, the contour representing the bony surface is slightly raised at the top of the contour, while that in the bare examples is smooth and continuous.




Figure 8.5: SAFT (a, c, e) and TFM (b, d, f) images resulting from FMCs captured with bare bone (a, b), non-contact tissue mimic (c, d) and contact tissue mimic (e, f).

The contact measurements, shown in parts (e) and (f), show large changes relative to the other two FMC reconstruction scenarios. In the SAFT example, (e), the centre of the surface is almost non-existent. In the TFM image, (f), on the other hand, the centre of the surface is of disproportionately high intensity, while the other parts of the surface show low levels of intensity. In addition to this, the TFM image, (f), displays a further

line of high intensity directly beneath the probe face. Finally, both images resulting from contact measurements showed noise beneath the centre of the bony surface.

In the second example, the differences between the bare bone inspection and that with the soft tissue mimic included were more pronounced. Comparing parts (b) and (d) of Figure 8.6, it can be seen that the thin surface profile was not replicated when the mimic was included. In addition to this, the level of noise increased significantly when the soft tissue was included in both SAFT and TFM images. The small, raised region of surface seen in the previous example was, once again, present in the images resulting from the inclusion of the tissue mimic. However, in this example, it was much more prominent, creating a distinct surface above that of the bone.




Figure 8.6: SAFT (a, c, e) and TFM (b, d, f) images resulting from FMCs captured with bare bone (a, b), non-contact tissue mimic (c, d) and contact tissue mimic (e, f).

The contact measurements presented similar characteristics to those in the previous example, with the centre of the surface being very weak in the SAFT images, while the opposite was true for the TFM images. As before, the TFM image displayed another surface close to the probe face but was, this time, less defined and of a higher intensity. In addition to this, there was a region of relatively high intensity beneath the bony surface. This was caused by the waves bouncing back from the probe face, making a return journey and, again, being received. These regions were, therefore,

caused by the fact that the probe face was closer to the surface of the bone than in the other inspection scenarios.

The surfaces extracted from the first example using the different surface extraction methods can be seen in Figure 8.7. The surfaces extracted from the bare surface had predictable characteristics, with the 'Radial' extraction method being sparser than the others. Additionally, this method resulted in a contour slightly further from the probe face than the others. The difference between the other two methods was minimal, given that only there was no strong soft tissue interfaces.

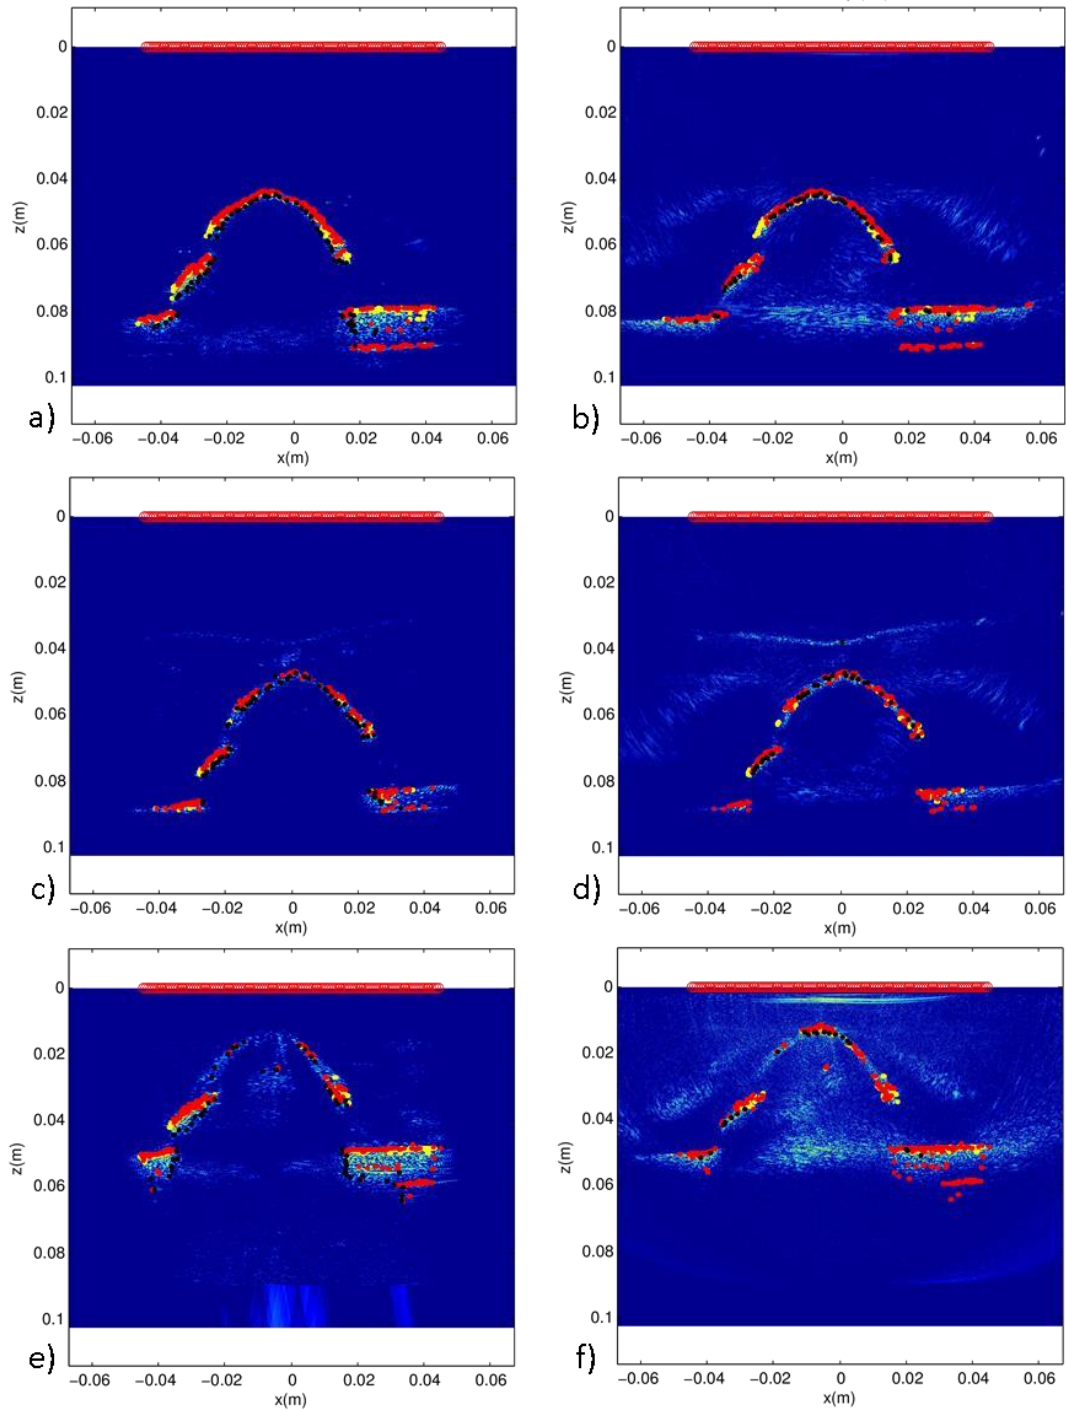


Figure 8.7: SAFT (a, c, e) and TFM (b, d, f) images resulting from FMCs captured with bare bone (a, b), non-contact tissue mimic (c, d) and contact tissue mimic (e, f). Also shown are the surface representations, extracted using the ‘Top’ (yellow), ‘Radial’ (black) and ‘New’ (red) methods.

The differences between the surface extraction methods were maintained when the soft tissue mimic was included, as can be seen in parts (c) and (d) of Figure 8.7. The main difference between the extracted surfaces from the bare and covered bony surfaces was an increase in the sparseness of the points. The accuracy of the extracted points was, however, maintained in the inspection with tissue mimic included. The contact inspection saw a significant reduction in coverage of the surface, with only certain areas being represented in the surface extraction.

The bare bone results of the second example did not provide the same density of extracted points and uniformity of coverage as was seen in the first example. In particular, there was a region from which almost no points were extracted, as can be seen in parts (a) and (b) of Figure 8.8. In saying this, the same characteristics of the different contour extraction methods were seen and the accuracy of the extracted points was also maintained.

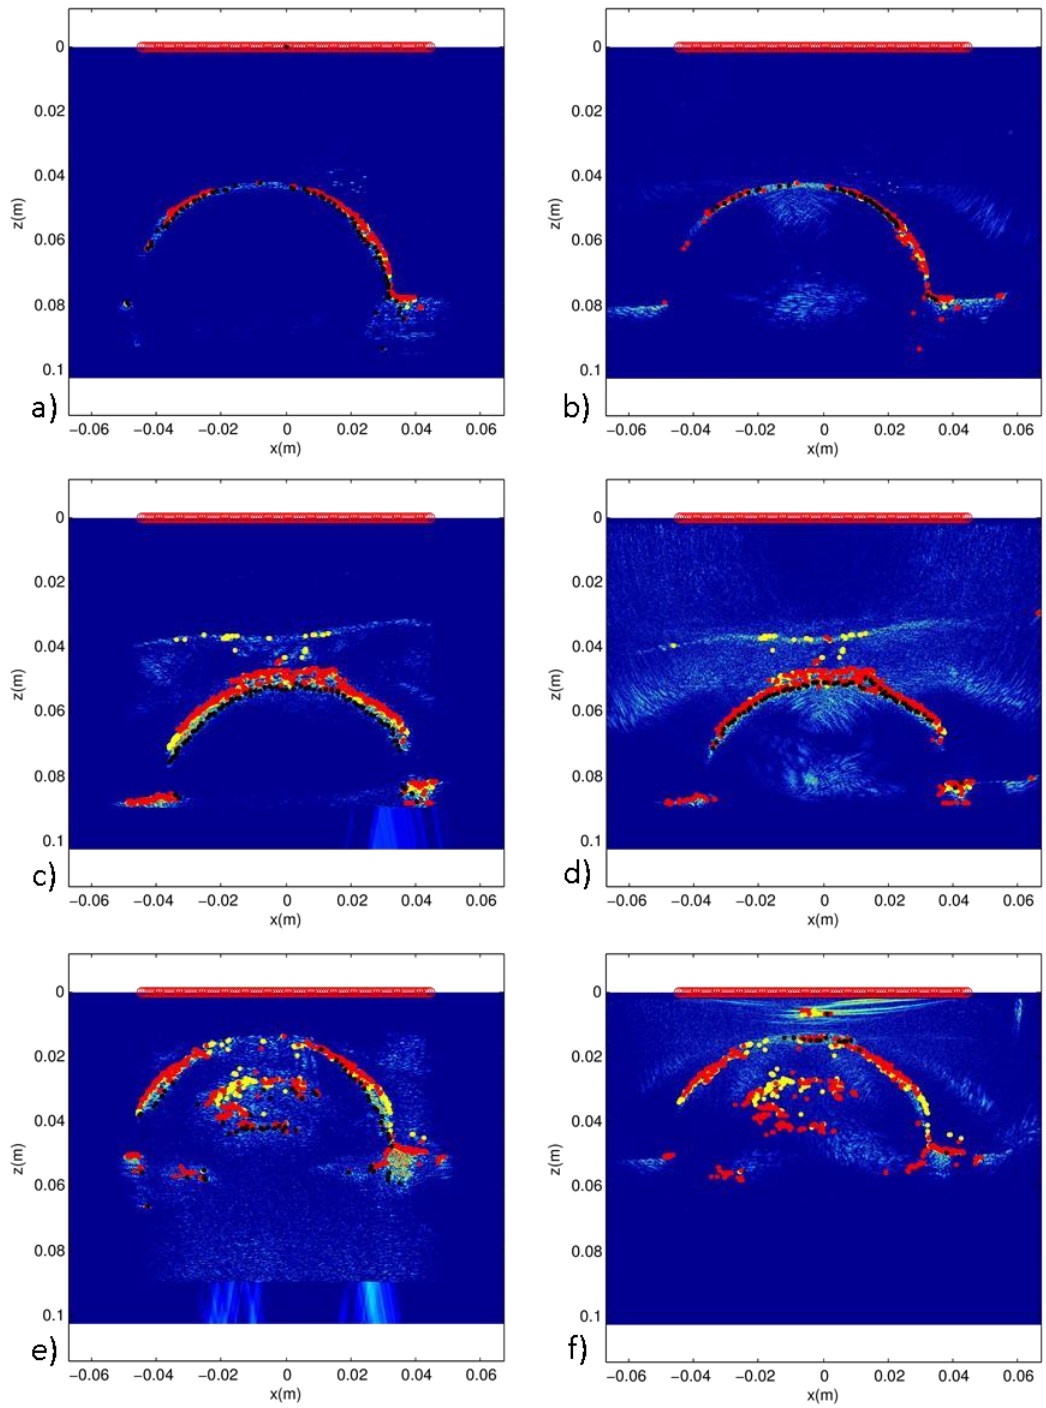


Figure 8.8: SAFT (a, c, e) and TFM (b, d, f) images resulting from FMCs captured with bare bone (a, b), non-contact tissue mimic (c, d) and contact tissue mimic (e, f). Also shown are the surface representations, extracted using the ‘Top’ (yellow), ‘New’ (black) and ‘Radial’ (red) methods.

The non-contact tissue mimic inspection showed a significant drop in performance relative to the previous example. In particular, the ‘Top’ extraction picked up points from the water-tissue mimic interface. While the ‘New’ method managed to avoid doing this, it could not separate the true bone surface and the additional surface, which can be seen in part (c) and (d) of Figure 8.8. The ‘Radial’ extraction method was not affected by either of these problems and retrieved the pixel closest to the bottom of the image.

The contour of extractions resulting from the contact inspection were very poor, as can be seen in parts (e) and (f) of Figure 8.8. In the SAFT and TFM images, regions of high intensity existed beneath the bone surface, which were, as described before, caused by return journeys of the waves. In addition to this, in the SAFT example, the centre of the surface was of low intensity and, as a result, no points were extracted in that region. The TFM example also had a number of points extracted from the region of high intensity that existed directly beneath the probe face.

When inspecting the composition of the tissue mimic material, it was found that the material appeared to be isotropic. As can be seen in Figure 8.9, the colour and opacity remained the same throughout the exposed width and depth. From this inspection, then, it was concluded that the scattering agent was uniformly distributed throughout the material, as the powder had not sunk to the bottom.



Figure 8.9: A quarter of the tissue mimic material, cut from the container using a scalpel, displaying an isotropic dispersion of powder.

Assessing the points of contact between the sample and the tissue mimic, it can be seen in Figure 8.10 that a reasonably good bond was made between the two materials. In saying this, small regions appeared to show a slight gap between the two materials. However, it is unknown what material would have filled this area during inspection and how large any other gaps might have been.



Figure 8.10: The partially exposed sample after removal of a quarter of the tissue mimic material, showing small gaps between the osseous tissue and the mimic.

8.4 Discussion

The main outcome of this experiment was that the soft tissue interface and scattering imposed by the soft tissue mimic did not make accurate bone surface representation and extraction to be difficult. While there was a slight degradation in the level of coverage, this did not result in a reduction in accuracy of the reconstruction. In saying this, there were a number of problems caused by the inclusion of the soft tissue mimic. One such problem was a thickening of the surface bone representation, which can be seen in parts (c) and (d) of Figure 8.6. However, it should be noted that this was not a direct comparison and, as such, might not have been due to the inclusion of the mimic. The probe angle could, for example, have been different from one inspection to the

next, which would have affected the retrieved bone surface. In evidence of this idea, the first example did not see such a change. Further to this, the thickening of the surface profile did not result in any obvious decrease in the accuracy of the surface extraction. In addition to this, there was a region of high intensity directly beneath the probe face in the TFM images of both examples for contact inspection. As can be seen in part (f) of both Figure 8.5 and Figure 8.6, there was a dip in the surface profile of the water-tissue mimic interface. As such, this could have resulted in a gap between the probe and the surface. Given that the surface would have been less than 1cm from the probe face, this would have caused disproportionately strong reflections. This is because the intensity of the beam that close to the element is of much higher intensity than elsewhere. This caused inaccurate surface extractions and highlights the problems that can be caused by contact imaging. An experienced operator with real time feedback would be required to ensure that the data captured was worthwhile. The requirement of real time feedback would complicate the system and increase cost. Additionally, the use of an operator would not only increase cost per scan, but would incur intra- and inter-operator variability. As such, robotic control, using a probe-to-skin standoff distance would be the most appropriate choice.

Another problem seen with the contact measurements was the double reflection of the bone surface, which caused regions of high intensity, as demonstrated in Figure 8.6 (e). When using contact measurements, the probe would often be close to the bone surface, meaning this problem would be common. While it could be solved by decreasing the inspection depth, this would severely limit the coverage of the full surface. Maintaining a standoff distance between the probe face and the surface of the skin would limit the likelihood of this situation.

Additionally, both non-contact examples showed an unwanted surface directly above the bony surface when the tissue mimic was included. This was likely caused by a gap between the two materials. As well as being displayed in the data, the dissection of the tissue mimic showed evidence of possible small gaps between the bone and tissue mimic. This was likely a result of the freeze-thaw process, which saw the material expand and contract. It is difficult to link this to real soft tissue-bone interfaces, as the material that filled the gap is unknown. However, based on the results of the previous

section, it would be reasonable to assume such effects would be unlikely in human knee joints.

The difference between the synthetic aperture methods was minimal relative to the results seen in previous chapters. In earlier studies, the side lobes associated with the TFM images were of such high intensity that the surface contours could not be accurately extracted. The side lobes seen in this study, however, were easily eliminated simply by thresholding. It should be mentioned, however, that only one small part of a knee joint was imaged, with similar probe standoffs and angles. Therefore, different surface shapes and probe-to-surface angles could cause intense side lobes. This study was, therefore, inconclusive in which imaging method would be appropriate, due to the lack of variation in the surface types.

The performance of the different surface contour extraction methods was as expected, with the 'Top' and 'New' methods producing very similar results. The only exception to this was in the second example when the soft tissue interface was extracted by the 'Top' method, but not by the other. While the contours extracted using the 'Radial' method did not extract the wrong surfaces, it was obvious that the points were closer to the bottom of the image, than to the probe face. It has been shown in previous chapters that it is preferable to extract the points on the contour closest to the probe face. As such, the 'New' method was most likely to provide an accurate surface extraction.

A further possible weakness in the experimental method was that the volume of scatterer defined in the reference publication (A. J. Weir et al., 2015) was very low, at 1% by weight. Other publications have used higher scatterer volumes of ~4% (Bude & Adler, 1995; Kendall & Faragher, 2007), which may have produced higher levels of noise. However, the materials employed herein have been tested for scattering properties and found acceptable (A. J. Weir et al., 2015).

Additionally, the phantom only included one layer of tissue mimic. This is unrealistic, in that multiple soft tissue types would be present between the probe face and the bone surface. Therefore, an improvement to the experimental setup would be to have additional layers, which could be achieved by preparing smaller batches and setting them additively. The individual batches could contain different proportions of the

scattering agent, which would provide differences in acoustic response. Moreover, these could be set at varying angles, to produce interfaces that were not perpendicular to the probe face. This is illustrated in Figure 8.11, which shows setting of four layers of soft tissue, each with a different density of scatterer.




Figure 8.11: Proposed setting procedure for achieving multiple layers of soft tissue mimic, each with different scatterer densities.

8.5 Conclusion

A tissue mimic material was employed to further investigate the effects of having an interface preceding the bony surface and to consider the consequences of scattering on surface extraction. It was found that the inclusion of a scattering agent had no obvious effect on the resulting images. Keeping in mind that the validity of the scattering effects of the material had been confirmed by other authors (A. J. Weir et al., 2015), this was considered to be a success. Further to this, the additional interface did not cause any problems in surface extraction, if the appropriate method was employed. However, only one interface and tissue type was employed.

The experimental method included contact measurements, in an effort to recreate the kind of inspections that take place both in medical ultrasound imaging and in related research (Barratt et al., 2006b, 2008a; Hanisch et al., 2015b). This exposed some of the weaknesses associated with contact imaging, including the need for an experienced

operator and for a decreased imaging depth. This provided experimental evidence to back the idea that robotic probe manipulation, providing a standoff distance between the probe face and the skin would offer the best solution in this project.

Therefore, while the results were positive and provided more evidence to suggest that FMC and synthetic aperture imaging methods could accurately image bone through soft tissue, there were weaknesses in the method, meaning further work is required in this area. The suggested improvements in the method would have made the sample more realistic, but it would be impossible to say whether or not such an apparatus would accurately recreate the real-life properties of a human knee joint. As such, the only way to truly test the ability of these methods would be to employ intact cadaveric specimens.

The outcome of both this and the previous chapters, then, was that the evidence collected implied that using FMC and synthetic aperture methods to image bony surfaces through soft tissue layers is, in fact, quite possible. However, there were weaknesses in the methods employed and, as such, the results would have to be confirmed. The only reliable method through which an unquestionable conclusion could be drawn on this topic would be to image intact knee joints, cadaveric or otherwise.

Chapter 9 Future Work, Practical Implementation and Conclusions

Four main hurdles must be overcome before a system can be considered fully functional. These are: i) the accurate reconstruction of complex bony surfaces, ii) the same task, but with the inclusion of soft tissues, iii) overcoming line of sight issues to fully reconstruct the intact knee joint and iv) the practical implementation of the system in a non-laboratory environment. While some of these four points have been addressed herein, others have not. This chapter, therefore, assesses the level of completion of each of these four topics and discusses possible improvements, solutions and

implementation methods. Finally, an overview of the work presented in the thesis is presented and conclusions are drawn.

9.1 Reconstruction of Complex Bony Surfaces

The most important and fundamental aspect of the system was to ensure that it was possible to accurately reconstruct complex bony surfaces using both synthetic aperture methods and robotic positioning. This was of particular importance for two reasons: firstly, synthetic aperture methods, while common in research in NDT, have rarely been used for surface imaging. As such, there is little evidence to suggest whether or not such a methods would be suitable for bone surface extraction. Secondly, in ultrasound bone imaging, both synthetic aperture methods and robotic positioning are rare and it was unknown whether accurate reconstruction would be possible.

Based on the work presented, it is believed that the first objective has been achieved. While imaging of a number of simple, metallic shapes exposed weaknesses in the system, the influence of these frailties was amplified due to the rudimentary nature of such samples. A number of complex surfaces which represent both the acoustic characteristics and shape of the bones within the human knee have been reconstructed to within an acceptable accuracy of under 1mm. However, there are a number of ways in which the results of this task could be improved.

Throughout this work, it was established that while TFM offered the most accurate surface representation, the associated side lobes could be of such intensity and coverage that highly inaccurate surface profiles could be extracted. While the use of a narrow aperture SAFT could alleviate this, this imposed a loss of resolution and coverage. A solution to this could be the creation of a side lobe reduction algorithm. This would minimise or eliminate the side lobes in an image, while retaining the positive aspects of TFM. While a number of groups have managed to reduce side lobes via data processing methods (Hunter et al., 2008; M. Li & Hayward, 2012; Prado, Higuti, & Kitano, 2013), these have not been designed specifically to reduce the side lobe levels. Therefore, a further algorithm could be created which would take

advantage of the fact that the method by which side lobes are formed within TFM is well understood.

For example, a form of recursive imaging could be employed, whereby an initial image is reconstructed and the position strongest true reflectors could be identified through global thresholding (eg. $>-2\text{dB}$). Using this information, the approximate corresponding samples within the A scans could be identified and set to zero. The reconstruction method could then be applied to the altered FMC data. This would result in an image which would not include the strongest reflectors and, importantly, the associated side lobes. The two scalar images resulting from this method could then be added together. This would limit the effects of the side lobes of the strongest reflectors and increase the intensity of the weaker, but still valid reflectors.

Alternatively, different probe parameters could be employed which would reduce the intensity of the side lobes formed. It has been shown that probe parameters influence the nature of side lobes heavily, with lower centre frequencies and smaller apertures contributing to reductions in side lobes. While both of these characteristics also have associated weaknesses, a balance could be found between the positive and negative aspects of the probe parameters. For example, it was shown in Chapter 7 that a 2.25MHz probe significantly reduced the side lobe intensity relative to that found with the 5MHz probe. While there was a loss of surface resolution, it was such that accurate surface extraction could still take place. This effect could be further explored *in silico* so that an optimal centre frequency could be defined. This could be achieved by simply testing a number of representative samples (such as those presented in Chapter 7) with a number of probe configurations. These would be based on real probe parameters, employing a range of centre frequencies. The resulting data would allow for an optimisation of centre frequency.

One of the reasons for inaccurate surface reconstructions was inaccurate TCP calibration. There were inherent problems with the calibration parts themselves, rather than the principle of the procedure. This could be remedied by accurate manufacture of the suggested parts, which would allow for more robust orientation calibration. Achieving this with traditional metal machining methods proved difficult and time consuming. Additionally, while the parts can be tested after completion, there is no

way to ensure a particular level of accuracy during manufacture. An alternative to traditional manufacture, therefore, could be to 3D print the parts, which would ensure they were made to a known level of accuracy, relative to the associated CAD drawings. In addition to this, the results presented herein could have been made more robust by repeat experiments. For each experiment type, only one sample was employed, due to time constraints. As such, it could be argued that the ability of system to reconstruct with accuracy consistently has not proved possible. This could have been mitigated by employing multiple samples and performing the same inspection scenario on each. Doing this would allow for a statistical analysis that would prove the repeatability of the system and provide validation of the results.

To conclude, while there are a number of methods by which the work presented herein could have been improved, the principle of accurate 3D surface reconstruction using robotic positioning and synthetic aperture methods has been proven. Future work using these methods would, therefore, focus on improvement rather than creation of a proof of concept.

9.2 Reconstruction of Complex Bony Surfaces with Preceding Soft Tissue Interfaces

The second major hurdle in completing a full proof of concept for the system was to show that the levels of accuracy achieved with bare bone imaging could be replicated when soft tissues were included. While this was investigated and the results were positive to an extent, a number of weaknesses were identified with the methods, bringing into question the validity of the results. There are a number of ways in which the current system could have been employed to improve its capability in this field without major changes to hardware and processing methods.

The first of these possibilities would be to improve the FE modelling technique presented in Section 7.2. As discussed previously, this could include employing a higher resolution image, allowing for more accurate material segmentation. Additionally, a comprehensive literature search could be employed to provide a more

robust definition of the soft tissues. Finally, a 3D model could be employed, which would allow for a higher level of accuracy and realism. To achieve this, material segmentation of CT or MRI data could be achieved using a specialist medical image processing package, such as Mimics (Materialise NV, Leuven, Belgium). This would allow for the production of a CAD model of the knee, with segmented materials, which would be used in PZFlex via the CAD importation tool. Doing this would not only allow for more accurate FMC simulation, but would allow for testing with different probe geometries and frequencies, including 2D arrays. Additionally, with simple repositioning of probe models being possible, different robot paths could be tested and compared. Further to this, the possibility of reconstruction tendon and ligaments could be tested using this method, which could be of great benefit to the capability of the system. While such implementations would require a very large finite element count, the solution times could be vastly reduced by use of the PZFlex Cloud. Simulation using FE could, therefore, be an important and powerful tool which would reduce the problems associated with using real specimens or subjects.

In addition to FE modelling, soft tissues could be modelled by physical soft tissue mimics. This has been carried out, but was simplistic in that only one layer of tissue was employed. This could be improved by including additional layers, each of which should have different material properties and be set at different angles. This could be further improved by carrying out material characterisation on the phantom materials and tailoring the constituent materials to recreate the properties of particular soft tissues. This method is, however, limited in use, in that only small bone samples could be employed and anatomical soft tissue shapes and dimensions could not be recreated. These limitations could, however, be overcome simply by acquiring a commercial, anatomically accurate, full leg phantom. While such products are available, such as the Blue Phantom Muskuloskeletal Model, the cost is excessive (~£4,000). If such a cost was within a budget, however, such an item would be a valuable tool, allowing for multiple inspection techniques to be tested, providing accurate results.

Despite the advantages of tissue substitutes, there is no guarantee that they will accurately simulate real tissue. As such, it is preferable that inspections be performed on real, intact legs. This could be achieved by using animal specimens, but these, again

may not accurately simulate human samples, due to differences in shape, scale and tissue properties. The most desirable option for determining the effect of soft tissue inclusion, then, is to employ human samples. There are two ways in which this could be achieved: firstly, fresh, frozen cadaveric specimens could be employed and, secondly, human subjects could be used.

Cadaveric specimens have been employed before in bone surface imaging as a means of acquiring accurate and realistic data (Barratt et al., 2008b; C. S. K. Chan et al., 2004; Foroughi et al., 2007; Talib et al., 2005; Vincent et al., 2004). This method, however, poses a number of challenges, in that it would require the department in which the work was carried out to have a licence to work with human tissues. This was not possible in this body of work, due to the immobility of certain pieces of equipment. However, if the necessary equipment could be brought together in the same department and ethical approval could be gained, this would provide the most accurate and realistic results.

Human subjects could also be employed, but this too would cause problems. While the immobility of the robots would not be a problem, as the subjects could be moved, human-robot interaction is tightly controlled due to safety concerns. As such, the industrial robots employed herein would be unsuitable for human interaction of any type. However, safety in human-robot interaction is becoming a popular research topic, with interaction necessary in many fields (Kosuge & Hirata, 2004; Kulić & Croft, 2005). This is particularly true in industrial robotics (Bicchi, Peshkin, & Colgate, 2008), as many tasks require both interaction and cooperation between humans and robots. One of the methods that have been employed to aid in such tasks is the development of intrinsically safe robots. These machines typically have high levels of compliance and are backdrivable, which provides safety around humans. An example of such manipulators is the KUKA-LBR series of lightweight robot arms (Bischoff et al., 2010), an example of which can be seen in Figure 9.1. This uses high performance actuators and integrated torque sensors to allow for sophisticated control algorithms (Bicchi et al., 2008).




Figure 9.1: The KUKA LBR iiwa 14 R820.

Such a robot could be employed in the task of scanning the knee of a human subject, as the safety of the subject would be guaranteed. In fact, a particular type of this robot has been tailored for biomedical applications and has been employed in other research efforts which require similar safeguards (Hennes, Bollue, Arenbeck, & Abel, 2013; Mönnich, Stein, Raczkowski, & Wörn, 2010). While scanning the full knee would be near impossible if the links of the robot were to remain outwith the water, simple scans could be performed on one side of the leg, with the subject sitting in a water bath.

The potential improvements described so far in this section could be implemented into the current system, as they were direct improvements on established practices. However, there are a number of published methods that were not included as part of the system which could have aided in bone imaging through soft tissues. There is a whole field of study within bone imaging which relates to image processing and contour extraction which was not been fully investigated during this project. This is discussed in Section 2.2.14. It is not yet known if any of these methods would improve performance, as the design thereof was not intended for this purpose. Importantly, however, these methods have yet to be investigated and, as such, could offer improvements in the extraction of bone surfaces.

9.3 Reconstruction of the Intact Knee Joint

While the problem of soft tissue imaging was addressed to an extent within the presented body of work, difficulties associated with line of sight with intact knee joints were not addressed at all. Two regions in particular will likely cause problems when attempting to image the entire surface: the condylar surfaces within the joint space and the region of the femur located behind the patella. As can be seen in Figure 9.2 (a), the joint surfaces between the tibia and femur has obvious line of sight issues. This is confirmed in part (c), which shows a pressure map of an element firing on the same sample. While ultrasonic energy can enter the joint space, there is no chance of reflection back to the probe face. The problems caused by the patella can be seen in Figure 9.2 (d), which shows how the pressure struggles to penetrate the cortical layer, blocking all ultrasonic energy from interacting with the region behind the patella.




Figure 9.2: Colour slices of the sagittal (a) and coronal (b) planes of the knee joint. Also shown are corresponding maximum pressure maps (c, d).

There are a number of ways in which the level of detriment caused by this could be investigated; with the first method being to employ a 3D FE model. Using the PZFlex Cloud, as described above, a 3D model could be investigated from a number of probe positions. The resulting images would give an indication of which regions were completely inaccessible, which would allow for different probe positions to be tested. As well as this, intact human or animal knee joints could be inspected, allowing for the same outcomes to be determined.

While particular probe positions might allow for greater access to certain regions, it would be impossible to overcome the physical limitations of ultrasonic waves. As such, there will be regions which are impossible to image using a fixed knee position. The methods mentioned so far will, therefore, be inadequate in reconstructing the entire

joint surface and alternatives will be necessary. One possible technique could be to scan the same subject more than once, with each scan being performed with the knee in a different pose. In scanning the knee in extension and then in flexion, surfaces that were previously inaccessible will become exposed. This effect can be seen in Figure 9.3, which shows the knee joint in extension in part (a) and in flexion in part (b). It is clear that different parts of the femoral condyles are exposed and can, therefore, be reconstructed.




Figure 9.3: Knee in extension (a) and flexion (b), showing the possible regions for inspection for movement analysis in red.

If this method was applied, it would be necessary to spatially relate the surface points obtained from the first scan to those of the second. This could be achieved by scanning parts of both the femur and tibia which would remain exposed at all times, despite any angle of flexion. Examples of such regions are shown in red in Figure 9.3, which would serve as references from which the necessary rotations and translations for the tibial and femoral point clouds could be calculated accurately. To do this, once the points extracted from each image had been rotated and translated by the robot positional data,

the points representing the tibia and femur would be separated. The points representing the associated reference region would then be matched with the corresponding set of points from the previous scan. This step could be achieved by providing an estimated flexion angle, shown as α in Figure 9.3 (b). Treating the knee joint as a simple hinge joint, this could serve as an elemental rotation around one axis, which would serve as a rough registration. From this point, ICP could be applied to match the second reference finely with the first. The same rotations and translations would then be applied to the surface points representing the epiphysis.

This method should, given an appropriately large flexion angle, provide enough access to the femoral condyles to accurately reconstruct the surfaces. Additionally, further scans could be included so that different regions would become exposed at every step. However, there are a number of potential problems with this method. Firstly, the patella will remain in place, blocking ultrasonic access to regions at the centre of the distal femur ie. the patellar notch. The regions that are blocked will, however, change and it is hoped that enough variation would be provided, so that all femoral surfaces would be exposed during some part of the inspection. It should also be noted, that parts of the surfaces of the tibial condyles would still be inaccessible, despite any angle of flexion. While the condyles would see different levels of exposure with different flexion angles, certain regions would remain inaccessible throughout. The hole-filling action of the point wrapping feature in Geomagic Wrap could compensate for this and recreate the surface with the necessary accuracy. However, if this is not possible, alternative methods would be necessary. In particular, SSMs (as mentioned in the literature review) could be employed to compensate for the sparsity of the resulting data.

9.4 Practical Implementation

In the event that intact human knee joints could be accurately reconstructed using the apparatus and methods described thus far, a system must then be produced to translate the success from the laboratory to a clinical environment. If the data capture and processing methods had been proven, the major task would be to provide accurate

probe manipulation, taking into account the necessities of patient safety, full spatial access to the knee, coupling media and the housing thereof. As yet, full immersion water baths have been employed. This would cause practical problems during implementation, in that all areas of the knee would have to be accessed, which would likely require that a six axis robot be at least partially submerged. With the robots employed herein and those suggested in the previous section, this would not be possible due to possible wearing of joints and for concerns for electrical safety.

One of the possible solutions to this would be to employ a fixed gantry, much like those employed in CT. Similar ultrasound tomography systems have been investigated before, but have focussed on smaller, more accessible regions, such as the breast (Birk, Zapf, Balzer, Ruitter, & Becker, 2012; Gemmeke & Ruitter, 2007; Gemmeke et al., 2010; Zapf, Schwarzenberg, & Ruitter, 2008). The movement of these systems is limited, usually only providing rotation around one axis, as the sample is small enough not to require further inspection and large transducer numbers are also usually exploited.

A more complex positioning system would, therefore, be required for imaging the entire knee joint. One possible way in which this could be achieved is presented in Figure 9.4, which shows a scanning gantry with four probes. The probes depicted have a square face, as they are assumed to be 2D arrays, but this is not necessarily this case, as 1D arrays should also be capable of full coverage. The leg would be placed in the centre of the system, so that all the probe faces would be angled towards the surface. Each probe could then scan in two directions, independent of the other probes. This would allow each probe to fully scan one side of the leg, meaning all four probes would allow for full capture, in the same vein as the cuboidal paths employed previously.




Figure 9.4: Possible ‘box’ design for ultrasound CT, with four probes (a), each with two axis of motion (b), allowing for full coverage using the cuboidal path employed herein.

Scan times could also be reduced in this setup by using multiple probes at the same time. As demonstrated in Figure 9.5, this could be achieved by employing two opposing probes, moving in opposite directions. In doing this, the leg of the subject would act as a barrier between the two probes and would ensure that transmissions from one probe would not be received by the other, nor would the transmissions from either probe interfere with those of the other. The scan time could be further reduced by employing all four probes at once, possibly using the method described above in combination with moving one pair of probes from bottom to top and the other from top to bottom.




Figure 9.5: Method of employing two probes at once, with the leg ensuring transmissions from one probe do not interfere with those of the other.

Other scanning mechanisms could be employed. For example, that shown in Figure 9.6 has more in common with the method employed in CT, with the probes creating a spiral pattern. In this configuration, the movement of each probe is dependent on the other. With the rotations and linear movement shown in parts (b) and (c) of Figure 9.6, this maintains an opposing position of each probe relative to the other. This, therefore, ensures that the transmissions and receptions of the probes do not interfere, as depicted in part (d).




Figure 9.6: Ring configuration of ultrasound computed tomography, with two probes with linked rotational and linear movement.

These configurations do not, however, account for the coupling medium. It has been assumed until this point that the leg of the subject would be fully submerged. The devices would, then, be housed within a water bath with dimensions slightly larger than the maximum dimensions of the device. While submersible ultrasound probes are common and would not present a problem, immersing linear actuators in water is rare. Many applications require splash proof actuators, but fully immersed applications are uncommon and, as such, alternatives might be necessary.

A possible alternative to fully immersing the device could be to employ a balloon within the device. As illustrated in Figure 9.7, the balloon could be attached to the corners of the device. Using a deformable material and filling it with enough water would provide the pressure to force a contact with the face of the probe. The balloon would deform as the probes moved, maintaining contact at all times. The contact

between the probe face and the balloon could be further ensured by attaching a curved face to the probe, using a low loss rubber, with an acoustic impedance comparable to water. This could also be achieved with an additional water-filled balloon encapsulating the face of the probe. Similar methods have been employed with wheel probes within NDT (Cawley & Adams, 1989; Freemantle, Hnakinson, & Brotherhood, 2004) and would allow better contact with the inner balloon. The coupling between the probe and the balloon could be achieved either by way of ultrasound gel or by a continuous spray of water on the outer balloon surface. The latter would be acceptable as many linear actuators are sold as ‘splash proof’.




Figure 9.7: Action of balloon coupling, with deformation allowing for continuous contact.

An alternative to a bespoke fixed gantry system is to employ a human-safe 6-axis robot. In this situation, the path of the robot is important, in that it must both guarantee the safety of the subject and ensure full coverage of the surfaces of the knee joint. The first and most simple way in which this could be achieved would be to use very basic paths composed of lines and arcs. This would require input of approximate leg dimensions and a simple registration of the leg position. Seeing as there might be regions which are inaccessible for a fixed robot system, two robots might be preferred. Such a configuration is shown in Figure 9.8 and could decrease scan time, as two probes could be employed at the same time, using the leg as a barrier between the two.

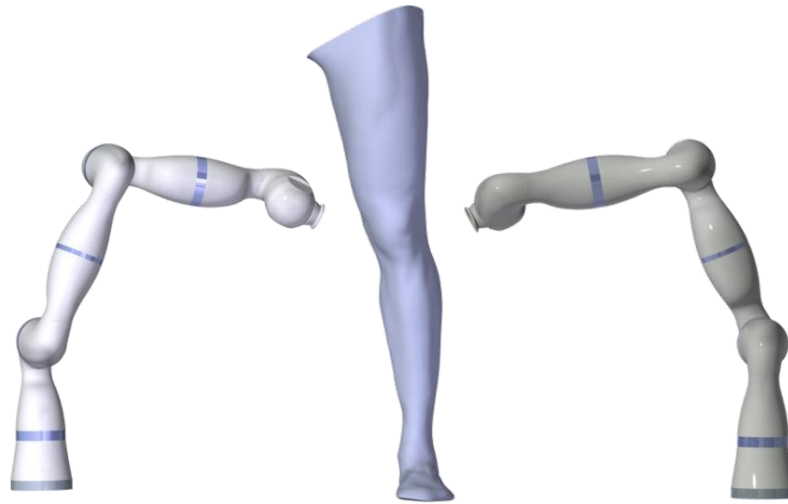


Figure 9.8: Possible double robot configuration allowing for full access to the knee joint.

In addition to a simple cuboidal or spherical path, patient-specific robot paths could be implemented. In the same way as laser scanners are employed to reverse engineer and generate part-specific paths for inspection within NDT (R. H. Brown et al., 2015), a laser scanner could be employed to generate a path that would be normal to the surface of the skin throughout the entire scan. Using similar path generation methods as described in Section 3.3.1.3, this would ensure both full coverage of the region of interest and the maintenance of a defined standoff between the probe face and the surface of the skin. Additionally, this method would allow for multiple scans with different levels of flexion, as described above. With the gantry systems, on the other hand, space would be limited and, as such, it would be difficult to perform additional scans with any other levels of flexion.

The reference model of the surface of the skin could be acquired in a number of ways. The most obvious method would be to employ a commercial scanner, such as the Faro Quantum. However, this would be time consuming and would be susceptible to inter- and intra-operator variability. A more robust method would be to mount a laser scanner on the end effector of the robot. Following a simple square path along the edges of the unfilled water bath, an accurate surface mesh of the skin of the patient could be acquired with high positional accuracy.

In addition to this method, complex, patient-specific paths could be defined using ultrasound data. A simple path, as described above could be employed initially, which would allow for an approximate 3D surface model to be created in the same way as described in section 5.2. From this, a patient-specific path could be created which would allow for normality and standoff to be maintained between the probe and the bone surface where possible. For this task, SSMs would be well suited, in that simple landmark data could be employed, which would return a relatively accurate, full 3D surface. This, in addition to position and orientation data relative to the robot coordinate system, would provide a good reference model with which to define the robot path.

Given the commercial availability of the robots, this method would appear preferable to the development of a bespoke gantry system. However, the issue of ultrasonic coupling has yet to be addressed. One solution to the need for coupling media would be to fully immerse the robots during the scan. While this would not be possible for most six-axis robots, there are some which are partially waterproofed. The KR AGILUS WP (KUKA Robotics, Germany), for example, has a protection rating of IP67 which allows for full immersion up to 1 metre for up to 30 minutes. If two robots were employed, the scan could be completed in under that time. As such, it could be possible to fill the tank after the laser scan, then drain it after the ultrasound scan.

There are also alternatives to this. For example, within NDT, free water jet systems have been employed to provide a means of coupling and could be used to couple the probe face to the skin. While these systems have been employed almost exclusively for through-transmission testing, they have been modified to allow for pulse echo applications (Cuevas & Hernández, 2014; Deutsch, 2000). However, it is unknown if such a configuration would allow for the complexities of FMC. Additionally, ensuring that the water column contains no bubbles is not a trivial task, requiring expert knowledge. As such, there is much cause to doubt the likelihood of success using a water jet coupling system.

An additional alternative found in NDT is bubbler systems, which have a small, water-filled chamber between the probe face and the inspection surface (Cuevas & Hernández, 2014). This ensures there is a constant coupling, but requires a level of

contact with the surface. This could be achieved by employing a patient specific robot path, generated using laser scanning, as described above. The high level of compliance offered by many lightweight robots would ensure such contact would not injure the subject. It should be mentioned that this method is not particular to bubbler systems and could be replicated with a probe encased in a balloon of water or a low loss rubber (Wagle & Kato, 2009). In this case, ultrasonic coupling could be ensured by applying a continual spray of water above the probe.

If the SSMs described in the previous section were successful in accounting for the high level of variability within the knee joints of OA sufferers, it would not be necessary to image the entire joint, as only landmarks would be necessary. If this was the case, it would be possible to have the subject sit or stand in a water bath and have a robot controlled probe dip into the water, immersing the probe, but not the joints of the robot. The probe mount needed to achieve this would constrict the movement of the probe in a similar manner to that seen with the probe mount used herein. However, the mount could be designed such that access was guaranteed to the regions in which the landmarks were located.

As of yet, only robotic probe manipulation has been considered. However, it is recognised that medical ultrasound inspections are currently carried out by human operators, which might imply that patients would be more comfortable with a person, rather than a robot. In particular, older people are less likely to accept robotic technology than the young (Broadbent, Stafford, & MacDonald, 2009). Therefore, the technique described directly above could be employed, but using a human operator to control the end effector of the robot. Lightweight robots are highly compliant and can be handled by users. They would, therefore, offer the possibility of an increase in the accuracy of the positional data over more ubiquitous stereoscopic optical tracking methods (Wiles, Thompson, & Frantz, 2004). This method would, however, be highly susceptible to intra- and inter-operator variability, both in terms of data quality and in inspection time.

When using any of the systems mentioned above, it would be of vital importance to ensure that the leg of the subject remained still during data capture. As it would be difficult if not impossible for a patient to remain still for tens of minutes, an

immobiliser would be necessary. As it would be necessary to maintain access to the knee during the data capture, the positioning of the immobiliser would have to take this into account. This could be accomplished by employing immobilisation straps as shown in Figure 9.9, above and below the knee. While this would not necessarily stop any movement of the muscles, the positioning of the straps should be such that the bones of the joint would not move during inspection.



Figure 9.9: Possible position of immobilisation straps that would ensure no movement of the joint during ultrasonic bone surface capture.

There are, therefore, numerous methods that could be employed to implement the system in a clinical environment. While human manipulation is a possibility that should not be overlooked, the negative aspects of the method make it a significantly less attractive technique. As such, robotically guided manipulation is preferred, of which two main areas were presented: gantry systems and six-axis robots. The six-axis systems offer far more possibilities in terms of probe paths, but the problems with coupling are lessened if gantry systems are employed. Therefore, the decision on which option to peruse would also have to take into consideration other factors including the availability of suitable six-axis robots, the consumption of time and resources in the

development of bespoke gantry systems and the capability of employing multiple probes.

9.5 Conclusions

A system has been developed which allows for accurate reconstruction of complex bony surfaces in three dimensions using ultrasound imaging. Using robotic positioning and synthetic aperture methods, techniques which are usually confined to NDT have been applied to the biomedical task of imaging the bony surfaces of the knee joint for planning and registration in robotic knee arthroplasty.

In order to test the system at a fundamental level, a number of metallic samples with simple geometries were imaged. This provided information about the weaknesses and strengths of the system in a number of different areas. An outcome of this work was the confirmation that TFM, while providing impressive surface resolution and coverage, was susceptible to high intensity side lobes. This made extraction of the surface profiles and, therefore, the reconstruction of the 3D surface point clouds inaccurate. By limiting the synthetic aperture width, SAFT offered an alternative to TFM which significantly reduced the impact of side lobes, but came with the cost of a reduction in both surface resolution and coverage.

In addition to this, this study exposed the detrimental effect that certain 3D surface features can have when using 1D arrays. In particular, sharp edges can lead to ghost surfaces, caused by the isotropic nature of such reflections. Additionally, curved surfaces can cause a distortion of the surface profile in the resulting image due to the reflections from different regions on the curve arriving at the receiver at different times. These problems essentially arise from the fact that 1D arrays are subject to three dimensional effects – a fact which is difficult to account for while reconstructing 2D images. Despite this, the observed effects were heavily amplified by the fact that the surfaces were polished and simple. Therefore, such effects would be unlikely to be as damaging in bone surface imaging.

In an effort to test the ability of the system to reconstruct complex surfaces, both a human-shaped epoxy resin and real bovine distal femur were imaged. Using a simple cuboidal scan path, the surfaces were fully reconstructed, with the mean errors of 0.828mm and 0.877mm for the composite and bovine samples, respectively. Echoing the results in Chapter 4, TFM, while performing well in some areas, produced high intensity side lobes that caused damaging levels of inaccuracy. Again, narrow aperture SAFT was found to produce the best results. Further to this, it was found that it was beneficial to favour the points in the surface profile which were closest to the probe face during surface profile extraction.

The success of Chapter 5 and Chapter 6 was found without the inclusion of soft tissue and the problems that can be encountered due to it. Therefore, investigations into the effect of soft tissue inclusion were made in Chapter 7 and Chapter 8. This began with the employment of FE models of sagittal and coronal CT slices of the knee joint, so that FMC with the correct soft tissues and scale could be simulated. With only tendon proving to cause any real problems, the study provided results which suggested that bone surface identification was possible using both FMC and synthetic aperture methods. However, due to simplified material definitions, the simulations may not have been accurate.

Therefore, to further test the effect of such tissues, a soft tissue mimic was applied to a bovine femoral condyle. While it must be mentioned that only one layer was employed, the inclusion of the mimic caused no problems in bone surface identification. The studies carried out suggested, therefore, that the inclusion of soft tissue would not seriously affect the system in accurately identifying the bony surfaces of the knee joint. However, it was concluded that more work would be required to produce a definite result.

In summary, then, of the work carried out in this thesis, it could be said that two of the major hurdles were addressed, while the third was not. The first step of reconstructing complex bony surfaces to sub-millimetre accuracy was achieved. The second obstacle of achieving the same, but with intact soft tissues was addressed, but not completed. Finally, the task of accurately reconstructing all the surfaces of the knee joint *in situ* was not tackled in any experiment or simulation.

The final chapter described and discussed possible ways in which the incomplete areas of the project could be fulfilled. In particular, a number of image processing and surface contour extraction methods were suggested that could help improve the accuracy of surface profile extraction with soft tissues. Further to this, a number of methods were described which could compensate for the line of sight problems faced when imaging the intact joint. Finally, a number of ways in which the system could be implemented in a clinical environment were presented. These included bespoke gantry systems, as well as commercially available 6-axis robotic manipulators. The suggestions in this chapter were broad and were based on successes either in this thesis or in the work of other groups. As such, it serves as evidence for high probability of success of this project given the correct resources.

With a likelihood of success in this project in mind, it should be mentioned that while this work has been aimed towards preoperative imaging for robotic knee arthroplasty, there are numerous other areas in which it could be of use. For example, it could be used to provide the surface meshes necessary for patient specific knee implants (Harrysson, Hosni, & Nayfeh, 2007). Also within the realm of robotic knee arthroplasty, some of the methods could be employed for bone surface registration. Currently, the physical bone is registered to the CAD model by touching a probe to the bony surfaces (Cobb & Pearle, 2013). This requires that the knee is incised, the patella is removed and many of the soft tissue structures are displaced. Therefore, when soft tissue balancing is carried out, the effect of the soft tissue is absent. If ultrasound imaging was used to perform this registration, all the soft tissues would be intact, providing a more accurate planning procedure. The system could also be used as a clinical assessment tool for patients who have undergone knee arthroplasty and for assessment of the development of osteoarthritis in general. Post-operative CT scans are undesirable due to both cost and the application of doses of ionising radiation to the patient – both of which could be circumvented with the use of an ultrasound imaging system. This would allow for more accurate assessment not only for the benefit of the patient and the associated treatment plan, but also for assessment of the clinical return from robotic arthroplasty systems, which has yet to be proven (Burnett & Barrack, 2013; Kamat et al., 2009). In addition to this, the system could be applied

to other joints for other procedures, providing a non-ionising method of producing accurate bone surface models. Finally, the system could be employed within NDT to provide a 3D surface profile of a part under inspection. This could be used to allow for correction for refractive effects in dual media, providing more accurate results.

To conclude, while the work presented in this thesis does not amount to a complete and functional system, evidence of proof of concept was delivered. Further to this, the problems likely to be encountered in future development were discussed and possible solutions considered. This study, therefore, serves as evidence for further investigation into this area, which would be of benefit not only to robotic knee arthroplasty procedures but other fields also.

References

- Abolmaesumi, P., Salcudean, S. E., Zhu, W.-H., Sirouspour, M., & DiMaio, S. P. (2002). Image-Guided Control of a Robot for. *IEEE TRANSACTIONS ON ROBOTICS AND AUTOMATION*, 18(1), 11–23.
- Adams, R. D., & Cawley, P. (1988). A review of defect types and nondestructive testing techniques for composites and bonded joints. *NDT International*, 21(4), 208–222. doi:10.1016/0308-9126(88)90333-1
- Adams, R. D., & Drinkwater, B. W. (1997). Nondestructive testing of adhesivelybonded joints. *NDT & E International*, 30(2), 93–98. doi:10.1016/S09638695(96)00050-3
- Albert, J. M. (2013). Radiation risk from CT: Implications for cancer screening. *American Journal of Roentgenology*, 201(July), 81–87. doi:10.2214/AJR.12.9226
- Anderson, J. J., Anthony, J. M., Zhang, Y., Wilson, P. W. F., Kelly-hayes, M., Wolf, P. A., ... Kannel, W. B. (1993). The effects of specific medical conditions on the functional limitations of elders in the framingham study. *American Journal of Public Health*, 84(3), 351–358.
- Angelsena, B. A. J., Tarp, H., Holmb, S., Kristoffersen, K., & Whittingham, T. A. (1995). Which transducer array is best ? *Europaen Journal of Ultrasound*, 2, 151–164.
- Arbeille, P., Poisson, G., Vieyres, P., Ayoub, J., Porcher, M., & Boulay, J. L. (2003). Echographic examination in isolated sites controlled from an expert center using a 2-D echograph guided by a teleoperated robotic arm. *Ultrasound in Medicine and Biology*, 29(7), 993–1000. doi:10.1016/S0301-5629(03)00063-2
- Arnau, A., & Soares, D. (2008). Fundamentals of piezoelectricity. In *Piezoelectric Transducers and Applications* (pp. 1–38). doi:10.1007/978-3-540-77508-9_1
- Arokoski, J. P. A., Jurvelin, J. S., Väättäin, U., & Helminen, H. J. (2000). Normal and pathological adaptations of articular cartilage to joint loading. *Scand J Med Sci.*, 10, 186–198.
- ArthritisResearchUK. (2013). *Osteoarthritis in general practice*. The Medical press.
- Azhari, H. (2010). APPENDIX A TYPICAL ACOUSTIC PROPERTIES. In *Basics of Biomedical Ultrasound for Engineers* (pp. 313–314). John Wiley & Sons, Inc.
- Azhari, H. (2010). *Basics of Biomedical Ultrasound for Engineers*. Wiley-IEEE Press.
- Bachmeier, C. J., March, L. M., Cross, M. J., Lapsley, H. M., Tribe, K. L., Courtenay, B. G., & Brooks, P. M. (2001). A comparison of outcomes in osteoarthritis patients undergoing total hip and knee replacement surgery. *Osteoarthritis and Cartilage / OARS, Osteoarthritis Research Society*, 9(2),

137–46. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11330253>

- Bae, W. C., Schumacher, B. L., & Sah, R. L. (2007). Indentation probing of human articular cartilage: Effect on chondrocyte viability. *Osteoarthritis and Cartilage / OARS, Osteoarthritis Research Society*, *15*(1), 9–18. doi:10.1016/j.joca.2006.06.007
- Bandini, a, Chicot, D., Berry, P., Decoopman, X., Pertuz, a, & Ojeda, D. (2013). Indentation size effect of cortical bones submitted to different soft tissue removals. *Journal of the Mechanical Behavior of Biomedical Materials*, *20*, 338–46. doi:10.1016/j.jmbbm.2013.02.011
- Bannouf, S., Robert, S., Casula, O., & Prada, C. (2013). Data set reduction for ultrasonic TFM imaging using the effective aperture approach and virtual sources. *Journal of Physics: Conference Series*, *457*, 012007. doi:10.1088/17426596/457/1/012007
- Barnett, S. B., Rott, H., Ter Haar, G. R., Ziskin, M. C., & Maedai, K. (1997). THE SENSITIVITY OF BIOLOGICAL TISSUE TO ULTRASOUND. *Ultrasound in Med. and Biol.*, *23*(6), 805–812.
- Barratt, D. C., Chan, C. S. K., Edwards, P. J., Penney, G. P., Slomczykowski, M., Carter, T. J., & Hawkes, D. J. (2008a). Instantiation and registration of statistical shape models of the femur and pelvis using 3D ultrasound imaging. *Medical Image Analysis*, *12*, 358–374. doi:10.1016/j.media.2007.12.006
- Barratt, D. C., Chan, C. S. K., Edwards, P. J., Penney, G. P., Slomczykowski, M., Carter, T. J., & Hawkes, D. J. (2008b). Instantiation and registration of statistical shape models of the femur and pelvis using 3D ultrasound imaging. *Medical Image Analysis*, *12*(3), 358–74. doi:10.1016/j.media.2007.12.006
- Barratt, D. C., Penney, G. P., Chan, C. S. K., Slomczykowski, M., Carter, T. J., Edwards, P. J., & Hawkes, D. J. (2006a). Self-Calibrating 3D-Ultrasound-Based Bone Registration for Minimally Invasive Orthopedic Surgery. *IEEE Transactions on Medical Imaging*, *25*(3), 312–323.
- Barratt, D. C., Penney, G. P., Chan, C. S. K., Slomczykowski, M., Carter, T. J., Edwards, P. J., & Hawkes, D. J. (2006b). Self-calibrating 3D-ultrasound-based bone registration for minimally invasive orthopedic surgery. *IEEE Transactions on Medical Imaging*, *25*(3), 312–23. doi:10.1109/TMI.2005.862736
- Bäthis, H., Perlick, L., Tingart, M., Lüring, C., Zurakowski, D., & Grifka, J. (2004). Alignment in total knee arthroplasty. *The Journal of Bone and Joint Surgery*, *86*(5), 682–687. doi:10.1302/0301-620X.86B5.14927
- Besl, P., & McKay, N. (1992). A Method for Registration of 3D Shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *14*(2).
- Bicchi, A., Peshkin, M. a, & Colgate, J. E. (2008). Safety for Physical Human–Robot Interaction. In *Handboof of Robotics* (pp. 1335–1348).

- Bilge, H., Karaman, M., & O'Donnell, M. (1996). Motion Estimation Using Common Spatial Frequencies In Synthetic Aperture Imaging. In *1996 IEEE ULTRASONICS SYMPOSIUM* (pp. 1551–1554).
- Birk, M., Zapf, M., Balzer, M., Ruiter, N., & Becker, J. (2012). A comprehensive comparison of GPU- and FPGA-based acceleration of reflection image reconstruction for 3D ultrasound computer tomography. *Journal of Real-Time Image Processing*, *9*(1), 159–170. doi:10.1007/s11554-012-0267-4
- Bischoff, R., Kurth, J., Schreiber, G., Koeppe, R., Albu-Schaeffer, A., Beyer, A., ... Hirzinger, G. (2010). The KUKA-DLR Lightweight Robot arm - a new reference platform for robotics research and manufacturing. *Robotics (ISR), 2010 41st International Symposium on and 2010 6th German Conference on Robotics (ROBOTIK)*, 1–8. doi:10.1108/01439910710774386
- Blomme, E., Bulcaen, D., & Declercq, F. (2002). Air-coupled ultrasonic NDE : experiments in the frequency range 750 kHz \pm 2 MHz. *NDT & E International*, *35*, 417–426.
- Brenner, D. J., & Hall, E. J. (2007). Computed Tomography — An Increasing Source of Radiation Exposure. *The New England Journal of Medicine*, *357*(22), 2277–2284.
- Broadbent, E., Stafford, R., & MacDonald, B. (2009). Acceptance of Healthcare Robots for the Older Population: Review and Future Directions. *International Journal of Social Robotics*, *1*(4), 319–330. doi:10.1007/s12369-009-0030-6
- Brown, C. P., Oloyede, a, Crawford, R. W., Thomas, G. E. R., Price, a J., & Gill, H. S. (2012). Acoustic, mechanical and near-infrared profiling of osteoarthritic progression in bovine joints. *Physics in Medicine and Biology*, *57*(2), 547–59. doi:10.1088/0031-9155/57/2/547
- Brown, R. H., Pierce, S. G., Collison, I., Dutton, B., Jackson, J., Lardner, T., ... Morozov, M. (2015). Automated Full Matrix Capture for Industrial Processes. *41st Annual Review of Progress in Quantitative Nondestructive Evaluation*, *34*, 1967–1976.
- Buckwalter, J. A., & Lohmander, S. (1994). Operative Treatment of Osteoarthritis. *The Journal of Bone and Joint Surgery*, *76-A*(9), 1405–1418.
- Buckwalter, J. A., & Mankin, H. J. (1997). Articular Cartilage. Part II: Degeneration and Osteoarthrosis, Repair, Regeneration, and Transplantation. *J Bone Joint Surg Br*, *79*, 612–632.
- Bude, R. O., & Adler, R. S. (1995). An easily made, low-cost, tissue-like ultrasound phantom material. *Journal of Clinical Ultrasound: JCU*, *23*(4), 271–273. doi:10.1002/jcu.1870230413
- Burnett, R. S. J., & Barrack, R. L. (2013). Computer-assisted total knee arthroplasty is currently of no proven clinical benefit: A systematic review knee. *Clinical Orthopaedics and Related Research*, *471*(1), 264–276. doi:10.1007/s11999-012-2528-8

- Caine, D. J., & Golightly, Y. M. (2011). Osteoarthritis as an outcome of paediatric sport: an epidemiological perspective. *British Journal of Sports Medicine*, 45(4), 298–303. doi:10.1136/bjsm.2010.081984
- Callahan, C. M., Drake, B. G., Heck, D. a., & Dittus, R. S. (1995). Patient outcomes following unicompartmental or bicompartamental knee arthroplasty. A metaanalysis. *The Journal of Arthroplasty*, 10(2), 141–50. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/7798094>
- Cawley, P., & Adams, R. D. (1989). Defect types and non-destructive testing techniques for composites and bonded joints. *Materials Science and Technology*, 5(5), 413–425. doi:10.1179/026708389790222366
- Chan, C. S. K., Barratt, D. C., Edwards, P. J., & Penney, G. P. (2004). Cadaver Validation of the Use of Ultrasound for 3D Model Instantiation of Bony Anatomy in Image Guided Orthopaedic Surgery. *Medical Image Computing and Computer-Assisted Intervention*, 3217, 397–404.
- Chan, V., & Perlas, A. (2011). *Atlas of Ultrasound-Guided Procedures in Interventional Pain Management*. (S. N. Narouze, Ed.). New York, NY: Springer New York. doi:10.1007/978-1-4419-1681-5
- Chanel, L., Nageotte, F., Vappou, J., Luo, J., & Cuvillon, L. (2014). Assistance to High Intensity Focused Ultrasound (HIFU) therapy : Real-time motion compensation using ultrafast ultrasound imaging ultrasound probe robotic arm. In *Surgteca* (pp. 1–4).
- Chatelain, P., Krupa, A., Navab, N., Chatelain, P., Krupa, A., Navab, N., ... Navab, N. (2015). Optimization of ultrasound image quality via visual servoing. In *IEEE Int. Conf. on Robotics and Automation*.
- Chen, Y., & Medioni, G. (1991). Object Modelling by Registration of Multiple Range Images. In *IEEE Int. Conf. on Robotics and Automation*.
- Chen, Y., Nguyen, M., & Yen, J. T. (2011). 7.5 MHz dual-layer transducer array for 3-D rectilinear imaging. *Ultrasonic Imaging*, 33(3), 205–16. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3229277&tool=pmc-entrez&rendertype=abstract>
- Chiao, R. Y., & Thomas, L. J. (1994). Analytic evaluation of sampled aperture ultrasonic imaging techniques for NDE. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 41(4), 484–493. doi:10.1109/58.294109
- Chiao, R. Y., Thomas, L. J., & Silverstein, S. D. (1997). Sparse array imaging with spatially-encoded transmits. *1997 IEEE Ultrasonics Symposium Proceedings. An International Symposium (Cat. No.97CH36118)*, 2(c), 1679–1682. doi:10.1109/ULTSYM.1997.663318
- Choong, P. F., Dowsey, M. M., & Stoney, J. D. (2009). Does accurate anatomical alignment result in better function and quality of life? Comparing conventional and computer-assisted total knee arthroplasty. *The Journal of Arthroplasty*, 24(4), 560–9. doi:10.1016/j.arth.2008.02.018

- Ciampi, F., Pujol, O., Ferr, J. M., & Radeva, P. (2012). Coronary Atherosclerotic Plaque Characterization By Intravascular Ultrasound. In *Ultrasound Imaging* (pp. 177–201). doi:10.1007/978-1-4614-1180-2
- Citak, M., Suero, E. M., Citak, M., Dunbar, N. J., Branch, S. H., Conditt, M. a, ... Pearle, A. D. (2012). Unicompartmental knee arthroplasty: Is robotic technology more accurate than conventional technique? *The Knee*, 4–7. doi:10.1016/j.knee.2012.11.001
- Citak, M., Suero, E. M., Citak, M., Dunbar, N. J., Branch, S. H., Conditt, M. a., ... Pearle, A. D. (2013). Unicompartmental knee arthroplasty: Is robotic technology more accurate than conventional technique? *Knee*, 20(4), 268–271. doi:10.1016/j.knee.2012.11.001
- Clark, W. W. (1991). Noise exposure from leisure activities: a review. *The Journal of the Acoustical Society of America*, 90(1), 175–81. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/1880286>
- Clarke, A. J., Evans, J. A., Truscott, J. G., Milner, R., & Smith, M. A. (1994). A phantom for quantitative ultrasound of trabecular bone. *Phys. Med. Biol.*, 39, 1677–1687.
- Cobb, J., & Pearle, A. (2013). Unicompartmental Knee Arthroplasty: Robotics. In *Knee Surgery using Computer Assisted Surgery and Robotics* (pp. 95–104). doi:10.1007/978-3-642-31430-8
- Cochran, S., Earnshaw, J., Bargain, B., & Kirk, K. (2002). Extending the Synthetic Aperture Focussing Algorithm to Deal With Flat and Curved Features in NDT. In *2002 IEEE ULTRASONICS SYMPOSIUM* (pp. 833–836).
- Conditt, M. A., & Roche, M. W. (2009). Minimally Invasive Robotic-Arm-Guided Unicompartmental Knee Arthroplasty. *The Journal of Bone and Joint Surgery*, 91(1), 63–68. doi:10.1097/01.BTK.0000026302.58597.68
- Cooper, I., Nicholson, I., Liaptsis, D., Wright, B., & Mineo, C. (2013). DEVELOPMENT OF A FAST INSPECTION SYSTEM FOR COMPLEX COMPOSITE STRUCTURE - THE INTACOM PROJECT. In *5th International Symposium on NDT in Aerospace* (pp. 13–15).
- Counter, S. A., Olofsson, A., Grahn, E. E. H. F., & Borg, E. (2005). MRI Acoustic Noise: Sound Pressure and Frequency Analysis I with Acoustic Dampers. *Journal of Magnetic Resonance Imaging*, 7(3), 606–611.
- Croft, P. (1990). Review of UK data on the rheumatic diseases - 3: osteoarthritis. *British Journal of Rheumatology*, 29, 391–395.
- Crowther, P. (2006). PRACTICAL EXPERIENCE OF PHASED ARRAY TECHNOLOGY FOR POWER STATION APPLICATIONS. *Insight*, 46(9), 525–528.
- Cuevas, E., & Hernández, S. (2014). Robot-based solution To Obtain an Automated , Integrated And Industrial Non-Destructive Inspection Process. In *6th International Symposium on NDT in Aerospace* (pp. 12–14).

- Cuevas, E., López, M., & García, M. (2012). Ultrasonic Techniques and Industrial Robots : Natural Evolution of Inspection Systems. In *4th NDT AEROSPACE*.
- Culjat, M. O., Goldenberg, D., Tewari, P., & Singh, R. S. (2010). A review of tissue substitutes for ultrasound imaging. *Ultrasound in Medicine & Biology*, *36*(6), 861–73. doi:10.1016/j.ultrasmedbio.2010.02.012
- Cunningham, L., Mulholland, A. J., Harvey, G., & Bird, C. (2012). Ultrasonic wave propagation in heterogeneous media Simulated data using PZFlex. *Proceedings of the Acoustics 2012 Nantes Conference*, (April), 1463–1468.
- Davies, B., Jakopec, M., Harris, S. J., Rodriguez y Baena, F., Barrett, A., Evangelidis, A., ... Cobb, J. (2006). Active-Constraint Robotics for Surgery. *Proceedings of the IEEE*, *94*(9), 1696–1704. doi:10.1109/JPROC.2006.880680
- Davies, J., Simonetti, F., Lowe, M., & Cawley, P. (2006). Review of Synthetically Focussed Guided Wave Imaging Techniques with Application to Defect Sizing, (figure 1), 1–9.
- Davies, M., Groom, C., Dunne, J., Xu, D., Owusu-Apenten, C., Epps, A., ... Gill, S. (2011). *Managing high value capital equipment in the NHS in England*. London.
- de González, A. B., & Darby, S. (2004). Risk of cancer from diagnostic X-rays. *The Lancet*, *363*, 1910. doi:10.1016/S0140-6736(04)16374-5
- Decking, J., Theis, C., Achenbach, T., Roth, E., Nafe, B., & Eckardt, A. (2004). Robotic total knee arthroplasty: the accuracy of CT-based component placement. *Acta Orthopaedica Scandinavica*, *75*(5), 573–9. doi:10.1080/00016470410001448
- Decking, R., Markmann, Y., Fuchs, J., Puhl, W., & Scharf, H.-P. (2005). Leg Axis After Computer-Navigated Total Knee Arthroplasty. *The Journal of Arthroplasty*, *20*(3), 282–288. doi:10.1016/j.arth.2004.09.047
- Denham, R. A., Jeffery, R. S., & Morris, R. W. (1991). Coronal alignment after total knee replacement. *The Journal of Bone and Joint Surgery*, *73*(5), 709–714.
- Deutsch, W. A. K. (2000). Automated Ultrasonic Inspection. In *WCNDT*. Dieppe, P., Basler, H. D., Chard, J., Croft, P., Dixon, J., Hurley, M., ... Raspe, H. (1999). Knee replacement surgery for osteoarthritis: effectiveness, practice variations, indications and possible determinants of utilization. *Rheumatology (Oxford, England)*, *38*(1), 73–83. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10334686>
- Dobie, G., Galbraith, W., Macleod, C., Summan, R., & Pierce, G. (2012). Synthetic Aperture Guided Wave Imaging using a Mobile Sensor Platform. In *2012 IEEE International Ultrasonics Symposium Proceedings* (pp. 1905–1908).
- Dobie, G., Summan, R., MacLeod, C., Pierce, G., & Galbraith, W. (2014). An Automated Miniature Robotic Vehicle Inspection System. In *AIP Conf. Proc.*
- Dobie, G., Summan, R., Pierce, S. G., Galbraith, W., & Hayward, G. (2011). A noncontact ultrasonic platform for structural inspection. *IEEE Sensors Journal*,

11(10), 2458–2468. doi:10.1109/JSEN.2011.2138131

- Donald, I., MacVicar, J., & Brown, J. (1958). Investigation of Abdominal Masses by Pulsed Ultrasound. *Lancet*, 11, 88–95.
- Drinkwater, B. W., & Wilcox, P. D. (2006). Ultrasonic arrays for non-destructive evaluation: A review. *NDT & E International*, 39(7), 525–541. doi:10.1016/j.ndteint.2006.03.006
- Dussik, K. T. (1941). Über die Möglichkeit, hochfrequente mechanische Schwingungen als diagnostisches Hilfsmittel zu verwenden. *Zeitschrift Für Die Gesamte Neurologie Und Psychiatrie*, 174(1), 153–168.
- Dziewierz, J., & Gachagan, A. (2013). Computationally Efficient Solution of Snell's Law of Refraction. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 60(6), 1256–1259.
- Dziewierz, J., Gachagan, A., Lord, N., & Mullholland, A. J. (2012). An applicationspecific design approach for 2D ultrasonic arrays . In *51st Annual Conference of the British Institute of Non-Destructive Testing 2012 (NDT 2012)*.
- Dziewierz, J., Lardner, T., & Gachagan, A. (2013). A design methodology for 2D sparse NDE arrays using an efficient implementation of refracted-ray TFM. In *2013 IEEE International Ultrasonics Symposium (IUS)* (pp. 136–138). Ieee. doi:10.1109/ULTSYM.2013.0035
- Elfring, R., Fuente, M. De, & Radermacher, K. (2009). Accuracy of Optical Localizers for Computer Aided Surgery, 328–330.
- El-Hawary, R., Roth, S. E., King, G. J. W., Chess, D. G., & Johnson, J. a. (2006). Visual Servoing in Medical Robotics: A Survey.Part 2 Tomographic Imaging Modalities - Techniques and Applications. *The International Journal of Medical Robotics + Computer Assisted Surgery: MRCAS*, 2(3), 251–255. doi:10.1002/rcs
- Ellams, D., Swanson, M., & Young, E. (2011). *8th Annual Report* (Vol. 1450).
- Erikson, K., Fry, F., & Jones, J. (1974). Ultrasound in Medicine - A Review. *IEEE TRANSACTIONS ON SONICS AND ULTRASONICS*, 21(3), 144–170.
- Fang, D. M., Ritter, M. a, & Davis, K. E. (2009). Coronal alignment in total knee arthroplasty: just how important is it? *The Journal of Arthroplasty*, 24(6 Suppl), 39–43. doi:10.1016/j.arth.2009.04.034
- Farley, J. M. (2004). BEST PRACTICE IN THE APPLICATION OF NDT – AN UPDATE. In *Proc. World Conf. NDT*.
- Firestone, F. A. (1946). The Supersonic Reflectoscope, an Instrument for Inspecting the Interior of Solid Parts by Means of Sound Waves. *The Journal of the Acoustical Society of America*, 17(3).
- Fishman, K., Ney, R., Chaney, L., Pizer, M., Vannier, W., Kuhlman, E., & Robertson, D. (1992). Three-dimensional Imaging. *Radiology*, 182(3), 899.

- Fledelius, H. C. (1997). Ultrasound in ophthalmology. *Ultrasound in Medicine & Biology*, 23(3), 365–375. doi:10.1016/S0301-5629(96)00213-X
- Fletcher, G., Gloth, T., Edelsbrumer, H., & Fu, P. (2006). Methods, Aparatus and Computer Program Products That Reconstruct Surfaces From Data Point Sets. USA. doi:10.1074/JBC.274.42.30033.(51)
- Fleute, M., & Lavall, S. (1998). Building a Complete Surface Model from Sparse Data Using Statistical Shape Models : Application to Computer Assisted Knee Surgery. *Lecture Notes in Computer Science*, 1496, 879–887.
- Foroughi, P., Bector, E., Swartz, M. J., Taylor, R. H., & Fichtinger, G. (2007). Ultrasound Bone Segmentation Using Dynamic Programming. *2007 IEEE Ultrasonics Symposium Proceedings*, 2523–2526. doi:10.1109/ULTSYM.2007.635
- Fradkin, L., Zernov, V., Elston, G., Taneja, R., Bell, I., Lines, D., ... Fitzgerald, P. (2013). Towards semi-automated non-destructive evaluation. *Journal of Physics: Conference Series*, 457, 012008. doi:10.1088/1742-6596/457/1/012008
- Fred, H. L. (2004). Drawbacks and Limitations of Computed Tomography. *Texas Heart Institute Journal*, 31(4), 345–348.
- Freemantle, R., Hnakinson, N., & Brotherhood, C. (2004). RAPID PHASED ARRAY ULTRASONIC IMAGING OF LARGE AREA COMPOSITE AEROSPACE STRUCTURES. In *WCNDT*.
- Friedman, L., Finlay, K., & Jurriaans, E. (2001). Ultrasound of the knee. *Skeletal Radiology*, 30(7), 361–77. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11499776>
- Frost, G. (2001). Inventing Schemes and Strategies: The Making and Selling of the Fessenden Oscillator. *Technology and Culture*, 43(2).
- Gachagan, A., McNab, A., & Reynolds, P. (2004). Analysis of ultrasonic wave propagation in metallic pipe structures using finite element modelling techniques. *IEEE Ultrasonics Symposium*. doi:10.1109/ULTSYM.2004.1417911
- Garcia, T., Hornof, W. J., & Insana, M. F. (2003). On the ultrasonic properties of tendon. *Ultrasound in Medicine and Biology*, 29(12), 1787–1797. doi:10.1016/S0301-5629(03)01069-X
- Gavrilov, L. R., Hand, J. W., Abel, P., & Cain, C. A. (1997). A Method of Reducing Grating Lobes Associated with an Ultrasound Linear Phased Array Intedned for Transrectal Thermotherapy. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Fr 1256 Equency Control*, 44, 1010–1017.
- Gee, A., Prager, R., Treece, G., & Berman, L. (2003). Engineering a freehand 3D ultrasound system. *Pattern Recognition Letters*, 24(4-5), 757–777. doi:10.1016/S0167-8655(02)00180-0

- Gemmeke, H., Berger, L., Birk, M., Gobel, G., Menshikov, A., Tcherniakhovski, D., ... Ruiter, N. V. (2010). Hardware setup for the next generation of 3D Ultrasound Computer Tomography. *IEEE Nuclear Science Symposium & Medical Imaging Conference*, 2449–2454. doi:10.1109/NSSMIC.2010.5874228
- Gemmeke, H., & Ruiter, N. V. (2007). 3D ultrasound computer tomography for medical imaging. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 580(2), 1057–1065. doi:10.1016/j.nima.2007.06.116
- Gill, S., Abolmaesumi, P., Fichtinger, G., Boisvert, J., Pichora, D., Borshneck, D., & Mousavi, P. (2012). Biomechanically constrained groupwise ultrasound to CT registration of the lumbar spine. *Medical Image Analysis*, 16(3), 662–674. doi:10.1016/j.media.2010.07.008
- Gilmour, T. P., & Subramanian, T. (2011). Three-dimensional reconstruction of transcranial ultrasound images obtained through the temporal bone window using a helmet-mounted mechanical beam-steering device. *Conference Proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference, 2011*, 413–6. doi:10.1109/IEMBS.2011.6090131
- Girardeau-Montaut, D. (2013). CloudCompare Manual.
- Gourdon, a., Poignet, P., Poisson, G., Vieyres, P., & Marche, P. (1999). A new robotic mechanism for medical application. *1999 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (Cat. No.99TH8399)*, 33–38. doi:10.1109/AIM.1999.803139
- Grandia, W. a., & Fortunko, C. M. (1995). NDE applications of air-coupled ultrasonic transducers. *1995 IEEE Ultrasonics Symposium. Proceedings. An International Symposium, 1*, 697–709. doi:10.1109/ULTSYM.1995.495667
- Gray, H. (1918). *Anatomy of the Human Body*. (W. H. Lewis, Ed.) (20th ed.). Philadelphia: Lea and Febiger.
- Greenway, B. (2012). Robot accuracy. *Industrial Robot: An International Journal*, 27(4).
- Gregg, E., Kriska, A., Salamone, L., Roberts, M., Anderson, S., Ferrel, R., ... Cauly, J. (1997). Relationships with Bone Mass , Osteoporosis and Fracture Risk. *Osteoporosis Int*, 7, 89–99.
- Grewe, M. G., & Gururaja, T. R. (1989). ACOUSTIC PROPERTIES OF PARTICLE/POLYMER COMPOSITES FOR TRANSDUCER BACKING APPLICATIONS. *1989 ULTRASONICS SYMPOSIUM*, 713–716.
- Griffin, D., & Galambos, R. (1941). The sensory Basis of Obstacle Avoidance by Flying Bats. *J Exp Zool*, 86.
- Grotenhuis, R., & Hong, A. (2012). Imaging the Weld Volume Via the Total Focus Method. In *4th International CANDU In-service Inspection Workshop and NDT in Canada 2012 Conference (Vol. Vi)*.

- Hacihaliloglu, I., Rasoulilian, A., & Rohling, R. N. (2014). Local Phase Tensor Features for 3-D Ultrasound to Statistical Shape + Pose Spine Model Registration. *IEEE Transactions on Medical Imaging*, 33(11), 2167–2179.
- Hagag, B., Abovitz, R., Kang, H., Schmitz, B., & Conditt, M. (2011). Surgical Robotics. In J. Rosen, B. Hannaford, & R. M. Satava (Eds.), *Surgical Robotics* (pp. 219–246). Boston, MA: Springer US. doi:10.1007/978-1-4419-1126-1
- Hailey, D. (2006). Open magnetic resonance imaging (MRI) scanners (abstract). *Issues Emerg Health Technol.*, 92.
- Halmshaw, R. (1991). *Non-Destructive Testing* (2nd ed.). London: Edward Arnold.
- Hanisch, C., Hsu, J., Noorman, E., & Radermacher, K. (2015a). LEVEL SET BASED SEGMENTATION OF THE DISTAL FEMUR FROM 3D ULTRASOUND VOLUME IMAGES. In *CAOS 2015 Proceedings*.
- Hanisch, C., Hsu, J., Noorman, E., & Radermacher, K. (2015b). MODEL BASED RECONSTRUCTION OF THE BONY KNEE ANATOMY FROM 3D ULTRASOUND IMAGES. In *CAOS 2015 Proceedings*.
- Harrysson, O. LA, Hosni, Y. A., & Nayfeh, J. F. (2007). Custom-designed orthopedic implants evaluated using finite element analysis of patient-specific computed tomography data: femoral-component case study. *BMC Musculoskeletal Disorders*, 8(1), 91. doi:10.1186/1471-2474-8-91
- Harvie, P., Sloan, K., & Beaver, R. J. (2012). Computer navigation vs conventional total knee arthroplasty: five-year functional results of a prospective randomized trial. *The Journal of Arthroplasty*, 27(5), 667–72.e1. doi:10.1016/j.arth.2011.08.009
- Hasenstab, A., & Osterloh, K. (2009). Defects in wood non destructive locating with low frequency Ultrasonic Echo Technique. *NDTCE*, 1–8.
- Hashemi, J., Chandrashekar, N., & Slauterbeck, J. (2005). The mechanical properties of the human patellar tendon are correlated to its mass density and are independent of sex. *Clinical Biomechanics*, 20(6), 645–652. doi:10.1016/j.clinbiomech.2005.02.008
- Hasse, W., & Maurer, A. (2004). Latest Developments on Industrial Ultrasonic Testing of Aircraft Components. In *WCNDT*.
- Heger, S., Portheine, F., Ohnsorge, J. A. K., Schkommodau, E., & Radermacher, K. (2005). User-interactive registration of bone with a-mode ultrasound. *IEEE Engineering in Medicine and Biology*, (April), 85–95.
- Heimann, T., & Meinzer, H.-P. (2009). Statistical shape models for 3D medical image segmentation: a review. *Medical Image Analysis*, 13(4), 543–63. doi:10.1016/j.media.2009.05.004
- Henckel, J., Richards, R., Lozhkin, K., Harris, S., Rodriguez y Baena, F. M., Barrett, a R. W., & Cobb, J. P. (2006). Very low-dose computed tomography for planning and outcome measurement in knee replacement. The imperial knee protocol. *The*

- Journal of Bone and Joint Surgery. British Volume*, 88(11), 1513– 8. doi:10.1302/0301-620X.88B11.17986
- Hennes, M., Bollue, K., Arenbeck, H., & Abel, D. (2013). Detection of Rehabilitation-Relevant Events during Endeffector Based Robot Assisted Rehabilitation of Upper Extremities. In *Converging Clinical and Engineering Research on Neurorehabilitation* (Vol. 1, pp. 765–768). doi:10.1007/978-3-64234546-3
- Herring, J. A., Laskin, S., & D, M. (1978). Unicompartamental tibiofemoral resurfacing arthroplasty. *The Journal of Bone and Joint Surgery*, 60-A(2), 182– 185.
- Hill, C. (1976). Ultrasonic Imaging. *J. Phys. E: Sci. Instrum*, 9(3).
- Hillger, W., Bühling, L., & Ilse, D. (2014). Scanners for Ultrasonic Imaging Systems. In *11th European Conference on Non-Destructive Testing* (pp. 2–8).
- Hindi, A., Peterson, C., & Barr, R. G. (2013). Artifacts in diagnostic ultrasound. *Reports in Medical Imaging*, 6, 29–48.
- Hoffmeister, B. K., Jones, C. I., Caldwell, G. J., & Kaste, S. C. (2006). Ultrasonic characterization of cancellous bone using apparent integrated backscatter. *Physics in Medicine and Biology*, 51(11), 2715–27. doi:10.1088/00319155/51/11/002
- Holler, P., & Schmitz, V. (1983). Products of Research and Development in NDT. *Nuclear Engineering and Design*, 76, 233–249.
- Holmes, C., Drinkwater, B. W., & Wilcox, P. D. (2005). Post-processing of the full matrix of ultrasonic transmit–receive array data for non-destructive evaluation. *NDT & E International*, 38(8), 701–711. doi:10.1016/j.ndteint.2005.04.002
- Holmes, C., Drinkwater, B. W., & Wilcox, P. D. (2008). Advanced post-processing for scanned ultrasonic arrays: application to defect detection and classification in non-destructive evaluation. *Ultrasonics*, 48(6-7), 636–42. doi:10.1016/j.ultras.2008.07.019
- Hoppe, H., DeRose, T., Dechamp, T., McDonald, J., & Stuetzle, W. (1992). Surface Reconstruction from Unorganized Points. *Computer Graphics*, 26(2).
- Hu, S., Hoffman, E. A., & Reinhardt, J. M. (2001). Automatic Lung Segmentation for Accurate Quantitation of Volumetric X-Ray CT Images. *IEEE Transactions on Medical Imaging*, 20(6), 490–498.
- Huang, N. F. R., Dowsey, M. M., Ee, E., Stoney, J. D., Babazadeh, S., & Choong, P. F. (2012). Coronal alignment correlates with outcome after total knee arthroplasty: five-year follow-up of a randomized controlled trial. *The Journal of Arthroplasty*, 27(9), 1737–41. doi:10.1016/j.arth.2012.03.058
- Hughes, S. (2001). Medical ultrasound imaging. *Physics Education*, 36(6), 468–475. doi:10.1088/0031-9120/36/6/304

- Humeida, Y., Wilcox, P. D., Todd, M. D., & Drinkwater, B. W. (2014). A probabilistic approach for the optimisation of ultrasonic array inspection techniques. *NDT & E International*, 68, 43–52. doi:10.1016/j.ndteint.2014.07.007
- Hunter, A. J., Drinkwater, B. W., & Wilcox, P. D. (2008). The wavenumber algorithm for full-matrix imaging using an ultrasonic array. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 55(11), 2450–62. doi:10.1109/TUFFC.952
- Hunter, A. J., Drinkwater, B. W., & Wilcox, P. D. (2010). Autofocusing ultrasonic imagery for non-destructive testing and evaluation of specimens with complicated geometries. *NDT & E International*, 43(2), 78–85. doi:10.1016/j.ndteint.2009.09.001
- Ill, T. C. K., Stulberg, D., Peshkin, M., Quaid, A., Lea, J., & Goswami, A. (1995). Total Knee Replacement. *IEEE Engineering in Medicine and Biology*, (May/June), 301–306.
- Insall, J., & Aglietti, P. (1980). A five to seven-year follow-up of unicondylar arthroplasty. *The Journal of Bone and Joint Surgery. American Volume*, 62(8), 1329–37. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/7440612>
- Jain, A. K., & Taylor, R. H. (2004). Understanding bone responses in b-mode ultrasound images and automatic bone surface extraction using a bayesian probabilistic framework. *Medical Imaging*, 5373, 131–142. doi:10.1117/12.535984
- Jaramaz, B. B., Hafez, M. A., & Digioia, A. M. (2006). Computer-Assisted Orthopaedic Surgery. *Proceedings of the IEEE*, 94(9), 1689–1695.
- Jensen, J. A., & Nikolov, S. I. (2002). Transverse flow imaging using synthetic aperture directional beamforming. In *2002 IEEE ULTRASONICS SYMPOSIUM* (Vol. 00, pp. 1523–1527).
- Jensen, J. A., & Nikolov, S. I. (2004). Directional Synthetic Aperture Flow Imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 51(9).
- Jensen, J. A., Nikolov, S. I., Gammelmark, K. L., & Pedersen, M. H. (2006). Synthetic aperture ultrasound imaging. *Ultrasonics*, 44 Suppl 1, e5–15. doi:10.1016/j.ultras.2006.07.017
- Jie Zhang, Drinkwater, B. W., & Wilcox, P. D. (2013). Comparison of ultrasonic array imaging algorithms for nondestructive evaluation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 60(8), 1732–45. doi:10.1109/TUFFC.2013.2754
- Jobst, M., & Connolly, G. D. (2010). Demonstration of the Application of the Total Focusing Method to the Inspection of Steel Welds. *10th European Conference on Non-Destructive Testing*, 1–11.
- Jordan, K. M. (2003). EULAR Recommendations 2003: an evidence based approach to the management of knee osteoarthritis: Report of a Task Force of the Standing

- Committee for International Clinical Studies Including Therapeutic Trials (ESCISIT). *Annals of the Rheumatic Diseases*, 62(12), 1145–1155. doi:10.1136/ard.2003.011742
- Juni, P. (2003). Population requirement for primary knee replacement surgery: a cross-sectional study. *Rheumatology*, 42(4), 516–521. doi:10.1093/rheumatology/keg196
- Jurvelin, J. S., Lyyra, T., Rasanen, T., & Kolmonen, P. (1995). Comparison of optical, needle probe and ultrasonic techniques for the measurement of articular cartilage thickness. *J. Biomechanics*, 28(2), 231–235.
- Kaiys, R., & Svilainis, L. (1995). Analysis of adaptive imaging algorithms for ultrasonic non-destructive testing, 33(1), 19–30.
- Kamat, Y. D., Aurakzai, K. M., Adhikari, A. R., Matthews, D., Kalairajah, Y., & Field, R. E. (2009). Does computer navigation in total knee arthroplasty improve patient outcome at midterm follow-up? *International Orthopaedics*, 33(6), 1567–1570. doi:10.1007/s00264-008-0690-0
- Kane, D., Balint, P. V., Sturrock, R., & Grassi, W. (2004). Musculoskeletal ultrasound - A state of the art review in rheumatology. Part 1: Current controversies and issues in the development of musculoskeletal ultrasound in rheumatology. *Rheumatology*, 43(7), 823–828. doi:10.1093/rheumatology/keh214
- Kane, D., Grassi, W., Sturrock, R., & Balint, P. V. (2004). A brief history of musculoskeletal ultrasound: “From bats and ships to babies and hips.” *Rheumatology*, 43(7), 931–933. doi:10.1093/rheumatology/keh004
- Kang, T., Kim, H., Son, T., & Choi, H. (2003). Design of Quadruped Walking and Climbing Robot. In *IEEE/RSJ Intl. Conference on Intelligent Robots and Systems*.
- Karaman, M., & O’Donnell, M. (1995). Synthetic aperture imaging for small scale systems. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 42(3), 429–442. doi:10.1109/58.384453
- Kendall, J. L., & Faragher, J. P. (2007). Ultrasound-guided central venous access: A homemade phantom for simulation. *Canadian Journal of Emergency Medicine*, 9(5), 371–373.
- Khaled, N. (2007). plane_line_intersect.
- Khallaghi, S. (2010). Biomechanically Constrained Groupwise Statistical Shape Model to Ultrasound Registration of the Lumbar Spine. In *Information Processing in Computer-Assisted Interventions* (pp. 47–54).
- Kholkin, a. L., Pertsev, N. a., & Goltsev, a. V. (2008). Piezoelectricity and crystal symmetry. In *Piezoelectric and Acoustic Materials for Transducer Applications* (pp. 17–38). doi:10.1007/978-0-387-76540-2_2
- Kjer, H. M., & Wilm, J. (2010). *Evaluation of surface registration algorithms for PET motion correction Bachelor thesis*. Retrieved from

<http://orbit.dtu.dk/getResource?recordId=262855&objectId=1&versionId=1>

- Kjer, M., & Wilm, J. (2012). ICP.
- Klauser, A. S., & Peetrons, P. (2009). Developments in musculoskeletal ultrasound and clinical applications. *Skeletal Radiology*, 1061–1071. doi:10.1007/s00256009-0782-y
- Klein, H.-M., Günther, R. W., Verlande, M., Schneider, W., Vorwerk, D., Kelch, J., & Hamm, M. (1992). 3D-surface reconstruction of intravascular ultrasound images using personal computer hardware and a motorized catheter control. *Cardiovascular and Interventional Radiology*, 15(2), 97–101. doi:10.1007/BF02734099
- Kohl, C., Krause, M., Maierhofer, C., & Wostmann, J. (2005). 2D- and 3Dvisualisation of NDT-data using data fusion technique. *Materials and Structures*, 38(283), 817–826. doi:10.1617/14293
- Kohl, C., & Streicher, D. (2006). Results of reconstructed and fused NDT-data measured in the laboratory and on-site at bridges. *Cement and Concrete Composites*, 28(4), 402–413. doi:10.1016/j.cemconcomp.2006.02.005
- Kohles, S. S., Mason, S. S., Adams, A. P., Berg, R. J., Blank, J., Gibson, F., ... Saha, A. K. (2012). Ultrasonic wave propagation assessment of native cartilage explants and hydrogel scaffolds for tissue engineering. *Int J Biomed Eng Technol.*, 10(3), 296–307. doi:10.1038/nature11130.Reduced
- Koizumi, N., Lee, D., Seo, J., Funamoto, T., Sugita, N., Yoshinaka, K., ... Mitsuishi, M. (2012). TECHNOLOGIZING AND DIGITALIZING OF MEDICAL PROFESSIONAL SKILLS FOR A NON-INVASIVE ULTRASOUND THERAGNOSTIC SYSTEM. In *Proceedings of the ASME/ISCIE 2012 International Symposium on Flexible Automation* (pp. 1–6).
- Koizumi, N., Nagoshi, M., Hashizume, H., & Mitsuishi, M. (2009). Construction Methodology for a Remote Ultrasound Diagnostic System. *IEEE TRANSACTIONS ON ROBOTICS*, 25(3), 522–538.
- Kosuge, K., & Hirata, Y. (2004). Human-robot interaction. In *IEEE International Conference on Robotic and Bioinformatics*. Retrieved from <http://robocare.istc.cnr.it/RC-Ws-1-proceedings.pdf#page=23>
- Kotha, S. P., DePaula, C. A., Mann, A. B., & Guzelsu, N. (2008). High frequency ultrasound prediction of mechanical properties of cortical bone with varying amount of mineral content. *Ultrasound in Medicine & Biology*, 34(4), 630–7. doi:10.1016/j.ultrasmedbio.2007.09.020
- Kowal, J., Amstutz, C., Langlotz, F., Talib, H., & Ballester, M. (2006). Automated Bone Contour Detection in Ultrasound B-mode Images for Minimally Invasive Registration on Computer-Assisted Surgery - an In Vitro Evaluation. *The International Journal of Medical Robotics + Computer Assisted Surgery: MRCAS*, 3, 314–348.

- Krause, M., B, M., Frielinghaus, R., Kretzschmar, F., Kroggel, O., Langenberg, K. J., ... Wollbold, F. (1997). Comparison of pulse-echo methods for testing concrete. *NDT & E International*, 30(4), 195–204.
- Krause, M., Mielentz, F., Milman, B., Müller, W., Schmitz, V., & Wiggenger, H. (2001). Ultrasonic imaging of concrete members using an array system. *NDT & E International*, 34(6), 403–408. doi:10.1016/S0963-8695(01)00007-X
- Krupa, A., & Chaumette, F. (2005). Control of an ultrasound probe by adaptive visual servoing. *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, 2007–2012*. doi:10.1109/IROS.2005.1545272
- Kryzstoforski, K., Krowicki, P., Swiatek-Najwer, E., Bedzinski, R., & Keppler, P. (2011). Noninvasive ultrasonic measuring system for bone geometry examination. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 7(January), 85–95. doi:10.1002/rcs
- Kujala, U. M., Kettunen, J., Paananen, H., Aalto, T., Battié, M. C., Impivaara, O., ... Sarna, S. (1995). Knee osteoarthritis in former runners, soccer players, weight lifters, and shooters. *Arthritis and Rheumatism*, 38(4), 539–46. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/7718008>
- KUKA. (2009). *KUKA Robot Sensor Interface 2.3*.
- KUKA. (2010). *KUKA System Software 5*.
- Kulić, D., & Croft, E. A. (2005). Safe planning for human-robot interaction. *Journal of Robotic Systems*, 22(7), 383–396. doi:10.1002/rob.20073
- Kurjak, M. D. (2000). Ultrasound scanning - Prof. Ian Donald (1910-1987). *European Journal of Obstetrics Gynecology and Reproductive Biology*, 90, 187–189. doi:10.1016/S0301-2115(00)00270-0
- Lambert, J., Antoine, P., Gens, G., Bimbard, F., Lacassagne, L., & Iakovleva, E. (n.d.). ANALYSIS OF MULTICORE CPU AND GPU TOWARD PARALLELIZATION OF TOTAL FOCUSING METHOD ULTRASOUND RECONSTRUCTION.
- Lambert, J., Antoine, P., Gens, G., Bimbard, F., Lacassagne, L., & Iakovleva, E. (2012). Performance Evaluation of Total Focusing Method on GPP and GPU. In *2012 Conference on Design and Architectures for Signal and Image Processing (DASIP)* (pp. 1–8).
- Leary, O., Tweedie, A., Leary, R. L. O., Gachagan, A., Holmes, C., Wilcox, P. D., & Drinkwater, B. W. (2007). Total Focussing Method for Volumetric Imaging in Immersion Non Destructive Evaluation. *2007 IEEE Ultrasonics Symposium*.
- Lee, M. A. (2007). Cross-National Research on Aging. *Research Highlights in the Demography and Economics of Aging*. doi:10.1038/nbt1352
- Lemons, R. a. (1974). Acoustic microscope—scanning version. *Applied Physics Letters*, 24(4), 163. doi:10.1063/1.1655136

- Li, C., Pain, D., Wilcox, P. D., & Drinkwater, B. W. (2013). Imaging composite material using ultrasonic arrays. *NDT & E International*, 53, 8–17. doi:10.1016/j.ndteint.2012.07.006
- Li, C., Xu, C., Gui, C., & Fox, M. D. (2005). Level Set Evolution without ReInitialization: A New Variational Formulation. *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, 1, 430–436. doi:10.1109/CVPR.2005.213
- Li, M., & Hayward, G. (2012). Ultrasound nondestructive evaluation (NDE) imaging with transducer arrays and adaptive processing. *Sensors (Basel, Switzerland)*, 12(1), 42–54. doi:10.3390/s120100042
- Liang, X. L. X., & Boppart, S. a. (2010). Biomechanical Properties of In Vivo Human Skin from Dynamic Optical Coherence Elastography. *IEEE Transactions on Biomedical Engineering*, 57(4), 953–959. doi:10.1109/TBME.2009.2033464
- Lines, D. I. A. (2003). The application of real-time medical ultrasonic techniques to NDT. In *Proceedings of NDT 2003* (pp. 189–192).
- Lines, D. I. A., Pettigrew, I. G., Kirk, K. J., Cochran, S., & Skamstad, J. A. (2006). Rapid Distributed data Collection and Processing With Arrays - the Next Step Beyond Full Waveform Capture. *9th Joint FAA/DoD/NASA Aging Aircraft Conference*, 1–11.
- Liu, T. ., & Wang, J. . (1994). A reliability approach to evaluating robot accuracy performance. *Mechanism and Machine Theory*, 29(1), 83–94. doi:10.1016/0094-114X(94)90022-1
- Lockwood, G. R., Talman, J. R., & Brunke, S. S. (1998). Real-time 3-D ultrasound imaging using sparse synthetic aperture beamforming. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 45(4), 980–8. doi:10.1109/58.710573
- Lotke, P. A., & L, E. M. (1977). Influence of Positioning Knee of Prosthesis in Total Knee Replacement. *Journal of Bone and Joint Surgery*, 59(1), 77–79.
- Lu, C., Deng, D., & Li, L. (2012). Ultrasonic Phased Array Inspection for Gas Pressure Welds Joint of High Speed Railway. *Przełąd Elektrotechniczny*, 88(1b), 173–176.
- Lu, Z., Xu, C., Zhao, X., Zhang, L., Wang, H., & Pan, Q. (2013). Ultrasonic Transmission Testing of Twin-Robots Coordinated Control. In *Proceedings of the 2013 IEEE International Conference on Mechatronics and Automation* (pp. 1256–1260).
- Magenat-thalman, N., & Schmid, J. (2008). MRI Bone Segmentation Using Deformable Models and Shape Priors. *Med Image Comput Comput Assist Interv.*, 11(1), 119–126.
- Malkani, A., Rand, J., Bryan, A., & Wallriches, S. (1995). Total Knee Arthroplasty with the Kinematic Condylar Prosthesis. A Ten Year Follow Up Study. *The Journal of Bone and Joint Surgery*, 77-A, 423–431.

- Manek, N. J., Hart, D., Spector, T. D., & MacGregor, A. J. (2003). The association of body mass index and osteoarthritis of the knee joint: an examination of genetic and environmental influences. *Arthritis and Rheumatism*, 48(4), 1024–9. doi:10.1002/art.10884
- Martín, C. J., Martínez, O., Octavio, A., & Godoy, G. (2008). Reduction of Grating Lobes in SAFT Images. In *2008 IEEE International Ultrasonics Symposium Proceedings* (pp. 721–724).
- Martinez, O., Akhnaq, M., Ullate, L. G., & Espinosa, F. M. De. (2003). A small 2D ultrasonic array for NDT applications. *NDT&E International*, 36, 57–63.
- Masuda, K., Kimura, E., Tateishi, N., & Ishihara, K. (2001). Three dimensional motion mechanism of ultrasound probe and its application for tele-echography system. *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No.01CH37180)*, 2, 1112–1116. doi:10.1109/IROS.2001.976317
- Masuda, K., Vilchis, a., Troccaz, J., & Cinquin, P. (2003). Safe control of ultrasound probe on abdomen by pulling cables for robotic tele-echography system. *Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE Cat. No.03CH37439)*, 4, 3676–3679. doi:10.1109/IEMBS.2003.1280953
- Mautz, R., & Tilch, S. (2011). Survey of optical indoor positioning systems. *2011 International Conference on Indoor Positioning and Indoor Navigation*, 1–7. doi:10.1109/IPIN.2011.6071925
- McCann, D. ., & Forde, M. . (2001). Review of NDT methods in the assessment of concrete and masonry structures. *NDT & E International*, 34(2), 71–84. doi:10.1016/S0963-8695(00)00032-3
- McDicken, W. N. (1981). *Diagnostic Ultrasonics: principles and Use of Instruments* (2nd ed.).
- Mcgilp, A., Dziewierz, J., Lardner, T., Gachagan, A., & Bird, C. (2014). Inspection of Complex Components using 2D Arrays and TFM. In *BINDT*.
- McNab, a., & Campbell, M. J. (1987). Ultrasonic phased arrays for nondestructive testing. *NDT International*, 20(6), 333–337. doi:10.1016/0308-9126(87)90290-2
- Mechlin, M., Thickman, D., Kressel, H. Y., Gefter, W., & Joseph, P. (1984). Magnetic resonance imaging of postoperative patients with metallic implants. *AJR. American Journal of Roentgenology*, 143(6), 1281–4. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/6333797>
- Menegaldo, L. L., Santos, M., Ferreira, G. a N., Siqueira, R. G., & Moscato, L. (2008). SIRUS: A mobile robot for floating production storage and offloading (FPSO) ship hull inspection. *International Workshop on Advanced Motion Control, AMC, 1*, 27–32. doi:10.1109/AMC.2008.4516036
- Merola, M., Chappuis, P., Escourbiac, F., Grattarola, M., Jeskanen, H., Schedler, B.,

- ... Zabernig, A. (2002). Non-destructive Testing of Divertor Components. *Fusion Engineering and Design*, 62, 141–146.
- Mineo, C., Herbert, D., Morozov, M., & Pierce, S. G. (2012). Robotic NonDestructive Inspection. In *BINDT 2012*.
- Mineo, C., Morozov, M., Pierce, G., Nicholson, I., & Cooper, I. (2013). ComputerAided Tool Path Generation for Robotic Non-Destructive Inspection. In *BINDT*.
- Mohr, F., Haase, W., & Ungerer, D. (2014). Automated Non-Destructive Examination of Complex Shapes. In *11th European Conference on NonDestructive Testing*.
- Moir, H. M., Jackson, J. C., & Windmill, J. F. C. (2013). Extremely high frequency sensitivity in a ' simple ' ear. *Biology Letters*, 9, 20130241. doi:10.1098/rsbl.2013.0241
- Mönnich, H., Stein, D., Raczekowsky, J., & Wörn, H. (2010). An automatic and complete self-calibration method for robotic guided laser ablation. *Robotics*, 1086–1087.
- Moreau, L., Drinkwater, B. W., & Wilcox, P. D. (2009). Transmission Cycles for Rapid Nondestructive Evaluation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 56(9), 1932–1944.
- Motesharei, A. (2014). Investigating the Biomechanical Outcomes of a RoboticAssisted Versus Conventional Unicompartmental Knee Arthroplasty.
- Mozes, A., & Chang, T. (2010). Three-dimensional A-mode ultrasound calibration and registration for robotic orthopaedic knee surgery. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 6(June 2009), 91–101. doi:10.1002/rcs
- Muir, H. (1977). Proteoglycans of cartilage. *J. Clin. Path.*, 12, 67–81.
- Muller, W., Schmitz, V., & Schafer, G. (1986). Reconstruction by the Synthetic Aperture Focussing Technique. *Nuclear Engineering and Design*, 94, 393–404.
- Nakadate, R., Solis, J., Takanishi, A., Minagawa, E., Sugawara, M., & Niki, K. (2011). Out-of-plane visual servoing method for tracking the carotid artery with a robot-assisted ultrasound diagnostic system. *Proceedings - IEEE International Conference on Robotics and Automation*, 5267–5272. doi:10.1109/ICRA.2011.5979594
- Nakadate, R., Solis, J., Takanishi, A., Sugawara, M., Niki, K., & Minagawa, E. (2010). Development of the Ultrasound Probe Holding Robot WTA-IRII and an Automated Scanning Algorithm based on Ultrasound Image Feedback. *ROMANSY 18 Robot Design, Dynamics and Control SE - 42*, 524, 359–366. doi:10.1007/978-3-7091-0277-0_42
- Netravali, N. a, Shen, F., Park, Y., & Bargar, W. L. (2013). A perspective on robotic assistance for knee arthroplasty. *Advances in Orthopedics*, 2013, 970703.

doi:10.1155/2013/970703

- Nikolov, S., Gammelmark, K., & Jensen, J. a. (1999). Recursive ultrasound imaging. *1999 IEEE Ultrasonics Symposium. Proceedings. International Symposium (Cat. No.99CH37027)*, 2, 1621–1625. doi:10.1109/ULTSYM.1999.849306
- Nilsson, P., Appelgren, A., Henrikson, P., & Runnemalm, A. (2012). Automatic Ultrasonic testing for Metal Deposition. In *18th World Conference on Nondestructive Testing* (pp. 16–20).
- Njiki, M., Bouaziz, S., Elouardi, A., Casula, O., & Roy, O. (2013). A multi-FPGA implementation of real-time reconstruction using Total Focusing Method. *2013 IEEE International Conference on Cyber Technology in Automation, Control and Intelligent Systems*, 468–473. doi:10.1109/CYBER.2013.6705491
- Njiki, M., Elouardi, A., Bouaziz, S., Casula, O., & Roy, O. (2013). A real-time implementation of the Total Focusing Method for rapid and precise diagnostic in non destructive evaluation. In *2013 IEEE 24th International Conference on Application-Specific Systems, Architectures and Processors* (pp. 245–248). Ieee. doi:10.1109/ASAP.2013.6567582
- Nock, L. F., & Trahey, G. E. (1992). Synthetic Receive Aperture Imaging with Phase Correction for Motion and for Tissue Inhomogeneities-Part I: Basic Principles. *IEEE Trans Ultrason Ferroelectr Freq Control*, 39(4), 489–485.
- Novotny, P. M., Stoll, J. A., Dupont, P. E., & Howe, R. D. (2007). Real-Time Visual Servoing of a Robot Using Three-Dimensional Ultrasound. *2007 IEEE International Conference on Robotics and Automation*, (April), 10–14.
- O'Donnell, M., & Thomas, L. J. (1992). Efficient Synthetic Aperture Imaging from a Circular Aperture with Possible Application to Catheter Based imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 39(3), 366–380.
- O'Leary, R. L., Hayward, G., Smillie, G., & Parr, A. C. . (2005). CUE Materials Database.
- Oralkan, O., Ergun, S., Johnson, J. A., Karaman, M., Demirci, U., Kaviani, K., ... Khuri-Yakub, B. T. (2002). Capacitive Micromachined Transducers: Next-Generation Arrays for Acoustic Imaging? *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 49(11), 1596–1610.
- Osmer, C. (2008). A gelatine-based ultrasound phantom. *Anaesthesia*, 63, 96–107.
- Osteoarthritis: national clinical guideline for care and management in adults*. (2008). London.
- Palani, P., Joseph, A., Jayakumar, T., & Raj, B. (1995). Ultrasonic velocity measurements for estimation of grain size in austenitic stainless steel. *NDT & E International*, 28(3), 179–185.
- Pearle, A. D., Kendoff, D., Stueber, V., Musahl, V., & Repicci, J. a. (2009). Perioperative management of unicompartmental knee arthroplasty using the

- MAKO robotic arm system (MAKOplasty). *American Journal of Orthopedics (Belle Mead, N.J.)*, 38(2 Suppl), 16–9. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/19340378>
- Pearle, A. D., Warren, R. F., & Rodeo, S. a. (2005). Basic science of articular cartilage and osteoarthritis. *Clinics in Sports Medicine*, 24(1), 1–12. doi:10.1016/j.csm.2004.08.007
- Peat, G., McCarney, R., & Croft, P. (2001). Knee pain and osteoarthritis in older adults: a review of community burden and current use of primary health care. *Annals of the Rheumatic Diseases*, 60(2), 91–7. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1753462&tool=pmc-entrez&rendertype=abstract>
- Peng, J., Peng, H., Zhang, Y., Gao, X., Peng, C., & Wang, Z. (2013). Study on the railway wheel ultrasonic inspection method using the full matrix capture. *2013 Far East Forum on Nondestructive Evaluation/Testing: New Technology and Application*, 118–124. doi:10.1109/FENDT.2013.6635540
- Peng, T., Peng, J., Peng, H., Wang, L., Gao, X., Luo, L., & Li, J. (2014). Efficiency improvement of total focusing imaging based on sparse array method. In *FENDT*. doi:10.1109/FENDT.2014.6928233
- Penney, G. P., Barratt, D. C., Chan, C. S. K., Slomczykowski, M., Carter, T. J., Edwards, P. J., & Hawkes, D. J. (2006). Cadaver validation of intensity-based ultrasound to CT registration. *Medical Image Analysis*, 10(3), 385–395. doi:10.1016/j.media.2006.01.003
- Prado, V. T., Higuti, R. T., & Kitano, C. (2013). COMPARISON BETWEEN ULTRASONIC SYNTHETIC APERTURE PLATE-LIKE STRUCTURES. In *COBEM*.
- Priester, A. M., Natarajan, S., & Culjat, M. O. (2013). Robotic ultrasound systems in medicine. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 60(3), 507–23. doi:10.1109/TUFFC.2013.2593
- Pritzker, K. P. H., Gay, S., Jimenez, S. a, Ostergaard, K., Pelletier, J.-P., Revell, P. a, ... van den Berg, W. B. (2006). Osteoarthritis cartilage histopathology: grading and staging. *Osteoarthritis and Cartilage / OARS, Osteoarthritis Research Society*, 14(1), 13–29. doi:10.1016/j.joca.2005.07.014
- Pua, E. C., Fronheiser, M. P., Noble, J. R., Light, E. D., Wolf, P. D., von Allmen, D., & Smith, S. W. (2006). 3-D ultrasound guidance of surgical robotics: a feasibility study. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 53(11), 1999–2008. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/18519222>
- Rajna, T., Herold, F., & Baylard, C. (2009). Reverse engineering process of cryostat components of Wendelstein 7-X. *Fusion Engineering and Design*, 84(7-11), 1558–1561. doi:10.1016/j.fusengdes.2009.03.008

- Ramadas, S. N., Leary, R. L. O., & Gachagan, A. (2009). Ultrasonic Sensor Design for NDE Application : Design Challenges & Considerations. In *Proceedings of the National Seminar and Exhibition on Non-Destructive Evaluation* (pp. 88– 91).
- Ramm, O. T. Von, Smith, S. W., & Pavy, H. G. (1991). High-speed Ultrasound Volumetric Imaging System- Part 11 : Parallel Processing and Image Display. *IEEE TRANSACTIONS ON ULTRASONICS*, 38(2), 109–115.
- Rand, J. A., & Ilstrup, D. (1991). Survivorship analysis of total knee arthroplasty. Cumulative rates of survival of 9200 total knee arthroplasties. *Am, J Bone Joint Surg*, 73(3), 397–409.
- Reed, S. C., & Gollish, J. (1997). The accuracy of femoral intramedullary guides in total knee arthroplasty. *The Journal of Arthroplasty*, 12(6), 677–82. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/9306219>
- Rennick, S. L., Fenton, T. W., & Foran, D. R. (2005). The effects of skeletal preparation techniques on DNA from human and non-human bone. *Journal of Forensic Sciences*, 50(5), 1016–9. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/16225205>
- Reynolds, P. (2005). Short Course: Finite Element Modelling for Ultrasonic Applications: Part 2: Wave Propagation. In *IEEE Ultrasonics Symposium*.
- Rigal, J., Escoffer, C., Querleux, B., Faivre, B., Agache, P., & Leveque, J. (1989). Assesment of the Aging of Human Sking by In Vivo Ultrasonic Imaging. *The Journal of Investigative Dermatology*, 93(5).
- Roche, M., O'Loughlin, P. F., Kendoff, D., Musahl, V., & Pearle, A. D. (2009). Robotic arm-assisted unicompartmental knee arthroplasty: preoperative planning and surgical technique. *American Journal of Orthopedics (Belle Mead, N.J.)*, 38(2 Suppl), 10–5. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/19340377>
- Romero, D., Mart, O., & Mart, C. J. (2009). Using GPUs for beamforming acceleration on SAFT imaging. In *2009 IEEE International Ultrasonics Symposium Proceedings* (pp. 1334–1337).
- Saarakkala, S., Laasanen, M. S., Jurvelin, J. S., & Töyräs, J. (2006). Quantitative ultrasound imaging detects degenerative changes in articular cartilage surface and subchondral bone. *Physics in Medicine and Biology*, 51(20), 5333–46. doi:10.1088/0031-9155/51/20/017
- Saijo, Y., Tanaka, M., Okawai, H., Sasaki, H., Nitta, S., & Dunn, F. (1997). Ultrasonic Tissue Characterisation of Infarcted Myocardium by Scanning Acoustic Microscopy. *Ultrasound in Med. and Biol.*, 23(1), 77–85.
- Salcudean, S., Bell, G., & Bachmann, S. (1999). Robot-assisted diagnostic ultrasound—design and feasibility experiments. *Proc. Int. Conf. in Medical Image Computing and Computer-Assisted Intervention* (), 1062–1072. Retrieved from <http://www.springerlink.com/index/x07201782px81x22.pdf>

- Sattar, T. P., & Brenner, A. A. (2006). 7-axis Arm for NDT of Surfaces with Complex & Unknown Geometry. In *Climbing and Walking Robots*.
- Sattar, T. P., Rodriguez, H. E. L., Shang, J., & Bridge, B. (2006). Automated NDT of Floating Production Storage Oil Tanks with a Swimming and Climbing Robot. In *Climbing and Walking Robots*.
- Sawyer-Glover, A., & Shellock, F. (2000). Pre-MRI procedure screening: recommendations and safety considerations for biomedical implants and devices. *Journal of Magnetic Resonance Imaging: JMRI*, 12(3), 510. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10992321>
- Schickert, M. (2002). Ultrasonic NDE of Concrete. In *2002 IEEE ULtr* (Vol. 00, pp. 739–748).
- Schickert, M. (2005). Progress in ultrasonic imaging of concrete. *Materials and Structures*, 38(283), 807–815. doi:10.1617/14298
- Schmidt, T., & Dutta, S. (2012). Automation in Production Integrated NDT Using Thermography. In *NDT in Aerospace*.
- Schumann, S., Puls, M., Ecker, T., Schwaegli, T., Stifter, J., Siebenrock, K., & Zheng, G. (2010). Determination of Pelvic Orientation from Ultrasound Images Using Patch-SSMs and a Hierarchical Speed of Sound Compensation Strategy. In *Information Processing in Computer-Assisted Interventions* (pp. 157–167).
- Schwabe, D. M., Maurer, A., & Koch, R. (2010). Ultrasonic Testing Machines with Robot Mechanics – A New Approach to CFRP Component Testing. In *Aerospace* (pp. 1–5).
- Schwarzschild, B. M. (1979). Scanning acoustic microscopy. *Physics Today*, 32(5), 20. doi:10.1063/1.2995549
- Scuderi, G. R., & Tria, A. J. (2010). *The Knee: a Comprehensive Review*. World Scientific Publishing Inc.
- Seedhom, B. B., Longton, E. B., Wright, V., & Dowson, D. (1972). Dimensions of the knee. Radiographic and autopsy study of sizes required by a knee prosthesis. *Annals of the Rheumatic Diseases*, 31(1), 54–58. doi:10.1136/ard.31.1.54
- Seo, C. H., & Yen, J. T. (2007). 256 x 256 array transducer with row-column addressing for 3-D imaging. *2007 IEEE ULTRASONICS SYMPOSIUM*, 2381–2384.
- Seydel, J. A. (1975). Methods Development for Non-Destructive Measurement of Bond Strength in Adhesively Bonded Structures. *Proceedings of the ARPA/AFML Review of Quantitative NDE*, (June 1974).
- Shellock, F. G. (2002). Magnetic resonance safety update 2002: implants and devices. *Journal of Magnetic Resonance Imaging: JMRI*, 16(5), 485–96. doi:10.1002/jmri.10196
- Shellock, F. G., & Crues, J. V. (1988). High-field-strength MR imaging and metallic biomedical implants: an ex vivo evaluation of deflection forces. *AJR. American*

- Journal of Roentgenology*, 151(2), 389–92. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/3260731>
- Shen, W. S. W., Gu, J., & Shen, Y. S. Y. (2005). Proposed wall climbing robot with permanent magnetic tracks for inspecting oil tanks. *IEEE International Conference Mechatronics and Automation, 2005*, 4(July), 2072–2077. doi:10.1109/ICMA.2005.1626882
- Shung, K. K., & Zipparo, M. (1996). Ultrasonic Arrays and Transducers. *IEEE ENGINEERING IN MEDICINE AND BIOLOGY*, (December).
- Silva, A. K. A., Silva, E. L., Egito, E. S. T., & Carriço, A. S. (2006). Safety concerns related to magnetic field exposure. *Radiation and Environmental Biophysics*, 45(4), 245–52. doi:10.1007/s00411-006-0065-0
- Siu, D., Rudan, J., Wevers, H. W., & Griffiths, P. (1996). Femoral articular shape and geometry. A three-dimensional computerized analysis of the knee. *The Journal of Arthroplasty*, 11(2), 166–73. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/8648311>
- Smith, J. (2011). *Functional Assessment to Compare Electromagnetic Navigated and Conventional Total Knee Arthroplasty*.
- Smith, S. W., Lee, W., Light, E. D., Yen, J. T., Wolf, P., & Idriss, S. (2002). Two Dimensional Arrays for 3-D Ultrasound Imaging. *2002 IEEE ULTRASONICS SYMPOSIUM*, 1545–1553.
- Sodickson, A., Andriole, K. P., Prevedello, L. M., & Nawfel, R. D. (2009). Recurrent CT, Cumulative Radiation Exposure, and Associated Radiation-induced Cancer Risk from CT of Adults. *Radiology*, 251(1), 175–184.
- Spies, M., & Rieder, H. (2010). Synthetic aperture focusing of ultrasonic inspection data to enhance the probability of detection of defects in strongly attenuating materials. *NDT & E International*, 43(5), 425–431. doi:10.1016/j.ndteint.2010.04.002
- Standring, S. (2008). *Gray's Anatomy* (14th ed.). Churchill Livingstone.
- Stepinski, T., & Lingvall, F. (2010). Synthetic aperture focusing techniques for ultrasonic imaging of solid objects., 438–441.
- Suh, J. K., Youn, I., & Fu, F. H. (2001). An in situ calibration of an ultrasound transducer: a potential application for an ultrasonic indentation test of articular cartilage. *Journal of Biomechanics*, 34(10), 1347–53. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11522315>
- Summan, R. (2012). *Positioning for mobile NDE inspection robots*. Retrieved from <http://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.581903>
- Sutcliffe, M., Weston, M., Dutton, B., Charlton, P., & Donne, K. (2012). Real-time full matrix capture for ultrasonic non-destructive testing with acceleration of post-processing through graphic hardware. *NDT & E International*, 51, 16–23. doi:10.1016/j.ndteint.2012.06.005

- Sutcliffe, M., Weston, M., Dutton, B., & Cooper, I. (2012). Real-time full matrix capture with auto-focussing of known geometry through dual layered media. In *BINDT* (pp. 1–8).
- Synnevag, J., Austeng, A., & Holm, S. (2007). Adaptive Beamforming Applied to Medical Ultrasound Imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, *54*(8), 1606–1613.
- Szabo, T. L. (2014). *Diagnostic Ultrasound Imaging: Inside Out* (2nd ed.). Elsevier.
- Talib, H., Rajamani, K., Kowal, J., Nolte, L.-P., Styner, M., & Ballester, M. a G. (2005). A comparison study assessing the feasibility of ultrasound-initialized deformable bone models. *Computer Aided Surgery: Official Journal of the International Society for Computer Aided Surgery*, *10*(5-6), 293–9.
doi:10.3109/10929080500379390
- Talmant, M., Foiret, J., Minonzio, J., Bossy, E., Grimal, Q., Laugier, P., ... Gl, C. (2011). *Bone Quantitative Ultrasound*. *Bone Quantitative Ultrasound*. doi:10.1007/978-94-007-0017-8
- Tapani, J., & Helmut, Y. (1994). Ultrasound Synthetic Aperture Imaging: Monostatic Approach. *IEEE TRANSACTIONS ON ULTRASONICS. FERROELEaRICS, AND FREQUENCY CONTROL*, *41*(3), 333–339.
- Tatarinov, A., Sarvazyan, N., & Sarvazyan, A. (2005). Use of multiple acoustic wave modes for assessment of long bones: model study. *Ultrasonics*, *43*(8), 672–80. doi:10.1016/j.ultras.2005.03.004
- Torres, P. M. B., & Goncalves, P. J. S. (2012). 3D FEMUR RECONSTRUCTION USING A ROBOTIZED ULTRASOUND PROBE. *The Fourth IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechatronics*, 884–888.
- Torres, P. M. B., Sanches, M. J., Goncalves, P. J. S., & Martins, J. M. M. (2012). Robotic 3D Ultrasound. In *Proceedings of RECPAD 2012, 18th Portuguese Conference on Pattern Recognition* (pp. 2–3).
- Töyräs, J., Lyyra-Laitinen, T., Niinimäki, M., Lindgren, R., Nieminen, M. ., Kiviranta, I., & Jurvelin, J. . (2001). Estimation of the Young's modulus of articular cartilage using an arthroscopic indentation instrument and ultrasonic measurement of tissue thickness. *Journal of Biomechanics*, *34*(2), 251–256. doi:10.1016/S0021-9290(00)00189-5
- Toyraas, J., Rieppo, J., Nieminen, M. T., Helminen, H. J., & Jurvelin, J. (1999). Characterization of enzymatically induced degradation of articular cartilage using high frequency ultrasound Characterization of enzymatically induced degradation of articular cartilage using high frequency ultrasound. *Phys Med Biol.*, *44*(11), 2733–2733.
- Trebackz, H., & Natali, a. (1999). Ultrasound velocity and attenuation in cancellous bone samples from lumbar vertebra and calcaneus. *Osteoporosis International*, *9*(2), 99–105. doi:10.1007/s001980050121

- Trots, I., Nowicki, A., Lewandowski, M., & Tasinkevych, Y. (2010). Synthetic Aperture Method in Ultrasound Imaging.
- Tweedie, A., O'Leary, R., Harvey, G., Gachagan, A., Holmes, C., Wilcox, P. D., & Drinkwater, B. W. (2007). Total Focussing Method for Volumetric Imaging in Immersion Non Destructive Evaluation. In *IEEE Ultrasonics Symposium* (pp. 1017–1020).
- Viceconti, M., Zannoni, C., Testi, D., & Cappello, A. (1999). CT data sets surface extraction for biomechanical modeling of long bones. *Computer Methods and Programs in Biomedicine*, *59*, 159–166. doi:10.1016/S0169-2607(98)00107-2
- Vieyres, P., Poisson, G., Courrèges, F., Smith-guerin, N., Novales, C., & Arbeille, P. (2006). A Tele-Operated Robotic System for Mobile Tele-Echography : the Otelo Project. In *Topics in Biomedical Engineering* (pp. 461–473).
- Vincent, D., Jocelyne, T., & Tonetti, J. (2004). Automatic delineation of the osseous interface in ultrasound images by information fusion. *Proceedings of the 7th International Conference on Image Fusion*, 862–867.
- Wagle, S., & Kato, H. (2009). Ultrasonic detection of fretting fatigue damage at bolt joints of aluminum alloy plates. *International Journal of Fatigue*, *31*(8-9), 1378–1385. doi:10.1016/j.ijfatigue.2009.03.017
- Wang, C., Gachagan, A., & Leary, R. L. O. (2011). High Intensity Focused Ultrasound Array Transducers for Sonochemistry Applications. In *IEEE International Ultrasonics Symposium*. doi:10.1109/ULTSYM.2011.0578
- Wang, H., Xu, C., & Xu, Y. (2013). Ultrasonic inspection of complex composites. *FENDT 2013 - Proceedings of 2013 Far East Forum on Nondestructive Evaluation/Testing: New Technology and Application*, 65–68. doi:10.1109/FENDT.2013.6635530
- Wang, Y. C., McPherson, K., Marsh, T., Gortmaker, S. L., & Brown, M. (2011). Health and economic burden of the projected obesity trends in the USA and the UK. *Lancet*, *378*(9793), 815–25. doi:10.1016/S0140-6736(11)60814-3
- Watson, S. J., Jones, A. L., Oatway, W. B., & Hughes, J. S. (2005). *Ionising Radiation Exposure of the UK Population : 2005 Review*. Chilton.
- Weidlinger-Associates. (2012). PZFlex User Manual. doi:10.1007/SpringerReference_28001
- Weill, R., & Shani, B. (1991). Assessment of Accuracy of Robots in Relation with Geometrical Tolerances in Robot Links. *Annals of the CIRP*, *40*(1), 395–399.
- Weir, A. J., Sayer, R., Wang, C., & Parks, S. (2015). A wall-less poly(vinyl alcohol) cryogel flow phantom with accurate scattering properties for transcranial Doppler ultrasound propagation channels analysis. In *IEEE Engineering in Medicine and Biology*.

- Weir, D. J., Moran, C. G., & Pinder, I. M. (1996). KINEMATIC CONDYLAR TOTAL KNEE Arthroplasty. *The Journal of Bone and Joint Surgery. British Volume*, 78-B, 907–911.
- Welch, J. N., Johnson, J. A., Bax, M. R., Badr, R., & Shahidi, R. (2000). A RealTime 3D Freehand Ultrasound System for Image-Guided Surgery. *2000 IEEE ULTRASONICS SYMPOSIUM*, 1601–1604.
- Wells, J. (2013). normnd.
- White, T. S. (2005). A Mobile Climbing Robot for High Precision Manufacture and Inspection of Aerostructures. *The International Journal of Robotics Research*, 24(2003), 589–598. doi:10.1177/0278364905055701
- Wilcox, P. D., Holmes, C., & Drinkwater, B. W. (2007). Advanced reflector characterization with ultrasonic phased arrays in NDE applications. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 54(8), 1541–50. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/17703657>
- Wiles, A. D., Thompson, D. G., & Frantz, D. D. (2004). Accuracy assessment and interpretation for optical tracking systems. *Medical Imaging 2004: Visualizations, Image-Guided Procedures and Display. Proceedings of SPIE*, 5367, 421–432. doi:10.1117/12.536128
- Wilhjelm, J. E., Pedersen, P. C., & Jacobsen, S. M. (2001). The influence of roughness, angle, range, and transducer type on the echo signal from planar interfaces. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 48(2), 511–21. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11370365>
- Williams, D. H., & Garbuz, D. S. (2010). Total knee arthroplasty : techniques and results. *BC Medical Journal*, 52(9), 447–454.
- Windolf, M., Götzen, N., & Morlock, M. (2008). Systematic accuracy and precision analysis of video motion capturing systems-exemplified on the Vicon-460 system. *Journal of Biomechanics*, 41(12), 2776–2780. doi:10.1016/j.jbiomech.2008.06.024
- Wojcik, G., Jr, J. M., & Carcione, L. (1999). Combined transducer and nonlinear tissue propagation simulations. *International Mechanical Engineering Congress & Exposition*. Retrieved from <http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Combined+transducer+and+nonlinear+tissue+propagation+simulations#0>
- Wojcik, G., Mould, J., Lizzi, F., Abboud, N., Ostromogilsky, M., & Vaughan, D. (1995). Non-linear modeling of therapeutic ultrasound. *1995 IEEE Ultrasonics Symposium*; ; 273-274; 1995; ; 1617-1622; 1995, 1617–1622.
- Woolhead, G. M., Donovan, J. L., & Dieppe, P. a. (2005). Outcomes of total knee replacement: a qualitative study. *Rheumatology (Oxford, England)*, 44(8), 1032–7. doi:10.1093/rheumatology/keh674

- Wright, N. C., Riggs, G. K., Lisse, J. R., & Chen, Z. (2008). Self-reported osteoarthritis, ethnicity, body mass index, and other associated risk factors in postmenopausal women—results from the Women’s Health Initiative. *Journal of the American Geriatrics Society*, 56(9), 1736–43. doi:10.1111/j.15325415.2008.01812.x
- Xu, N., Li, Y., & Zhou, Z. G. (2013). Ultrasonic Nondestructive Evaluation of Composite Components Using Advanced Total Focusing Method. *Advanced Materials Research*, 652-654, 1353–1356. doi:10.4028/www.scientific.net/AMR.652-654.1353
- Yan, C. X. B., Goulet, B., Pelletier, J., Chen, S. J. S., Tampieri, D., & Collins, D. L. (2011). Towards accurate, robust and practical ultrasound-CT registration of vertebrae for image-guided spine surgery. *International Journal of Computer Assisted Radiology and Surgery*, 6(4), 523–537. doi:10.1007/s11548-010-05362
- Yao, J. Q., & Seedhom, B. B. (1999). Ultrasonic measurement of the thickness of human articular cartilage in situ. *Rheumatology (Oxford, England)*, 38(12), 1269–71. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11477292>
- Yao, W., Abolmaesumi, P., Greenspan, M., & Ellis, R. E. (2005). An estimation/correction algorithm for detecting bone edges in CT images. *IEEE Transactions on Medical Imaging*, 24(8), 997–1010. doi:10.1109/TMI.2005.850541
- Yen, J. T., Steinberg, J. P., & Smith, S. W. (2000). Sparse 2-D array design for real time rectilinear volumetric imaging. *IEEE TRANSACTIONS ON ULTRASONICS, FERROELECTRICS, AND FREQUENCY CONTROL*, 47(1), 93–110.
- Young, K., & Pickin, C. G. (2000). Accuracy assessment of the modern industrial robot. *Industrial Robot: An International Journal*, 27(6), 427–436. doi:10.1108/01439910010378851
- Zapf, M., Schwarzenberg, G. F., & Ruitter, N. V. (2008). High Throughput SAFT for an Experimental USCT System as MATLAB Implementation with Use of SIMD CPU Instructions. *Proc. SPIE*, (692010), 692010–692010–11. doi:10.1117/12.770443
- Zhang, J., Hunter, A., Drinkwater, B. W., & Wilcox, P. D. (2012). Monte Carlo inversion of ultrasonic array data to map anisotropic weld properties. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 59(11), 2487–97. doi:10.1109/TUFFC.2012.2481
- Zhang, Y., Rohling, R., & Pai, D. K. (2002). Direct surface extraction from 3D freehand ultrasound images. *IEEE Visualization, 2002. VIS 2002.*, 45–52. doi:10.1109/VISUAL.2002.1183755

Appendix A

Calibration Part Design

A1. Calibration Part Design

In order for the position of a probe to be known to the robot controller, a calibration must be undertaken, allowing the probe position to be related to the robot. KUKA robot controllers allow for two levels of tool calibration: TCP and orientation calibration. The TCP is the origin of the local tool coordinate system, while the orientation is the definition of the axes relative to it. Calibration allows these parameters to be defined relative to the coordinates of the robot and, therefore, calculate the position and orientation of the tool based on known joint angle changes.

Orientation calibration is the upper level, with TCP calibration a prerequisite of use.

The ‘tool’ in this case was the ultrasound probe. As such, the orientation of the tool would be of great importance, given that 2D images would be placed in 3D space based on the position and orientation data provided by the robot controller.

While a number of methods of calibration exist (KUKA, 2010), the most suitable were found to be the XYZ 4 point method and the ABC 2 point method. The first stage of calibration was TCP definition, which required a choice of TCP that would relate well to the probe and the data it would retrieve. With this in mind, the centre point of the array was chosen. This can be seen in Figure A.1, which shows the three tool axes and the TCP as the origin. Calibration of the TCP using the 2 point method requires that the TCP be in contact with a static reference point, which is usually in the form of a spike. Bringing the probe face into contact with a sharp object would likely cause damage and, as such, the probe could not be used in the TCP calibration. To remedy this, a spike was made such that the point was located at the TCP, relative to the robot. This can be seen in Figure A.2, which shows the spike screwed into the separator




Figure A.1: the axial definitions of the tool, relative to the probe, with the TCP located at the origin. This configuration placed the 2D imaging plane in the X-Y plane (positive X) for simple reconstruction.




Figure A.2: the TCP calibration spike in place on the separator. The points of these spikes represented represented the centre of the origin of the array coordinate system in the two probe configurations.

To accurately calculate the desired length of the calibration spike, a Faro Quantum touch probe (Faro Technologies, FL, USA) was employed. Measurements were made at points along the two surfaces shown in Figure A.3, which correspond to the desired length of the calibration spike.




Figure A.3: Two surfaces from which planes were constructed, allowing the desired length of the TCP calibration spike to be determined.

Once the point representations of both surfaces had been acquired, these were loaded into Matlab 2014a (Mathworks, MA, USA). Each of the groups of surface points was treated as a plane, the distance between which would be equivalent to the desired length of the calibration spike. To calculate this distance, the Matlab function ‘plane_line_intersect’ (Khaled, 2007) was employed. This function requires the vector normal of one of the planes, any point on that plane, and two points from the intersecting line. In order to calculate the vector normal of the plane, the Matlab function ‘normnd’ (Wells, 2013) was employed. This function calculates the normal to a plane represented by a group of 3D coordinates by way of Principal Component Analysis (Hoppe, DeRose, Dechamp, McDonald, & Stuetzle, 1992) – a method that not only computes the normal of a group of points, but confirms whether or not it is even possible to fit a normal. The first point needed to define the intersecting line in the ‘plane_line_intersect’ function, any point was chosen from one of the planes. The second point required was produced by adding the normal associated with that plane to the first point. The plane with which this line intersected was the opposing plane.

This can be seen in Figure A.4, which shows all the measured surface points on the probe face plate in red, those on the rear face of the probe holder in blue, an intersecting line from one of the points and the intersecting point in yellow. The transparent panels are an attempt to better represent the planes.




Figure A.4: The measured surface points for the probe face (red) and the back of the probe holder (blue). Also shown is the intersecting line from one of the points on the probe face and the point of intersection (yellow).

The described procedure was repeated for every point on both of the planes, producing 21 distances. A distance of 6mm was subtracted from each of these results to account for the 3mm radius of the protective sphere covering the touch probe, in which the touch probe sat in the centre. This yielded a mean distance of 43.22mm, a mean difference between distances of 0.05mm and a maximum difference of 0.14mm. As the difference between the calculated distances was so small and there was an agreement between these results and those found with callipers, the mean distance of 43.22mm was accepted as the desired distance of the spike.

It should, however, be mentioned that the probe has a face plate covering the elements of the array, acting as both a protective cover and a matching layer. The TCP was not located on the face of the probe, but rather on the face of the elements of the array. Based on typical matching layer material properties (O'Leary, Hayward, Smillie, & Parr, 2005), the face plate thickness was found to be 0.125mm.

Subtracting this from the first distance yielded a desired spike length of 43.10mm.

Once the spike had been manufactured, the length thereof was tested for accuracy. This was achieved using a similar procedure to that described previously. The technique differed in that instead of finding the distance between two planes, the required measurement was that between a plane and a single point.

However, unlike the previous method where the mean of all measurement points was employed, the maximum distance was chosen. The reason for this is illustrated in Figure A.5, which shows a number of scenarios where the touch probe is in contact with the spike tip. However, in each of these scenarios, the measurement point of the touch probe is in a different position. As the Faro arm is a hand held measurement device, it is very likely that the centre of the sphere would not be normal to the spike at every measurement, as it is in the middle scenario. As can be seen in Figure A.5, when the sphere is off centre, the measurement point dips below that at the normal. Therefore, if the assumption is made that all measurements were made in contact with the tip of the spike, it is reasonable to say that the maximum distance calculated will be the closest to being normal to the tip and, in turn, the most accurate measurement of spike length.

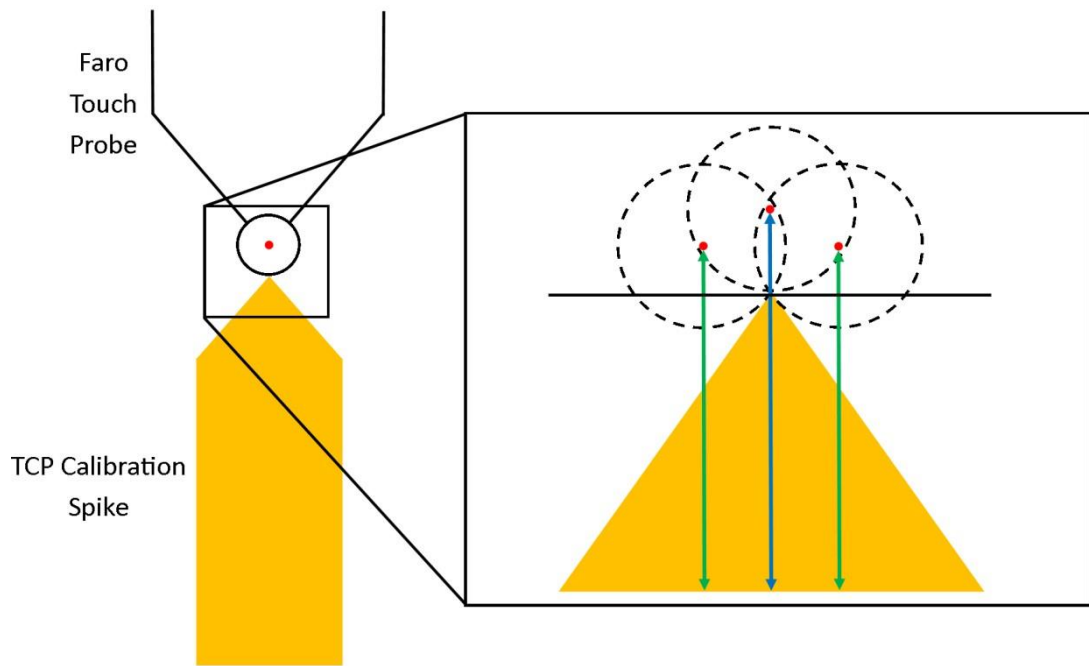


Figure A.5: A number of scenarios in which the touch probe is in contact with the tip, but the Faro measurement point is in a different location. The measurement point is shown as a red dot, the differences in position are illustrated using green and blue length lines and the associated spherical protector positions are denoted by dashed lines.

The TCP calibration spike length was measured as 42.97mm, with the maximum difference between that length and any of the other recordings being 1.56mm. This can be seen in Figure A.6, which shows all the measurement points bunched together in green, while the chosen, most accurate reading is displayed in red. While the actual length of the spike was 0.13mm short of the desired length, it was accepted, as manufacturing constraints meant that it was unlikely that the accuracy could be improved upon.

distance between the plane and the tip of the spike is 42.9685mm




Figure A.6: The measured surface points for the flat part of the separator (blue), the tip positions (green) and the most accurate tip position recording (red). Also shown is the intersecting line from the chosen tip position measurement and the point of intersection (black).

With the TCP calibration spike completed, it was possible to carry out the four point XYZ calibration procedure, which defined the physical TCP relative to the robot. This consisted of touching the calibration spike and the static reference spike by manually jogging the robot, as can be seen in Figure A.7. It is advantageous to vary the poses used as much as possible, providing a more accurate and robust calibration. However, this was difficult to achieve, as the tip of the calibration spike was not as sharp as desired. The result of the spike being blunt was that the X direction of the calibration spike could not vary too much from the Z direction of the world frame. Therefore, while the calibration procedure may produce a low error, inaccuracies in positional reconstruction by the robot controller could be high. According to Kevin Amos (project manager at Phoenix Systems and former training manager at KUKA Robotics), acceptable calibration error is any value under 0.8mm. This calibration error is a value produced by the robot controller, representing the positional difference

between each of the four calibration points. The calibration error for the TCP calibration of Tool 1 was 0.774mm, while that of the Tool 2 was 0.562mm. While the former value was very close to the prescribed limit, it was still below it and, as such, was accepted.




Figure A.7: the XYZ 4 point calibration procedure showing four varying points of contact, with the final ensuring the X axis of the tool is parallel to the Z axis of the WORLD reference frame.

As well as the TCP calibration, orientation calibration was also required to allow for correct positioning of captured data. The two point ABC method was employed for this, which defines the tool axes relative to those of the world reference frame. The first part of this technique requires that the static reference spike come into contact with a point on the negative X axis of the tool. However, as can be seen in Figure A.1, this would be located at some point inside the probe. As such, another spike was manufactured to the same specifications as the TCP calibration spike, but to a shorter

length, as can be seen in Figure A.8 (a) and (b). Once placed in the same position as the TCP calibration spike was, it fulfilled the requirement of providing a point on the negative X axis. The second part of the two point ABC method requires that the reference spike come into contact with a point on the positive XY plane, which can be seen in Figure A.1. This was achieved using the spikes shown in parts (c) and (d) of Figure A.8.




Figure A.8: Orientation calibration spikes for negative X direction calibration for Tool 1 (a) and Tool 2 (b) and positive XY plane for Tool 1 (c) and Tool 2 (d).

Appendix B

Composite Human Distal Femur Surface Reconstruction

B1. Results Tables




Figure B.1: The standard deviation achieved for data compared with the laserscanned reference mesh for the cuboidal scan path. ‘Smoothing’ refers to the application of a smoothing function to the 2D points forming contours in each image, ‘nearest neighbour’ denotes the application of the outlier elimination algorithm and

‘sampling’ indicates whether or not the mesh itself or a densely sampled point cloud representing the mesh was employed.




Figure B.2: The maximum error achieved for data compared with the laser-scanned reference mesh for the cuboidal scan path. The parameter descriptions correspond to those used in Figure B.1.




Figure B.3: The mean error achieved for data compared with the CAD reference mesh for the cuboidal scan path. The parameter descriptions correspond to those used in Figure B.1.




Figure B.4: The standard deviation achieved for data compared with the CAD reference mesh for the cuboidal scan path. The parameter descriptions correspond to those used in Figure B.1.




Figure B.5: The maximum error achieved for data compared with the CAD reference mesh for the cuboidal scan path. The parameter descriptions correspond to those used in Figure B.1.




Figure B.6: The mean error achieved for compared data from the hemispherical scan path. The parameter descriptions correspond to those used in Figure B.1.

Appendix C

Bovine Distal Femur Surface Reconstruction

C1. Results Tables




Figure C.1: Comparison of the mean error attained using different robot velocities.

The compared point clouds employed the 'Bottom' contour extraction method, did not use a single pass robot path and had the nearest neighbour point elimination algorithm applied.




Figure C.2: Comparison of mean error, for the four processing methods based on changes in contour extraction method and in the application of smoothing functions. The compared point clouds employed a 1% path velocity and no single top side pass.




Figure C.3: Comparison of the outlier elimination (denoted as nearest neighbour) and the use of the single top side pass attained using different contour extraction methods. The compared point clouds employed no smoothing function, used the ‘Top’ contour extraction method and employed a 1% velocity.




Figure C.4: Comparison of the mean error of the mesh attained using different TCP velocities. The compared meshes employed the ‘Top’ extraction method, did not have the nearest neighbour point elimination algorithm applied and did not use a single top side pass.




Figure C.5: Comparison of the mean error attained with both smoothed and unsmoothed meshes. The compared meshes employed the full top side pass, did not

have the nearest neighbour point elimination algorithm applied and employed a 1% TCP velocity.




Figure C.6: Comparison of the standard deviation attained using different contour extraction methods at different velocities. The compared meshes employed no smoothing function, did not use the single top side pass and had the nearest neighbour point elimination algorithm applied.




Figure C.7: Comparison of the outlier elimination (denoted as nearest neighbour) and the use of the single top side pass attained using different contour extraction methods. The compared meshes employed no smoothing function, used the contour extraction method described in chapter 4 and employed a 1% velocity.