

# **Quantitative Assessment of Patient Suitability for Magnetic Resonance-guided High Intensity Focused Ultrasound Therapy**

*Ngo Fung Daniel Lam*

A dissertation submitted in partial fulfillment  
of the requirements for the degree of  
**Doctor of Philosophy**  
of  
**The Institute of Cancer Research.**

Division of Radiotherapy and Imaging, Joint Department of Physics  
The Institute of Cancer Research and  
The Royal Marsden NHS Foundation Trust

January 5, 2021

I, Ngo Fung Daniel Lam, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Signed:

Ngo Fung Daniel Lam

January 5, 2021

# Abstract

Clinical assessment of a patient's suitability for magnetic resonance-guided high intensity focused ultrasound (MRgHIFU) therapy currently involves subjective judgements based on available diagnostic images and prior clinical experience. The presence of organs at risk and acoustic obstructions, such as bone and air, and target depth are taken into account. A quantitative method of assessing suitability from images available at referral may minimise the number of patients incorrectly offered, or denied, treatment. A workflow for this assessment is developed herein for pelvic tumour patients.

Novel workflow components include identification of each patient's ideal treatment angle, assessment of the percentage tumour volume that can be covered using standard 'treatment cells' defined in the MRgHIFU control software, and assessment of the percentage tumour volume that can be treated (i.e. receive a cytotoxic thermal dose). Volunteer and patient image datasets, with the subjects lying both supine ('referral imaging') and in an oblique supine decubitus (treatment) position, were used for methodology development and testing. A method of identifying a subject's ideal treatment angle using predicted tumour coverage was developed. These angles were compared with clinically-used treatment angles. Practical methods for assessment of tumour coverage from referral imaging have been developed and their predictive capability quantified. Tumour treatability in treatment image datasets was analysed using the k-Wave acousto-thermal simulation package.

Calculated ideal treatment angles were within  $5\pm 2^\circ$  of clinical treatment angles. Predictions of tumour coverage derived from referral images agreed with those from treatment images within  $12\pm 7\%$  (range: 4-21%). Refinements to the

tumour coverage method improved computational speed by a factor of 7 on average (from  $19.7 \pm 8.8$  to  $2.8 \pm 2.0$  hours). Tumour treatability was  $32 \pm 14\%$  (range: 15-50%) less than tumour coverage with ablated tissue volumes lying  $9.3 \pm 1.6$ mm shallower than the geometric focus, suggesting tumour coverage overestimation. Despite limitations, the developed methods show significant promise.

# Acknowledgements

I would like to thank my supervisors, Ian Rivens, Gail ter Haar, Emma Harris, for assisting me with my scientific endeavours, ranging from reading my drafts to providing an experienced opinion on my scientific results. I would like to thank members of the ICR MRI Physics unit, in particular Matthew Blackledge, Simon Doran, and Matthew Orton, for assisting me with extracting segmentation data from OsiriX and Horos software. I would like to thank Ari Partanen and Mathieu Allard from Profound Medical for providing technical assistance and details with the Sonalleve<sup>®</sup> MRgHIFU system. I would like to thank Sharon Giles and Nandita de Souza, other members of The Royal Marsden Hospital's MR unit and therapy ultrasound team, and the patients and volunteers whose data was critical for the development of this PhD project.

# Contents

<b>1</b>	<b>Introduction</b>	<b>31</b>
1.1	Clinical Use of High Intensity Focused Ultrasound . . . . .	31
1.2	Current Clinical Workflow for MRgHIFU . . . . .	32
1.3	Hypothesis . . . . .	34
1.4	Background . . . . .	34
1.5	Prospective Patient Suitability Workflow . . . . .	35
1.6	Thesis Structure . . . . .	38
<b>2</b>	<b>Background</b>	<b>39</b>
2.1	Introduction . . . . .	39
2.2	Physical Principles of HIFU . . . . .	39
2.3	Drawbacks of HIFU . . . . .	43
2.4	Image Segmentation . . . . .	44
2.5	Screening for High Intensity Focused Ultrasound . . . . .	48
2.5.1	Patient Positioning . . . . .	49
2.5.2	Sonication Placement . . . . .	50
2.6	Simulation of Acoustic Propagation . . . . .	52
2.6.1	Finite-Difference Methods . . . . .	52
2.6.2	Pseudo-spectral and k-space Methods . . . . .	55
2.6.3	Angular Spectrum Methods . . . . .	57
2.6.4	Summary of Acoustic Simulation Methods . . . . .	60
2.7	Simulation of Thermal Bioeffects . . . . .	62
2.7.1	Finite-Difference Methods . . . . .	62

2.7.2	Pseudo-spectral and k-space Methods . . . . .	63
2.7.3	Monte Carlo Methods . . . . .	65
2.7.4	Summary of Thermal Simulation Methods . . . . .	66
<b>3</b>	<b>Clinical Studies</b>	<b>68</b>
3.1	Introduction . . . . .	68
3.2	The Sonalleve <sup>®</sup> V2 MRgHIFU System . . . . .	68
3.3	Magnetic Resonance (MR) Image Acquisition . . . . .	71
3.3.1	Common MR Scanning Protocol . . . . .	71
3.3.2	Volunteer Image Acquisition . . . . .	71
3.3.3	Patient Image Acquisition . . . . .	72
<b>4</b>	<b>Image Registration</b>	<b>75</b>
4.1	Introduction . . . . .	75
4.2	Method . . . . .	75
4.2.1	Evaluation of Uncertainties in Image Registration . . . . .	76
4.3	Results . . . . .	77
4.3.1	Evaluation of Uncertainties in Image Registration . . . . .	77
4.4	Discussion and Conclusion . . . . .	77
<b>5</b>	<b>Image Segmentation</b>	<b>78</b>
5.1	Introduction . . . . .	78
5.2	Method . . . . .	79
5.2.1	Body Outline . . . . .	79
5.2.2	Organs at Risk . . . . .	82
5.2.3	Acoustic Obstructions . . . . .	82
5.2.4	Volunteers . . . . .	85
5.2.5	Patients . . . . .	85
5.2.6	Target Tumour . . . . .	85
5.2.7	Evaluation of Segmentation Quality . . . . .	88
5.3	Results . . . . .	90
5.3.1	Target Tumour Misalignment . . . . .	90

5.3.2	Evaluation of Segmentation Quality . . . . .	90
5.4	Discussion . . . . .	91
5.4.1	Target Tumour Misalignment . . . . .	91
5.4.2	Manual Segmentation Validation . . . . .	92
5.4.3	Automatic Segmentation Validation . . . . .	92
5.5	Conclusion . . . . .	92
<b>6</b>	<b>Predicting Target Coverage from Referral Imaging</b>	<b>93</b>
6.1	Introduction . . . . .	93
6.2	Method . . . . .	95
6.2.1	Overview . . . . .	95
6.2.2	Target Coverage Analysis . . . . .	97
6.2.3	Uncertainty of Tumour Coverage Analysis due to Quantisation	106
6.2.4	Cell Packing Density . . . . .	107
6.3	Results . . . . .	107
6.3.1	Gel-Pad Thickness and Membrane Bowing Distance . . . . .	107
6.3.2	Volunteer Target Coverage . . . . .	110
6.3.3	Patient Tumour Coverage . . . . .	112
6.3.4	Uncertainty of Tumour Coverage Analysis . . . . .	115
6.3.5	Cell Packing Density . . . . .	116
6.4	Discussion . . . . .	117
6.4.1	Gel-Pad Thickness and Membrane Bowing Distance . . . . .	117
6.4.2	Volunteer Target Coverage . . . . .	117
6.4.3	Patient Target Coverage . . . . .	119
6.4.4	Cell Packing Density . . . . .	122
6.4.5	Limitations of the Study . . . . .	122
6.5	Conclusion . . . . .	123
<b>7</b>	<b>Improved Method for Calculating Patient Tumour Coverage</b>	<b>125</b>
7.1	Introduction . . . . .	125
7.2	Method . . . . .	127
7.2.1	Overview . . . . .	127



7.2.2	Comparison with the Device-First Method . . . . .	129
7.2.3	CPU versus GPU Implementations of the Ray-triangle Intersection Algorithm . . . . .	136
7.3	Results . . . . .	137
7.3.1	Comparison with the Device-First Method . . . . .	137
7.3.2	CPU versus GPU Implementations of the Ray-triangle Intersection Algorithm . . . . .	138
7.4	Discussion . . . . .	138
7.5	Conclusions . . . . .	139
<b>8</b>	<b>Identification of the Ideal Treatment Angle</b>	<b>141</b>
8.1	Introduction . . . . .	141
8.2	Method . . . . .	143
8.2.1	Calculation of the Starting Treatment Angle . . . . .	143
8.2.2	Identifying the Ideal Treatment Angle . . . . .	148
8.3	Results . . . . .	153
8.3.1	Starting Treatment Angle . . . . .	153
8.3.2	Ideal Treatment Angle . . . . .	153
8.4	Discussion . . . . .	155
8.4.1	Comparison between the Starting and Ideal Treatment Angles	156
8.4.2	Comparison between the Ideal and Clinical Treatment Angles	157
8.4.3	Comparison between the Starting and Clinical Treatment Angles . . . . .	157
8.4.4	Precision Required in Patient Positioning . . . . .	158
8.4.5	Limitations of this Study . . . . .	158
8.5	Conclusions . . . . .	159
<b>9</b>	<b>Evaluating Patient Treatability</b>	<b>160</b>
9.1	Introduction . . . . .	160
9.2	Development of Simulation Methods . . . . .	161
9.2.1	Overview . . . . .	161
9.2.2	Inputs . . . . .	164

9.2.3	Simulation Grid Extraction . . . . .	164
9.2.4	Medium Generation . . . . .	166
9.2.5	Acoustic Simulation . . . . .	169
9.2.6	Thermal Simulation . . . . .	180
9.2.7	Calculation of Tumour Treatability . . . . .	182
9.3	Results . . . . .	185
9.3.1	Acoustic Simulation of a Patient Sonication . . . . .	185
9.3.2	Thermal Simulation . . . . .	189
9.3.3	Tumour Treatability . . . . .	193
9.4	Discussion . . . . .	196
9.4.1	Acoustic Simulation of a Patient Sonication . . . . .	196
9.4.2	Thermal Simulation . . . . .	197
9.4.3	Tumour Treatability . . . . .	199
9.4.4	Limitations of this Study . . . . .	202
9.5	Conclusions . . . . .	203
<b>10</b>	<b>Overall Conclusions and Future Work</b>	<b>205</b>
10.1	Overall Conclusions . . . . .	205
10.2	Future Work . . . . .	207
10.2.1	Treatability Estimation . . . . .	207
10.2.2	Structural Deformation . . . . .	209
10.2.3	Robustness . . . . .	209
10.2.4	Clinical Translation . . . . .	210
	<b>Appendices</b>	<b>212</b>
	<b>A Volunteer and Patient Data</b>	<b>212</b>
	<b>B Standard Operating Procedure for Image Registration</b>	<b>215</b>
	<b>C Investigation of the Ideal Sample Square Size</b>	<b>217</b>

<b>D Paper: Prediction of pelvic tumour coverage by magnetic resonance-guided high-intensity focused ultrasound (MRgHIFU) from referral imaging</b>	<b>219</b>
<b>Bibliography</b>	<b>234</b>

# List of Figures

- 1.1 Proposed workflow for the assessment of patient treatability. The input is a referral image dataset (RID), in which the patient is imaged supine on a diagnostic scanner. Step 1: important anatomical tissues are identified and segmented. Step 2: an ideal treatment position is calculated (represented by the isocentre line, which goes through the predicted transducer home position and magnetic isocentre) by iterative assessment of (Step 3) the percentage of the tumour volume covered, which is the percentage of the tumour that can be reached by the HIFU focus. Step 4: when the ideal treatment angle is found, the percentage of the tumour volume treatable, which is the percentage of the tumour that has received a thermal dose  $\geq 240$  CEM43, is calculated using acoustic and thermal simulations. From this, a clinical decision on whether to pursue patient treatment with MRgHIFU is made. . . . . 37
- 3.1 Schematic of the Sonalleve<sup>®</sup> V2 MRgHIFU system showing: LEFT - a subject lying on the MR bed, compressing the acoustic-coupling gel-pad and bowing the acoustic membrane which seals the oil bath. The membrane coil bows with the membrane. Ideally, the target tissue is centred directly above the transducer's home position, and the centre of the membrane/gel pad, and below the magnetic isocentre. RIGHT- a coronal view of the MRgHIFU couch showing the transducer's home position below the centre of the membrane. 69

5.1	Body masking workflow, depicted for a representative volunteer treatment dataset and described in the text. Step 1: Otsu thresholding using a 512-bin histogram. Step 2: Connected components labelling and extraction of the largest component, assumed to be the body. Step 3: Anterior, left and right edges of the body segment are outlined, and Step 4: the array is zero-padded by at least 50 voxels at each side, in order to allow for processes in Step 5 to function. Step 5: morphological dilation, flood-filling and erosion. Step 6: the padding is removed, and the result is the body mask. . . . .	81
5.2	Extracorporeal air segmentation workflow. Input Fat and Water treatment datasets are individually normalised and then summed together (Step 1). Step 2: an intensity histogram is generated from the image, and the threshold is found at the lowest-intensity trough. Step 3: the image is thresholded so that voxels with intensities below the threshold are marked. Step 4: Using connected-components labelling, individual connected segments are identified. Segments that are not connected to the anterior, right or left edges of the image dataset are removed. Step 5: morphological flood-filling is used to fill holes in the remaining segments, and the output result is the air segmentation. . . . .	84
5.3	Target tumour segmentation from T2wLFOV images (a representative axial slice from a patient). The green line is the manually drawn tumour contour. . . . .	86
5.4	Quantitative data (a,b) and a representative image (c) of the agreement between treatment and registered-referral tumour segments. (a): Mean Surface-Surface Distance (MSSD) and Hausdorff distance is depicted. (b): Dice Similarity Coefficient. (c): The blue region represents the G29 treatment tumour surface, and the red region represents the G29 registered-referral tumour surface. . . .	90

5.5	Automatic segmentation of the volunteer treatment image datasets results in spurious segments between the subject and the gel-pad. Green represents the manually-segmented validation slice, magenta represents the automatic segmentation, and white denotes the overlap between the two segments. . . . .	91
6.1	Schematic of developmental methodology used in this study. The target (tumour) volume coverage estimated from a subject's referral image dataset (top row), and that calculated from the treatment image dataset (bottom row), were compared in order to assess the accuracy of predicting target (tumour) coverage from referral datasets. Step 1: The referral dataset is rotated into the same orientation as the treatment dataset using affine registration both to allow comparison with the treatment dataset and to simulate the rotation in Step 2 of the prospective patient workflow (Figure 1.1). In step 2, segmentation of relevant tissues (e.g. bony acoustic obstructions) was performed to identify tissues that could obstruct acoustic access to the target. Step 3: target volume coverage was calculated for the rotated referral dataset and the treatment dataset, and then compared to assess the accuracy of the methodology. . . . .	96
6.2	Ray intersection with segments, represented as binary (0 and 1) arrays, were checked as follows. Rays are represented as a set of 'ray-points' (red stars) that form a line. Scores at each ray-point were linearly interpolated from the values of the surrounding voxels. Each ray-point's score is multiplied by the interval between ray-points, and summed to generate a path length of intersection for the entire ray. . . . .	98

6.3	Practical restrictions applied to the transducer’s translation capabilities (solid red lines) for volunteer datasets only. (a) For treatment datasets (TID), the left-right translation was limited by the extent of acoustic coupling between the volunteer’s skin and the gel pad. Registered-referral datasets shared these left-right restrictions. (b) For registered-referral datasets (RRID), the transducer’s inferior-superior translation was restricted by the extent of pelvic bone and the requirement for a full body outline within the image. Treatment datasets shared these inferior-superior restrictions. . . . .	99
6.4	Method used to predict transducer’s anterior-posterior home position. The magnetic isocentre in the treatment dataset is known because the registered-referral image dataset (RRID) had been registered to the treatment dataset. A line was drawn downwards from the treatment image dataset (TID) isocentre and intersected the skin at the skin point. From this skin point, the home position was calculated using the average compressed gel-pad thickness, the average membrane bowing distance, and the calibrated distance between undeformed membrane and home position of 67.5 mm (see Figure 3.1). . . . .	101
6.5	Method for quantifying target volume covered within a dataset (volunteer treatment dataset in this example). A regular 3D grid of potentially accessible points was created (blue crosses) within the target: soft tissue (volunteers) or tumour (patients). For each transducer position and tilt, the acoustic beam was checked for intersection with any acoustic obstructions (green contours) or organs at risk. If no obstruction exists, an 8 mm treatment cell was created around the focus (yellow ellipse). Grid points within a treatment cell were marked as ‘covered’ (red crosses). . . . .	104
6.6	The coarse cell packing (red) leaves more of the target volume (black) in between cells uncovered than the fine cell packing (green).107	

6.7	Body Mass Index (BMI), weight, angle and height versus the compressed gel-pad thickness close to the isocentre for volunteer treatment datasets. Uncertainties for weight, height and BMI were estimated to be $\pm 1^\circ$ , $\pm 1$ kg, $\pm 4$ mm and 2% respectively. . . . .	108
6.8	Body Mass Index (BMI), weight, angle and height versus the membrane bowing distance close to the isocentre for volunteer treatment datasets. Uncertainties for angle, weight, height, and BMI were estimated to be $\pm 1^\circ$ , $\pm 1$ kg, $\pm 4$ mm and 2% respectively. . .	109
6.9	Representative axial anatomical images of volunteers, with the treatment image (green) overlaid on the registered-referral image (magenta). The soft tissue volume covered as predicted from the registered-referral dataset only (red), that calculated from the treatment dataset only (cyan) and that identified by both datasets (white) were overlaid on the anatomical data. For volunteers, the target volume was defined as the soft tissue volume covered in the treatment dataset (white plus blue volumes). The percentage target volume covered is the white volume divided by the white plus blue volume represented as a percentage. . . . .	111
6.10	Left: Anatomical overlay between registered-referral (red) and treatment (grayscale) datasets, with the isocentre at the crosshair centre. The isocentre line goes vertically through the isocentre. The contents of the green rectangle is shown magnified on the right. Right: The anterior-posterior distance between the registered-referral skin point and the treatment skin point, directly below the isocentre, was measured to be 12.2 mm. . . . .	112



6.11 Representative axial anatomical images of patients, with the treatment (green) and registered-referral anatomy (magenta) overlaid. The entire treatment dataset tumour segment outline (orange), the treatment-only covered volume (cyan), the registered-referral-only covered volume (red) and the overlap between the treatment and registered-referral covered volumes (white) were overlaid on the anatomical images. The overall treatment tumour covered volume is white plus blue; the overall registered-referral tumour covered volume is white plus red. For simpler visual interpretation, only the treatment tumour outline is shown (orange line). . . . . 113

6.12 Bar charts depicting the percentage of the target volume covered for volunteers (a) and the percentage of the tumour volume covered for patients (b). (a) For volunteers, target volume was defined as the soft tissue volume covered as assessed from the treatment images. The percentage of the target volume that was predicted to be covered from registered-referral datasets was plotted, for each volunteer, for each treatment position (Steep (blue) or Shallow (orange)). (b) For patients, the percentage of the registered-referral tumour segment volume predicted to be covered from registered-referral datasets (red), are compared to the percentage of the treatment tumour segment volume calculated to be covered from treatment datasets (cyan). For simpler visual identification, the numbers above each pair of bars represents the difference in % Tumour Volume Covered predicted from registered-referral datasets and that calculated from treatment datasets. . . . . 114

6.13 Sagittal anatomical image of patient G3 at the tumour centroid, overlaid with acoustic obstructions (AOs) and organs at risk (OARs, combined red overlay), the tumour (green) and the virtual transducer (blue transducer, orange acoustic beam). Notably, the extracorporeal air was present inferior to and posterior of the tumour. The virtual transducer is focused on the isocentre, but offset in the Inferior-Superior direction in order to draw attention to the relative position of the extracorporeal air and the tumour. . . . . 115

6.14 Spatial distribution of reachable treatment cell density for patient G24 in the axial plane. Point colours correspond to the number of treatment cell centres found within a 1 mm<sup>3</sup> cube centred at the point. The maximum cell packing density was 4 cells per mm<sup>3</sup>. The tumour is overlaid as a mesh segment (red), and an 8 mm diameter treatment cell (orange) is overlaid for comparison. . . . . 116

7.1 Workflow for the target-first methodology, demonstrated using a representative patient treatment dataset. Step 1: Input Dixon images are segmented to give acoustic obstructions (green), organs at risk (blue, OARs), the tumour (red) and the body outline (grey). The segments are then transformed into triangular surface meshes. In step 2, candidate cell positions (black crosses) are generated within and around the tumour mesh. In step 3, the tumour coverage is calculated by back-projecting plausible transducer positions (only one shown, for ease of visual interpretation) from each candidate cell position. Each transducer position is then tested for whether the beam would intersect OARs and acoustic obstructions, and whether placing the transducer in that position would breach the specified limits of translation for the MRgHIFU system. In step 4, if the current iteration is not the final iteration, then the reachable and unreachable cell positions are identified. Exclusion meshes were derived such that the next iteration of cell points would be generated only outside of the exclusion meshes, and within and around the remaining uncovered tumour volume. If the current iteration was the final iteration, then the percentage of the tumour volume covered is output. . . . . 128

7.2 An illustrative example of a segment, where the '1' boxes are the segment voxels and the '0' boxes are the background voxels. DFM is device-first method. If the allowed path length of intersection was set to 0 mm, the red line demonstrates the effective surface of the segment that rays may not pass without being marked as intersecting. . . . . 130

7.3 Workflow for Step 3 in the target-first methodology. Target grid points (blue crosses) were generated during the meshing stage. Step 1: Each candidate cell centre (black cross) generated vectors pointing towards the posterior with a maximum angulation of  $10^\circ$  off the anterior-posterior axis. Step 2: For each vector, the associated transducer position was calculated, and a virtual transducer generated such that the focus was at the cell centre and the transducer beam axis was aligned with the vector, but in the opposite direction. Step 3: The transducer acoustic beam was checked for obstruction with OAR or acoustic obstruction meshes, and the transducer position derived in Step 2 was assessed for whether device limits must be breached in order to place the transducer there. If the acoustic beam is not obstructed and the transducer position was within specification, all target grid points within an 8 mm treatment cell around the transducer geometric focus was marked as covered (green crosses). The total number of covered target grid points multiplied by the product of the spacing in between target grid points was the covered tumour volume. . . . . 134

7.4 Computational time required when using the GPU implementation (blue) or the MATLAB-native CPU implementation (red) of the ray-triangle intersection algorithm when assessing whether increasing numbers of candidate cell centres are reachable. . . . . 138

8.1 Workflow for calculating the starting treatment angle. Step 1: Tumour, organs at risk and bone (an acoustic obstruction) are segmented and meshed. Step 2: A set of vector rays, each representing patient treatment angles, are generated within 3D space. All rays originated from the tumour centroid, are in an axial plane and are angled at most 40° away from the vertical. Step 3: rays that intersect any organs at risk and acoustic obstructions, and have an origin-to-skin intersection distance greater than the MRgHIFU focal length (140 mm), are removed to improve computational performance. Step 4: For all remaining rays, the acoustic window is estimated by fitting a cone (originating at the tumour centroid, with the cone axis coincident with the ray). The treatment angle associated with the ray that has the largest cone is used as the starting treatment angle. . . . . 145

8.2 A visual comparison of the mesh surfaces of two meshing methods applied to the same example segment, represented here by a two dimensional binary pixel array. Each pixel in the segment array is either '1' or '0', representing the segment tissue/material, or everything else, respectively. The mesh surface generated by the device-first meshing method (red, described in Chapter 7), which was used for OARs and bones, is compared to that generated by the marching cubes isosurface meshing method (cyan), which was used for the body outline. . . . . 147

8.3 Workflow for identifying the ideal treatment angle. Inputs are the body outline, tumour, organ-at-risk (OAR) and acoustic obstruction (i.e. bone) meshes, as well as the starting treatment angle. For ease of visual interpretation, the body outline is not shown except in Step 2, where it is used. Step 1: The ray representing the isocentre line of the starting treatment angle ('starting ray') is generated. Rays are generated perpendicular to the inferior-superior axis around the starting ray, each ray representing the isocentre line of a particular treatment angle. Step 2: for each ray, the transducer home position is identified using the techniques developed in Section 6.2.2.5. Each ray's intersection with the body mask mesh is found and the intersection point defined as the skin point. From the skin point, estimates of the compressed gel-pad thickness (9.8 mm) and the membrane bowing distance (10.0 mm) are used to obtain the estimated unbowed membrane, from which the transducer's home position is known to be a fixed distance (67.5 mm for the Sonalleve V2) posterior. The patient's meshes are positioned and rotated such that the tumour centroid, the skin point, the magnetic isocentre and the transducer's home position all lie on a single 'isocentre line' (red). Step 3: for each treatment angle, the tumour coverage is calculated. Ultimately, tumour coverage is plotted against treatment angle. . . . . 150

8.4 The variation of tumour coverage with treatment angle, for all patients. The starting treatment angle (red vertical), in which the approximate acoustic window was maximal, the ideal treatment angle (blue dashed vertical), in which the tumour coverage is maximal, and the clinical treatment angle (black vertical), which was used in treatment, are shown. Tumour volume coverage at angles 5° steeper and shallower than the ideal treatment angle were calculated, and the lower of the two was depicted by the horizontal pink line. For patients G27 and G29, the starting and ideal treatment angles are the same. For patient G3, tumour coverage was identical for at all angles tested. . . . . 154

8.5 Difference in estimated tumour volume covered between the starting treatment angle (red arrow), and the ideal treatment angle (blue arrow), for patient G8<sub>2</sub>. (a): Tumour volume is represented as discrete points; those covered by both the starting and ideal angles are grey, by the starting angle only were red, and by the ideal angle only were blue. The tumour itself is represented as the red mesh in (a) and is grey in the other subplots. (b): An axial view of the tumour within the pelvis showing covered volumes relative to the acoustic obstructions and organs-at-risk (green meshes). Views of the covered tumour volumes from (c) the ideal angle (19° from supine) and (d) the starting angle (29° from supine). . . . . 155

9.1 Schematic for the methodology used to quantify patient treatability. Inputs are the patient images (representative treatment image slices shown here), the deepest and shallowest reachable treatment cells (magenta), the covered tumour volume (yellow) and the entire tumour volume (red). The transducer (blue) was positioned and angled such that the geometric focus was placed at the reachable treatment cell centre. In step 1, the simulation grid was extracted from the larger dataset. In step 2, each voxel of the extracted region was assigned acoustic and thermal properties based on voxel intensity. A density map is displayed, with white denoting the highest density and black denoting the lowest density in the image. In step 3, an acoustic simulation is performed in order to estimate the acoustic pressure field in the tissue. The pressure field, overlaid on the density map and tumour, shows regions of high (yellow) and low (blue) pressure. In step 4, a thermal simulation is performed to identify the ablated tissue resulting from acoustic energy absorption and heat transfer. In step 5, the ablated tissue volumes (cyan) from sonicating the deepest and shallowest reachable treatment cells were used to estimate the maximum treatable depth and the proportion of tumour treatable. . . . . 163



9.2 Extraction of the simulation grid from a patient image dataset, using the known transducer position and angulation required to target the intended treatment cell. (a): An example cross-section of the original image dataset is shown, with the tumour (red), covered tumour (yellow), shallowest reachable treatment cell (magenta) and associated transducer position and angulation (blue beam, orange beam axis) overlaid. (b): The original image dataset is rotated such that the transducer beam axis is made vertical, and the Y-axis is defined as parallel to that. (c): The simulation grid is extracted from the rotated dataset. The coordinate systems used to refer to directions within (a) the original image dataset (patient orientation-based directions), (b) the rotated and (c) the extracted grid are shown below the images. . . . . 165

9.3 Comparison of the pressure amplitude simulated by k-Wave and that calculated by O’Neil’s solution at the spatial resolution of 551  $\mu\text{m}$  (2.2 grid points per minimum wavelength) along the beam axis (top) and across the beam axis through the transducer focus (bottom). k-Wave underestimated the maximum focal peak by 9 MPa (33%). . . . . 175

9.4 Comparison of the pressure amplitude simulated by k-Wave and that calculated by O’Neil’s solution at the spatial resolution of 119  $\mu\text{m}$  (10.2 grid points per minimum wavelength) along the beam axis (top) and across the beam axis through the transducer focus (bottom). k-Wave underestimated the maximum focal peak by 2 MPa (7%). . . . . 176

9.5 Validation of the linear simulation against the O’Neil analytical solution as a function of spatial resolution. (a): the overall (solid black) and normalised (dashed) L2 relative error between the k-Wave and O’Neil pressures along the beam axis are shown. The red line, at 10%, is the accepted uncertainty for pressure measured by hydrophones [84]. (b): the overall (solid black) and normalised (dashed) L2 relative error between the k-Wave and O’Neil pressures across the beam axis are shown. The red line, at 10%, is the accepted uncertainty for pressure measured by hydrophones [84]. (c): the normalisation factor that the k-Wave pressures are multiplied by in order to match the O’Neil pressure at peak pressures, at each spatial resolution. (d): the computational time required for each simulation at each spatial resolution. . . . . 177

9.6 Schematic diagram illustrating the calculation of the maximum treatable depth. (a): if the deepest ablated tissue volume was less than the treatment cell volume, a line was drawn between the centroids of the deepest and shallowest ablated tissue volumes. The cell-equivalent point was found by interpolating the ablated tissue volume along this line until the cell-equivalent point, where it was the same volume as the treatment cell (84 mm<sup>3</sup> for a 4 mm cell), is found. The maximum treatable depth is defined to be half a treatment cell length (5 mm for a 4 mm cell) anterior of the cell-equivalent point. (b): if the deepest ablated tissue volume was greater than the treatment cell volume, an equivalent volume ellipsoid, with the same ratio of diameter to length as the treatment cell was created and with the same volume as the deepest ablated tissue volume, is centred at the deepest ablated tissue centroid. The maximum treatable depth was defined to be half the ellipsoid length anterior of the deepest ablated tissue centroid. . . . . 184

9.7	Four representative examples of the acoustic pressure field (see colour bar, only pressure values >10% of the focal peak pressure are shown for clarity), generated with no electronic steering applied, targeting both deepest and shallowest reachable treatment cells (magenta). All acoustic simulations were performed at source acoustic power of 300 W. The acoustic field is overlaid on the patients' density maps (lighter means material is denser, bone is white and oil is black). Tumour is shown in red. The cross-section is the X-Y slice in which contains the peak focal pressure. The focal peak is offset both axially and laterally from the intended treatment cell. . . . .	187
9.8	Focal peak pressure is plotted against path length in tissue. The dotted line is a linear regression. . . . .	188
9.9	Cross-sections of the ablated tissue (cyan) when sonicating the deepest and shallowest reachable treatment cell (magenta, centered at the geometric focus of the transducer but shown in images with full cross-section for scale, 4 mm at the widest diameter and 10 mm maximum length) are shown overlaid on MR images of patient anatomy. Thermal diffusion was simulated for a total heating time of 16 seconds (standard for 4 mm cells) and a cooling time of 30 seconds. The X-Y cross-sections shown are those with the largest ablated tissue cross-sectional area. For all patients and both target positions, the ablated tissue is offset from the treatment cell as quantified in Table 9.4. The treatment cells and ablated tissue are positioned relative to the transducer (blue) in the X-Y plane as shown in the large subfigure. . . . .	190

9.10	The path length in tissue is plotted against (top) ablated tissue volume and (bottom) the offset between the transducer geometric focus and the ablated tissue centroid in the Y-axis, with positive offset meaning the centroid is closer to the transducer than the geometric focus. Dotted lines are linear regression. The Y-offset associated with G8 <sub>2</sub> shallowest target point, which was estimated from the Y-offsets of other patients and target points, is not included in the Y-offset plot. . . . .	192
9.11	Cross-sections of the treatable volumes (cyan) for each patient are shown overlaid on MR images of patient anatomy with the tumour coverage (yellow) and the tumour itself (red). The treatable tumour volume is a subset of the covered tumour volume, which is a subset of the tumour. The cross-sections shown here are from slices in the same plane as the cell-equivalent point (which had been interpolated or extrapolated from the centroids of the deepest and shallowest ablated tissue volumes). The transducer (130 mm aperture width) is shown for scale. . . . .	194
9.12	The tumour-skin distance is plotted against the percentage difference between the tumour treatability and the tumour coverage. . .	195
A.1	Representative axial slices of volunteers, obtained at the magnetic isocentre. Only the referral, steep and shallow treatment images were used in this project. . . . .	213
A.2	Representative axial slices of patients, obtained at the tumour centroid. The tumour is highlighted in red. . . . .	214
C.1	Graphs showing (a): the absolute difference between treatment and registered-Referral virtual transducer zero-position for each volunteer and treatment position tested and (b): the mean absolute difference in zero-position over all tested volunteers and treatment positions. The red line denotes the square side length with the minimum mean absolute difference, which is 15 mm. . . . .	218

# List of Tables

2.1	Summary of acoustic simulation methods. . . . .	61
2.2	Summary of thermal simulation methodology. . . . .	67
3.1	Details of treatment cells and associated safety margins with each cell. . . . .	70
3.2	Details of volunteers from whom imaging data was acquired. Positive rotation of the pelvis denotes rotation anti-clockwise from the perspective of the volunteer. . . . .	73
3.3	Details of gynaecological tumour patients from whom imaging data was acquired. Positive rotation of the pelvis denotes rotating anti-clockwise from the perspective of the volunteer. . . . .	74
4.1	Inter-observer (top) and intra-observer (bottom) uncertainty for image registration. MDCP is mean Euclidean distance between corresponding points. SD is standard deviation. . . . .	77
5.1	Manual segmentation validation results. SD is standard deviation. DSC is Dice Similarity Coefficient. MCCD is mean contour-contour distance. . . . .	90
5.2	Automatic segmentation validation results. DSC is Dice Similarity Coefficient. MCCD is mean contour-contour distance. . . . .	91
6.1	Table of measured (nominal) gel-pad thickness, and measured membrane-bowing distance for patients. . . . .	109

7.1	Spacing between candidate cell centres for each iteration of the target-first method. . . . .	132
7.2	Comparison of the percentage tumour covered required between the device-first method (DFM) and the target-first method (TFM). .	137
7.3	Comparison of computational time required by the original device-first method (DFM), and by the refined target-first method (TFM). .	137
9.1	Acoustic and thermal properties of different tissues and materials. Sources for these properties are referenced. . . . .	169
9.2	Results for the simulations involving the modification, or lack thereof, of bone sound speed. . . . .	172
9.3	Peak focal pressure and path length in tissue travelled by the acoustic beam for each patient and target treatment cell. All acoustic simulations were performed using a source acoustic power of 300 W. . . . .	186
9.4	Differences between the intended treatment cell and the ablated tissue. For context, a 4 mm treatment cell is 4 mm in diameter (across X and Z), 10 mm in length (along Y), and 84 mm <sup>3</sup> in volume. SD is standard deviation. Thermal diffusion was simulated for a total heating time of 16 seconds (standard for 4 mm cells) and a cooling time of 30 seconds. . . . .	191
9.5	Details of patient tumours. . . . .	193
9.6	Tumour treatability versus tumour coverage. . . . .	194
9.7	Clinical data on patient treatments. . . . .	196

## Chapter 1

# Introduction

### 1.1 Clinical Use of High Intensity Focused Ultrasound

The aim of high-intensity focused ultrasound (HIFU) ablation is to induce coagulative cell necrosis by increasing the temperature within a localised region to greater than 55°C, for at least 1 second [1]. For the Sonalleve<sup>®</sup> V2 MRgHIFU system, the total sonication time required for ablating a 'treatment cell' (a standard volumetric region defined in the Sonalleve control software, generated by electronically steering the transducer focus in concentric circles) ranged from 16 seconds for a 2 mm-diameter cell to 56 seconds for a 16-mm diameter cell [2]. HIFU ablation has been FDA-approved and/or CE-marked for various diseases, including uterine fibroids [3–5], metastatic bone tumour pain palliation [6, 7], prostate cancer [8–11], and essential tremor [12]. HIFU therapy of various conditions including, but not limited to, breast and liver cancer [13, 14], thyroid nodules [15], hypertension [16], Parkinson's disease and neuropathic pain palliation [17], have either been approved by regulators outside the United States or are in clinical trials. In the United Kingdom, a clinical trial for the HIFU ablation of recurrent gynaecological cancers is being conducted (NCT02714621) [18].

HIFU therapy of pelvic tumours is particularly challenging because of the depth of the tumours within the body. HIFU systems can only treat targets within the focal length constraints of their transducers, and identifying acoustic access

which is free from obstruction by acoustically opaque tissues, such as gas and bone, and from organs at risk is challenging [18]. Failure to correctly identify suitable patients for magnetic resonance-guided HIFU (MRgHIFU) therapy could deprive them of their only treatment option, while failure to identify patients who cannot be treated wastes patient time and hospital resources on screening sessions (explained below). Patients must therefore be carefully assessed prior to being accepted for treatment.

## **1.2 Current Clinical Workflow for MRgHIFU**

Currently, the clinical evaluation process relies heavily on experience and opinion. The process for MRgHIFU treatment using the Sonalleve<sup>®</sup> V2 (Profound Medical, Mississauga, Canada) is as follows. Patients are referred to the MRgHIFU clinic on the basis of diagnostic imaging or follow-up imaging after unsuccessful prior treatment [18, 19]. Such 'referral image datasets' are usually obtained with the patient lying supine. From clinical information, including the referral dataset, only patients deemed by clinical judgement to be potential candidates for HIFU treatment are invited for screening [18, 20]. At screening, patients are asked to lie on a 15 mm thick rubber absorber, representing an acoustic-coupling gel pad and designed to prevent accidental sonication. The absorber sits on top of a membrane covering the transducer, which lies in an oil-bath below the MRgHIFU couch. The patient is asked to lie in one or two 'best guess' treatment positions on the MRgHIFU couch, and are imaged with treatment conditions being mimicked as closely as possible to produce 'screening image datasets'. The 'best guess' positions are designed to maximise acoustic access to the tumour, and are identified by the treatment team using prior clinical experience and subjective interpretation of the referral image dataset. A dummy treatment is performed by clinicians on the screening image dataset in order to estimate how much of the tumour could be reached by the transducer focus, and thus of the potential for a successful treatment. Suitable patients, those for whom clinical trial eligibility criteria are fulfilled, including satisfying a minimum threshold of tumour volume that could be reached, are invited back for treatment. In a previous metastatic



bone pain palliation trial, 16 of 37 patients (43%) initially considered for treatment were found at screening not to satisfy eligibility criteria because of disease that could not be targeted, for reasons that included tumour accessibility [19]. In a pilot planning study assessing MRgHIFU for the treatment of recurrent gynaecological tumours, 9 of 20 eligible patients (45%) who underwent screening imaging were subsequently found to be untreatable because <50% tumour coverage could be achieved without risk of damage to surrounding structures [18]. These two studies suggest that, for abdominal pelvic tumours, more than 40% of patients may be overestimated from clinical judgement as being suitable for MRgHIFU at the referral stage. During the treatment session, patient imaging is acquired before sonication, during sonication and after sonication. Imaging obtained before sonication are used to plan the positions of the individual treatment cells. Low-power sonications ('test shots') are used to calibrate the position of the transducer focus with the region of observed temperature increase and to assist the clinician in identifying suitable sonication power settings. MR thermometry [21, 22] is obtained during treatment, and is used to monitor the patient during sonication to ensure that organs at risk and healthy tissue do not undergo what the clinical team judges to be excessive heating, as well as ensuring that the cell-killing (ablative) temperature ( $\geq 55^{\circ}\text{C}$  for  $\geq 1$  second) is reached within the tumour target during sonication. After sonication, the resulting thermal dose, which is calculated by the Sapareto-Dewey equation [23], is checked to ensure that the cell-killing threshold ( $\geq 240$  cumulative equivalent minutes at  $43^{\circ}\text{C}$  (CEM43) [23, 24]) is achieved within the tumour target. The thermal dose is used to confirm ablation because the time at which the tissue crosses the ablation temperature threshold of  $55^{\circ}\text{C}$  must be interpolated from low-frame-rate MR-derived temperature maps, a process which introduces uncertainties. Furthermore, the temperature is shown only in the Sonalleve<sup>®</sup> software during sonication, whereas the thermal dose is accumulated over the entire course of the MRgHIFU treatment. The patient is also imaged after sonication in order to identify changes in tissue as a result of the treatment.

## 1.3 Hypothesis

The principle hypothesis of this PhD is that a workflow could be developed, that could predict patient suitability for MRgHIFU treatment from referral imaging, and improve patient selection for MRgHIFU screening. Given the limited time involved in the PhD, aspects of the workflow were individually developed and tested. A prospective workflow for assessing patient suitability for MRgHIFU treatment from referral imaging has been proposed below.

The motivation behind this research is to reduce the subjectivity involved in making suitability judgements that results from clinical experience and personal preferences, in order to minimise the number of patients incorrectly denied treatment, and maximise the number who would benefit from a screening scan. A quantitative estimation of patient treatability could improve clinical decision-making regarding whether the patient is suitable for MRgHIFU therapy. In the literature, quantitative screening software for skin lesions reduced the clinician error rate from 37% to 13% [25]. In the long-term, with sophisticated screening methods, it may even be possible to obviate the need for a screening visit, and reducing the load on the resources of a busy clinical MR department.

## 1.4 Background

Patient screening aims to identify whether a patient satisfies eligibility criteria for a treatment or a clinical trial, such as whether  $\geq 50\%$  of the tumour can be reached [18], whereas treatment planning aims to maximise the treated tumour volume whilst minimising harm to healthy tissue and treatment time. Hence, patient screening comprised the identification of the ideal treatment position, because the proportion of tumour that could be reached (tumour coverage) and treated (tumour treatability) is expected to vary with position, and the estimation of the tumour coverage and treatability. On the other hand, treatment planning, which happens immediately before and during treatment, comprises optimisation of sonication placement, duration and power using imaging data acquired in the treatment position which is usually already decided beforehand. Methodology for

patient positioning and sonication placement, which may be used to assess how much of the tumour could be reached by HIFU, are discussed in Section 2.5. In this section, previously published methodology for patient screening is described. Though there appears to be a wide body of literature on HIFU treatment planning [26–32], there appears to be far less research published on patient screening for MRgHIFU treatment suitability.

The current clinical screening method has several downsides. Firstly, it relies on clinician experience and judgement for both patient positioning and treatment cell placement, which means that screening results may vary between different HIFU treatment clinics. Secondly, no methods are used to investigate whether sufficient energy could be delivered to treatment cells in order to ablate the tissue within. When assessing patient suitability, treatment cells are assumed by clinicians to retain their shape and size regardless of the acoustic path.

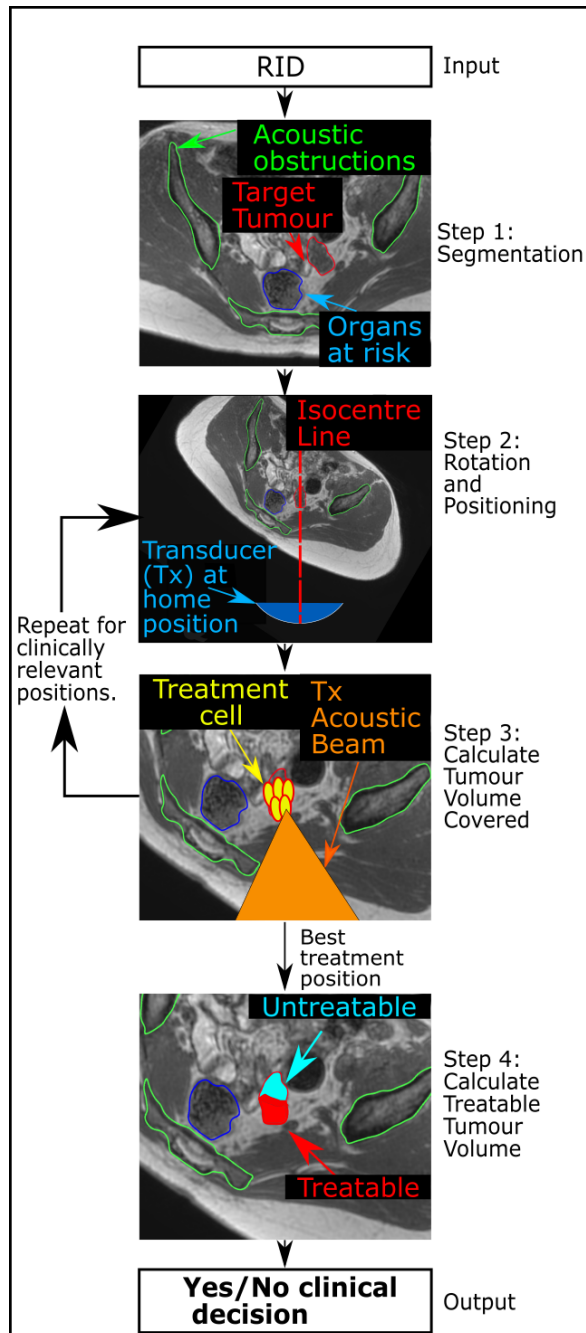
Currently, quantitative screening methods for assessing the patient suitability for the HIFU therapy of uterine fibroids [33, 34] have been developed. The authors examined a large number of fibroid treatments (422 [34] and 240 treatments [33]), and developed linear equations to quantify a score denoting relative patient treatability based on factors such as subcutaneous fat thickness and MR image intensity ratios. Thresholds for treatment success were determined by analysing receiver operating characteristic (ROC) curves. The overall positive predictive value (i.e. the probability that the patient is treatable given that they are predicted treatable) is 0.9 [33]. These models are not patient specific. However, these models assume that the fibroids are easily reachable by the HIFU focus and do not take into account the presence of organs at risk or bone, which would be a major factor when targeting pelvic region tumours. Furthermore, these approaches had been tested for targets within the uterus only and with the patient lying prone. It was not used to explore whether patients could be treated in other orientations.

## **1.5 Prospective Patient Suitability Workflow**

The general aims of the project are to develop aspects of a prospective workflow, shown in Figure 1.1, designed for assessing whether patients are suitable for

MRgHIFU therapy from their referral imaging.

The prospective patient suitability workflow, designed to predict the proportion of a patient's target tumour that could be treated using only their referral imaging, comprises three steps. In Step 1, key anatomical components such as bone, organs at risk (OARs), and the target tumour, are segmented from the input referral image dataset ('RID'). Next, Steps 2 and 3 are performed iteratively: first, the image is then rotated into a potential treatment position, which is selected from a set of promising treatment positions. Secondly, the percentage of tumour volume that can be reached by the HIFU focus (% tumour volume covered) is calculated for that potential treatment position. In Step 4, using the patient orientation that maximises the target volume covered (i.e. the ideal patient orientation), acoustic and thermal modelling are used to estimate the treatable tumour volume, defined as the percentage of the tumour volume that receives a thermal dose  $\geq 240$  CEM43, in order to facilitate a quantitative clinical decision as to whether a patient should proceed to treatment or not.



**Figure 1.1:** Proposed workflow for the assessment of patient treatability. The input is a referral image dataset (RID), in which the patient is imaged supine on a diagnostic scanner. Step 1: important anatomical tissues are identified and segmented. Step 2: an ideal treatment position is calculated (represented by the isocentre line, which goes through the predicted transducer home position and magnetic isocentre) by iterative assessment of (Step 3) the percentage of the tumour volume covered, which is the percentage of the tumour that can be reached by the HIFU focus. Step 4: when the ideal treatment angle is found, the percentage of the tumour volume treatable, which is the percentage of the tumour that has received a thermal dose  $\geq 240$  CEM43, is calculated using acoustic and thermal simulations. From this, a clinical decision on whether to pursue patient treatment with MRgHIFU is made.

## **1.6 Thesis Structure**

The thesis develops and tests methodology to deliver steps 2, 3 and 4 in the prospective patient suitability workflow. Chapter 2 describes the current state of the art in image segmentation, patient positioning, treatment cell placement, and acousto-thermal simulation methodologies, and fundamental methods in those fields that had been used in this project. Chapter 3 describes the process of volunteer and patient image acquisition, and the clinical MRgHIFU system. Chapters 4 and 5 describes the methods and validation of image registration and image segmentation respectively. Chapter 6 describes the development and testing of a methodology for predicting tumour coverage from referral imaging, by comparing the tumour coverage predicted from referral imaging to that from treatment imaging. Chapter 7 describes a refinement of the methodology used to predict tumour coverage in Chapter 6, with the aim of improving computational speed without sacrificing accuracy. Chapter 8 describes a method for identifying the ideal treatment angle, using the tumour coverage methodology developed in Chapter 7. Finally, Chapter 9 assesses patient treatability by simulating acoustic propagation through complex tissue geometry and the resulting thermal bioeffects in order to assess the treatable tumour volume.

## Chapter 2

# Background

### 2.1 Introduction

A review of the physics of HIFU as a therapeutic modality is presented here, as are some clinical drawbacks. The prospective patient workflow comprises steps involving image segmentation (Step 1), patient positioning and sonication placement (Step 2 and 3) and acousto-thermal simulation (Step 4). Literature describing the state-of-the-art for these topics is presented in this chapter.

### 2.2 Physical Principles of HIFU

HIFU is a therapeutic modality which uses focused ultrasound to induce coagulative necrosis by adiabatic heating within a localised region of tissue, such as a target tumour, whilst preserving as much healthy tissue around the target as possible. Localised ablation is achieved by focusing the ultrasound beam. This concentrates the acoustic intensity at the focus without an increase in source pressure. Hence, in homogeneous tissues, temperatures at the focus rise faster than those in the pre- and post-focal regions. HIFU treatments can be ultrasound [14] or magnetic resonance-guided [18]. Systems deliver ultrasound from a transducer placed within a human cavity [29] or noninvasively from an extracorporeal transducer [1]. This PhD project focuses on MRgHIFU therapy using an extracorporeal transducer system.

Acoustic waves are longitudinal waves that compress and rarefact the medium through which they travel. Instantaneous acoustic intensity is defined as

the product of instantaneous acoustic pressure and particle velocity. In a plane wave approximation, pressure is related to intensity using equation 2.1:

$$I = \frac{P^2}{2\rho c} \quad (2.1)$$

where  $I$  is acoustic intensity,  $P$  is the acoustic pressure amplitude,  $\rho$  is medium density and  $c$  is the medium speed of sound.

For a HIFU treatment, in order to reach the target, the acoustic waves propagate through oil, within which the HIFU transducer sits, an acoustic-coupling liquid or gel, subcutaneous fat and muscle. Acoustic waves are partially reflected and partially transmitted when propagating from one medium to another. The proportion of acoustic energy that is reflected or transmitted is dependent on the acoustic impedances  $Z$ , defined as a product of the density  $\rho$  and speed of sound  $c$  of the medium, of the two media [35]. Generally, the greater the difference between the acoustic impedance of the two media, the greater the reflected proportion of energy.

Acoustic energy is attenuated as it propagates through a medium, as a result of viscous absorption (collisions between molecules induced by the compression and rarefaction of the acoustic pressure wave) and scattering by small inhomogeneities in the medium [1]. Acoustic attenuation follows the frequency-dependent power law [1] in equation 2.2:

$$I = I_0 e^{-\alpha x} \quad (2.2)$$

where:

$$\begin{aligned} \alpha &= \alpha_0 f^\gamma \\ \alpha &= \alpha_a + \alpha_s \end{aligned} \quad (2.3)$$

where the attenuation coefficient  $\alpha$ , the absorption coefficient  $\alpha_a$ , the scattering coefficient  $\alpha_s$  and the frequency-dependent exponent  $\gamma$  are material-dependent properties,  $f$  is the acoustic frequency, and  $x$  is distance that the acoustic wave



travels through the medium.  $I$  and  $I_0$  are the acoustic intensities at distance  $x$  away from the source, and at source, respectively [1]. On average, within human soft tissues, approximately 75% of attenuation is absorption [36]. Medium sound speed, density and attenuation coefficient are known to vary with temperature [37, 38].

Density, and therefore speed of sound, increases in a compressed medium and decreases in a rarefacted medium, which distorts the original waveform, leading to nonlinear propagation. This results in a gradual "shocking" of the acoustic wave, in essence transferring energy from the fundamental to higher harmonic frequency components. The extent to which this occurs depends on the medium's nonlinearity parameter  $B/A$ , the incident pressure amplitude, and the distance travelled by the acoustic wave. As acoustic energy absorption increases with frequency, higher harmonic content results in greater attenuation and absorption of acoustic energy, resulting in more heating. Higher frequency components of the acoustic wave are diffracted less, generating smaller sidelobes and therefore energy is retained along the beam path, which may also contribute to increased heating at the focus. Nonlinear effects may therefore increase the energy deposited, the size of the generated lesion, and the maximum temperature at the focus [39].

The temperature change resulting from the absorption of acoustic energy within tissue is described by the Pennes bioheat equation [40] (equation 2.4):

$$c\rho \frac{\partial T}{\partial t} = \nabla \cdot (K \cdot \nabla T) + \rho_b c_b V (T - T_b) + Q \quad (2.4)$$

where

$$Q = \sum_{f_0}^{f_{max}} \alpha_a(f) I(f) \quad (2.5)$$

where  $T$  is temperature,  $c$  is the specific heat capacity of the medium,  $\rho$  is the density of the medium,  $K$  is the thermal conductivity,  $\rho_b$  is the blood density,  $c_b$  is the blood specific heat capacity,  $V$  is the blood perfusion rate,  $T_b$  is the arterial blood temperature, and  $Q$  is the volumetric thermal energy deposited into the

medium by absorption of acoustic energy. The absorption coefficient  $\alpha_a(f)$  is a function of frequency, as shown in equation 2.3, and  $I(f)$  represents the acoustic intensity for each harmonic frequency component. Frequencies range from the fundamental frequency  $f_0$  to the highest harmonic frequency  $f_{max}$

Proton resonance frequency shift thermometry is the current MR-based clinical standard for noninvasive in-vivo temperature monitoring in non-adipose tissue [41]. The standard deviation in measured temperature is approximately 1 °C[42]. In MRgHIFU, thermometry data is obtained more slowly than real-time (in 3-second intervals [19]), so the ablation temperature threshold of 55 °C for 1 second must be interpolated, introducing uncertainties. An alternative quantity – thermal dose – is used clinically to assess whether the target was successfully ablated. Thermal dose [23] is a concept derived from the field of clinical hyperthermia, a therapeutic modality in which the target site is heated to 40-43 °C for 1 hour or longer [43]. Because hyperthermia treatments vary in the temperature to which the target site is heated, the duration of heating, and the number of treatment fractions, comparisons between results of different hyperthermia trials were difficult. Furthermore, in the case of a failure to reach the planned target temperature at the tumour, a method for quantifying the 'dose' given to the patient resulting from the treatment attempt was needed. Thermal dose  $D$  combines the temperature at the target and the time spent at this temperature into a single quantity, using equation 2.6:

$$D = \int_0^{t_{final}} R^{T-43} dt \quad (2.6)$$

where:

$$R = \begin{cases} 0.5, & \text{if } T \geq 43^\circ\text{C} \\ 0.25, & \text{if } 37^\circ\text{C} \leq T < 43^\circ\text{C} \\ 0, & \text{otherwise} \end{cases} \quad (2.7)$$

where  $D$  is the thermal dose in units of cumulative equivalent time at 43 °C (CEM43), as defined by Sapareto and Dewey [23],  $T$  is temperature in degrees Celsius, the heating occurs from time  $t = 0$  to time  $t = t_{final}$ . A threshold of 240 cumulative equivalent minutes at 43 °C (CEM43) has been shown to exist for com-

plete tissue coagulation [24].

## 2.3 Drawbacks of HIFU

Clinical applications of HIFU, both ultrasound- and MR-guided, have been documented, demonstrating both the potential and the drawbacks of this therapeutic modality. In 2004, Wu et al. published a paper describing follow up of multi-centre treatments of solid carcinoma by extracorporeal ultrasound-guided HIFU. Target sites treated included the liver, bone, breast, kidney, pancreas and lung [14]. However, complications were noted with the 1038 patients treated. They included mild local pain within 1 week of treatment (20-30% of patients, with 5-10% patients prescribed oral analgesics), skin burns (10-20% of patients, attributed by Wu et al. to clinical inexperience), low grade fever (up to 38.5°C for 5-10% of patients, with several patients noted to experience fevers of up to 39.5°C). Adverse effects with a lower incidence rate include hepatic abscesses (6/474 liver carcinoma patients treated), bone infection (4/153 bone metastases patients treated), bowel perforation (4/153 patients, attributed to abdominal adhesion induced by pre-HIFU surgery which caused the HIFU focus to drift off-target), bone fracture (4/153 bone metastases patients treated) and nerve damage (4/153 bone metastases patients treated). In another study, published in 2009, Li et al. recorded adverse effects from 59 patients treated for liver carcinoma using ultrasound-guided HIFU [44]. They noted that all patients suffered burns along the path of the ultrasound beam, with 8 (14%) patients suffering first-degree burns, 48 (81%) second-degree burns and 3 (5%) third degree burns, and pain (52 (88%) requiring non-steroidal analgesics or no pain-relief, and the remainder requiring morphine). Out of 59 patients treated, 48 showed elevated levels of glutamic-pyruvic transaminase or glutamic oxaloacetic transaminase (up to 80 units per litre), which are biomarkers of liver damage (threshold being 40 units per litre). Ten (17%) patients experienced skin numbness post-treatment and six (10%) patients experienced mild haematuria. Damage to the gallbladder and bile duct was observed in 2 (3%) and 5 (8%) patients respectively. The incidence of adverse effects was noted by Li et al. to be greater than that observed in other studies,

and was attributed to clinical inexperience and the clinical team's lack of access to real-time temperature monitoring of the target. They further noted that the presence of acoustically reflective material in the acoustic beam path, such as bone and air within the intestinal tract, resulted in decreased therapeutic efficacy and increased risk to healthy tissue. According to these authors, proper case selection was the most important factor for successful treatment of liver carcinoma cases with HIFU. These results suggest that patient screening is important in identifying patients who would receive the greatest clinical benefit at the lowest risk for healthy tissue damage, and that treatment planning is important for minimising both treatment time (up to 8 hours for a 40 x 40 x 40 mm lesion [44]) and harm to healthy tissue.

## **2.4 Image Segmentation**

Image segmentation was required for this PhD in order to identify and delineate tissues of interest, such as bone, organs at risk and the target tumour, within the input images. The target tumour needs to be identified in order to assess whether MRgHIFU can treat it. Organs at risk (OARs) need to be identified, since the ultrasound may cause damage to them. Acoustic obstructions, such as bone and air, need to be identified since they are opaque to ultrasound, and their presence could affect whether the target tumour can be reached by the focused beam.

In the literature, manual segmentation is described as the gold standard for validation of image segmentation, despite being subject to intra- and interobserver variability [45–51]. The degree of intra- and interobserver variability is routinely quantified by examining the Dice Similarity Coefficient (DSC), a measure of volume agreement, between segments extracted at different times by the same observer and by different observers, respectively [45, 46, 48]. In a study of prostate segmentation, three different observers manually segmented the prostate, a soft tissue target surrounded by muscle and adipose tissue, for 50 MR datasets. The interobserver variability was quantified by the median DSC value between all observers, which was 0.87 [45]. In another study on the segmentation of femora and innominate bones, two observers manually segmented

eight patient datasets, of which four (randomly chosen) had been re-segmented two weeks later by the same observers. The mean intraobserver DSC values for femora and innominate bones was 0.970 and 0.963 respectively, and the mean inter-observer DSC values for femora and innominate bones was 0.971 and 0.965 respectively [46].

Automatic image segmentation methods were investigated briefly in this study as a possible option for reducing the time involved in manual image segmentation. Automatic segmentation methods reported in the literature include image intensity-based methods, such as thresholding and feature-space clustering. These methods take advantage of similarities in image intensity and spatial position in order to segment regions of interest from the image. Thresholding involves segmentation of an image into regions based on whether individual pixel or voxel intensities are above or below certain intensity thresholds. The thresholds themselves can be defined in many ways [52, 53]. Thresholding has been used to segment adipose tissue from whole-body MRI imaging [53], and to segment adipose tissue and muscle in the leg [50, 54]. Edge-based segmentation methods, which identify discontinuities in image intensity gradients to mark edges, can also be used [55]. The assumption of tissue homogeneity, i.e. that image pixels or voxels of the same tissue and material are spatially proximal, could be used in order to group neighbouring pixels or voxels into the same segment despite some intensity variation (which can be expected from image noise and artefacts) [55, 56]. Binary image operations have been developed [57] and used as intermediate steps after an initial estimation of the segment was obtained by another process. These include connected-components labelling [57] (a process in which sets of connected segment pixels or voxels are identified, and each is labelled as being separate), morphological dilation and erosion [58] (in which pixels or voxels are added or removed from the edge of the segments, respectively), and flood-filling [58] (in which non-segment pixels or voxels surrounded by segment pixels or voxels are detected and re-assigned as segment pixels or voxels). These image operations are designed to remove spurious parts of the segment or fill in holes in the segment, both of which arise from image intensity heterogeneity and image

noise. A published method for segmenting subcutaneous adipose tissue around the pelvis region from MR images used the aforementioned methods in combination [53]. However, both image intensity and edge-based segmentation methods are vulnerable to image artefacts, particularly image noise, partial volume effects, susceptibility artefacts and intensity level heterogeneity. These methods are therefore usually used in pre-processing steps before more advanced segmentation methods are applied [46].

More advanced image segmentation methods use prior information, usually derived from a set of manually-delineated segments, to develop a model to perform automatic segmentation. For the segmentation of bone from MR imaging, the literature suggests that atlas-based methods and statistical shape models provide the best results [55]. Using these methods, hip joints and the pelvic bones have been segmented, and when compared to manual segmentations, have agreed with a DSC greater than 0.92 [46]. In comparison, the acceptable DSC threshold was 0.8 for the segmentation of head and neck organs for radiotherapy [59]. The downside to atlas-based methods or statistical shape models is that a large number of manual segments (for example, 49 patients [45], 28 patients [46]) is required in order to generate the atlas or model. Atlas and model segmentation appears suitable for bone segmentation, because the shape of bone is consistent between different patients, so prior shape information can easily be translated into a segmentation for a specific patient. However, soft tissue organs at risk such as the bladder and rectum vary in position, size and shape, which can present challenges to atlas- or model-based segmentation. Bladder and rectal volumes have been recorded as varying by up to  $\pm 30\%$  [60]. Rectal positions have been reported to vary by up to 5 mm in the Anterior-Posterior direction, up to 3 mm Left-Right, and by less than 6 mm in the Inferior-Superior direction [61]. Atlas-based segmentation of the rectum resulted in a median DSC of 0.77 over 30 patients [62], which is below the threshold of acceptability mentioned above.

Semi-automatic methods, which attempt to balance the accuracy associated with manual segmentation with the speed associated with automatic segmenta-

tion, have been published. They include active contour and level set methods, in which an initial contour is manually drawn within, or outside, the object to be segmented, and the contour grows or shrinks, respectively, to align with the boundaries of the object. Ma et al. proposed segmenting pelvic organs at risk, including the vagina, bladder and rectum, from axial T2-weighted MR images using geometric deformable models [51]. Initial contours are manually drawn within the relevant organs. The contours then expand outwards, at a rate determined by the similarity between the image at the contour points and the mean and standard deviation of image intensity surrounded by the contour, which are assumed to be similar within a segment. This expansion is balanced against the influence of prior shape information and the smoothness of the contour curve, which prevents the contour leaking out of the tissue boundary. Results appeared encouraging, with the average symmetric absolute surface distance (a measure of the average spatial distance between the surfaces of the semi-automatically generated segment and the manually generated validation segment) of 0.38 mm for a patient dataset with a voxel resolution of  $0.69 \times 0.69 \times 5.40 \text{ mm}^3$  and 0.50 mm for a patient dataset with a voxel resolution of  $1.00 \times 1.00 \times 5.40 \text{ mm}^3$ . The disadvantage of these models is that they require careful parameter tuning to allow the contour to ignore image noise and, at the same time, prevent the contour from leaking out due to blurred image boundaries resulting from partial volume effects. Furthermore, these methods are sensitive to the initial contour. Another semi-automatic segmentation method involves contour interpolation, in which contours are generated manually on parallel, non-adjacent 2D slices of a 3D volumetric image, and for the slices between those, contours are generated automatically by interpolation. Schenk et al. proposed such a method for liver parenchyma segmentation from CT images [63]. In their method, for the interactive contouring, the user placed seed points on the boundary of the tissue to contour. After the first seed point was placed, the software calculated the minimum cost pixel-to-pixel path between the most recent seed point and the second-most recent seed point. When interactive contouring was only performed for one slice per 12 mm (corresponding to 1 in 3 slices for 4 patients, 1 in 6 slices for 1 patient), the DSC was  $\geq 0.98$

and time spent was reduced by at least 62% compared to validation segments obtained by interactive contouring of every slice. The manually contoured slices were transformed into distance maps, in which each slice pixel was assigned a number representing distance away from the contour. The distance maps were interpolated in the slices between the manually contoured slices. Contours were then derived from the interpolated distance maps. However, because of this, the choice of which slices to manually segment is important. If the boundary of the object to be segmented changes rapidly in shape and size along the axis perpendicular to the slice plane, more manually segmented slices are required so that the interpolated contours fit more accurately to the object boundary.

The body outline is in strong contrast with the surrounding extracorporeal air, and the extracorporeal air contrasts strongly with the gel-pad, oil-bath and body outline. They could therefore be segmented with intensity-based automatic image segmentation techniques, as was done by Kullberg et al [53]. Organs at risk have relatively low contrast against surrounding tissue, and the outline of bone is blurred due to partial volume effects, so a semi-automatic segmentation method based on slice interpolation was chosen in order to obviate the parameter tuning required for active contour and level set methods, whilst saving time compared to entirely manual segmentation.

## **2.5 Screening for High Intensity Focused Ultrasound**

Patient screening is used to identify whether a patient satisfies eligibility criteria for a treatment or a clinical trial. A criterion may, for example, be whether  $\geq 50\%$  of the tumour can be reached [18]. Treatment planning, on the other hand, is designed to maximise the treated tumour volume whilst minimising harm to healthy tissue and treatment time. Hence, patient screening involves the identification of the ideal treatment position, because the proportion of tumour that may be reached (tumour coverage) will vary with position, and determination of the maximum tumour volume that may be covered. On the other hand, treatment planning, which



is carried out immediately before and during treatment, involves optimisation of sonication placement, duration and power using imaging data acquired in the treatment position which has usually been decided beforehand. Similar methodologies are used by both processes, in particular the optimisation of sonication placement in order to estimate the proportion of tumour that could be reached (used in patient screening) or ablated (used in treatment planning), as discussed later in Section 2.5. The following section introduces published methodologies regarding patient positioning and sonication placement, which are relevant for the current study on patient screening. For context, the current clinical workflow for patients, including patient screening, is described in Section 1.2.

### **2.5.1 Patient Positioning**

Patient positioning is important in determining the proportion of tumour reachable by the HIFU focus [18, 19] and is therefore important in the process of patient screening. However, patient positioning does not appear to have been systematically studied in the literature. In 1996, McGough et al. published a ray-tracing method in which patients were positioned with respect to an ultrasound phased-array transducer by maximising the number of transducer elements whose output could access the geometric focus of the transducer, which was placed by clinicians at or around the tumour centroid. A transducer element was able to contribute to the focus if a ray linking it and the geometric focus was not obstructed by bone and/or air. Images from a single patient were only shown to demonstrate qualitatively the capability of their methodology. Quantitative results and a larger patient cohort to assess the quality of their positioning methodology would have improved their analysis. Nevertheless, using ray-tracing to identify the ideal patient positioning appears to have been an innovative idea. In 2016, Scherrer et al. [27, 28] proposed a clinical workflow for MRgHIFU ablation therapy, which included identification of the patient position relative to the MRgHIFU transducer. They proposed that the patient position be decided on the basis of the maximisation of an acoustic window (approximated to be a cone) under the constraint that no OARs or acoustic obstructions, such as bone or air, lay within the cone.

Whether the position with the largest acoustic window would result in the greatest tumour coverage, or whether positions with smaller acoustic windows produced the maximum tumour coverage, was not assessed. In a conference presentation dealing with the simulation of kidney ablation by HIFU [64], Abbas et al. positioned the transducer relative to the patient tissue volumes by manually identifying the position with the "shortest acoustic path within the tissue". No further details were provided. In conclusion, although patient positioning methods are mentioned in the literature, it appears that little work has been done on assessing the quality of the positioning method.

## **2.5.2 Sonication Placement**

The Sonalleve<sup>®</sup> MRgHIFU system electronically steers the transducer focus in concentric circles to generate treatment cells. This allows the ablation of larger volumes of tissue with less time [42, 65]. Several standard protocols for generating differently-sized treatment cells are built into the Sonalleve<sup>®</sup> control software, with larger treatment cells being associated with increased sonication time. Sonication placement is the process of deciding where to place these treatment cells within the anatomical target. This is necessary for patient screening and in treatment. Methods for sonication placement have been published, but similarly to patient positioning, there appears to be no publicly available comparisons between them. In 2005, Fedewa et al. proposed a method of choosing sonication sites for ultrasound-guided transrectal HIFU therapy of the prostate. The transducer was positioned within the tip of a probe that was inserted into the rectum, and could rotate around the probe axis. Sonication sites were generated according to five parameters that clinicians would choose: maximum overlap between simulated lesions at depth, minimum and maximum overlap between simulated lesions in directions transverse to the depth, the maximum angle that the transducer could rotate around the probe axis, and the maximum volume fraction of the simulated lesion that exceeded the target boundary. The lesions were generated by non-linear simulations. In 2011, Sannholm proposed a methodology for sonication site placement for the treatment of uterine fibroids by extracorporeal MR-guided

HIFU [66]. The method was based on optimisation, with the aim of minimising treatment time whilst maximising treatment volume. A treatment cell defined the ablated tissue following each individual sonication. The uterine fibroid target was discretised into a set of regularly-spaced grid points within the target, with the spacing chosen by the clinician and with each point representing a unit of volume equal to the multiplicative product of the grid spacing. If a point was surrounded by any treatment cell, it would be marked as covered and the associated unit of volume would be added to the count of total volume coverage. Because the clinical targets were uterine fibroids, no OARs or acoustic obstructions were expected between the transducer and the target, so Sannholm's work did not take them into consideration. In 2016, Scherrer et al. [28] proposed placing treatment cells in layers parallel to the MRgHIFU bed surface, within which treatment cells would be centred at the corners of an equilateral triangle tiling. No detail was provided beyond this theoretical sketch, and no results were shown. In 2018, Williamson et al. [67] proposed a treatment cell placement strategy developed from the bubble meshing algorithm [68], an approach that uses repulsive "forces" between treatment cells, and between the target boundary and the treatment cells, to distribute treatment cells within the target volume. Treatment cell orientation was determined by the presence of acoustic obstructions (in this case, the ribcage). If an acoustic obstruction was detected, a 'torque' was applied to the treatment cell, and the optimisation would rotate the treatment cell to align towards the intercostal space. All treatment cells used were 3 mm diameter by 7 mm length ellipsoids. For five test targets, Williamson et al.'s strategy resulted in an increase of  $8\pm 2\%$  (mean  $\pm$  standard deviation) in target tissue covered by the treatment cells and a  $16\pm 5$  decrease in healthy tissue exposure compared to a raster approach, where treatment cells are packed in parallel layers with their centroids in a regular grid, which Williamson et al. claim is the current clinical approach. However, the new strategy had not been tested on patient data. Furthermore, when optimising the placement of treatment cells, the bubble packing strategy did not take into account whether the treatment cells could be reached. Therefore, the strategy appears to be aimed at reducing the overlap between individual

treatment cells, and hence maximising the target volume treated for a given period of time (a goal of treatment planning), rather than identifying the maximum achievable tumour coverage (a goal of patient screening).

## **2.6 Simulation of Acoustic Propagation**

Acoustic propagation simulations allow the estimation of the acoustic energy deposition during sonication of target tissues. This is input into thermal simulations to generate temperature and thermal dose maps, which are of clinical interest. Within the human body, muscle, fat and bone have different acoustic properties [69] and are distributed without particular symmetry, which means that simplification of the simulation into two dimensions cannot be performed. Furthermore, at the acoustic powers involved in HIFU (40W to 300W source acoustic power for the Sonalleve<sup>®</sup> V2), the acoustic wave propagates nonlinearly [70]. This results in increased energy deposition at the focus [70, 71]. Acoustic simulation methods that are capable of simulating nonlinear acoustic propagation from a focused, phased array transducer through a general heterogeneous medium include finite-difference time-domain (FDTD) methods [72–76], the hybrid angular spectrum (HAS) method [77, 78] with nonlinear corrections [79], and pseudo-spectral and k-space (PS) methods [80–84]. Each of these methods is reviewed in greater detail below. The pros and cons of the methods are summarised in Table 2.1. Finite-element methods have also been used in acoustic simulation [85] – however, because patient MR and CT imaging data takes the form of 3D arrays, the patient data needs to be meshed before finite element methods can be used. This introduces uncertainties, and raises the problem of interpenetrating meshes because adipose and muscular tissue are interwoven within the body. Hence, finite-element methods were not explored further.

### **2.6.1 Finite-Difference Methods**

Finite-difference time-domain (FDTD) methods approximate the spatial and temporal derivatives of a governing equation in the form of linear sums and differences between spatial and temporal points. An early FDTD method for the sim-

ulation of nonlinear acoustics was published by Lee in 1993 [86, 87]. Lee used the Khokhlov–Zabozotskaya–Kuznetsov (KZK) equation [88, 89] as the governing equation. The KZK equation assumes a directed acoustic beam source and provides inaccurate results for the pressure field more than  $20^\circ$  off-axis and for the region close (within several source radii) to the source [86]. It is therefore used for simulating weakly focused beams where the region of interest is paraxial. Lee’s code assumed a homogeneous medium, and only simulated diffraction, nonlinear propagation and thermoviscous absorption. In 1996, Cleveland, Hamilton and Blackstock [90] published an extension of Lee’s code to include arbitrary relaxation processes and layered heterogeneity in the medium. Unfortunately, the limitations of the KZK equation suggest that it cannot be used to simulate clinically used HIFU transducers, which are strongly focused (where the ratio between the transducer aperture diameter and the focal length, a.k.a the F-number, is approximately 1 [73, 77]). In 1999, Hallaj and Cleveland [72] used FDTD methods to simulate nonlinear acoustics within a thermoviscous fluid. The model equation used for acoustic propagation was the Westervelt equation [91, 92]. The source frequency was 1.0 MHz and source pressures ranged from 1 MPa to 10 MPa. In 2001, this methodology was further developed by Hallaj, Cleveland and Hynynen to simulate the thermo-acoustic lensing effect [73], caused by the variation in medium sound-speed with temperature [37]. Hallaj et al. state that they lacked data on the temperature variation of other acoustic properties, such as medium density, attenuation coefficient or nonlinearity. In their FDTD method, acoustic and thermal simulation steps were performed in alternating steps, with the acoustic simulation outputting the pressure field to derive the heat source input to the thermal simulation. This outputs the temperature field that was used to update the medium sound speed for the acoustic simulation. Thermal simulations were performed using FDTD methods with the Pennes bioheat equation [40] as the governing equation. Transducers with F-numbers ranging from 0.8 to 1.3 were tested, with the maximum transducer power simulated being 140 W. Source frequencies ranged from 1.0 to 1.5 MHz. Hallaj et al. worked on 2D, azimuthally symmetric acoustic simulations with a radial and axial spatial resolution of 0.1

mm and with a time resolution of 10 ns for the acoustic simulations, corresponding to the Courant–Friedrichs–Lewy number (CFL) of 0.16. The CFL is a quantity representative of numerical stability for time-stepping simulation schemes, and relates the maximum velocity in the simulation  $u$ , the spatial resolution  $\delta x$  and the time resolution  $\delta t$  and is given by:  $CFL = u\delta t / \delta x$ . The azimuthally symmetric nature of their simulation means that it would have limited clinical applicability given the heterogeneity and lack of symmetry in the human body. Furthermore, Hallaj et al. did not present validation data for their method. In 2009, Pinton et al. [76] developed a nonlinear, attenuating full wave ultrasound simulation. To prevent acoustic waves from reflecting from the boundaries of the simulation grid, perfectly matched layers (PMLs) were implemented on all sides of the simulation grid [93]. In the PMLs, wave behaviour was governed by a separate set of non-physical equations, which only attenuated components of the wave travelling normally to the PML boundary. Pinton et al. stated that, for three-dimensional simulations involving more than  $10^9$  spatial points and more than  $10^{13}$  spatio-temporal points, approximately 32 hours of computation on a 56-processor cluster was required. Their code was validated against experimental data obtained using a membrane hydrophone in a water bath, with pressure features that were above 0.3% of the pressure peak being visually indistinguishable between the simulation and experimental data. In a follow-up paper published in 2011, they validated their model against an experimental setup representing transcranial ablation therapy [71]. The source frequency was 1.0 MHz, and the transducer was programmed to emit a 20-cycle enveloped pulse with a time-averaged peak intensity of  $20\text{W}/\text{cm}^2$ . The acoustic simulation grid was  $896 \times 896 \times 726$  points, with spatial resolution of  $154 \mu\text{m}$  and a time resolution of 20 ns, corresponding to a CFL of 0.2. The number of time steps to be simulated was not provided by Pinton et al, but by considering the number of time steps required to travel from one end of the computational grid to the other along the beam axis, at the minimum sound-speed (1540 m/s) in the medium, the estimated number of time steps per simulation was 3240. The computational time required for the simulation was not provided by Pinton et al, but the number of spatio-temporal points

to process were similar to those mentioned in the validation study, suggesting a similar computation time of 32 hours on a 56-processor cluster.

Finite-difference methods appear unsuitable for large-scale (on the order of 10 cm) and high-frequency (on the order of 1 MHz or higher) simulations, because a high spatial resolution (at least 10 points per shortest wavelength) is required to accurately model short wavelengths [82]. This increases the memory and computational time requirements for the simulation. Due to the high spatial resolution requirement, high temporal resolution is also required for numerical stability, as represented by the CFL, which increases the number of time steps required to complete a simulation, and therefore further increases the computation time required.

## **2.6.2 Pseudo-spectral and k-space Methods**

Pseudo-spectral and k-space (PS) methods have been proposed as a faster alternative to FDTD [82]. Instead of discretising the governing equation as in FDTD, in pseudo-spectral methods the acoustic pressure at each time step is approximated as a Fourier series. This can be done using the computationally efficient Fast Fourier transform (FFT) algorithm. The FFT is also used to calculate spatial derivatives in the governing equation. Because the pressure is represented as a sum of sinusoids, the minimum spatial resolution requirement is two grid points for the shortest wavelength. The time derivative is approximated using a finite difference in pseudo-spectral time domain methods. In 'k-space' methods, the time derivative is approximated using a k-space time propagator, which allows larger time steps to be taken for similar numerical accuracy and stability when compared to using the finite difference [82, 94, 95]. According to Tabei et al., a k-space method with a spatial resolution of 3 points per shortest wavelength generates the same error as a FDTD method with a spatial resolution of 14 points per shortest wavelength [94].

PS methods have been used to simulate acoustic propagation in patients at clinically-relevant settings for HIFU ablation [64, 83]. Open source PS simulation packages include k-Wave [80, 81, 96] and openPSTD [97]. k-Wave was originally

developed as a simulation package for modelling photoacoustic waves, and has been expanded to include modelling of nonlinear acoustic propagation through heterogeneous media with power-law frequency-dependent attenuation. It has been validated against linear analytical solutions [98] and experimental data [84], and has been used in simulations of HIFU [64, 83]. openPSTD is software developed to assess acoustic sound propagation in an urban environment, and as of the date of thesis submission, only supports acoustic propagation in two dimensions within a space composed of rectangular subdomains [97]. A limitation of k-Wave is that, because of the way k-Wave models the power-law frequency-dependent attenuation within its governing equations [81, 96] the power-law attenuation exponent must be fixed across the propagation medium, even though all other acoustic parameters can vary spatially. In 2018, a simulation study by Suomi et al., on the efficacy of HIFU ablation of the kidney, modelled nonlinear acoustic propagation from an annular transducer (operating frequency 0.95 MHz) in patient kidneys [83]. Suomi et al. used a computational grid of 1200 x 1200 x 1200 points, with an isotropic spatial resolution of 185  $\mu\text{m}$ , and simulated 31876 time steps with a time resolution of 8.15 ns, corresponding to a CFL of 0.18. Each simulation required approximately 50 hours on a 400 core cluster with a memory requirement of 200 GB. In the same year, Abbas et al. presented simulations which examined the effect of tissue geometry on HIFU ablation of kidney tumours [64]. Unlike Suomi et al, Abbas et al. performed linear simulations. No reason was given as to this decision. The computational grid was 864 x 864 x 864 points, with an isotropic spatial resolution of 243.6  $\mu\text{m}$ , and 26,867 time steps were simulated with a time resolution of 9 ns. Each simulation required approximately 53 hours on a 144-core cluster to complete. In 2019, a study by Martin et al. reported validation of k-Wave against experimental data [84]. Their work involved simulating acoustic propagation from a single-element bowl transducer (aperture diameter 64 mm, focal length 98 mm, operating frequency 1.1 MHz) within a water tank, with glycerol inserts placed in the beam path to introduce medium heterogeneity. Martin et al. used a computational grid consisting of 2048 x 864 x 864 points, with an isotropic spatial resolution of 100  $\mu\text{m}$ , and simulated



11467 time steps with a time resolution of 15 ns, corresponding to a CFL of 0.23. Simulated spatial-peak temporal-peak positive and negative pressure were lower than hydrophone measurements by 11% and 10% respectively, with hydrophone measurements having estimated uncertainty of 10%. The simulated focus was displaced from the measured focus by at most 1.4 mm. The simulation required 9.5 hours to simulate on a 864-core cluster with a total memory requirement of 1.3 TB. In an effort to reduce the computational requirements of PS methods, Grisey et al. proposed a scheme in which simulations were performed in multiple layers [99, 100], with each layer increasing in spatial resolution and decreasing in field-of-view size as they approached the focus, and with the output of one layer being Fourier-upsampled and used as the input for the next layer. However, no results had been provided as to how this would improve computation time. Evidence that the multilayer methodology converged to a single-domain simulation had not been presented. Grisey et al. also did not present specific information relating to the construction of the multilayer model, such as overlap between layers, the increase in spatial resolution per layer, and how the field-of-view size shrunk between layers.

### **2.6.3 Angular Spectrum Methods**

The angular spectrum method takes an initial pressure distribution within a plane, and propagates it in the direction perpendicular to that plane. Propagation to the next plane is performed by decomposing the pressure wave distribution in the initial plane into a weighted sum of 2D plane waves with different wavevectors, multiplying each 2D plane wave by a phase term related to the expected phase shift of propagation to the next plane, and then summing the modified 2D plane waves together to generate the pressure distribution in that next plane [101]. The decomposition and summation process can be performed quickly using the computationally efficient FFT algorithm. The angular spectrum method requires a method of generating the initial pressure distribution. The Rayleigh-Sommerfeld integral [101] has been suggested for this purpose [77], as has the fast nearfield method [102–104]. The disadvantage of the angular spectrum method is that

it assumes linear propagation, and that it can only simulate propagation through layered media in which the interfaces between different materials are all parallel to the initial plane. In 2016, Wang and Zhou [78] validated a modified angular spectrum method, capable of simulating acoustic attenuation and nonlinearity through layered media, against experimental data obtained by needle hydrophone in a water tank. The validation simulations were conducted with a concave source (operating frequency 2.08 MHz, source pressure 480 kPa) emitting into a two-layer medium (3 cm water, 5 cm tissue). The total computational grid was 768 x 768 x 212 points, with the transverse spatial resolution being 238  $\mu\text{m}$ , axial spatial resolution in the water layer being 356  $\mu\text{m}$  and that in the tissue layer being 390  $\mu\text{m}$ . The simulated peak negative pressure was greater than the measured value by 7%. Simulated peak positive pressures for the first three harmonics visually matched measured values. Computation time was approximately 50 minutes on a 4-core desktop PC.

The hybrid angular spectrum (HAS) method was developed by Vyas and Christensen in 2012 [77]. It incorporates modifications to simulate acoustic propagation through heterogeneous and attenuating media. In their implementation, only the first reflection from medium interfaces was simulated. Vyas and Christensen compared the HAS method and the FDTD method in a simulation of linear acoustic propagation through a 3D breast model. For both methods, the pressure distribution was initialised using a Rayleigh-Sommerfeld integral, with the source frequency being 1 MHz. The computing grid was 301 x 300 x 300 points, with isotropic spatial resolution of 150  $\mu\text{m}$ . The HAS method required 46 seconds to complete the simulation on a laptop, whilst the FDTD method required 467 minutes. However, Vyas and Christensen's implementation was only capable of simulating linear acoustic propagation. In 2015, Schwenke et al. [26] implemented a variant of the HAS method, with the ability to simulate nonlinear acoustic propagation [79], as a component of their treatment planning software. They claim to have validated it, but have not shown data. In 2018, Johnson et al. performed a validation study of Vyas and Christensen's linear HAS method [105]. Acoustic propagation and the thermal bioeffects from a MRgHIFU trans-

ducer sonicating a homogenous tissue phantom were modelled, with acoustic simulation performed using the HAS method and thermal simulations performed using a previously-validated FDTD method. Simulation results were then compared to MR thermometry data. The initial pressure distribution from the source transducer, a 256-element transducer with operating frequency of 0.94 MHz and acoustic powers of 6.3 and 7.9 W, was generated from the Rayleigh-Sommerfeld integral. It appears that Johnson et al. believed the effects of nonlinearity were negligible at the acoustic powers used. The simulation grid was 647 x 343 x 280 points, with an isotropic spatial resolution of 250  $\mu\text{m}$ . Johnson et al. state that the average computation time for a single simulation was 32 seconds. However, no description of the computing system they used was provided.

Although the angular spectrum method and its variants are computationally efficient in comparison to the FDTD and PS methods, they have several limitations. Firstly, because angular spectrum methods do not involve time stepping, they implicitly simulate acoustic propagation at pseudo-steady state conditions. This does not appear to be a problem for the simulation of continuous waves, as used in clinical HIFU ablation [65, 106]. Secondly, because the initial pressure distribution must be planar, transducers without a principal propagation direction, such as hemispherical transducers for transcranial HIFU, cannot be simulated using these methods. Thirdly, the HAS method as described by Vyas and Christensen, which appears the most clinically relevant [26, 105, 107], only incorporates transmission and the first backwards reflection. It appears possible for the code to be extended to allow for more reflections. Fourthly, for the HAS method as implemented by Vyas and Christensen, mirror sources (which arise from the inherent periodicity implied in the Fourier transform) could interfere with the wave propagation, resulting in ripple artefacts [108]. To reduce this interference, the transverse field-of-view can be made larger so that the boundaries are further away [78], absorbing boundary conditions could be applied to the transverse grid boundaries [109], or by removing the decomposed plane waves with wavevectors that are angled away from the propagation direction past a depth-dependent critical angle when propagating from one plane to the next [108].

Open-source software that use the angular spectrum method include Aber-sim [110], which can simulate nonlinear acoustic propagation through a homogeneous medium, FOCUS [111], which can simulate nonlinear acoustic propagation through layered media with all interfaces parallel to the source plane, and CREANUIS [112, 113], which was developed for ultrasound harmonic imaging and simulates nonlinear acoustic propagation through a medium which is homogeneous in sound-speed, density and attenuation coefficient but fully heterogeneous in nonlinearity parameter.

#### **2.6.4 Summary of Acoustic Simulation Methods**

A summary of the acoustic simulation methods discussed here is given in Table 2.1. A pseudo-spectral and k-space simulation package (k-Wave [80]) was selected for use in simulating acoustic propagation for this project. The FDTD and nonlinear HAS codes capable of simulating acoustic propagation through nonlinear, heterogeneous media are not publicly available. Furthermore, k-Wave had been validated [84, 98] and applied to simulations of HIFU ablation in patient data [64, 83].

**Table 2.1:** Summary of acoustic simulation methods.

<b>Methods</b>	<b>Finite-Difference (FDTD)</b>	<b>Time Domain</b>	<b>Pseudo-spectral and k-space</b>	<b>Angular Spectrum</b>
<b>Recommended Spatial Resolution</b>	8-10 points per shortest wavelength [114]		At least two points per shortest wavelength at the highest harmonic frequency [82, 95, 114]	4 points per shortest wavelength in the axial direction [78], 8 points per shortest wavelength (at the fundamental frequency) transverse [115]; for the hybrid angular spectrum method, Vyas and Christensen used 10 points per wavelength [77] in validation
<b>Advantages</b>	Can easily incorporate physical phenomena such as acoustic property variation with temperature [73]		For the same accuracy as FDTD simulation, the spatial resolution (and computation time) can be reduced [94]	Computationally rapid compared to FDTD (faster by up to 2 orders of magnitude compared to FDTD [77])
<b>Disadvantages</b>	Extremely high computational requirements [76]		k-Wave requires the power-law attenuation exponent be fixed [96]	Only directional propagation can be simulated. The hybrid angular spectrum only simulates the first reflection from medium interfaces [77]. No time stepping means that effects of dynamically-varying source pressure with time cannot be simulated. Relatively new method without much validation work published.
<b>Software Availability</b>	Open source KZK-based methods are available [86, 87, 90, 116], but only function for weakly focused acoustic beams and layered media		k-Wave [80] and openPSTD [97]	Only angular spectrum methods are available (CREANUIS [112], FOCUS [111], Abersim [110]), and these are limited in functionality.

## 2.7 Simulation of Thermal Bioeffects

Thermal simulation is used to predict the temperature change and thermal dose delivered to a target when acoustic energy is deposited. Since the human body is heterogeneous, without underlying symmetries, the simulation methodology must be capable of simulating heat transfer through a heterogeneous medium. The acoustic energy deposition is also without underlying symmetry, so the simulation methodology must be capable of simulating heat transfer from a general heat source. Much of the published literature appears to use the Pennes bio-heat equation [40], discussed above in Section 2.2, as the governing equation. Because thermal diffusion occurs on longer time scales than that for acoustic propagation [73], its simulations tend to use coarser spatial and/or temporal resolutions (e.g. same isotropic spatial resolution of 0.1 mm but increased time step from 10 ns (acoustic) to 0.1 seconds (thermal) [73]). For thermal simulations, time-steps are on the order of 0.1 seconds [73, 105]. Thermal simulation times are on the order of seconds [105] to minutes [64] on workstation hardware.

Because patient MR and CT image data is obtained in the form of 3D metrics, thermal simulation methodologies that can be directly applied to gridded data were explored. These included finite-difference time-domain [64, 71, 78, 83, 105, 117–120], pseudo-spectral and k-space [80, 121], and Monte Carlo methods [122–124]. Finite element methods have also been used for simulating heat transfer [85, 125], but as mentioned before, they require meshing of the medium which contributes to uncertainty and introduces the problem of interpenetrating meshes. In the literature, descriptions of the thermal simulations appear to be sparse compared to that for acoustic simulations.

### 2.7.1 Finite-Difference Methods

Garnier et al. developed an FDTD method for simulating tissue temperature evolution with time [118] using the bioheat equation as the governing equation. Temperatures were not allowed to exceed 100°C and the perfusion rate of tissue was set to zero upon necrosis (after exposure to  $\geq 340$  CEM43 thermal dose). The increase in tissue attenuation coefficient with temperature was incorporated into

the thermal simulation. The peak difference between the measured and simulated temperatures was approximately 5°C, with thermocouple uncertainty being up to  $\pm 9^\circ\text{C}$ . Garnier et al. used a time step of 0.02 seconds (a total of 1000 time steps were simulated for a 20 second-long exposure), but did not mention the spatial grid size, resolution or the computational time required for the FDTD bioheat simulation. In 2010, Dillenseger and Esneault used this FDTD simulation methodology to validate a separate pseudo-spectral time-domain method for computing the bioheat equation [121]. For a grid size of 256 x 256 x 64 points, with a spatial resolution of 0.4 mm, and a total number of 200 time steps being simulated with a time resolution of 0.1 seconds (corresponding to a total simulation duration of 20 seconds), the computational time for the FDTD simulation was 62 seconds on a 4-core desktop PC.

In 2018, Johnson et al. [105] used a three-dimensional FDTD method to perform thermal simulations as part of the validation for an acoustic simulation method in tissue phantoms. The FDTD method had previously been implemented as part of a method for speeding up MR thermometry [126]. In that process, validation had been performed by quantifying the root-mean-square error between the simulated temperature and the measured temperature (measured using proton resonance frequency shift thermometry) within a heated region of interest, with the resulting root-mean-square error being 0.8°C. Johnson et al. performed simulations with the following settings: the computational grid and spatial resolution was the same as those used in the acoustic simulation: 647 x 343 x 280 grid points, with a spatial resolution of 250  $\mu\text{m}$ . A total of 227 time steps was simulated with a time resolution of 0.08 sec (corresponding to a total simulation duration of 18.16 seconds). Blood perfusion was assumed to be negligible. The average computing time required for thermal simulations was 35.3 sec. However, no description of the computing system was provided by Johnson et al.

## **2.7.2 Pseudo-spectral and k-space Methods**

In 2010, Dillenseger and Esneault developed a pseudo-spectral time-domain method to compute the bioheat equation for a homogeneous medium [121],

where the derivative property of the Fourier transform was used to simplify the spatial derivative in the bioheat equation. Similarly to a previous simulation study performed by their group [118], they incorporated the change in tissue attenuation coefficient with temperature. Simulations were performed in a computational grid of  $256 \times 256 \times 64$  points with a spatial resolution of 0.4 mm. The code was implemented in C++ and run on a 4-core desktop PC. A total of 10 time steps, with time resolution being 2 seconds, was simulated, for a total duration of 20 seconds. Results were visually indistinguishable from the FDTD method developed by their group [118], and when compared against in-vivo animal experiment data obtained from thermocouple measurements, the peak difference between the measured and simulated temperatures was approximately  $5^{\circ}\text{C}$ , with thermocouple uncertainty being up to  $\pm 9^{\circ}\text{C}$ . The computational time required was 31 seconds compared to the 62 seconds required by an FDTD method that Dillenseger and Esneault's group implemented previously [118]. The biggest issue with Dillenseger and Esneault's method was that it was limited to simulating bioheat transfer in a homogeneous medium. Furthermore, Dillenseger and Esneault noted that the improvement of the pseudo-spectral method compared to the FDTD method was not as great as they expected, which they attribute to the computational complexity of the Fast Fourier Transform (FFT). It appears that for smaller computational grids, the extra processing required by the FFT requires means that the time saved is relatively short compared to FDTD methods. Dillenseger et al. note that computational speedup could be achieved for both FDTD and pseudo-spectral methods using GPU implementation of the software.

The k-Wave software package [80] also includes a thermal diffusion simulation component based on pseudo-spectral and k-space methods [127], with the bioheat equation as the governing equation. The divergence and gradient terms in the bioheat equation are calculated separately using Fourier transforms. The time derivative was calculated using a k-space propagator. In the k-Wave documentation [128], for a homogeneous medium, the thermal simulation agreed with an exact solution, derived using Green's functions [129], to machine precision.



### 2.7.3 Monte Carlo Methods

Monte Carlo methods set up an ensemble of simultaneous and random processes, with the end result of the processes being collected together and reported as the simulation output. Therefore, Monte Carlo methods can easily be parallelised. The only paper found on the topic was written in 2002 by Deng and Liu who developed a Monte Carlo method to simulate bioheat transfer using the bioheat equation as the governing equation [123]. They discretised the bioheat equation to extract a probability model from it. At a particular instant in time  $t = N\delta t$ , where  $\delta t$  is the time resolution and  $N$  the number of time steps, individual particles are created within the simulation domain grid and moved as governed by the extracted probability model. Each particle is only allowed at most  $N$  moves. If the particle runs out of steps, a temperature based on its initial temperature is registered; if the particle hits the boundary of the domain before running out of steps, depending on chance, it may be terminated and a temperature dependent on the boundary conditions is recorded, or it may be sent to another grid point. The temperature of a single point on the domain grid was calculated from the total tally of all particles created from that point, divided by the number of particles created from that point. Validation was performed by comparison with an analytical solution in one dimension, with the difference between the two being less than  $0.5^{\circ}\text{C}$  at any point. For a one-dimensional system consisting of 31 grid points, with 10,000 particles created and measured at a single location at time  $t = 2,000$  sec, required approximately 14 seconds of computational time. Unfortunately, neither the time step size, nor the computer system within which the simulation was performed, were mentioned. Deng and Liu mention that the nature of Monte Carlo simulation allows for easy parallelisation and therefore, easy speedup, if optimised hardware is used. Furthermore, the methodology allows the simulation of temperatures at particular spatio-temporal points of interest, without requiring the calculation of other spatio-temporal points beforehand, which could save time. However, Deng and Liu noted that computing accuracy at the domain boundaries is worse than that at the interior domain points, and that even for a

one-dimensional system, the method was time-consuming. It appears that if the desired output is the evolution of a temperature map over time, which is required for the identification of thermal dose contours, then another method may be more ideal. The authors used this method to evaluate the cooling effect of large blood vessels on temperature distributions during hyperthermia [130] and to correlate skin surface temperature and heat flux with pathophysiology [122]. Monte Carlo methods appear unpopular compared to finite-difference methods.

#### **2.7.4 Summary of Thermal Simulation Methods**

A summary of the thermal simulation methods explored is given in Table 2.2. The pseudo-spectral and k-space method was chosen as the thermal simulation method for this project. From acoustic simulation, grids on the order of 1000 x 1000 x 1000 points were expected; for such grids, pseudo-spectral methods were expected to be faster than FDTD methods. To identify where tissue ablation occurs, thermal dose must be evaluated for each grid point, suggesting that grid-based methods would be superior to Monte Carlo methods. Lastly, k-Wave is publicly available, whereas to the author's knowledge, the FDTD codes used in the literature are not.

**Table 2.2:** Summary of thermal simulation methodology.

<b>Methods</b>	<b>Finite-Difference Time Domain (FDTD)</b>	<b>Pseudo-spectral and k-space</b>	<b>Monte Carlo</b>
<b>Advantages</b>	Used widely in the literature [64, 71, 78, 83, 105, 119, 120], and therefore well understood. Available in open-source software FDTD [131]	Faster than FDTD for large simulation grids [121]. Available in open-source software k-Wave [80].	Can be parallelised easily for huge speed up. Fast if only a small number of spatio-temporal time points are to be examined [123].
<b>Disadvantages</b>	May be slower than pseudo-spectral methods for large grids [121]	May be slower than FDTD for small grids [121]. Not as common as FDTD in the literature.	Very slow if non-parallelised, and if more than several spatio-temporal points are examined [123]. Very little literature on this method.

## Chapter 3

# Clinical Studies


### 3.1 Introduction

A volunteer study was performed before imaging data from a recent clinical trial, MRgHIFU for Recurrent Gynaecological Cancer (NCT02714621), was made available. Image data from both the volunteer and the patient study were used over the course of the PhD. The purpose of the volunteer study was to provide human data similar to that which would be provided by patients. The methodology for predicting target (tumour) coverage from referral images (Chapter 6) was developed on the volunteer data. The patient data, which was obtained later, was used to test the methodology that had been developed.

### 3.2 The Sonalleve<sup>®</sup> V2 MRgHIFU System

The MRgHIFU system used at The Royal Marsden Hospital (RMH) for the clinical trial was the Sonalleve<sup>®</sup> V2 (Profound Medical, Mississauga, Canada). This contains a 256-element phased-array transducer (outer diameter 138 mm, inner diameter 44 mm, focal length 140 mm, with the 256 6-mm diameter circular transducer elements arranged pseudorandomly within this annulus). During treatment, the Sonalleve<sup>®</sup> bed replaces the imaging couch in the bore of the MR scanner. The transducer is mounted in an oil bath, filled with medicinal-grade white oil, on a robotic positioner with 3 linear and 2 rotational motion capabilities, and faces the patient through a thin (50  $\mu\text{m}$  thick) acoustically transparent membrane. The transducer's home position (posterior white cross in the left diagram in Figure 3.1)

is calibrated to always be 140 mm below the magnetic isocentre, and the undeformed membrane-to-isocentre distance is calibrated to be 72.5 mm. Acoustic coupling is achieved using a degassed-water wetted gel-pad (either 15 or 40 mm thick), which is placed on top of the membrane, and on which the patient is positioned for treatment. Due to the patient's body weight, the gel-pad compresses and the acoustic membrane bows when the patient is placed on top of the gel-pad. From its home position, the transducer can translate in 50  $\mu\text{m}$  steps by up to: 72.5 mm left or right and inferior or superior, and 34 mm towards the patient and 33 mm away. The transducer can angle up to  $10^\circ$  away from the perpendicular in the Left-Right and Inferior-Superior directions. It emits acoustic waves at a frequency of 1.22 MHz and, during therapy, outputs acoustic power ranging from 40W to 300W.



**Figure 3.1:** Schematic of the Sonalleve<sup>®</sup> V2 MRgHIFU system showing: LEFT - a subject lying on the MR bed, compressing the acoustic-coupling gel-pad and bowing the acoustic membrane which seals the oil bath. The membrane coil bows with the membrane. Ideally, the target tissue is centred directly above the transducer's home position, and the centre of the membrane/gel pad, and below the magnetic isocentre. RIGHT- a coronal view of the MRgHIFU couch showing the transducer's home position below the centre of the membrane.

The Sonalleve<sup>®</sup> transducer focus can be steered electronically to ablate standardised, approximately ellipsoidal volumes called ‘treatment cells’, resulting in ablation of large volumes of tissue whilst reducing total treatment time [42, 65]. This is done by rapidly steering the focus between particular ‘trajectory points’ that lie on the focal-plane, with larger treatment cells having more trajectory points, and taking advantage of heat diffusion to ablate tissue between the trajectory points. The precise positions of the trajectory points are derived from the configuration file. Each trajectory point is sonicated for 50 ms [42]. The transducer is switched off when the focal position is steered to another trajectory point, a process requiring <10 ms [42]. The Sonalleve compensates for the loss in focal peak intensity when electronically steering off-axis by increasing the source power (up to 20% at 8 mm lateral off-axis according to Philips internal documents). Each treatment cell has a cylindrical safety margin aligned with and centred on the cell associated with it, within which organs at risk should not be present[132], although in practice, clinician judgement is exercised. The range of treatment cells[18, 133] and their associated details are summarised in Table 3.1, with numbers extracted from the Sonalleve<sup>®</sup> configuration file.

Treatment Cell					
Diameter (mm)	16	14	12	8	4
Length (mm)	34.88	31.5	27.96	21.84	10
Cell Safety Margin					
Diameter (mm)	50	45	40	30	20
Height (mm)	80	72.5	65	50	35
Sonication Time (sec)	56	45	36	20	16

**Table 3.1:** Details of treatment cells and associated safety margins with each cell.

## **3.3 Magnetic Resonance (MR) Image Acquisition**

### **3.3.1 Common MR Scanning Protocol**

All subjects (all female, five volunteers, five patients) were scanned on a 3.0T Philips Achieva<sup>®</sup> MR scanner (Amsterdam, Netherlands), using a two-point Dixon sequence [134] (TE1/TE2 = 1.186 (out-of-phase) / 2.372 (in-phase) ms, TR = 3.62 ms, number of echoes = 2, flip angle = 10°); patients were further imaged using, amongst others, a T2w Large Field-of-View (T2wLFOV) sequence.

"Referral Image Datasets" were acquired with subjects lying supine on the standard MR bed using SENSE XL torso coils (Philips, Netherlands) wrapped around the pelvis. "Treatment Image Datasets" were acquired with subjects lying oblique supine decubitus on a gel-pad covering the acoustic window of the Sonalleve<sup>®</sup> V2 MRgHIFU bed, using two Sonalleve<sup>®</sup> coils – one built into the bed, and an external pelvic coil. All subjects were positioned by an experienced radiographer who was involved in the MRgHIFU treatments at the RMH. Each referral and treatment image dataset included in-phase, out-of-phase, water-only ('Water') and fat-only ('Fat') 3D images. Treatment angles were measured using ITK-Snap 3.6.0 software [135] (University of Pennsylvania, USA), by manually drawing a line between the axial-plane positions of the left and right ischial spines, and finding the angle between this and a horizontal line.

Volunteer images were acquired before the first patient was treated. Cohort-specific imaging information is given in Section 3.3.2 for volunteers and Section 3.3.3 for patients.

### **3.3.2 Volunteer Image Acquisition**

Five female volunteers (details in Table 3.2) were scanned with ethics approval from The Royal Marsden and the Institute of Cancer Research Committee for Clinical Research (CCR1406). Volunteer Body Mass Index (BMI), height and weight data was recorded. Uncertainties for weight, height and BMI were estimated to be  $\pm 1$  kg,  $\pm 4$  mm and 2% respectively. The volunteers did not have tumours.

In addition to the supine referral dataset described above, each volunteer was scanned in two "treatment" positions which were deemed to be plausible from experience treating patients with pelvic bone pain using MRgHIFU [18, 19]: 1) a nominally "Steep" (17-29°) treatment position and 2) a nominally "Shallow" (8-17°) treatment position. The exact angle was dependent on the volunteer's size and shape, as this affected how they could be fitted into the bore of the MR scanner. Each volunteer thus generated two treatment imaging datasets. The volunteers, wearing thin trousers, were rotated anticlockwise from the subject's perspective, with their left buttock roughly centered over the Sonalleve<sup>®</sup> acoustic window and with their right side elevated using angled foam pads placed under the Sonalleve<sup>®</sup> bed cushion. The volunteer pelvises were scanned from the L5-sacrum disc down to the inferior-most point of the ischial tuberosity in the axial direction. Fields-of-view were chosen to include the entire body outline in the axial slices. 15 mm-thick gel-pads were used to provide acoustic coupling between the skin and the Sonalleve<sup>®</sup> acoustic window. Voxel resolution for referral and treatment datasets was approximately  $0.78 \times 0.78 \times 1.50 \text{ mm}^3$ .

Lastly, volunteers were also scanned "Supine" (0-3°) to examine the effect of geometric deformation associated purely with changing from the diagnostic MR bed to the therapy bed. Treatment angles are recorded in Table 3.2. Representative examples of volunteer images are shown in Appendix A Figure A.1. Due to time limitations, no investigation was carried out with Supine treatment datasets.

### **3.3.3 Patient Image Acquisition**

After volunteer image acquisition, five patient datasets were acquired, as part of a recurrent gynaecological tumour clinical trial (NCT02714621, REC: 15/WM/0470)[18]. Patient data is recorded in Table 3.3. Weight data had been collected as part of the clinical trial data (estimated uncertainty  $\pm 1 \text{ kg}$ ), but height data (and therefore BMI) had not. To generate treatment imaging datasets, patients were oriented into a clinically judged treatment position, with the tumour as close to the magnetic isocentre as possible. Because pre-treatment diagnostic referral images were not available, the earliest (Day-7) follow-up supine images



Details of Volunteers					
Volunteer	1	2	3	4	5
Age (years)	28	44	29	27	36
Body Mass Index	20.2	26.4	23.5	23.8	20.9
Height (cm)	165	165	170	160	168
Weight (kg)	55	72	68	61	59
Pelvic Tilt from Supine (degrees)					
Steep	23	19	17	24	29
Shallow	17	12	8	13	16
Supine	1	0	0	1	3

**Table 3.2:** Details of volunteers from whom imaging data was acquired. Positive rotation of the pelvis denotes rotation anti-clockwise from the perspective of the volunteer.

were used as referral image datasets; this time point was chosen to minimise anatomical changes between the two imaging datasets. 15 mm-thick gel-pads were used for all patients, except patient G24. For patient G24, a 40 mm-thick gel-pad was manually cut out to provide a degassed-water-filled recess, into which the patient was lowered. For both referral and treatment datasets, patient Dixon image data were acquired after gadolinium contrast injection for improved contrast, and were acquired with a Field-of-View (FoV) of  $288 \times 288 \times 133$  voxels and voxel size  $0.87 \times 0.87 \times 1.50$  mm<sup>3</sup>. Treatment angles were measured from treatment datasets as described in Section 3.3.2. Representative examples of patient images are shown in Appendix A Figure A.2.

As part of a separate study, tumours were segmented from patient T2-weighted Large Field-of-View (T2wLFOV) datasets (TE = 90 ms, TR = 3620.4 ms, number of echoes = 16, flip angle = 90°, FoV  $672 \times 672 \times 40$  voxels, voxel size  $0.45 \times 0.45 \times 4.5$  mm<sup>3</sup>) obtained immediately pre-treatment. The clinical purpose of this sequence was to provide a wide overview of the anatomy, allowing clinicians to detect issues that could affect the quality of the treatment. Segmentation was performed by an experienced radiographer, who positioned the subjects, using in-house software (Adept v0.2, The Institute of Cancer Research,

UK) [18] (see Section 5.2.6). These segments were used to define the target tumour volume for each patient.

Details of Patients					
Patient	G3	G8 <sub>2</sub>	G24	G27	G29
Age (years)	74	59	64	53	72
Weight(kg)	61	61	42	76	57
Treatment Angle (degrees)	9	24	-6	33	16

**Table 3.3:** Details of gynaecological tumour patients from whom imaging data was acquired. Positive rotation of the pelvis denotes rotating anti-clockwise from the perspective of the volunteer.

Since the volunteer study began before the first patient was treated, the range of patient treatment angles was not known. The upper and lower bounds of the patient treatment angles were found to exceed the upper and lower bounds of the volunteer 'treatment' angles by 4 and 2 degrees, respectively. Patient mean weight (mean: 59 kg, with standard deviation: 12 kg, henceforth formatted  $59 \pm 12$  kg) matched average volunteer weight ( $63 \pm 7$  kg), but because height information was not captured for patients, the similarity in height and BMI between patients and volunteers was unknown.

## Chapter 4

# Image Registration

### 4.1 Introduction

Image registration is the process of aligning two image datasets, to allow for comparison between them. In this study, registration involved rotation of the referral dataset into the orientation used in the treatment dataset, thus allowing the target coverage predicted from the registered-referral dataset to be compared with that calculated from the treatment dataset. Registration was also used to align the tumour segments (see Section 5.2.6), which had been segmented from T2wLFOV datasets, with the Dixon image datasets that later underwent further processing. In order to reduce time spent in manual segmentation, pelvic bones segmented from volunteer referral images were also aligned with treatment datasets. In light of the limited volunteer and patient data available here, a manual image registration method was selected in order to ensure good registration quality.

### 4.2 Method

Each subject's referral image dataset was registered to the corresponding treatment image dataset by aligning 10 or more pelvic bone landmark points. These were manually placed throughout the pelvis using Horos v2.4.0, v3.0.1 (open-source, <https://horosproject.org/>) [136] and OsiriX Lite 10.0.4 (Pixmeo, Switzerland) [137]. The Standard Operating Procedure is attached in Appendix B. Horos and OsiriX calculated the required affine transformation [138, 139] and applied it to the referral dataset to create a registered-referral dataset. The registered-

referral dataset was ensured to be in the same coordinate system as the treatment dataset.

For patients, tumours had been segmented from T2wLFOV datasets (see Section 5.2.6). Registration from T2wLFOV datasets to the corresponding Dixon datasets was also performed to align tumour segments to Dixon imaging for referral and treatment datasets. Tumour segments were then further transformed to align with registered-referral datasets.

An external python environment (Python 3.6.2, NumPy 1.13.1 [140, 141], SciPy 1.3.1 [142], SimpleITK 1.0.1 [143]) was created and code was written for extracting data from Horos/OsiriX [144] transforming it into MATLAB format. The code was written by Dr Matthew Blackledge.

#### **4.2.1 Evaluation of Uncertainties in Image Registration**

As there is no ground-truth data available to determine the registration accuracy, it was assumed to be linked to the observer's ability to repeatedly place registration points on identical features in paired referral and treatment images. The Mean Euclidean Distance between Corresponding Points (MDCP) was calculated for paired, post-registration landmark points in the registered-referral and treatment datasets to represent the registration uncertainty [145].

It was assumed that the quality of the registration of volunteer data is indicative of the quality of the registration of patient data. To examine inter-observer variability, three observers (all medical physicists) each registered volunteer V1's referral dataset to its Steep treatment dataset three times, with at least a three day interval between repeats. To examine intra-observer variability, one observer (the thesis author) registered three volunteer referral datasets to their corresponding Steep treatment dataset, repeating each registration thrice with at least a three day interval between repeats. Registration validation used the Steep treatment dataset as the fixed image dataset because, as the most steeply angled, it was assumed to vary the most from the supine referral dataset, and hence to be the most difficult treatment dataset to register the referral dataset to.

## 4.3 Results

### 4.3.1 Evaluation of Uncertainties in Image Registration

Inter-observer uncertainty in image registration for volunteer V1			
MDCP	Observer 1	Observer 2	Observer 3
Repeat 1 (mm)	1.00	1.27	1.40
Repeat 2 (mm)	1.43	1.41	1.22
Repeat 3 (mm)	1.32	0.95	0.98
Mean±SD (mm)	1.25±0.18	1.21±0.19	1.20±0.17
Intra-observer uncertainty in image registration			
MDCP	Volunteer 1	Volunteer 2	Volunteer 3
Repeat 1 (mm)	1.00	1.55	1.18
Repeat 2 (mm)	1.43	1.37	1.29
Repeat 3 (mm)	1.32	1.28	0.98
Mean±SD (mm)	1.25±0.18	1.40±0.11	1.15±0.13

**Table 4.1:** Inter-observer (top) and intra-observer (bottom) uncertainty for image registration. MDCP is mean Euclidean distance between corresponding points. SD is standard deviation.

Results for the assessment of image registration uncertainty are recorded in Table 4.1. Between three observers, the mean distance between corresponding points for the referral imaging dataset for one volunteer, registered to one of their treatment imaging datasets, was on average  $1.2 \pm 0.2$  mm. For one observer, the mean distance between corresponding points for the referral imaging dataset for three volunteers registered to their Steep treatment imaging datasets was on average  $1.3 \pm 0.2$  mm. Single-factor ANOVA showed no statistically significant difference in registration uncertainty between observers ( $p > 0.95$ ), or in registration uncertainty between different Volunteers for one observer ( $p > 0.28$ ). For comparison, a dataset voxel was approximately  $0.75 \times 0.75 \times 1.50$  mm<sup>3</sup>.

## 4.4 Discussion and Conclusion

Mean post-registration misalignment was found to be less than the axial slice width of the Dixon image datasets, in line with results from the literature [146]. Hence, the introduced registration methodology was used throughout this study.

## Chapter 5

# Image Segmentation

### 5.1 Introduction

Image segmentation was performed in order to extract regions of interest from the image datasets. These included important tissues, such as bony anatomy, which act as acoustic obstructions; organs at risk (OARs), which would be damaged if ultrasound was transmitted through them, and patient tumours, which are the ablation target. Extracorporeal air, an acoustic obstruction, is also segmented. The presence of OARs and acoustic obstructions in the beam path prevents safe and effective sonication of the target, and hence is important in determining the extent of tumour coverage (Chapter 6). The body outline was segmented to assist with other segmentation processes, such as that of extracorporeal air, as well as identification of the transducer home position relative to the tissue volume of interest (see Chapters 6 and 8). Because medical image segmentation (step 1 in the workflow) is a well-studied field [55, 147], little time was spent developing new methodologies for image segmentation, in order to allow focus on other steps in the prospective patient workflow.

OARs were segmented manually due to low contrast with surrounding tissue. Bones were manually segmented because partial volume effects reduced the contrast between cortical bone and surrounding tissue. Extracorporeal air and the body outline were automatically segmented. Segments take the form of a three-dimensional binary array of voxels, the size of which matches the size of the image dataset from which the segments were extracted. A binary array is one

where each voxel in the array can only have one of two values: '1' or '0'. For all segments, the '1' voxels represent the segmented tissue or material, whilst the '0' voxels represent everything else.

## 5.2 Method

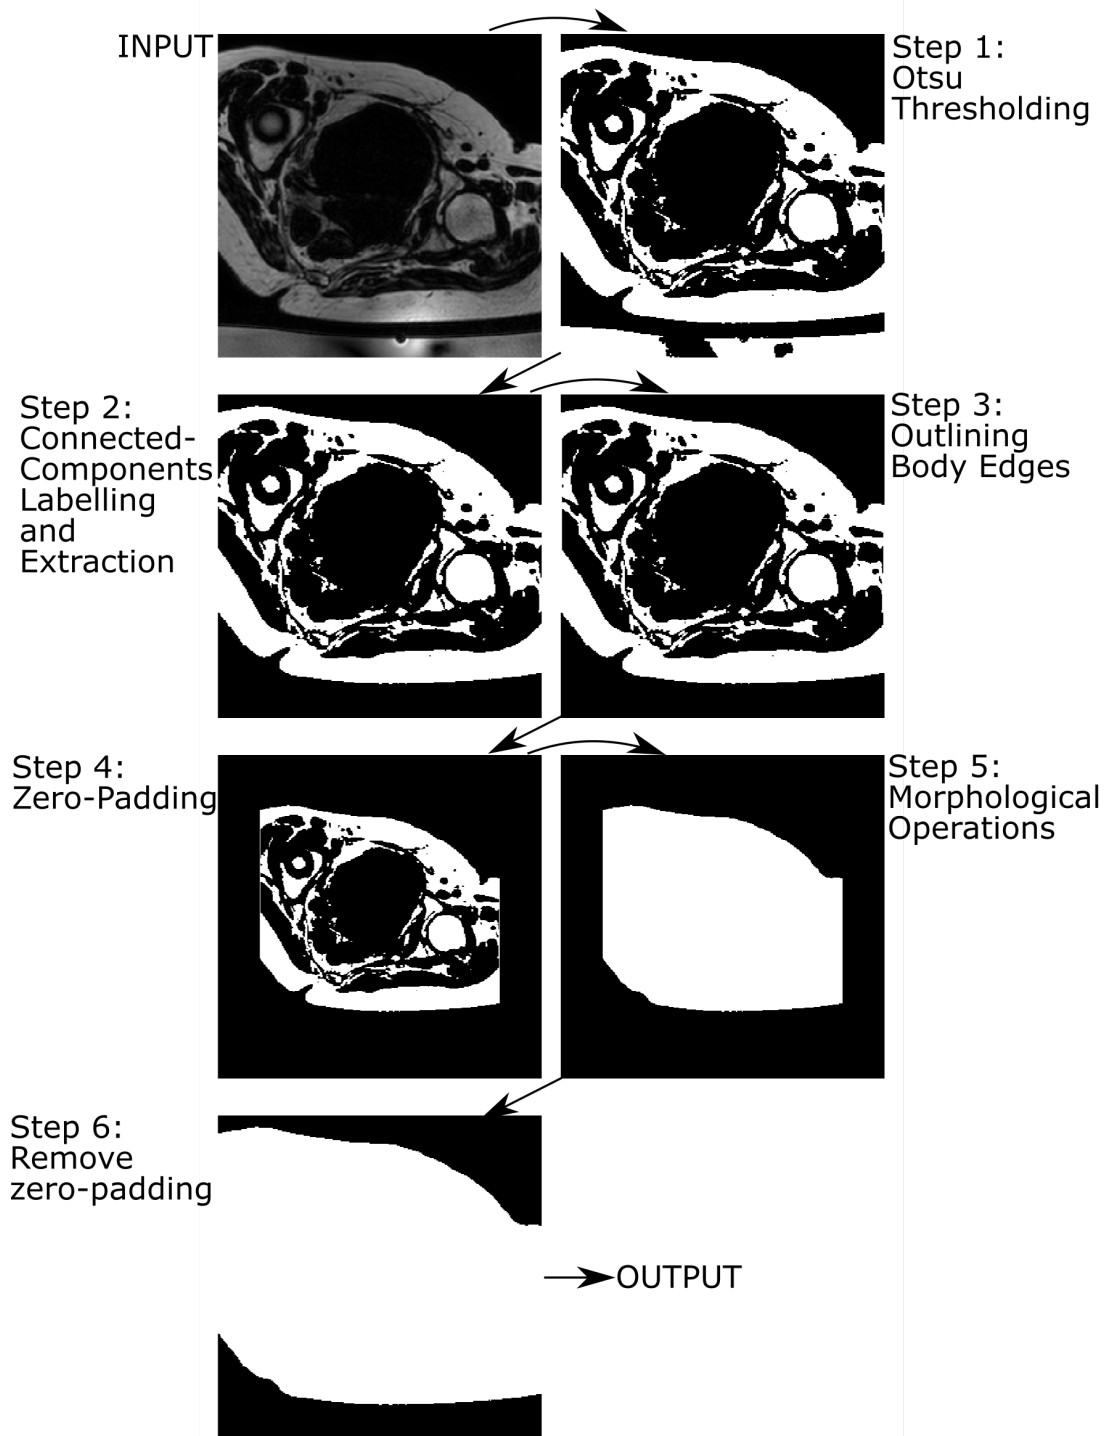
Manual segmentation of bony anatomy and OARs was performed by semi-automatic contouring (using contour interpolation between manually contoured slices) of in-phase Dixon image datasets using Horos, OsiriX and 3D Slicer v4.10.2 (open-source, <https://www.slicer.org>, Harvard University, Boston, USA) [148][149]. The software was chosen since the scripts required to extract contour data from them already existed within the ICR. Contour data was extracted from Horos and OsiriX in DICOM RTStruct format, and read into MATLAB, via Python 2.7 [144] and MATLAB scripts written at the ICR by Matthew Blackledge, Matthew Orton and Simon Doran. Contour data was extracted from Slicer3D using the open-source SlicerRT package[150]. Automatic segmentation of extracorporeal air and the body outline was performed as described below.

### 5.2.1 Body Outline

In order to delineate the body outline, and to separate the tissue volume from the gel-pad and oil bath in the treatment datasets, the body outline was segmented automatically from the registered-referral and treatment datasets. This was accomplished as described in Figure 5.1. In order to obtain the body outline from treatment datasets, the input image used was the Fat dataset (see Section 3.3.1); from registered-referral datasets, the input image was the in-phase dataset. In step 1, the input image was Otsu thresholded [52] using an 512-bin intensity histogram. Otsu thresholding is a method for automatically determining an intensity threshold by minimising a weighted sum of the intensity variance below the threshold and that above the threshold, with the weights being the probability of being below and above the threshold. Voxels below the threshold were marked '0' and those above were marked '1'. Because the gel-pad is water-based, it has relatively low intensity in the Fat images. The thresholding process therefore al-

lowed removal of the gel-pad from treatment datasets. For the registered-referral datasets, there is no gel-pad or oil-bath, so the Otsu thresholding separated the tissue from air. In step 2, in order to separate body tissue voxels from other structures such as noise or oil-bath, connected-components labelling [57] was performed. This process identified regions of '1' voxels that were connected only to each other, and each region was uniquely labelled. In order to fill holes that had resulted from the thresholding, morphological flood-filling [58] of each region was performed. The largest (by number of voxels) region, which was assumed to be the body, was extracted and the remaining regions discarded. In step 3, the body region was prepared for hole-filling. Each axial slice of the extracted body region was examined to determine whether it reached to the leftmost and rightmost column and anterior-most row. Where it did, for those columns, the topmost and bottommost extents of the body region were identified, and column voxels between them were marked as body segment voxels; for the anterior-most row, the leftmost and rightmost extent of the extracted body region was identified, and row voxels between these extents were marked as body segment voxels. This process allows the subsequent morphological operations to function properly in step 5. In step 4, in order to allow the morphological operations to function, the entire array was zero-padded by at least 50 voxels. This meant that the array was placed within a larger array, such that the original array was entirely surrounded by a border of '0' voxels at least 50 voxels thick. In step 5, for each axial slice, the body segment is morphologically dilated to smooth cracks in the thresholding, the interior is flood-filled to remove interior holes, and the segment is morphologically eroded to counter the effect of the earlier dilation. The structuring element used for dilation and erosion was a 2D disk with radius of 16 voxels. In step 6, the array is unpadded to return to the same dimensions as the body mask segment.





**Figure 5.1:** Body masking workflow, depicted for a representative volunteer treatment dataset and described in the text. Step 1: Otsu thresholding using a 512-bin histogram. Step 2: Connected components labelling and extraction of the largest component, assumed to be the body. Step 3: Anterior, left and right edges of the body segment are outlined, and Step 4: the array is zero-padded by at least 50 voxels at each side, in order to allow for processes in Step 5 to function. Step 5: morphological dilation, flood-filling and erosion. Step 6: the padding is removed, and the result is the body mask.

### **5.2.2 Organs at Risk**

Organs at risk (OARs) are healthy tissues that, if damaged, could lead to severely decreased quality of life. Hence, their presence within the beam path restricts acoustic access to the target, and thus affects the patient's treatability. Their segmentation is therefore necessary for the accurate assessment of target coverage. Relevant organs at risk in the pelvic region include the uterus, bladder, the small intestine (particularly the ileum), the colon and the rectum. Some patients had undergone pelvic exenteration prior to undergoing MRgHIFU therapy, so not all the relevant organs were present in all patients.

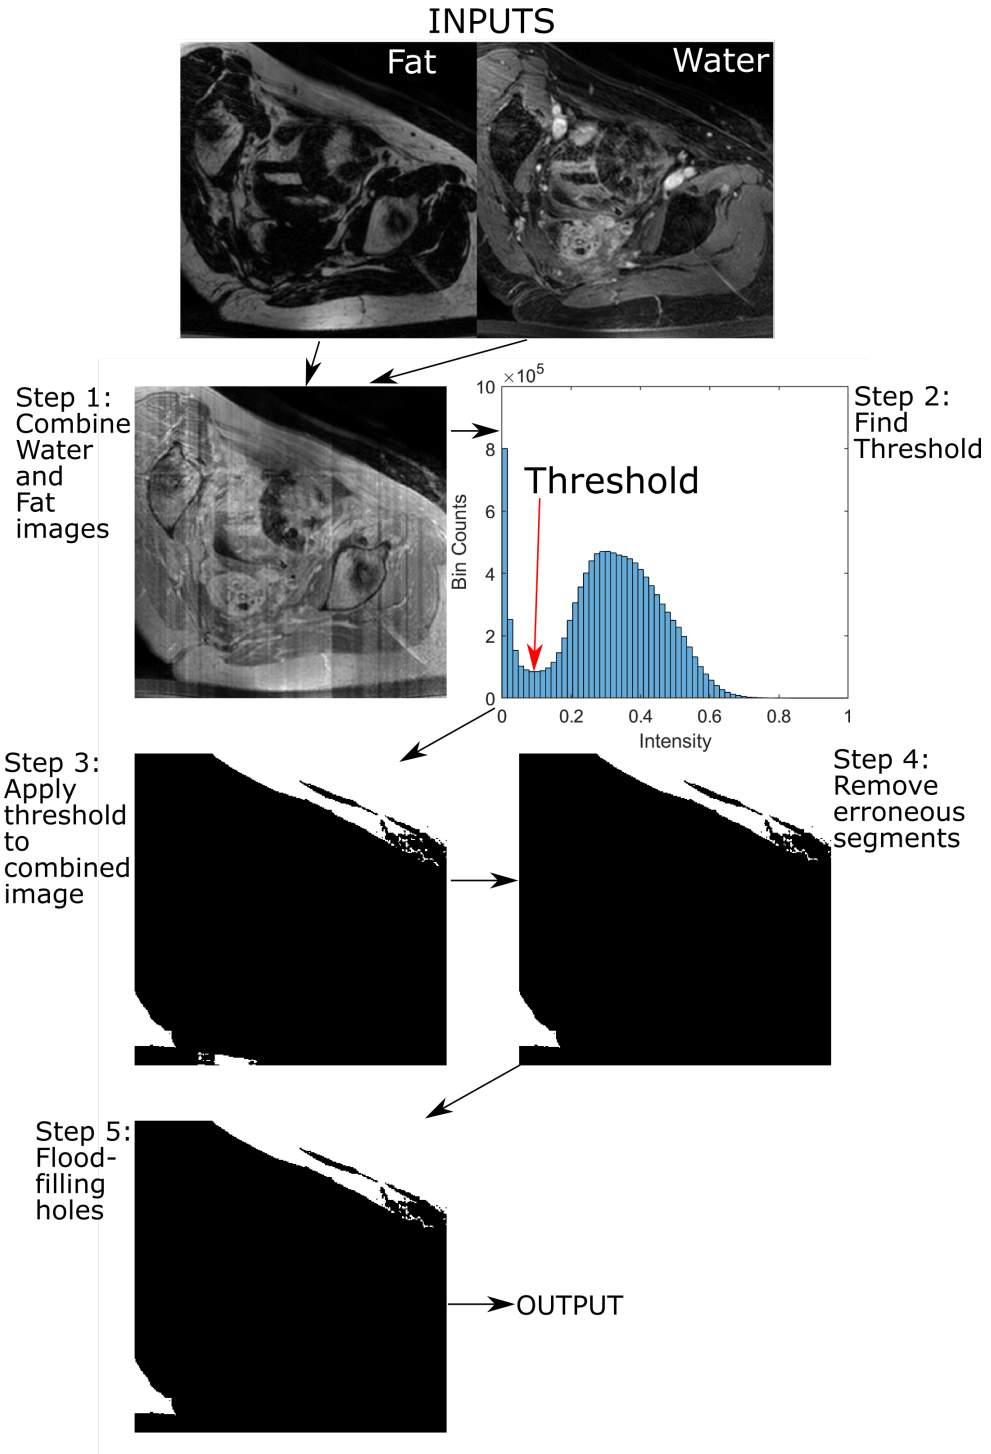
Volunteers have not undergone the clinical preparation, such as dieting and bowel-filling, that some patients will have undergone in an effort to protect OARs [18]. For this reason and for simplicity in this feasibility study, neither bowel gas nor OARs were segmented from volunteer data.

### **5.2.3 Acoustic Obstructions**

Tissues and materials that are acoustically opaque, such as bone, trapped air within the intergluteal cleft, and extracorporeal air act as acoustic obstructions. Because of their acoustic opacity, ultrasound cannot propagate through them. Segmentation of acoustic obstructions is therefore necessary for the accurate assessment of target coverage. For all subjects (volunteers and patients), bone was segmented manually from treatment and referral datasets, using the same method as used for OARs (see Section 5.2.2). In volunteers, it was assumed that the intergluteal cleft would be filled with acoustic-coupling gel in a real clinical procedure as part of clinical preparations, and hence, this was not treated as an acoustic obstruction. However, in patients, the intergluteal cleft was sometimes observed to contain air, and was therefore manually contoured and included as part of the extracorporeal air segment.

Extracorporeal air was segmented only for treatment datasets. When estimating target coverage from referral images, clinical preparations were assumed to ensure perfect acoustic coupling between the subject and the transducer. Extracorporeal air was not segmented for volunteers, because the trousers worn by

volunteers during image acquisition prevented skin-to-gel-pad acoustic coupling. Instead, volunteer acoustic coupling limits in the left-right direction were manually identified, as described in Section 6.2.2.4. For patients, extracorporeal air was segmented using an automatic segmentation method inspired by Kullberg et al. [53], which is shown in Figure 5.2 and described here. In step 1, input Fat and Water treatment dataset images were each normalised until their maximum intensity was 0.5, and then summed to increase contrast between air and the tissue, gel-pad and oil-bath. In step 2, a 64-bin intensity histogram was generated from the image, and the lowest intensity trough was identified. The intensity corresponding to the centre of that bin was taken as the threshold. This threshold was chosen because the extracorporeal air was assumed to have the lowest voxel intensity in the image (alongside cortical bone) and in preliminary experiments, this threshold performed well for multiple patients. In step 3, the image from step 1 was thresholded to extract voxels of low intensity. To remove segments representing cortical bone and bowel gas, the thresholded image was masked by the complement of the body outline, the segmentation of which is described in Section 5.2.1. In step 4, connected-components labelling (6-connectivity) split the thresholded image into disparate segments, each of which were tested for whether it was connected to the right, anterior or left sides of the dataset. Any segment that was not connected to any of these was assumed to be spurious. For example, a susceptibility artefact around the Sonalleve bed coil resulted in localised intensity variations in the image which passed the thresholding. These segments would be removed in this step. Finally, in step 5, any holes in the segments that resulted from image noise were filled using 3D morphological flood-filling, and the result was output as the extracorporeal air segment.



**Figure 5.2:** Extracorporeal air segmentation workflow. Input Fat and Water treatment datasets are individually normalised and then summed together (Step 1). Step 2: an intensity histogram is generated from the image, and the threshold is found at the lowest-intensity trough. Step 3: the image is thresholded so that voxels with intensities below the threshold are marked. Step 4: Using connected-components labelling, individual connected segments are identified. Segments that are not connected to the anterior, right or left edges of the image dataset are removed. Step 5: morphological flood-filling is used to fill holes in the remaining segments, and the output result is the air segmentation.

#### **5.2.4 Volunteers**

For volunteers, bone was manually segmented as an acoustic obstruction. Pelvic bones were segmented from referral imaging. The referral pelvic bone segments were aligned with the corresponding treatment datasets in order to reduce the burden of manual contouring. Femora were manually segmented separately from referral and treatment datasets, because of the likelihood of them being articulated differently between datasets (unlike the more rigid pelvis). To further reduce time-consuming manual delineation, the femurs were segmented only from the femur head to the axial slice at which the pelvic bones ended, and for treatment datasets, only the femur closest to the gel-pad was segmented.


#### **5.2.5 Patients**

Bone and the intergluteal cleft were manually segmented, and extracorporeal air was automatically segmented, as acoustic obstructions. All observed OARs were manually segmented. Because the treatment region was considerably smaller for patients than for volunteers, in order to reduce time-consuming manual delineation, only acoustic obstructions and OARs within the region-of-interest (the axial slices containing the tumour plus a margin of 10 axial slices in the inferior and superior directions) were segmented. The pelvic bones at the greater sciatic notch were always segmented, because the notch defines the superior edge of the sciatic foramen through which the acoustic beam is expected to sonicate the tumour. For pelvic bone in particular, the region of the greater sciatic notch was segmented regardless of whether it was within the region-of-interest, in order to enforce a constraint on the travel and angulation of a virtual transducer. For treatment datasets, the extracorporeal air was segmented manually.

#### **5.2.6 Target Tumour**

The target tumour for each patient was segmented by Sharon Giles, who carefully contoured the tumour outlines on T2wLFOV axial slices for both the referral and treatment datasets, using in-house software (Adept v0.2, The Institute of Cancer Research, UK) [18]. Image data from other sequences were referenced dur-

ing contouring. Treatment contours were used to guide segmentation in referral datasets. As the target tumour was segmented by an expert, it was assumed that the contour was accurate. Healthy volunteers had no tumours to segment.



**Figure 5.3:** Target tumour segmentation from T2wLFOV images (a representative axial slice from a patient). The green line is the manually drawn tumour contour.

The tumour segment output from Adept was represented as a set of 3D coordinate points. A MATLAB script was written to extract the coordinate data. As the T2wLFOV slice thickness was approximately 10 times the voxel in-plane side length, the interpolation involved in aligning the T2wLFOV image array with the corresponding Dixon image array was assumed to introduce large errors. The re-

referral and treatment tumour coordinates were therefore meshed using the Marching Cubes isosurface generation algorithm [151] to generate a triangulated surface mesh, and the coordinates of the surface mesh were aligned with the referral and treatment Dixon image dataset by using the registration described in Section 4.2. Once in the Dixon image coordinate space, a set of points with the same spacing as the Dixon image voxels was generated within the tumour surface meshes to represent the interior tumour volume. This was done by first obtaining an isosurface triangulated mesh using the marching cubes algorithm [151][152] with an isovalue of 0.5, chosen in order to place the isosurface midway between the centres of segment and non-segment voxels. Then, a regular grid of points with the same spacing as the Dixon image voxels was generated, and only those points that were within the mesh were kept [153]. These points were designated as the tumour coordinates. Referral tumour coordinates were further aligned with the registered-referral Dixon dataset by using the transformation derived from the referral-to-treatment registration (see Section 4.2).

### 5.2.6.1 Target Tumour Misalignment

Treatment and registered-referral tumour segments were compared to assess the misalignment of tumours which could not be explained by registration, and which was therefore attributed to body deformation and clinical preparation. To quantify the extent of the misalignment, the mean surface-to-surface distance (MSSD), which quantifies the mean distance between the two tumour surfaces and hence the mean disagreement in shape [154], was calculated as defined in equation 5.1:

$$MSSD = \frac{d(A, B) + d(B, A)}{2} \quad (5.1)$$

where

$$d(A, B) = \frac{1}{|A|} \sum_{a \in A} \min_{b \in B} \|a - b\| \quad (5.2)$$

where  $A$  and  $B$  are the two tumour surfaces which are represented by discrete points in the sets  $A = \{a_1, a_2, \dots, a_p\}$  and  $B = \{b_1, b_2, \dots, b_q\}$ ,  $\|a - b\|$  denotes the Euclidean norm between the points  $a \in A$  and  $b \in B$ , and  $|A|$  denotes the number

of points in the set  $A$ .

The Hausdorff distance, which represents the greatest disagreement between the two tumour surfaces and hence quantifies the worst-case shape difference, was also calculated using equation 5.3 [155]:

$$H = \max(h(A,B), h(B,A)) \quad (5.3)$$

where

$$h(A,B) = \max_{a \in A} \min_{b \in B} \|a - b\| \quad (5.4)$$

### 5.2.7 Evaluation of Segmentation Quality

Automatic segmentation quality for the body outline and for extracorporeal air was quantified by comparing randomly selected, automatically segmented image slices with corresponding manually segmented validation slices (Body: five slices per dataset, from three ‘steep’ treatment and two ‘shallow’ registered-referral datasets; Air: five slices per dataset from three patient treatment datasets). For body outline segmentation, each validation slice was at least 10 axial slices away from any other and no validation slice was created from slices with incomplete information (due to the registration process). In order to determine the ability of the segmentation to determine acoustic coupling between patient and transducer, only the extracorporeal air segments around a region of interest, where the body and gel-pad interfaced, were assessed. The patient treatment datasets from which validation slices were generated were visually ensured to contain some extracorporeal air in this region of interest. A patient treatment dataset with no extracorporeal air within this region of interest was also tested. The slices were positioned with one approximately at the target tumour centroid, two at the inferior and superior extents of the target tumour, and two more at most 10 slices (15 mm) away from the target tumour. It was assumed that body outline segmentation quality for volunteers would be indicative of the body outline segmentation quality for patients.

The assumption that the referral pelvic bone segments could be used to auto-



matically segment the pelvic bones in the treatment dataset was similarly tested against manual contouring performed on the treatment dataset (five slices per dataset, four datasets). The segmentation quality of the volunteer bony pelvis and femora was taken to be indicative of the segmentation quality for all manually segmented tissues. Each validation slice was at least 10 (15 mm) or 5 slices (7.5 mm) away from any others in the same treatment dataset for pelvis and femur respectively, and was at least 15 slices away from the superior-most and inferior-most slices of the treatment dataset, in order to avoid incomplete post-registration data.

For both manual and automatic segmentation, one observer manually produced all validation slices. The quality metrics that were used to quantify agreement between the segment slices and the validation slices were the Dice Similarity Coefficient (DSC) and the Mean Contour-to-Contour Distance (MCCD). The DSC quantified the area agreement between the segment slices and the validation slices, and was calculated for each validation slice. It is described by equation 5.5:

$$DSC = \frac{2T \cap V}{|T| + |V|} \quad (5.5)$$

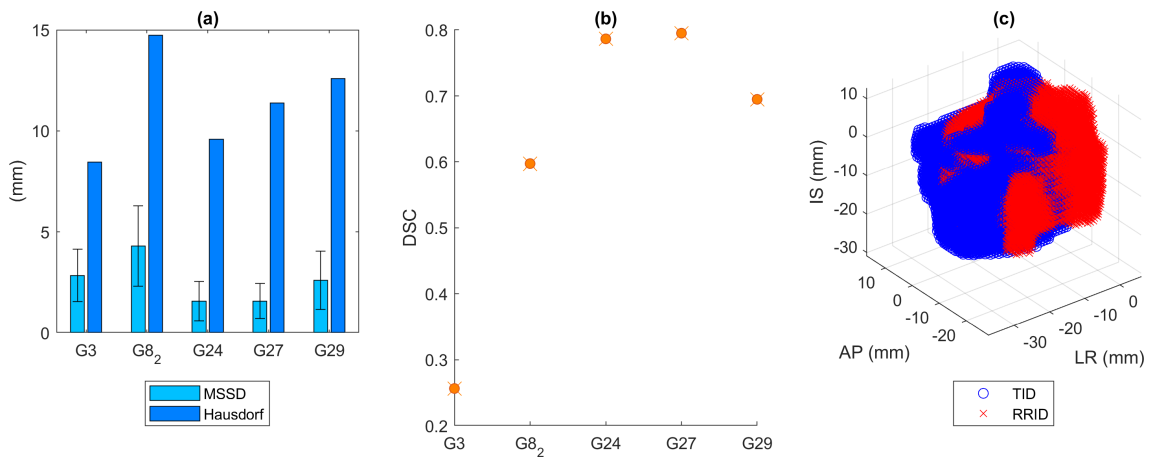
where  $T$  is the set of voxels that make up the automatically/manually generated segment in an image slice, and  $V$  is the set of voxels that make up the segment in the corresponding validation slice.

The MCCD was used to quantify the mean distance between the automatically/manually generated segment boundary in an image slice and the validation segment boundary in the corresponding validation slice. It was calculated for each validation slice as defined in equation 5.1, with the sets  $A$  and  $B$  representing points on the boundaries of the segments within the image slice and within the validation slice, respectively.

## 5.3 Results

### 5.3.1 Target Tumour Misalignment

Quantitative and visual results for the misalignment between treatment and registered-referral tumours are shown in Figure 5.4.



**Figure 5.4:** Quantitative data (a,b) and a representative image (c) of the agreement between treatment and registered-referral tumour segments. (a): Mean Surface-Surface Distance (MSSD) and Hausdorff distance is depicted. (b): Dice Similarity Coefficient. (c): The blue region represents the G29 treatment tumour surface, and the red region represents the G29 registered-referral tumour surface.

### 5.3.2 Evaluation of Segmentation Quality

Manual segmentation quality is shown in Table 5.1.


Manual Segmentation Validation		
	Pelvis	Femur
DSC (mean±SD)	0.93±0.01	0.96±0.01
MCCD (mm, mean±SD)	0.76±0.10	0.53±0.12

**Table 5.1:** Manual segmentation validation results. SD is standard deviation. DSC is Dice Similarity Coefficient. MCCD is mean contour-contour distance.

Automatic segmentation quality is shown in Table 5.2. For extracorporeal air segmentation, the edge-case (no extracorporeal air) patient treatment dataset resulted in a DSC of 1 and an undefined MCCD. A representative example of the extracorporeal air segmentation is shown in Figure 5.5.

Automatic Segmentation Validation		
	Extracorporeal Air	Body Outline
DSC(mean±SD)	0.89±0.06	0.991±0.003
MCCD(mm, mean±SD)	0.1±0.2	0.88±0.25

**Table 5.2:** Automatic segmentation validation results. DSC is Dice Similarity Coefficient. MCCD is mean contour-contour distance.



**Figure 5.5:** Automatic segmentation of the volunteer treatment image datasets results in spurious segments between the subject and the gel-pad. Green represents the manually-segmented validation slice, magenta represents the automatic segmentation, and white denotes the overlap between the two segments.

## 5.4 Discussion

### 5.4.1 Target Tumour Misalignment

Figure 5.4(a) places Hausdorff distance, a measure of the maximum disagreement between treatment and registered-referral tumour surfaces, at approximately 10 mm, while the mean MSSD is approximately 3 mm. The relatively large Hausdorff distance suggests that deformation of the tumour between treatment and referral datasets involves shape change, although the relatively small MSSD suggests that the shape change is relatively limited. The low DSC for G3 is probably due to the tumour being smaller (approximately 1000 mm<sup>3</sup>) than the others (approximately 10,000-40,000 mm<sup>3</sup>), so any mismatch in position between tumours results in a larger volume disagreement as a percentage of the total tumour volume.

### **5.4.2 Manual Segmentation Validation**

The manual segmentation quality for volunteer bones (DSC > 0.93, MCCD < 0.76 mm) demonstrates that this segmentation method is fit for purpose. In literature, a DSC of 0.8 is considered acceptable [156][157][59][158]. Given that volunteer and patient datasets were both manually segmented by the same observer using similar processes, the segmentation quality for volunteer bones was assumed to be indicative of the quality of the manual segmentations of acoustic obstructions and OARs from patient datasets as well.

### **5.4.3 Automatic Segmentation Validation**

The automatic extracorporeal air segmentation method was fit for purpose when applied to patients, even in the edge case where there was no extracorporeal air to segment. From visual examination, in volunteer datasets, the automatic air segmentation misidentified skin and gel-pad as air. This probably was due to volunteers wearing thin trousers when they were scanned, which created a thin separation layer between the patient and the gel-pad, which led to the misidentification of the trouser layer and the susceptibility artifact close to the bed coil as extracorporeal air. This could result in underestimations of accessible volume. Thus, the extracorporeal air segmentation algorithm can only be said to be fit for purpose when applied to the patient datasets, in which skin is directly coupled with the gel-pad.

## **5.5 Conclusion**

The image segmentation methodology presented here showed results that were acceptable and in line with literature. The segments generated here were used to predict target (tumour) coverage (Chapter 6) and to identify the ideal treatment position (Chapter 8).

## Chapter 6

# Predicting Target Coverage from Referral Imaging

## 6.1 Introduction

Target coverage is defined here as the proportion of the target volume that can be reached by the Sonalleve<sup>®</sup> focus. Targets include, but are not limited to, malignant tumours. Currently, target coverage is estimated manually and is subject to clinical opinion and prior experience. Here, target coverage is assumed to be a proxy for the volume of tumour that could be ablated. The quantitative prediction of target (tumour) coverage from referral imaging was a component of the prospective patient suitability workflow for MRgHIFU therapy (see Figure 1.1 Step 3).

As far as the author knows, previous literature on this topic is sparse to non-existent. Methods for MRgHIFU screening are discussed in Chapter 2. In 2018, Abbas et al. [159] described rotation of supine patient diagnostic CT datasets, before simulating ultrasound propagation through the rotated dataset. The shortest acoustic path had been identified manually, and the transducer had been positioned relative to the tissue volume such that the beam axis would align with the identified acoustic path. However, they did not quantitatively evaluate whether this was a valid approach to take. Various methods for predicting tumour coverage, or treatability, from treatment imaging have been developed as part of proposed treatment planning systems. In 2004, Seip et al. proposed an automated treat-

ment planning system for intracavity prostate ablation [160], based on a library of lesions associated with treatment parameters. This had been developed and tested by Fedewa et al. in 2005 [29]. The lesion library had been generated using nonlinear simulations in homogeneous media. A set of treatment sites had been generated within and around the target according to user-set parameters. Fedewa's method was specifically for intracavity HIFU, and assumed no acoustic obstructions or organs at risk (OARs) between the transducer and the tumour. In 2011, Sannholm [66] developed a combinatorial optimisation method of packing treatment cells within a target volume. It also assumed no OARs or acoustic obstruction between the transducer and the target. In 2018, Williamson et al. [67] proposed an alternative treatment cell placement strategy based around bubble-packing [68], and measured the resulting tumour coverage. Cells had been initialised randomly within the target volume, with their positions optimised. The presence of acoustic obstructions (in this case the ribcage) had been accounted for only by rotating the cells such that they would align towards the intercostal space.

In this chapter, a novel methodology for calculating target (tumour) coverage is described. Using this methodology, a feasibility study was performed to determine whether it is possible to predict tumour coverage accurately from referral imaging, by comparison with predictions made from subjects lying in treatment orientation. The methodology was first developed from volunteer imaging data, because patient data was unavailable at the start of this project (see Chapter 3). The novel methodology was designed to take into account the presence of acoustic obstructions, OARs, and the expected deformation of soft tissue caused by reorientation of the patient from supine to treatment orientation. Patient imaging data, which became available as a concurrent clinical trial (and this project) progressed, was used to test the accuracy of the method. Some of the research presented here has been published [161] (included in Appendix D).

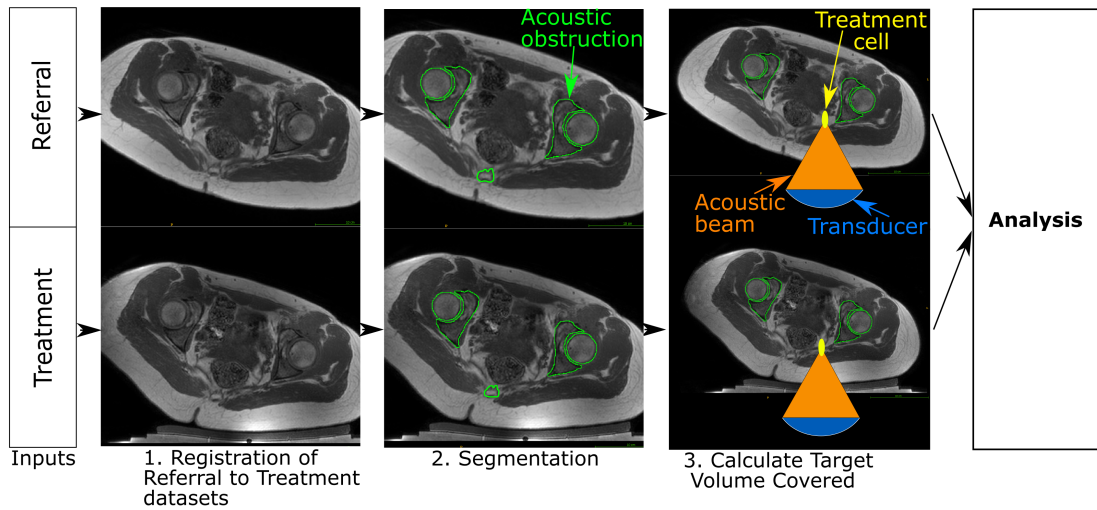
## 6.2 Method

### 6.2.1 Overview

To evaluate the methodology developed for calculating target coverage, estimations of target coverage from referral imaging and treatment imaging obtained for each subject (volunteer or patient) were compared. Here, the referral image dataset is the expected input into the prospective patient workflow and is used to predict target coverage. The treatment image dataset depicts the volunteer or patient positioned in a plausible or actual treatment position, respectively, on the MRgHIFU bed. The treatment dataset is used to calculate the ground-truth target coverage.

The workflow for this study is shown in Figure 6.1. Subject referral and treatment datasets were acquired following the procedures described in Chapter 3. As the treatment position is known from the treatment dataset, the referral dataset was oriented into the known treatment position to compare the predicted target coverage with the actual target coverage. This was achieved by an affine registration of the referral dataset to the treatment dataset (Step 1 in Figure 6.1), as described in Chapter 4. If the methodology is applied to patient screening from referral images, the treatment dataset would not be available, so the treatment orientation would have to be determined in another way. In this chapter, however, it is the accuracy of the methodology itself which is of interest. Segmentation of the acoustic obstructions and OARs (Step 2 in Figure 6.1) from both datasets was performed to identify tissues that could prevent target coverage, using the methods described in Chapter 5. This was followed by calculating target coverage (Step 3 in Figure 6.1) and comparing the results for predictions from referral datasets with those from treatment datasets, the methods for which are described in this chapter.

At the start of the project, clinical trial data were not available. The methodology was therefore developed using volunteer images, with the goal of testing it on anticipated clinical datasets. As a result of significant anatomical differences between volunteers and patients, certain adaptations were necessary. Firstly, vol-



**Figure 6.1:** Schematic of developmental methodology used in this study. The target (tumour) volume coverage estimated from a subject’s referral image dataset (top row), and that calculated from the treatment image dataset (bottom row), were compared in order to assess the accuracy of predicting target (tumour) coverage from referral datasets. Step 1: The referral dataset is rotated into the same orientation as the treatment dataset using affine registration both to allow comparison with the treatment dataset and to simulate the rotation in Step 2 of the prospective patient workflow (Figure 1.1). In step 2, segmentation of relevant tissues (e.g. bony acoustic obstructions) was performed to identify tissues that could obstruct acoustic access to the target. Step 3: target volume coverage was calculated for the rotated referral dataset and the treatment dataset, and then compared to assess the accuracy of the methodology.

unteers lacked target tumours. This could have been addressed by the creation of dummy tumours, but in the absence of an obvious method for defining the size, shape and position of dummy tumours in an unbiased and clinically relevant way, all the soft tissue in the pelvis was defined as “target tissue”. Secondly, while patients undergo dietary and physical bowel preparation prior to treatment in order to minimise the risk of damage to OARs, such as the rectum and bowels, volunteers were not required to do so. As a result, OAR tissues were ignored in the volunteer datasets. While these two limitations present challenges, they do not prevent like-for-like comparison between the referral and treatment target coverage predictions.



## 6.2.2 Target Coverage Analysis

### 6.2.2.1 Overview

Datasets from 5 volunteers, comprising pseudo-referral datasets and pseudo-treatment datasets, were available for the development of the methodology. Each volunteer had been placed in two different plausible treatment positions to generate two treatment image datasets, and in the supine position to generate a referral imaging dataset. The methodology was subsequently tested on 5 patients who had undergone ablative MRgHIFU treatment for recurrent gynaecological tumours. Volunteer and patient images were acquired as described in Section 3.3. As seen in Figure 6.1, after registering the referral dataset to match the treatment position in the treatment dataset (see Section 4), and segmenting important tissues and materials, such as the tumour, bony anatomy, OARs, extracorporeal air and the body mask, (see Section 5), the target coverage was calculated for treatment and registered-referral datasets as described. An MRgHIFU transducer was simulated and positioned relative to the treatment and registered-referral datasets. Practical and clinically-relevant transducer translation restrictions were applied to the virtual transducer to reduce computational speed and, for volunteers, to account for limited acoustic coupling since extracorporeal air was not segmented (see Section 5.2.3). For each subject (volunteer or patient), the target (tumour) volume that could be covered by treatment cells was calculated for the subject's treatment dataset and predicted from the subject's registered-referral dataset by exhaustively translating and tilting the virtual transducer, and the two volumes were then compared quantitatively. The details of this process are presented below.

### 6.2.2.2 MRgHIFU System

The MRgHIFU transducer, based on the system described in Section 3.2, was simulated in MATLAB R2018b. The virtual transducer consisted of 256 points that represented the centre of each transducer element. The acoustic beam was modelled as a set of 256 'rays' traced from each transducer element to the geometric focal point, where each 'ray' was represented as equally-spaced points

on a straight line from the centre of a transducer element to the geometric focus. This was done because the OAR and acoustic obstruction segments had not been meshed, in order to avoid errors associated with poor-quality meshing. As seen in Figure 6.2, each 'ray' was checked for intersection by summing the score associated with each ray-point and multiplying the result by the separation between consecutive ray-points, which was 0.2 mm. This was chosen to provide a compromise between computational speed and accurate discretisation of the 'ray'. A 'ray' was counted as intersecting if the ray's total path length of intersection was greater than 0 mm. If no 'rays' were intersecting any OAR and acoustic obstruction segments, the transducer was assumed to produce a perfect 8 mm treatment cell, i.e. an 8 x 21.84 mm ellipsoid [65][162], centred at the geometric focus with its long-axis aligned to the beam axis. An 8 mm cell was chosen for this study because it was the most commonly used treatment cell in a previous bone metastases pain palliation trial conducted at the ICR [19].




**Figure 6.2:** Ray intersection with segments, represented as binary (0 and 1) arrays, were checked as follows. Rays are represented as a set of 'ray-points' (red stars) that form a line. Scores at each ray-point were linearly interpolated from the values of the surrounding voxels. Each ray-point's score is multiplied by the interval between ray-points, and summed to generate a path length of intersection for the entire ray.

### 6.2.2.3 Volunteer Targets and Patient Tumours

For volunteers (who had no tumours), all soft tissue within the body was chosen to be the target. This decision was made instead of creating test targets within the volunteer body, because the choices of test target shape, size and position would affect test target coverage, and hence the test target coverage results would be more indicative of the parameters of the test target than of the accuracy of the methodology. For patients, the targets were patient tumours, which were manually identified and segmented as described in Section 5.2.6.

### 6.2.2.4 Practical and Clinically-relevant Restrictions on Translation



**Figure 6.3:** Practical restrictions applied to the transducer's translation capabilities (solid red lines) for volunteer datasets only. (a) For treatment datasets (TID), the left-right translation was limited by the extent of acoustic coupling between the volunteer's skin and the gel pad. Registered-referral datasets shared these left-right restrictions. (b) For registered-referral datasets (RRID), the transducer's inferior-superior translation was restricted by the extent of pelvic bone and the requirement for a full body outline within the image. Treatment datasets shared these inferior-superior restrictions.

Practical and clinically-relevant restrictions on transducer translation were applied to volunteers and patients. Translation restrictions represented limits, beyond which none of the simulated transducer elements were allowed to extend. For volunteers, because the target was all soft tissue within the body, including soft tissue in the incomplete slices at the superior and inferior extents of registered-referral datasets, restrictions were placed in order to prevent the transducer from translating beyond the Superior limit (Figure 6.3b) so that the incomplete slices would not be a source of error. The Inferior limit was placed


at the inferior-most extent of the pelvis, since when the volunteer study was performed, patient data was not available and thus gynaecological tumours had been assumed to be positioned within the inferior-superior extents of the pelvic bones. Because the automatic air segmentation algorithm does not function for volunteer imaging (see Section 5.2.3), the extent of acoustic coupling (skin-to-gel-pad and gel-pad-to-membrane) was manually identified in the left-right direction (Figure 6.3a), and the transducer was restricted from translating past those restrictions.

For patients, in order to reduce computational time, the total number of translation steps were reduced by placing limits on the extent of translation. The left-right and inferior-superior transducer translation limits were derived using knowledge of the transducer focal length, the transducer's maximum angulation, and trigonometry. The maximum left-right and inferior-superior translations that allowed the transducer to focus onto the tumour at maximum angulation were calculated. The automatic air segmentation (Section 5.2.3) was used to account for the limited extent of acoustic coupling in treatment datasets. As with volunteers, the transducer elements were restricted from extending into regions with incomplete slices in the registered-referral imaging dataset. An additional translation restriction in the anterior-posterior direction was implemented, which would reduce the anterior-posterior translation so that the transducer focus could only be translated between the following two limits: (1) the anterior-most extent of the tumour and (2) the posterior-most extent of the tumour plus half the treatment cell length (10.92 mm, see Table 3.1).

#### 6.2.2.5 Positioning the Simulated MRgHIFU Transducer

In this study, the treatment position was known from the treatment dataset. In the treatment datasets, the isocentre, and hence the transducer's home position (Section 3.2), was known. In the registered-referral dataset, because the treatment position is the same, the transducer's home position left-right and inferior-superior coordinates were the same as those identified in the treatment dataset. However, to mimic the prospective workflow (Figure 1.1), the anterior-posterior coordinate had to be estimated from data within the registered-referral dataset.

The method for doing so is shown in Figure 6.4. Briefly, it was assumed that: i) the gel-pad would be most compressed and the membrane most bowed at the isocentre line, and ii) after soft tissue deformation resulting from the reorientation into the treatment position, the isocentre-to-skin point distance would remain the same.



**Figure 6.4:** Method used to predict transducer's anterior-posterior home position. The magnetic isocentre in the treatment dataset is known because the registered-referral image dataset (RRID) had been registered to the treatment dataset. A line was drawn downwards from the treatment image dataset (TID) isocentre and intersected the skin at the skin point. From this skin point, the home position was calculated using the average compressed gel-pad thickness, the average membrane bowing distance, and the calibrated distance between undeformed membrane and home position of 67.5 mm (see Figure 3.1).

After identifying the skin point, an estimate of the compressed gel-pad thickness and the membrane bowing distance were required in order to obtain the home position (see Figure 6.4). In volunteer image datasets, the membrane was observed to bow towards the transducer due to the volunteer's body weight. Fur-

thermore, the gel-pad appeared thinner close to the isocentre, probably because the subject's body weight was compressing the gel-pad. Estimation of the membrane bowing distance and the actual thickness of the gel-pad at the isocentre would allow for a more accurate estimate of the registered-referral home position, and hence a more accurate prediction of tumour coverage.

Where possible, the gel-pad thickness and membrane bowing distances were measured for all volunteer Steep and Shallow treatment datasets. The compressed gel-pad thickness was measured in a region of interest close to (at random spots within 12 mm left or right and within 45 mm inferior or superior) the left-right inferior-superior magnetic isocentre from the in-phase Dixon datasets of each treatment dataset, using ruler tools in ITK-Snap, with five repeat measurements per treatment dataset. The choice of 12 mm in left-right and 45 mm anterior-posterior was made after observing that, for at least one volunteer, the compressed gel-pad thickness did not vary with distance from the isocentre as much in the inferior-superior direction as in the left-right direction. The membrane bowing distance was measured from the same in-phase Dixon datasets and using the same software, by comparing the anterior-posterior position of the membrane edges, which were pinned at the edge of the acoustic window and hence fixed at 72.5 mm below the isocentre (Section 3.2), with the posterior-most position of the membrane. The membrane bowing distance directly below the isocentre could not be measured because of a susceptibility artefact there, centred around the Sonalleve<sup>®</sup> bed receive coil.

In order to reproduce what would happen with the prospective workflow, and because patients and volunteers were observed to have similar weights (see Tables 3.3 and 3.2), the average membrane bowing distance and average compressed gel-pad thickness for patients were assumed to be the same as that calculated for volunteers. An exception was made for patient G24, who had been treated using a customised gel-pad, so their gel-pad thickness was independently measured and used for this positioning method. Each patient's actual gel-pad thickness and membrane bowing distance were measured in a similar way to that described for volunteers, with the exception that for G8<sub>2</sub> and G27,


gel-pad thickness and membrane bowing was measured on T2wLFOV instead of in-phase Dixon datasets, because the Dixon dataset field-of-view was too small. Furthermore, because the tumour was the region of interest, the gel-pad thickness and membrane bowing was measured in a smaller region posterior to the tumour centroid (within  $\pm 12$  mm in left-right and  $\pm 13.5$  mm in inferior-superior), with five repeats for each patient. The actual measured gel-pad thickness and membrane bowing distances for patients were compared to the volunteer-derived estimates. Tumour coverage for patients was calculated using both the volunteer-derived estimates and the actual measured quantities, and compared.

#### **6.2.2.5.1 A Modification to the Positioning Method**

A modification of the positioning technique was pursued. This was motivated by the assumption that the reliance of the original technique on a single point, the position of the skin on the isocentre line, made it vulnerable to outliers. The modified technique involved identifying the mean anterior-posterior position of the skin within a square region around the isocentre line, and using that average as the Anterior-Posterior position of the skin point, from which the registered-referral home position could be derived. An investigation was performed in order to identify the ideal size of the square region to sample (see Appendix C), resulting in a 15 x 15 mm square being chosen. Differences in volunteer target coverage calculated using the original (single point) and the modified (sampled square) technique were compared using a paired Student's t-test.

#### **6.2.2.6 Calculating Target Coverage**

For volunteers, a regular grid of target points [66], one per image voxel, was created in the soft tissue (see Figure 6); for patients, this grid was only created within the tumour. Each target grid point represented the voxel surrounding it, and would record whether it was covered or not. The transducer acoustic beam had been discretised into 256 rays, linking the centre of a transducer element to the focus. Each ray was itself discretised into regularly spaced (0.2 mm) points along its length, and each was tested for intersection with acoustic obstructions or



**Figure 6.5:** Method for quantifying target volume covered within a dataset (volunteer treatment dataset in this example). A regular 3D grid of potentially accessible points was created (blue crosses) within the target: soft tissue (volunteers) or tumour (patients). For each transducer position and tilt, the acoustic beam was checked for intersection with any acoustic obstructions (green contours) or organs at risk. If no obstruction exists, an 8 mm treatment cell was created around the focus (yellow ellipse). Grid points within a treatment cell were marked as 'covered' (red crosses).

OARs. If no point intersected acoustic obstructions or OARs, an 8-mm treatment cell was drawn around the focal point, and all grid points within this were marked as covered (Figure 6.5). This was repeated as the transducer was exhaustively translated and tilted within the Left-Right and Inferior-Superior limits specified in Section 6.2.2.4.



For volunteers, the target is large, so the transducer was translated in relatively large steps (4 mm) in order to sweep the focus through it. For patients, the tumour was relatively small, so smaller (2 mm) translation steps were required in order to accurately determine the tumour coverage. The choice of 4 mm translation steps for volunteers was made because that was the radius of the treatment cell. The choice of 2 mm, which was slightly longer than the longest diagonal across a patient image voxel (approximately 1.94 mm), was made as a compromise between increased computational time and improved accuracy in determining target coverage. At each translation step, the transducer was tilted. Because of the missing data (target tissue, acoustic obstructions and OARs) in the superior and inferior extents of the registered-referral images (see Figure 6.3), which were a result of aligning the referral dataset to the treatment dataset, the virtual transducer was restricted to only tilting in the left-right direction, by up to  $\pm 10^\circ$  (device specifications, Section 3.2) in  $2.5^\circ$  steps for both patients and volunteers. This was repeated with the transducer in different positions and angles, until all tilt angles and all translated positions were checked. The total number of target (tumour) grid points covered, multiplied by the image voxel volume, was used to quantify target (tumour) volume covered.

For volunteers, the accuracy of the methodology was quantified by comparing target volume coverage calculated from the treatment dataset with that from the registered-referral dataset, as given in equation 6.1, where CV is the covered target (tumour) volume. In effect, the soft tissue volume covered in the treatment dataset becomes the target volume for the registered-referral dataset, allowing calculation of the percentage target volume covered ( $TVC_{vol}$ ).

$$TVC_{vol} = 100\% \times \frac{(CV_{RegisteredReferral} \cap CV_{Treatment})}{(CV_{Treatment})} \quad (6.1)$$

For patients, the accuracy of the methodology was quantified by examining the difference between the percentage target (tumour) volumes covered ( $TVC_{pat}$ , calculated using equation 6.2) for both treatment and registered-referral datasets.

$$\begin{aligned}
TVC_{pat,Treatment} &= 100\% \times \frac{(CV_{Treatment})}{(TV_{Treatment})} \\
TVC_{pat,RegisteredReferral} &= 100\% \times \frac{(CV_{RegisteredReferral})}{(TV_{RegisteredReferral})}
\end{aligned} \tag{6.2}$$

where TV is the total tumour volume. The computation was performed on a 20-core cluster, with each core clocked at 2.6 GHz and 12.8 GB of RAM per core.

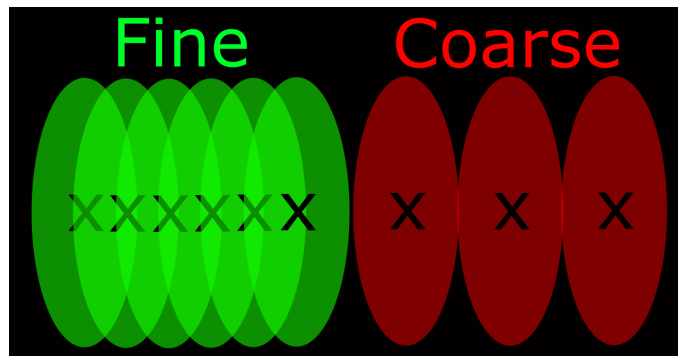
### 6.2.3 Uncertainty of Tumour Coverage Analysis due to Quantisation

In the process of calculating target (tumour) coverage, the target (tumour) volumes were discretised into target grid points. These target grid points are an imperfect representation of the original target (tumour) volume, and hence were a source of uncertainty for the target volume coverage results. The uncertainties related to the discretisation of the target (tumour) volume are assessed as described below.

Target (tumour) coverage was assessed by covering target grid points within the target (tumour) with treatment cells. To assess the effect discretisation has on the assessment of tumour coverage, the target volume covered by a single treatment cell was compared to the analytical volume of that treatment cell. Any absolute difference between the two was taken to be the uncertainty associated with calculating target coverage. A target grid of points was created, with the grid spacing being that of patient image voxel dimensions (0.86 x 0.86 x 1.5 mm). A single treatment cell (an 8 mm diameter, 21.84 mm length ellipsoid) was placed randomly within this target grid, such that no part of the treatment cell extended beyond the extents of the target grid. The covered volume calculated was compared against the treatment cell's analytical volume. This was repeated 100 times, and the mean difference between the analytical and computed results, as a percentage of the total treatment cell volume, was defined as the discretisation uncertainty.

## 6.2.4 Cell Packing Density

The methodology described in this chapter uses the exhaustive translation and angulation of the transducer to place treatment cells, and then assesses the total target (tumour) coverage within these treatment cells. Some translation and tilt combinations result in the treatment cells being placed closer together than the translation step size (4 mm or 2 mm for volunteers or patients, respectively) may suggest. Quantifying the cell packing density effectively identifies how closely cells are placed together in the target (tumour) volume, rather than having to derive it from translation and angulation increments. This is important because closer cell placement results in greater target (tumour) coverage, since more of the tumour volume between treatment cells is covered, as seen in Figure 6.6.



**Figure 6.6:** The coarse cell packing (red) leaves more of the target volume (black) in between cells uncovered than the fine cell packing (green).

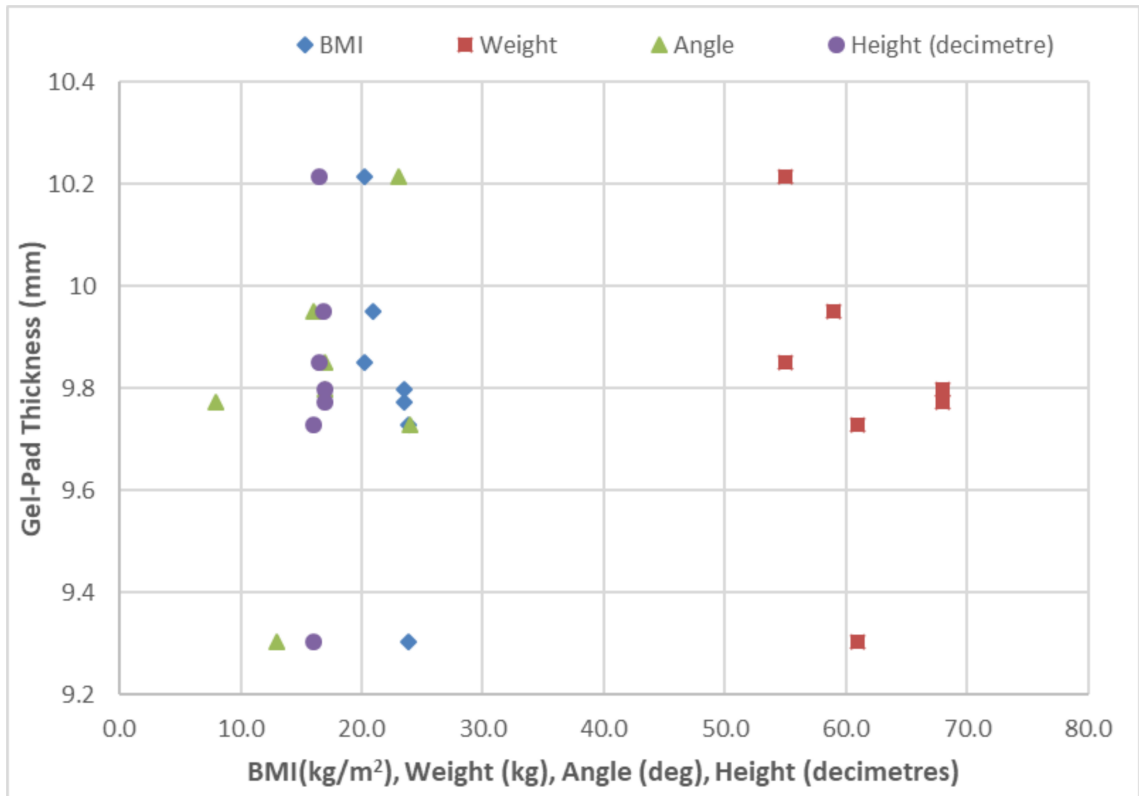
A study of the cell packing density was performed for patient G24 (chosen arbitrarily). The virtual transducer was translated and tilted using the translation and tilting settings used for calculating  $TVC_{pat}$  (Section 6.2.2.6). The treatment cell density was calculated by counting the number of treatment cell centres within a  $1 \text{ mm}^3$  region.

## 6.3 Results

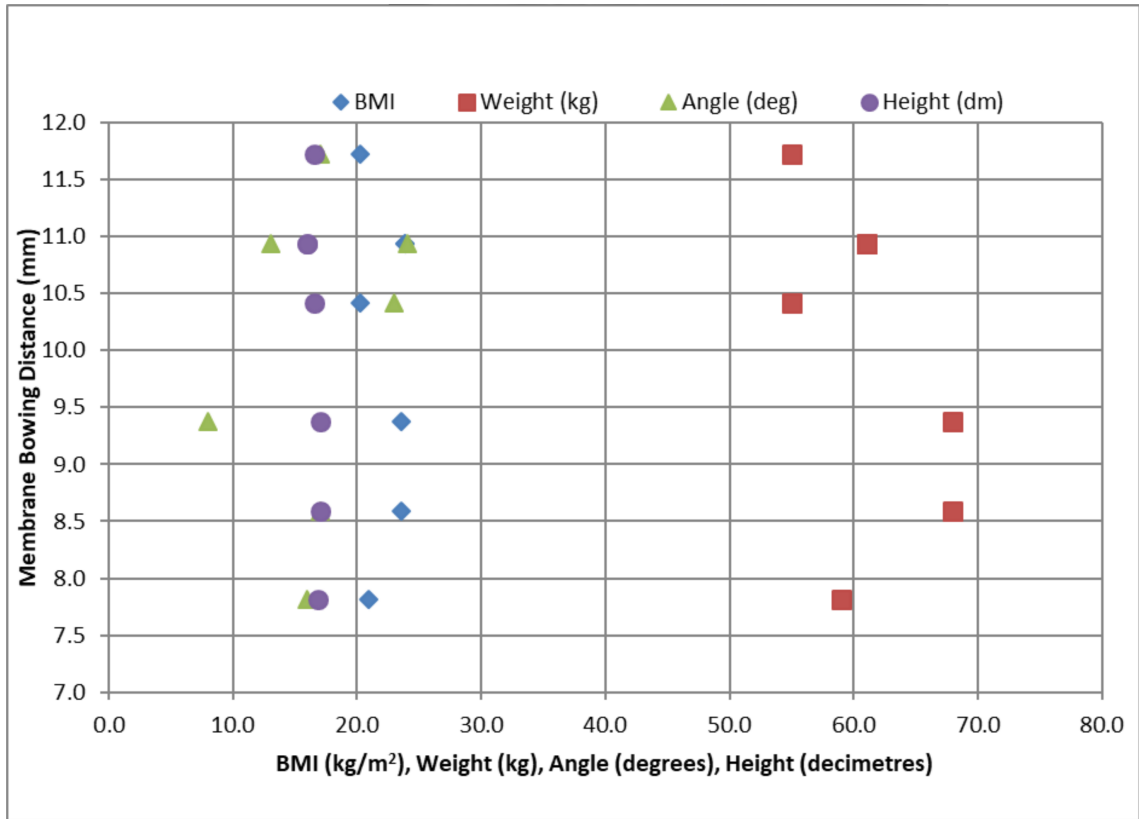
### 6.3.1 Gel-Pad Thickness and Membrane Bowing Distance

Of all ten volunteer treatment datasets, only seven captured the full extent of the gel-pad thickness and the membrane bowing at the isocentre. For volun-

teers, averaging over all measured datasets, the membrane bowing distance was  $10.0 \pm 1.3$  mm, and the gel-pad was found to be  $9.8 \pm 0.3$  mm thick at the isocentre. The gel-pad thickness and membrane bowing distance were plotted as a function of BMI and weight to assess the correlation between the variables (see Figures 6.7, 6.8).



**Figure 6.7:** Body Mass Index (BMI), weight, angle and height versus the compressed gel-pad thickness close to the isocentre for volunteer treatment datasets. Uncertainties for weight, height and BMI were estimated to be  $\pm 1^\circ$ ,  $\pm 1$  kg,  $\pm 4$  mm and 2% respectively.



**Figure 6.8:** Body Mass Index (BMI), weight, angle and height versus the membrane bowing distance close to the isocentre for volunteer treatment datasets. Uncertainties for angle, weight, height, and BMI were estimated to be  $\pm 1^\circ$ ,  $\pm 1$  kg,  $\pm 4$  mm and 2% respectively.

**Table 6.1:** Table of measured (nominal) gel-pad thickness, and measured membrane-bowing distance for patients.


Patient	G3	G8 <sub>2</sub>	G24	G27	G29
Gel-pad Thickness (mm)	12.3	8.0	5.3	10.9	8.6
(Nominal)	(15)	(15)	(40)	(15)	(15)
Membrane Bowing (mm)	5.0	10.0	4.0	10.0	9.0

## 6.3.2 Volunteer Target Coverage

For each volunteer in each of their two treatment positions (Steep and Shallow, see Section 3.3.2), the  $TVC_{vol}$  is shown in Figure 6.12 (top). On average, over all volunteers and treatment positions, the tumour volume covered predicted from the registered-referral dataset is within 9% (range: 2 to 22 %) that calculated from the treatment dataset. Representative axial images for each volunteer, in each treatment position, are shown in Figure 6.9.

### 6.3.2.1 A Modification to the Positioning Technique


The absolute difference in volunteer target coverage between the positioning techniques (one based on a single-point, the other based on the sample square region of interest) was  $0.4 \pm 0.2\%$ , corresponding to an absolute difference in volume of approximately  $1100 \pm 700 \text{ mm}^3$ . A paired t-test showed no statistically significant difference ( $p > 0.83$ ) between the two results.



**Figure 6.9:** Representative axial anatomical images of volunteers, with the treatment image (green) overlaid on the registered-referral image (magenta). The soft tissue volume covered as predicted from the registered-referral dataset only (red), that calculated from the treatment dataset only (cyan) and that identified by both datasets (white) were overlaid on the anatomical data. For volunteers, the target volume was defined as the soft tissue volume covered in the treatment dataset (white plus blue volumes). The percentage target volume covered is the white volume divided by the white plus blue volume represented as a percentage.

Because V2 Shallow had the worst agreement between registered-referral and treatment datasets (approximately 83%, Figure 6.12), this was investigated further. It was found that the skin point below the isocentre differed by 12.2 mm

between the treatment and registered-referral datasets (grayscale and red respectively, see Figure 6.10), meaning that for the registered-referral image, the transducer home position should have been placed 12.2 mm posterior to where it had been estimated to be.



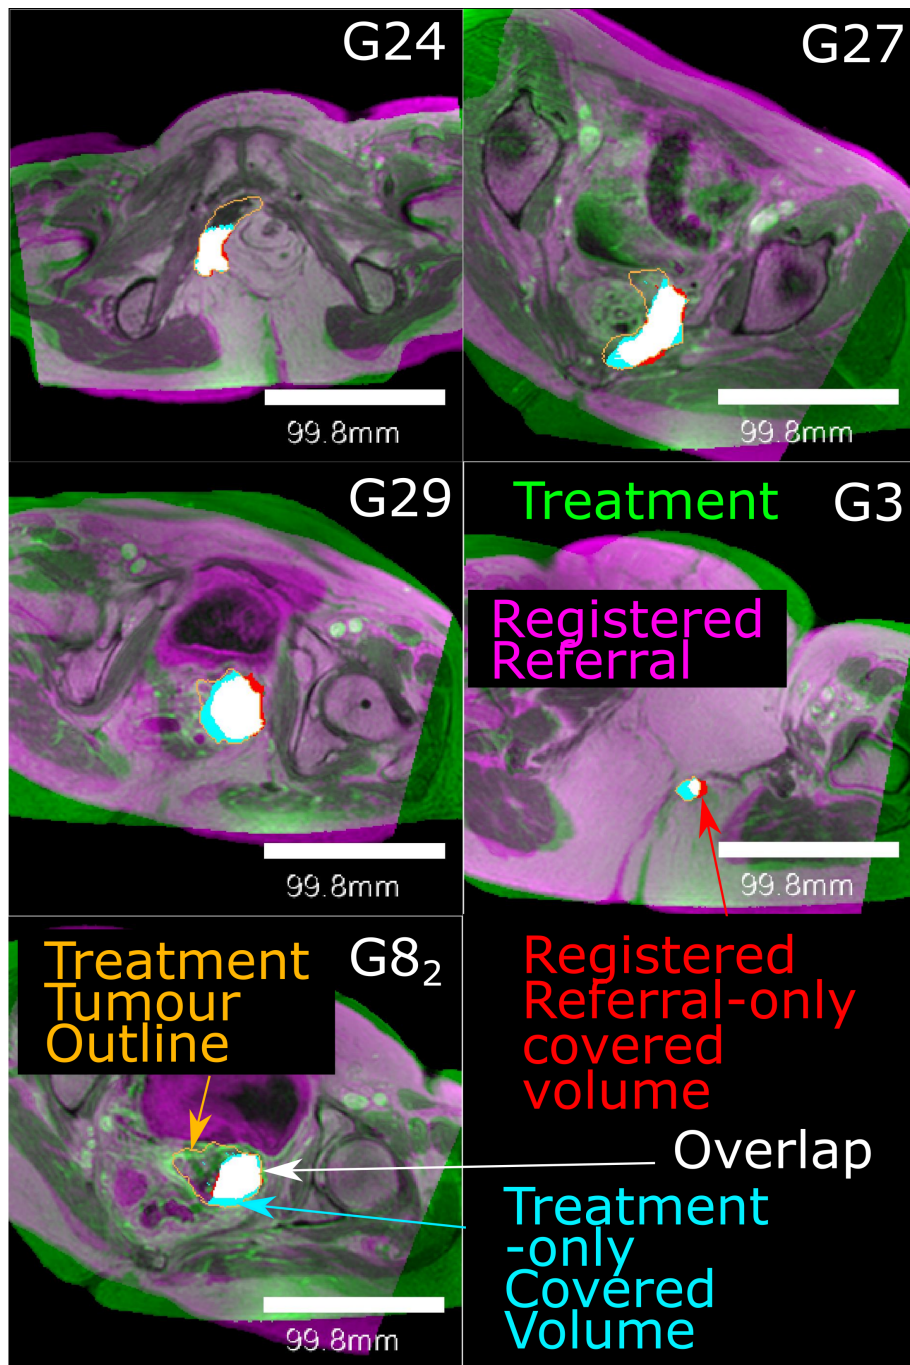
**Figure 6.10:** Left: Anatomical overlay between registered-referral (red) and treatment (grayscale) datasets, with the isocentre at the crosshair centre. The isocentre line goes vertically through the isocentre. The contents of the green rectangle is shown magnified on the right. Right: The anterior-posterior distance between the registered-referral skin point and the treatment skin point, directly below the isocentre, was measured to be 12.2 mm.

### 6.3.3 Patient Tumour Coverage

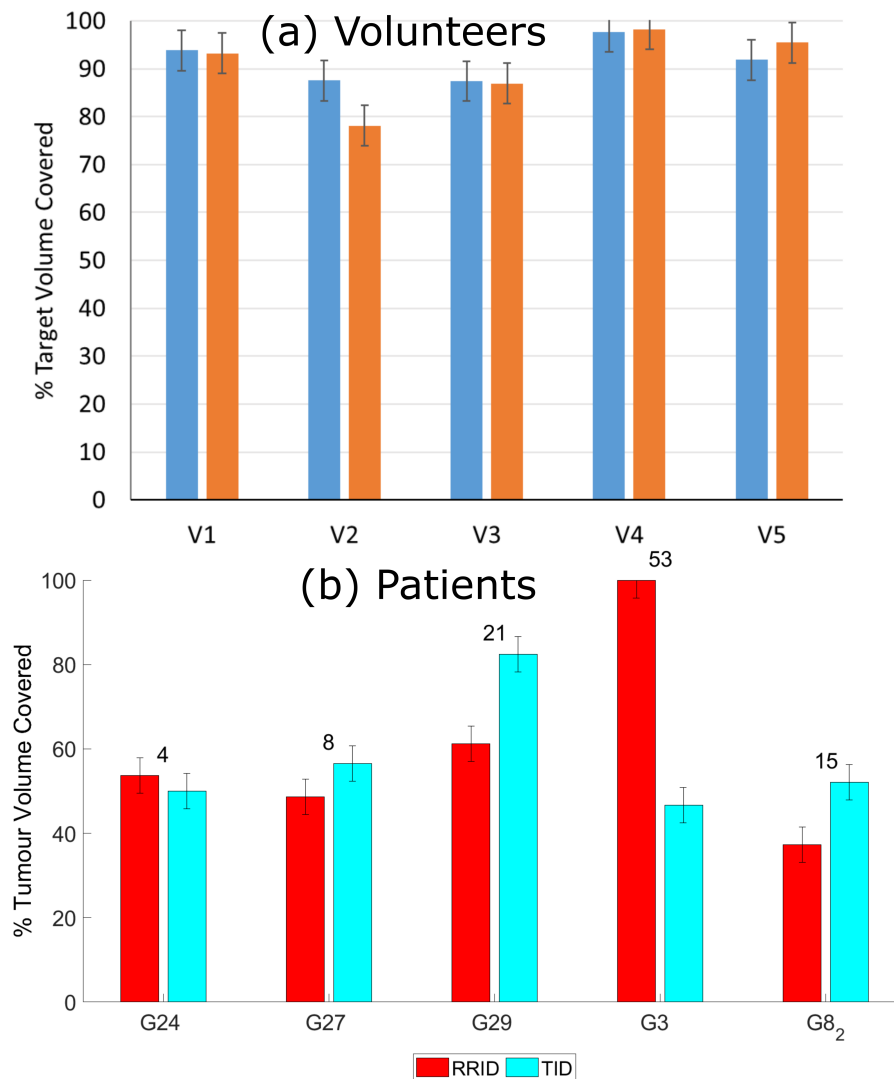
Because the difference between the referral positioning techniques explored in the volunteer study was not statistically significant, the patient study that followed was conducted using the positioning technique that used the skin point directly below the magnetic isocentre. The  $TVC_{pat}$  for each patient's treatment and referral datasets are shown in Figure 6.12 (bottom). On average, the registered-referral  $TVC_{pat}$  predicts the treatment  $TVC_{pat}$  to within 20% (range: 4 to 53%). If patient G3 is excluded, the mean percentage difference between registered-referral and treatment  $TVC_{pat}$  is 12%. Representative axial images for each patient are shown in Figure 6.11. The computational time required was  $19 \pm 10$  hours (mean  $\pm$  standard deviation, range: 3-32 hours).

For patient registered-referral datasets, patient tumour coverage was recalculated using the measured values for compressed gel-pad thickness and membrane bowing distance, instead of the volunteer-derived estimates. The maximum difference in tumour coverage between the two sets of results was 0.2%.






**Figure 6.11:** Representative axial anatomical images of patients, with the treatment (green) and registered-referral anatomy (magenta) overlaid. The entire treatment dataset tumour segment outline (orange), the treatment-only covered volume (cyan), the registered-referral-only covered volume (red) and the overlap between the treatment and registered-referral covered volumes (white) were overlaid on the anatomical images. The overall treatment tumour covered volume is white plus blue; the overall registered-referral tumour covered volume is white plus red. For simpler visual interpretation, only the treatment tumour outline is shown (orange line).



**Figure 6.12:** Bar charts depicting the percentage of the target volume covered for volunteers (a) and the percentage of the tumour volume covered for patients (b). (a) For volunteers, target volume was defined as the soft tissue volume covered as assessed from the treatment images. The percentage of the target volume that was predicted to be covered from registered-referral datasets was plotted, for each volunteer, for each treatment position (Steep (blue) or Shallow (orange)). (b) For patients, the percentage of the registered-referral tumour segment volume predicted to be covered from registered-referral datasets (red), are compared to the percentage of the treatment tumour segment volume calculated to be covered from treatment datasets (cyan). For simpler visual identification, the numbers above each pair of bars represents the difference in % Tumour Volume Covered predicted from registered-referral datasets and that calculated from treatment datasets.

Patient G3 showed the greatest disagreement between the registered-referral and treatment  $TVC_{pat}$ . The patient's treatment dataset was examined

more closely, with a sagittal slice being shown in Figure 6.13. The beam path is close to the extracorporeal air in the inferior-superior direction, restricting access to the tumour for some positions.



**Figure 6.13:** Sagittal anatomical image of patient G3 at the tumour centroid, overlaid with acoustic obstructions (AOs) and organs at risk (OARs, combined red overlay), the tumour (green) and the virtual transducer (blue transducer, orange acoustic beam). Notably, the extracorporeal air was present inferior to and posterior of the tumour. The virtual transducer is focused on the isocentre, but offset in the Inferior-Superior direction in order to draw attention to the relative position of the extracorporeal air and the tumour.

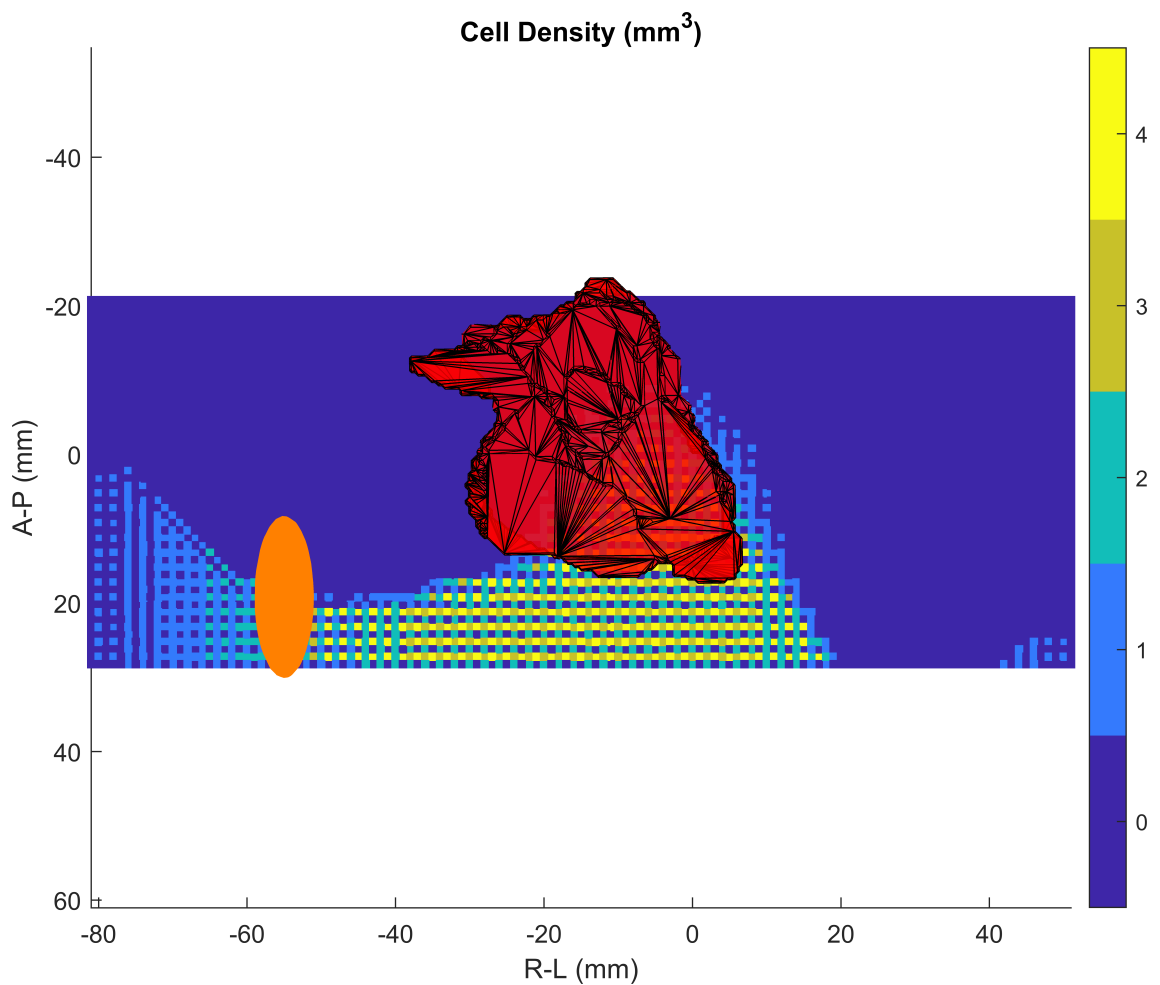
### 6.3.4 Uncertainty of Tumour Coverage Analysis

The uncertainty associated with the discretisation of the target (tumour) into grid points was estimated to be  $1.7 \pm 0.8\%$ . In total, the uncertainty associated with

volunteer and patient target coverage was calculated to be approximately 4.3%. The major contributions to uncertainty were the uncertainty involved in manual image segmentation, estimated to be 4% (Section 5.3) and the discretisation uncertainty at 1.7%.

### 6.3.5 Cell Packing Density

The maximum cell packing density observed was 4 cell centres per  $\text{mm}^3$ . The spatial distribution of the cell packing density for patient G24 is shown in Figure 6.14.



**Figure 6.14:** Spatial distribution of reachable treatment cell density for patient G24 in the axial plane. Point colours correspond to the number of treatment cell centres found within a  $1 \text{ mm}^3$  cube centred at the point. The maximum cell packing density was 4 cells per  $\text{mm}^3$ . The tumour is overlaid as a mesh segment (red), and an 8 mm diameter treatment cell (orange) is overlaid for comparison.

## 6.4 Discussion

### 6.4.1 Gel-Pad Thickness and Membrane Bowing Distance

Compressed gel-pad thickness and membrane bowing distance varied minimally between volunteers, with standard deviation in gel-pad thickness of 0.3 mm and in membrane bowing distance of 1.3 mm. This supported its use in patient data. The scatter plots (Figures 6.7 and 6.8) suggest that membrane bowing and gel-pad thickness do not vary with BMI, height or weight.

The compressed gel-pad thickness and membrane bowing distances measured from the treatment datasets of some patients (Table 6.1) varied from the averages derived from volunteer data (9.8 and 10.0 mm, respectively). In particular, the actual patient membrane bowing distances of 5 mm and 4 mm measured for patients G3 and G24, respectively, were significantly different. To evaluate the effect of this, the TVC was recalculated with the actual gel-pad thickness and membrane bowing distance for all patients. The maximum difference that resulted from using the average membrane bowing and gel-pad thickness, rather than the actual measured values, was 0.3% (patient G24), suggesting that other factors have a greater impact on the difference between the  $TVC_{pat}$  predicted from the referral dataset and that calculated from the treatment dataset.

### 6.4.2 Volunteer Target Coverage

For volunteers, the results show that the average target coverage agreement between treatment imaging datasets and registered-referral imaging datasets is 91% (range: 78 - 98%), corresponding to a mean difference of 9%. These results suggest that the techniques used for positioning the transducer home position in the registered-referral imaging datasets were sufficient to proceed to testing with patient data. The worst disagreement, 78%, for Volunteer 2 tilted at a 12° treatment angle, was probably related to the inaccurate placement of the registered-referral transducer home position, caused by a wrong assumption that the distance between the isocentre and the skin point directly below the isocentre (see Figure 6.4) would remain constant. As a consequence, the transducer focus was

predicted to reach 12 mm deeper into the volunteer than it actually could (Figure 6.10). In practice, because of the deformation of the soft tissue layers between the isocentre and the skin, the distance between the isocentre and the skin point had probably changed between the registered-referral and treatment datasets. The next worst disagreements were associated with Volunteer 2 in the Steep position, and Volunteer 3 in both treatment positions. For these cases, similarly to Figure 6.10, the home position for the registered-referral dataset was placed approximately 6 mm closer to the volunteer than was achieved in the treatment dataset, which resulting in an overestimate of the depth that the MRgHIFU focus could reach. Again, this was probably related to the assumption that the isocentre-to-skin point distance remained constant between the registered-referral and treatment datasets.

The registered-referral imaging dataset and treatment imaging dataset target coverage assessments differ only in the choice of transducer home position for the registered-referral imaging dataset, as discussed above, and in the femur segmentation. Another source of the difference in target coverage between the registered-referral and treatment datasets could have been related to the articulation of the femur. A change in articulation with respect to the pelvis between the referral and treatment datasets could have contributed to the differences in the target volume coverage predicted by the registered-referral dataset versus that calculated from the treatment dataset. This is because the femur is connected to the pelvis at the approximate inferior-superior midpoint of the pelvic bones. Therefore, a change in the position of the femur relative to the pelvis could change access to pelvic soft tissue, the target tissue for volunteers. Since the angulation is restricted to tilting left-right only, and the transducer movement is restricted to prevent translation beyond the inferior-superior extents of the pelvis, differences in femur segmentation were judged to have only a small effect. Another source of uncertainty was the discretisation of the target volume, which was found to contribute 1.7% to the difference in target covered volumes.

The modified transducer positioning technique, based on sampling within a square region of interest around the isocentre line, provided no statistically signif-

icant difference to the positioning method that used the skin point directly below the isocentre (see Figure 6.4). In light of this, the single skin-point transducer positioning technique was used when analysing patient data.

### 6.4.3 Patient Target Coverage

The goal of the work described in this chapter was to develop a methodology to quantitatively assess the tumour coverage from referral imaging, as opposed to the current clinical practice of qualitative assessment, and to assess the feasibility of the methodology. Agreement in tumour coverage between referral and treatment imaging datasets to within 20% was judged as acceptable [59]. From the results, quantitative prediction of tumour coverage from referral imaging appears feasible. Despite the simplicity of the technique used to account for the expected body deformation that results from reorientation from supine into a treatment position, the  $TVC_{pat}$  predicted from the registered referral imaging dataset and that calculated from the treatment imaging dataset had a mean difference of  $12 \pm 7\%$  (range: 4-21%), excluding what is believed to be an outlier for whom the difference was 53% (patient G3). The outlier is explored in greater detail below. In the context of current clinical studies, where more than 40% of referred patients ultimately fail screening, these results are encouraging [18][19]. The small cohort involved in this study (5 volunteers, 5 patients) represents lower than expected patient recruitment for the clinical trial. However, other published studies have also involved small patient cohorts, e.g. a transcranial simulation study involved 5 patients [74], a simulation study for kidney ablation examined 4 patients [24], and in various therapeutic feasibility studies, between 10 to 13 patients' data was available [163][164][165]. In addition, the capability of an automatic geometric optimisation technique for the packing of HIFU treatment cells has been demonstrated using test objects and the publicly available dataset of a single volunteer [67]. Results from these small-cohort feasibility studies also demonstrate high variance in results. For example, in the transcranial simulation study, simulation results differed from measured data by up to  $40 \pm 13\%$  [74]. The results here indicate a step towards the long-term objective of widespread quantitative analysis of

patient suitability for MRgHIFU therapy, with the aim of improving clinical decision-making and minimising the impact on patient and hospital time and resources.

The outlier, with the worst prediction by far (53% overprediction), was patient G3. The main reason for the poor results with this patient was thought to be the assumption of perfect acoustic coupling between the patient and the gel-pad when predicting  $TVC_{pat}$  from the referral dataset. In practice, the treatment imaging dataset showed that the peripheral position of the tumour was obstructed by extracorporeal air between the patient and gel-pad (see Figure 6.13). This highlights a possible advantage of the prospective patient suitability workflow, in which this methodology is a core component. Having established that greater tumour coverage could have been achieved by examining referral imaging, clinicians could have pursued more aggressive clinical preparation in order to increase tumour coverage, for example by using a customised, degassed-water-filled recessed gel-pad as was done for patient G24. This was not done in practice for patient G3, since the clinical team had limited experience with treating abdominal soft tissue tumours and had not expected such difficulty during treatment.

Two reasons why results for patients appear, in general, worse than those for volunteers, are discussed below. First, volunteer target volumes were approximately 10 times larger ( $300 \pm 100 \text{ cm}^3$ ) than patient targets ( $20 \pm 10 \text{ cm}^3$ ) whilst target grid spacing for volunteers (approximately  $0.78 \times 0.78 \times 1.50 \text{ mm}$ ) and that for patients ( $0.868 \times 0.868 \times 1.50 \text{ mm}$ ) are similar. Hence, a single target point being covered would have a larger proportional effect on percentage target volume covered for smaller target volumes than for larger target volumes. Secondly, volunteer OARs were ignored, whilst patient OARs changed shape, size and position between referral and treatment datasets, due to patient orientation from supine to treatment position, and clinical preparation such as pre-treatment dieting and bowel-preparation. This could have affected tumour coverage. For example, bladder and rectal volumes have been seen to vary by up to  $\pm 30\%$  [60], and rectum positions by up to 5 mm in the Anterior-Posterior direction, up to 3 mm left-right, and by less than one slice thickness in the Inferior-Superior direction [61].

Given that the 'referral' imaging in this study was actually the 7-day follow-



up imaging, post-ablation tumour volume changes were proposed as a possible source of error. However, in the literature, it has been suggested that superficial malignant tumour volume does not decrease by a statistically significant amount [166] even 6-8 weeks after ablation. This suggests that any post-ablation tumour volume change in 7 days is negligible. Tumour growth in the time between referral and treatment (maximum 1 month for NHS England [167]) was explored as a possible source of error if the software were to be used clinically. Cervical tumours have a mean metabolic tumour volume doubling time of approximately 300 days [168][169], suggesting that prediction error related to tumour growth would be minimal.

The threshold path length of intersection, below which rays were assumed not to have intersected, was proposed as a possible source of uncertainty. The path length of intersection was increased to 0.2 and 1.0 mm, and  $TVC_{pat,TID}$  for patient G3, who was chosen because of the acoustic beam's proximity to the extracorporeal air (Figure 6.13), was calculated. The resulting  $TVC_{pat,TID}$  was increased to 51% and 71% respectively. This suggests that the path length of intersection, which essentially defines a "margin of acceptance" around the OAR and acoustic obstruction segments, can play an important part in determining the tumour coverage. This implies that tumour coverage is sensitive to the path length of intersection, and by extension, the segmentation quality.

Computational time ranged from 3 hours to 32 hours for different patients. This was related to the Left-Right, Anterior-Posterior and Inferior-Superior dimensions of each patient tumour. The tumour of patient G3 had Left-Right, Anterior-Posterior and Inferior-Superior dimensions of 13.0 mm, 11.3 mm, and 12.0 mm respectively, with time elapsed during computation being 3 hours; for patient G27, the tumour dimensions were 51.2 mm, 54.7 mm and 31.5 mm respectively, with time elapsed during computation being 32 hours. The current computational time requirement (an average of 20 hours on a 20-core cluster node) is thought to be unrealistic in a clinical setting. Computational time could be improved with better software engineering, such as rewriting the code in a compiled language, and by employing specialised computational hardware such as graphics processing units

(GPUs) or field-programmable gate arrays (FPGAs). Refinements of the methodology for calculating patient tumour coverage, aimed at improving computational speed without sacrificing accuracy, are described in Chapter 7.

#### **6.4.4 Cell Packing Density**

Due to a combination of translation and tilting, treatment cells were placed in regions distant from the tumour. Furthermore, up to 4 treatment cell centres were placed within some 1 mm<sup>3</sup> regions, which appears redundant given that the target grid points within the tumour had the same spacing as image voxel dimensions (0.868 x 0.868 x 1.50 mm for patients). The cell density decreased as the transducer translated towards the anterior, likely because the acoustic beam began to intersect OARs and acoustic obstructions. The same phenomenon was observed for cell density near the left and right extremes of the image (Figure 6.14), because the transducer was forbidden from translating beyond a certain point (see Section 6.2.2.4), although it was allowed to tilt in the Left-Right direction by up to 10° off axis.

These results suggest that, if refinements of the methodology were pursued with the goal of reducing computational speed, several refinements to the placement of treatment cell centres could be made. One possible refinement would be to generate treatment cell centres close to and within the tumour. Another refinement would be to place treatment cell centres more evenly, in order to reduce the redundancy in their placement involved in the current methodology and hence reduce computational time.

#### **6.4.5 Limitations of the Study**

This study focused on tumour coverage rather than tumour treatability, because preliminary experiments suggested that calculating tumour coverage required less time than calculating tumour treatability. Hence, tumour coverage was planned to be used to identify the ideal treatment orientation. However, attenuation and refraction from the pre-focal fat and muscle may lead to less thermal ablation than expected from the tumour coverage. The region of heating may also be offset from the geometric focus. This is the subject of a later study (Chapter

9).

Another major limitation of the study is the small number of volunteers (n=5, two 'treatment' positions each) and patients (n=5) involved, which makes deriving a definitive conclusion about the quality of the methodology difficult. With a larger sample set, patient outliers could be clearly distinguished. Furthermore, this study is limited to predicting pelvic soft tissue tumour coverage in oblique supine decubitus treatment positions. For other tumour sites and other treatment positions, for example liver carcinomas treated in the prone position, it is likely that the prospective patient suitability workflow would need to be adapted. While the precise results of this study are only applicable to the specific diagnostic MR bed and MRgHIFU couch used, the core principles are expected to be applicable to other HIFU devices.

Deformation of the patient was not simulated – instead, only the simple assumption that the distance between the isocentre and the skin point directly posterior to it would be the identical in both the rotated referral dataset and the deformed treatment datasets, was applied. The work here is assumed to apply for the patients going through the prospective patient workflow, in which the patient referral dataset would be rotated such that the isocentre, the tumour centroid and the skin point would be collinear. Full simulation of the deformation of the soft tissues between the gel pad and the target for acousto-thermal modelling purposes may increase the predictive capability of the prospective patient suitability workflow.

## **6.5 Conclusion**

Novel methodology for predicting the MRgHIFU target coverage from supine (MR) referral imaging has been developed using 10 volunteer datasets and retrospectively applied to 5 patient datasets. The difference between the target coverage computed using referral versus treatment image datasets was within  $12\pm 7\%$  on average (range: 4-21%), when one patient, in which inadequate acoustic coupling was achieved during treatment, was excluded from the analysis. Alternative methods to account for soft tissue deformation should be explored. Despite the

relatively small cohort size, the focus on pelvic tumours, and the limited range of patient positions and MRgHIFU equipment on which the methodology was devised and tested, these results suggest that the development of a quantitative screening workflow, which should obviate the need for patient suitability to be assessed using solely clinical judgement based on operator experience, should be feasible.

## Chapter 7

# Improved Method for Calculating Patient Tumour Coverage

## 7.1 Introduction

The 'device-first' method described in Chapter 6 was developed using volunteer data, for which there was no tumour target. As we have seen, in volunteers, all soft tissue in the pelvic region was defined as the "target" tissue. Because this is a large target, a method based around exhaustive angulation and translation of the transducer was developed. The computational time required scaled with the number of transducer translation and tilting steps. For smaller targets, such as patient tumours, a higher step resolution (i.e. smaller translation and tilting steps) was required to ensure extensive tumour coverage. Some transducer position and angle combinations, particularly those with the transducer at the left-right or inferior-superior translational extent, require processing even though they are unlikely to contribute to tumour coverage. A problem arose from the representation of the acoustic beam. In this device-first method, when checking for acoustic beam intersection with organs at risk (OARs) or acoustic obstructions, the beam was represented as a set of rays reaching from each transducer element to the transducer focus, with each being represented as a collection of collinear points instead of a line, and each point being tested for intersection with acoustic obstructions or OARs. To avoid under counting intersections with acoustic obstructions or OARs, the distance between consecutive ray-point was required to be

less than the smallest image voxel dimension, and was set to 0.2 mm. Each ray (140 mm long) therefore consisted of 700 ray-points. There are 256 rays from the transducer elements to the geometric focus. The device-first method also introduces uncertainty from the interpolation of patient data into a new grid when rotating the referral imaging dataset, as required for the prospective workflow (Figure 1.1 Step 2). The computing time complexity of this method scales with the number of translation steps, the number of tilt increments per translation step, and the number of points per ray.

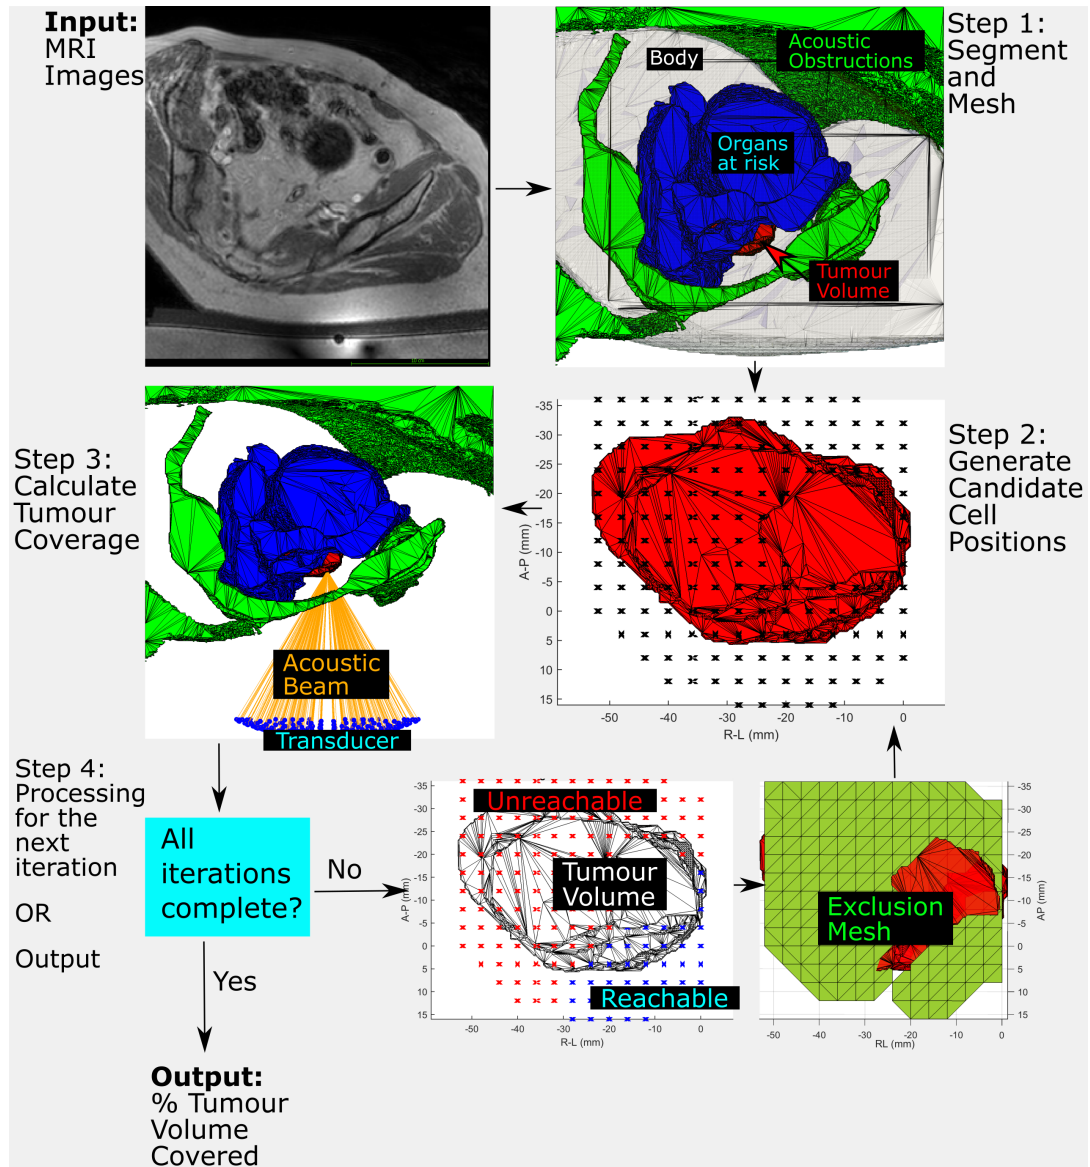
In this chapter, refinements to the method for calculating tumour coverage were investigated, with the aim of improving computational speed without sacrificing accuracy. A new 'target-first' method was developed, in which the problem was approached from the 'other end', as shown in Figure 7.1, and was tested on patient treatment images. Instead of scanning through device parameters to establish where the focus would cover part of the tumour, points representing candidate treatment cell positions were generated within the tumour, and the transducer beam interrogated to see whether it could reach this point from allowed positions and angles. Acoustic beam intersection checking could be sped up by representing the acoustic beam as a set of vector rays from the transducer elements to the geometric focus, and the OAR and acoustic obstruction segments as triangulated meshes. This allowed use of the Möller–Trumbore ray-triangle intersection algorithm [170] that takes advantage of GPU parallelisation [171]. An iterative, coarse-to-fine cell packing approach was adopted – with cells initially being packed loosely, and the packing tightening around the remaining uncovered tumour volume with each iteration. This improved computational speed by reducing the redundant processing of very closely packed treatment cells. The tumour coverage and the computational time required for the target-first method is compared here to that needed for the device-first method.

## 7.2 Method

### 7.2.1 Overview

Figure 7.1 shows an overview of the target-first method, with details given in the Sections below. Patient treatment images were used to test the workflow, firstly because the MRgHIFU transducer home position could be easily located (140 mm posterior to the magnetic isocentre, see Figure 3.1), and secondly because the automatic segmentation of extracorporeal air, a large acoustic obstruction, was only available for these images.

In Step 1, the image datasets are segmented to extract important OARs and acoustic obstructions using the methods described in Chapter 5. Each segment is then transformed into a triangular surface mesh. In Step 2, candidate treatment cell centres are generated within, and around, the tumour mesh. In Step 3, each candidate cell centre is assessed for whether it could be reached, by back-projecting virtual transducers at angles that could theoretically reach that point up to  $10^\circ$  away from the anterior-posterior axis. Each virtual transducer is tested for breaching device limits, or for obstruction of its acoustic beam by OARs or acoustic obstructions. If the acoustic beam is unobstructed and the transducer has not exceeded device limits, then the candidate cell centre is marked as reachable. For each reachable cell centre, 8 mm diameter treatment cells are drawn around them, and the tumour volume within the cells is marked as being covered by the MRgHIFU focus. In Step 4, in order to reduce computational time in processing cell centres in unreachable and already-covered regions, 'exclusion meshes', within which treatment cell centres generated in the next iteration of Step 2 would be removed, are generated from unreachable and redundantly reachable treatment cell centres. Steps 2 to 4 were repeated until step 3 had been completed for the final iteration. With every iteration, Step 2 generated candidate cell centres more finely around the tumour volume that remained uncovered by the MRgHIFU focus than the previous iteration. After Step 3 in the final iteration, the percentage of the tumour volume covered was calculated and output.



**Figure 7.1:** Workflow for the target-first methodology, demonstrated using a representative patient treatment dataset. Step 1: Input Dixon images are segmented to give acoustic obstructions (green), organs at risk (blue, OARs), the tumour (red) and the body outline (grey). The segments are then transformed into triangular surface meshes. In step 2, candidate cell positions (black crosses) are generated within and around the tumour mesh. In step 3, the tumour coverage is calculated by back-projecting plausible transducer positions (only one shown, for ease of visual interpretation) from each candidate cell position. Each transducer position is then tested for whether the beam would intersect OARs and acoustic obstructions, and whether placing the transducer in that position would breach the specified limits of translation for the MRgHIFU system. In step 4, if the current iteration is not the final iteration, then the reachable and unreachable cell positions are identified. Exclusion meshes were derived such that the next iteration of cell points would be generated only outside of the exclusion meshes, and within and around the remaining uncovered tumour volume. If the current iteration was the final iteration, then the percentage of the tumour volume covered is output.




## **7.2.2 Comparison with the Device-First Method**

In order to assess the performance of the target-first method, it was compared against the device-first method (Chapter 6). The meshing and cell packing of the target-first method was designed to be as similar to the device-first method as possible. The resulting tumour coverage calculated by the two methods, and the computation time required, was compared for all patient treatment image datasets.

### **7.2.2.1 Meshing**

In the device-first method, OAR and acoustic obstruction segments remained as binary 3D arrays. A binary array is one in which each voxel in the array can only have one of two values: '1' or '0'. For all segments, the '1' voxels represent the segmented tissue or material, whilst the '0' voxels represent everything else. Intersection between the acoustic beam and the OARs or acoustic obstructions was checked by discretising the acoustic beam into ray-points, the value of each point being interpolated from the values of the voxels surrounding each point. If for each ray, the allowed path length of intersection is 0 mm, i.e. no intersection of any part of the beam with any acoustic obstruction or OAR is allowed, this enforces a strict requirement for the ray-points to be at the centre of, or surrounded by, zero-value voxels. Therefore, the device-first method effectively assumes a segment surface boundary as shown in Figure 7.2.



**Figure 7.2:** An illustrative example of a segment, where the '1' boxes are the segment voxels and the '0' boxes are the background voxels. DFM is device-first method. If the allowed path length of intersection was set to 0 mm, the red line demonstrates the effective surface of the segment that rays may not pass without being marked as intersecting.

The target-first method was developed on the basis of using the Möller–Trumbore ray-triangle intersection algorithm [170] to speed up intersection checking. This algorithm required the OARs and acoustic obstructions to be represented as triangulated surface meshes. To replicate the segment surface boundary seen in the device-first method in the target-first method, the meshing strategy used for each segment array was performed as follows. The segment array was transformed into sets of coordinate points, where each coordinate was the centre of a '1' voxel. The coordinate points of the '0' voxels neighbouring the '1' voxels were found. Alpha-shape meshing [172], with the alpha radius set to be the smallest alpha-shape that resulted in enclosure of all points, was used to transform the coordinate points into a triangulated surface mesh. In order to remove spurious triangles, meshes were filtered to remove triangles with any edge length

exceeding the distance  $D$  (given by equation 7.1):

$$D = 2\sqrt{a^2 + b^2 + c^2} \quad (7.1)$$

where  $a$ ,  $b$  and  $c$  are the voxel dimensions.

As in the device-first method, a target grid was generated within the tumour mesh, the purpose of which was to keep track of what part of the tumour target had been covered by the MRgHIFU focus. The target grid was generated with spacing equivalent to image voxel dimensions (0.868 x 0.868 x 1.50 mm), in order to match the target grid spacing used in the device-first method. The discretisation error was assumed to be the same as that found for the device-first method (see Section 6.2.3). In order to reduce meshing time and improve computational speed, and because the transducer would be positioned posterior to the isocentre, the parts of the extracorporeal air segment that was anterior to the magnetic isocentre were removed. Preliminary work suggested that this could allow a native MATLAB implementation of the ray-triangle intersection algorithm to work five times faster.

#### 7.2.2.2 Cell Packing

To replicate the cell densities observed in the device-first method, cells were packed in a cuboidal manner, with a constant Left-Right, Anterior-Posterior and Inferior-Superior spacing between candidate cell centres. The spacings used for each iteration are shown in Table 7.1. The Inferior-Superior spacing was constant at 2 mm, because in the device-first method, transducers were limited to tilting in the Left-Right direction and could only be translated by 2 mm in the Inferior-Superior direction. The initial iteration's (Iteration 1) Left-Right and Anterior-Posterior spacing of 4 mm was chosen since this was the radius of the chosen 8 mm treatment cell. The final iteration's (Iteration 3) Left-Right and Anterior-Posterior spacing of 1 mm was chosen since patient target grid points were spaced by 0.868 mm in the Left-Right and Anterior-Posterior directions, and was chosen to reduce the redundancy involved in treatment cell placement in the device-first method. Iteration 2's Left-Right and Anterior-Posterior spacing was

chosen as a midpoint between the initial and final iteration settings.

**Table 7.1:** Spacing between candidate cell centres for each iteration of the target-first method.

Iteration No.	Left-Right (mm)	Anterior-Posterior (mm)	Inferior-Superior (mm)
1	4	4	2
2	2	2	2
3	1	1	2

For the first iteration, candidate cell centres were generated within a cuboidal region, with the sides of the cuboid defined by the Left-Right and Inferior-Superior extents of the tumour, the anterior extent of the tumour, and the posterior extent of the tumour plus half the treatment cell length. For all other iterations, in order to reduce computation time as cell centres are generated with increasingly fine spacing, the candidate cell centres were generated within a cuboidal region defined by the Left-Right and Inferior-Superior extents of the remaining uncovered tumour volume, the anterior extent of the remaining uncovered tumour volume, and the posterior extent of the remaining uncovered tumour volume plus half the treatment cell length. The remaining uncovered tumour volume was used to reduce the total number of cell points generated, compared to when the entire tumour volume was used. Any candidate cell centres within the exclusion mesh, which had been derived from the previous iteration, were then removed.

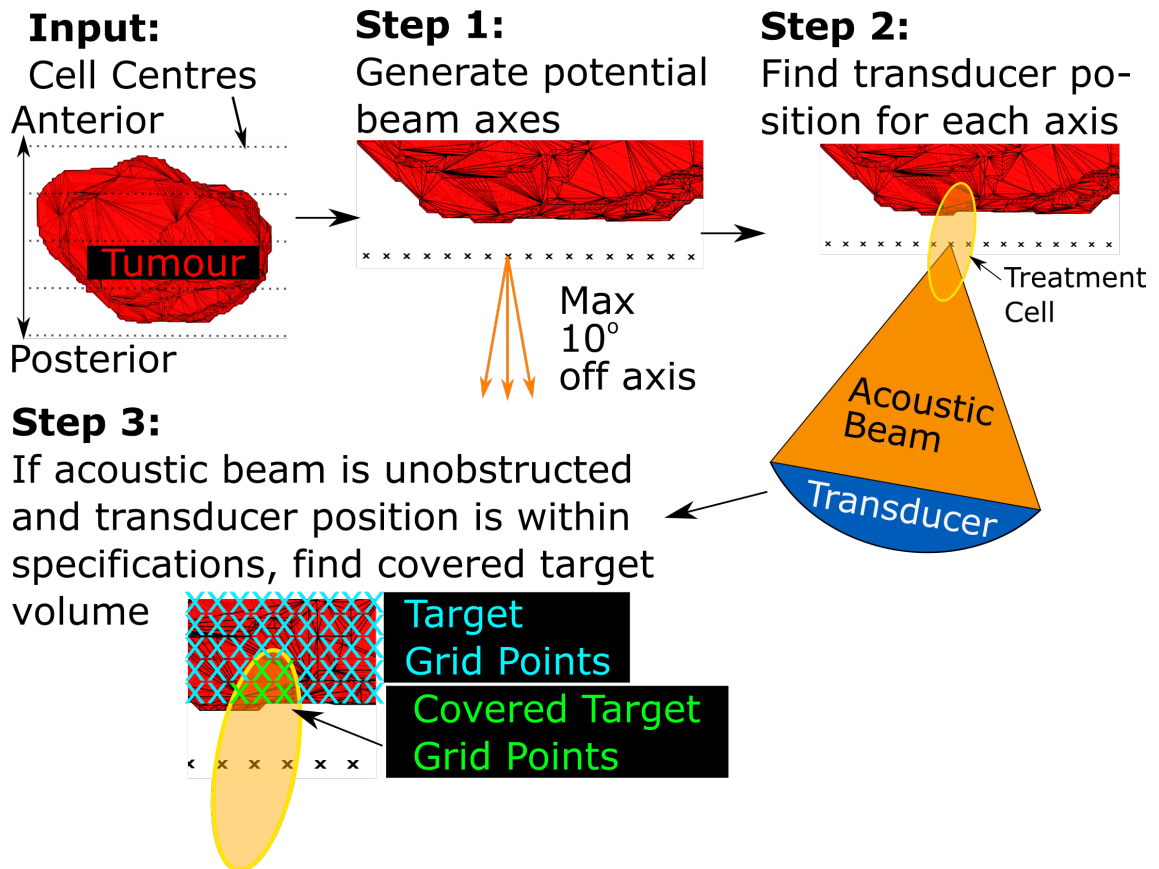
### 7.2.2.3 Calculating Tumour Coverage

The workflow for calculating tumour coverage from the candidate cell centres is shown in Figure 7.3. A virtual MRgHIFU transducer was created, with the same specifications as the virtual MRgHIFU transducer used in the device-first method (see Section 6.2.2.2). Similarly to the device-first method, the virtual transducer was only allowed to angle in the Left-Right direction, with angular increments of  $2.5^\circ$  up to a maximum angle of  $10^\circ$  away from the Anterior-Posterior axis. From the MRgHIFU home position, the transducer was allowed to move up to 72.5 mm away from the home position in the left, right, inferior and superior directions, up to 34 mm in the anterior direction and 33 mm in the posterior direction. Because

the input images were treatment datasets, the home position is known to be 140 mm posterior of the magnetic isocentre. Unlike in the device-first method, however, the acoustic beam was now represented as a series of lines connecting the geometric focus with each of the transducer elements. Intersection of the acoustic beam with OAR or acoustic obstruction meshes was assessed using a ray-triangle intersection algorithm [170], which was implemented for GPUs [173]. This choice was made in order to speed up intersection checking.

In step 1, a set of vectors were generated from each of the candidate cell centres. These vectors represent beam axes, and were generated based on the allowed angulation of the virtual transducer. Hence, the vectors were generated with angulation in the Left-Right direction only, with each vector being separated from other vectors by  $2.5^\circ$ , and with the maximum angulation away from the Anterior-Posterior axis being  $10^\circ$ . In step 2, for each beam axis, the associated transducer position was derived by assuming that the geometric focus was at the cell centre, and then 'back-projecting' the transducer position. In step 3, the transducer position was first checked to see if it exceeded the translation limits away from the home position defined in the device specifications above. If the transducer position was within device specifications, the acoustic beam was tested for obstruction with any OAR or acoustic obstruction meshes using the ray-triangle intersection algorithm [170]. The acoustic beam was obstructed if any line within the acoustic beam intersected an OAR or acoustic obstruction mesh. If the acoustic beam was not obstructed, and the transducer position fell within device specifications, then an 8 mm diameter treatment cell (the same used in the device-first method) was drawn around the transducer geometric focus and all target grid points within the treatment cell were marked as covered. This was repeated for all candidate cell centres within an iteration.

After all iterations were complete, the total number of covered target grid points, multiplied by the product of the spacing between target grid points, was the total covered tumour volume. Percentage tumour volume covered (%TVC) was calculated (defined in equation 6.2 in Chapter 6).



**Figure 7.3:** Workflow for Step 3 in the target-first methodology. Target grid points (blue crosses) were generated during the meshing stage. Step 1: Each candidate cell centre (black cross) generated vectors pointing towards the posterior with a maximum angulation of  $10^\circ$  off the anterior-posterior axis. Step 2: For each vector, the associated transducer position was calculated, and a virtual transducer generated such that the focus was at the cell centre and the transducer beam axis was aligned with the vector, but in the opposite direction. Step 3: The transducer acoustic beam was checked for obstruction with OAR or acoustic obstruction meshes, and the transducer position derived in Step 2 was assessed for whether device limits must be breached in order to place the transducer there. If the acoustic beam is not obstructed and the transducer position was within specification, all target grid points within an 8 mm treatment cell around the transducer geometric focus was marked as covered (green crosses). The total number of covered target grid points multiplied by the product of the spacing in between target grid points was the covered tumour volume.

#### 7.2.2.4 Processing for the Next Iteration

For each iteration except the final one, results were processed in order to improve computational speed in the following iteration, in which candidate cell centres would be generated with finer spacing, thus packing treatment cells more closely.

Firstly, all target grid points that remained uncovered were identified. The next iteration of candidate cell centres were generated from the Left-Right, Inferior-Superior and Anterior-Posterior extents of these uncovered points (see Section 7.2.2.2), in order to reduce the total number of cell centres being generated and assessed, and hence improve computation time. Secondly, inaccessible regions and 'redundantly accessible' regions were identified, so that the candidate treatment cell centres generated in the next iteration within those regions were removed. Treatment cells within inaccessible regions were automatically assumed inaccessible. Treatment cells within redundantly accessible regions were identified as covering regions of the tumour that had already been covered in the previous iterations. The removal of candidate treatment cell centres reduced the total number of cell centres to assess, and therefore improved computation time. Inaccessible and redundantly accessible regions were generated from candidate cell centres that were found to be inaccessible or redundantly accessible, respectively, and were identified as follows:

1. From the first iteration to the current iteration inclusive, the treatment cell centres that could not be reached (inaccessible), and those that could be reached (accessible), were identified.
2. For each individual iteration, inaccessible and accessible cell centres were transformed into inaccessible and redundantly accessible meshes. This iteration-by-iteration meshing allowed the finely sampled later iterations to retain detail in the mesh.
  - Inaccessible cell centres were meshed using the alpha-shape algorithm [172], with the alpha-radius set to be  $\sqrt{A^2 + B^2 + C^2}$  where A, B and C are the Left-Right, Inferior-Superior and Anterior-Posterior cell packing spacing for the iteration.
  - Redundantly accessible cell centres were defined as the accessible cell centres that were further than 1 alpha-radius away from all inaccessible cell centres. Redundantly accessible cell centres were meshed in

the same way as for the inaccessible cell centres, to create a redundantly accessible mesh.

3. Collectively, the inaccessible and redundantly accessible meshes are termed as the 'exclusion mesh'.

As shown in the workflow (Figure 7.1), cell centres generated (see Section 7.2.2.2) in the next iteration that lie within the exclusion mesh were removed before target coverage was assessed for that iteration.

#### 7.2.2.5 Time Comparison

The required computational time for the device-first method to calculate target coverage (Section 6.2.2.6) was compared with that for the target-first method (Figure 7.1 Steps 2-4 until all iterations were complete). The device-first method was run on a 20-core cluster node (2.6 GHz and 12.8 GB per core). The target-first method was run on a 7-core cluster node (2.6 GHz and 12.8 GB per core) with a single NVIDIA Tesla V100 16GB GPU. The difference in hardware arose from university policy; only a maximum of 7 CPU cores could be requested per GPU used.

### 7.2.3 CPU versus GPU Implementations of the Ray-triangle Intersection Algorithm

In order to assess the expected improvement in computational time that arises from implementing the ray-triangle intersection algorithm [170] on a GPU [173], a comparison with an implementation that used MATLAB-native CPU routines [174] was performed. For this comparison, the treatment dataset of patient G8<sub>2</sub> was arbitrarily selected. The difference in computational speed between the GPU and native implementations was assessed for six sets, each set holding a different number of candidate cell centres (10, 50, 100, 250, 500, 1000). The candidate cell centres were generated randomly around the magnetic isocentre. Each cell centre had a transducer placed directly posterior to it. The computational time required to assess the reachability of all the cell centres using the vectorised implementation was compared to that using the GPU implementation. The vec-



torised implementation was performed on a single 2.6 GHz core; the GPU implementation performed on a Tesla V100 16GB GPU. Uncertainty in computational time was assessed by repeating computation for the 10, 50, 100 cell centre sets until three repeats were obtained, and then calculating the standard deviation in computational time as a fraction of the mean computational time.

## 7.3 Results

### 7.3.1 Comparison with the Device-First Method

The mean absolute difference between the device-first method and the target-first method in %TVC is  $0.31 \pm 0.08\%$  (range: 0.18-0.38%). Results are recorded in Table 7.2.

**Table 7.2:** Comparison of the percentage tumour covered required between the device-first method (DFM) and the target-first method (TFM).

Patient	%TVC (DFM)	%TVC (TFM)
G3	46.6	47.0
G8 <sub>2</sub>	52.1	52.5
G24	50.0	50.3
G27	56.6	56.4
G29	82.5	82.1

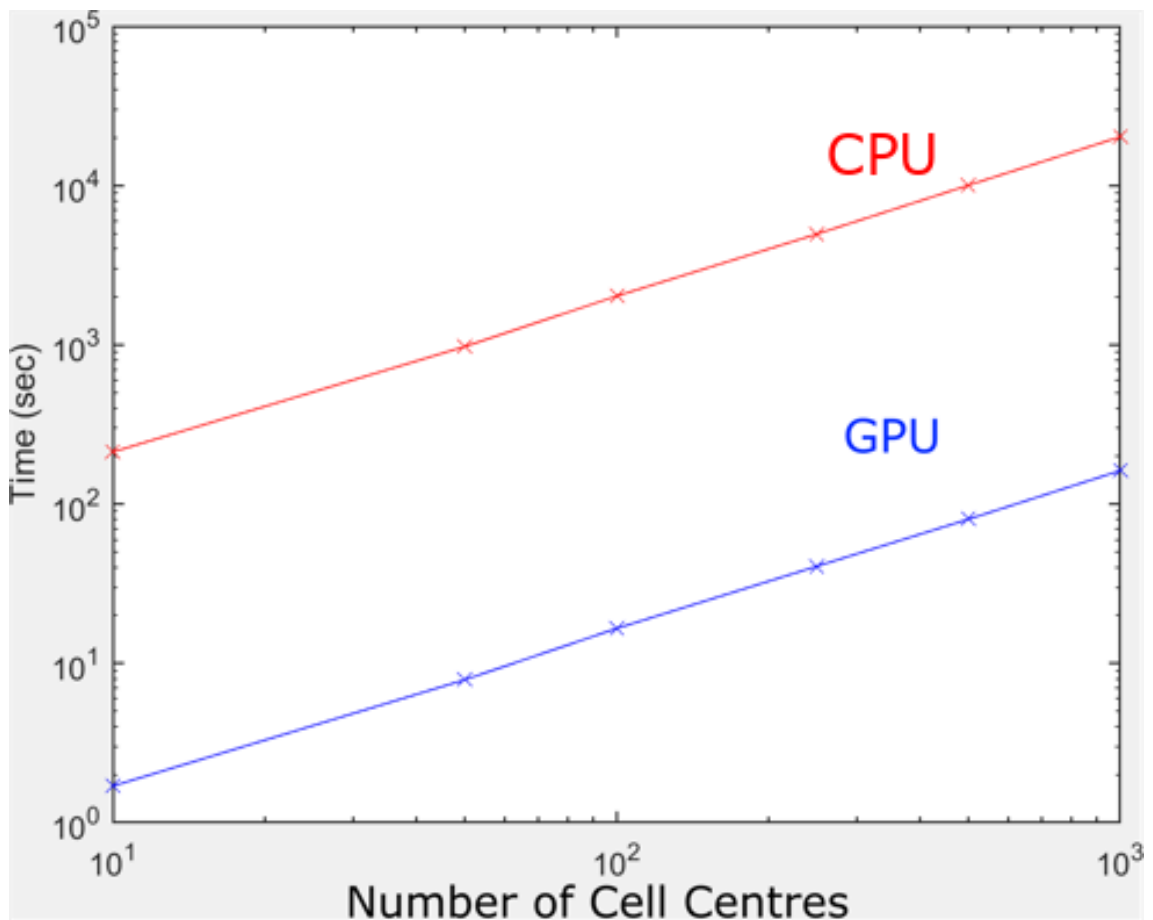
Refinements to the methodology and the use of more sophisticated computing equipment resulted in a decreased computational time for the target-first method, which had a mean of  $2.8 \pm 2.0$  hours (range: 0.3-5.5 hours), compared to the device-first method, which required  $19.7 \pm 8.8$  hours (range: 4-29 hours). On average, the target-first method was 7 times faster than the device-first method. Computational times are shown in Table 7.3.

**Table 7.3:** Comparison of computational time required by the original device-first method (DFM), and by the refined target-first method (TFM).

Patient	Time (DFM) (hours)	Time (TFM) (hours)
G3	4.2	0.3
G8 <sub>2</sub>	24.7	5.5
G24	15.7	0.7
G27	28.6	4.5
G29	25.2	3.2

### 7.3.2 CPU versus GPU Implementations of the Ray-triangle Intersection Algorithm

The computational time required by the GPU implementation, and that required by the vectorised implementation, for different numbers of cell centres is shown in Figure 7.4. The GPU implementation is approximately 120 times faster than the vectorised implementation for each set of cell centres tested. The uncertainty in computational time was calculated to be approximately 1%.



**Figure 7.4:** Computational time required when using the GPU implementation (blue) or the MATLAB-native CPU implementation (red) of the ray-triangle intersection algorithm when assessing whether increasing numbers of candidate cell centres are reachable.

## 7.4 Discussion

The maximum difference between the tumour coverage estimated using the refined target-first method and that using the original device-first method is less

than 0.4%, which is less than the 1.7% discretisation error (Section 6.3.4). This is also less than the 4% associated with image segmentation uncertainty. This demonstrates that, given the same conditions, the methods yield the same tumour coverage.

The computational speedup associated with using the GPU implementation instead of the vectorised implementation was a factor of 120, so even though the vectorised implementation could be parallelised across 20 CPU cores (as was the case in the device-first method), an estimated speedup by a factor of 6 could be predicted. This is shown in the computational time results (Table 7.3), with a mean speedup of approximately 7 times (range: 4.5 times for G8<sub>2</sub> to 23.6 times for G24). The variation in speedup may be due to many factors. Increasing the size and shape complexity of the OARs and acoustic obstructions result in their meshes consisting of a greater number of triangles, and therefore more time may be required for intersection checking. Another reason for the variation in speedup amongst patients may be related to the position of the tumour. The back-projected transducer position was checked initially for whether it was within device specifications, which was relatively computationally efficient compared to the acoustic beam intersection checking. Hence, if a tumour was positioned far from the home position, candidate cell centres could be judged to be inaccessible based on exceeding device specifications, without computationally intensive intersection checking. Computational overhead, such as transferring data from RAM memory to GPU memory and back, could have reduced the overall computational speedup. Computational processing that did not involve ray-triangle intersection checking, such as creation of the exclusion mesh and reducing the number of cell centres, was done on the CPU, and therefore may have reduced the overall computational speedup.

## 7.5 Conclusions

The target first method yields the same results as the device-first method, with a mean difference of 0.3% in estimates of tumour coverage, with the former method being seven times faster than the latter. The target-first method can therefore be

substituted for the device-first method in any prospective workflow.

## Chapter 8

# Identification of the Ideal Treatment Angle

### 8.1 Introduction

In order to deliver the best possible HIFU treatment, it is critical to position the patient in the position that maximises acoustic access and thus energy delivery to the tumour target [18, 19]. Selection of this ideal treatment position is an important part of the proposed prospective patient workflow (Figure 1.1 Step 2) for determining a patient's suitability for MRgHIFU treatment.

Methods for identifying the ideal treatment positions for ultrasound therapy have appeared in literature. As mentioned in Section 1.4 and 2.5.1, McGough et al. [120] defined the ideal patient position for ultrasound hyperthermia as that in which the greatest number of transducer array elements had unobstructed acoustic access to prostate cancer targets. In 2016, Scherrer et al. [28] proposed that the ideal treatment position for HIFU ablation was the position in which the angle of a cone emitted from the tumour centroid was maximised, under the constraint that no organs at risk (OARs), such as rectum, or acoustic obstructions, such as bone, lay within, or intersected, the cone. Scherrer et al. showed no results for this. In essence, they assumed that maximising the cone angle, which served as a proxy for the acoustic window, would maximise tumour coverage (the percentage of the tumour that could be covered by the MRgHIFU focus) and tumour treatability (the percentage of the tumour that could be ablated). In 2018, Abbas

et al. [159] proposed – again, without justification or validation – that the ideal patient position was that which gave the shortest acoustic path within the tissue.

In this chapter, a patient's treatment position is represented by the 'treatment angle', which is defined as the angle at which the patient is rotated away from supine around their body axis when positioned for treatment. Here, the ideal treatment angle is defined as the treatment angle at which tumour coverage is maximised. In the proposed prospective patient workflow, after the ideal treatment angle has been identified, the patient is expected to be placed at that angle with the tumour centroid on the isocentre line (the line going through the magnetic isocentre and transducer home position). The next steps are to calculate tumour coverage and treatability. The aim of the work described in this chapter was to develop a method for identifying the ideal treatment angle. Furthermore, because of the challenges associated with achieving precise three-dimensional patient placement in clinical practice, the tumour coverage at a range of angles (steeper and shallower than the ideal) was examined, in order to assess the acceptable margin of error when placing the patient.

This aim was achieved by first implementing the method suggested by Scherrer et al. to rapidly identify an initial 'starting' treatment angle from patient referral images. In order to investigate whether this starting treatment angle gave the maximum tumour coverage, the starting treatment angle, and angles steeper and shallower than that, were transformed into treatment positions and the tumour coverage was calculated for each position. In order to identify the ideal treatment angle, and to assess the effect of imprecise patient positioning, tumour coverage was plotted as a function of angle. Finally, for each patient, the predicted (starting and ideal) treatment angles were compared to the clinical treatment angle, which represented a compromise between the clinically judged 'optimal' treatment angle and the practical limitations that prevented the placement of the patient at that angle.

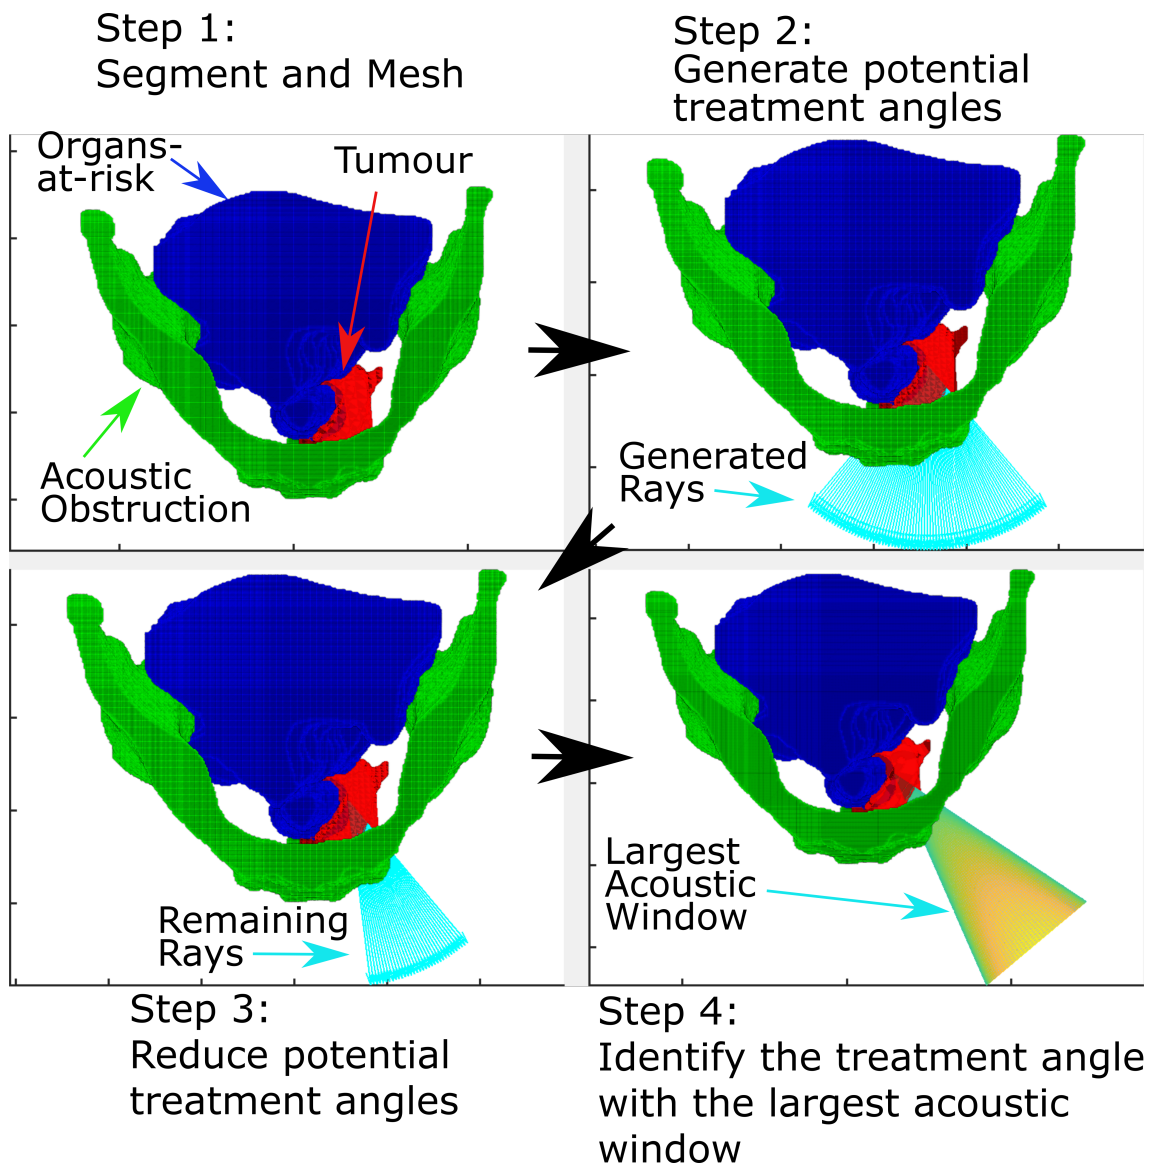
## 8.2 Method

### 8.2.1 Calculation of the Starting Treatment Angle

The starting treatment angle provides a starting point from which the ideal treatment angle can be identified. The methodology for calculating this starting treatment angle is outlined in Figure 8.1. In step 1, organs at risk (OARs), bone (an acoustic obstruction), and the tumour were segmented from input patient 'referral' images, using the methods described in the Image Segmentation chapter (Chapter 5). This was necessary because ultrasound propagation through OARs is considered unsafe, and bone is acoustically opaque. The tumour was segmented as the target. The body outline was segmented for use in the various processes required to calculate the ideal treatment angle, including the calculation of tumour coverage. As in Section 7.2.2.1, these segments were meshed in order to allow the use of rapid geometric intersection algorithms, as described later. In step 2, vector rays representing a set of treatment angles and positions, originating from the tumour centroid and all lying within an axial plane perpendicular to the inferior-superior axis, were generated up to a maximum angle away from the patient anterior-posterior axis. One of these rays represents the starting treatment angle. Details are given in Section 8.2.1.2 below. In step 3, any rays that intersect any OARs or bones were removed, as ultrasound cannot safely propagate through OARs or bone to the tumour. Intersections were checked using a ray-triangle intersection algorithm [170], as was done in a previous study (Section 7.2.2.3). Rays for which the tumour centroid-to-skin distance was greater than the MRgHIFU transducer focal length (140 mm) were removed, because in those treatment positions the focus cannot reach the tumour. Furthermore, reducing the total number of rays to process improves the computational speed of step 4. In step 4, acoustic windows were estimated for all rays in order to identify that with the largest acoustic window, which was defined as the starting treatment angle. Estimation of the acoustic window size for each treatment angle was carried out, because it is faster to calculate the acoustic window (on the order of minutes for each treatment angle) than the tumour coverage (on the order of hours to days

for each treatment angle) or treatability (on the order of days to weeks), according to preliminary experiments. The acoustic window was approximated using Scherrer et al.'s method [28]: a cone, with its origin at the tumour centroid (same as the ray origin) and with the ray being the cone axis, was created for each ray. Then, for each ray, the cone was widened symmetrically around the axis until it met an OAR or bone segment. This determined its maximum size. Intersection between the cone and OARs or bones was checked using a cone-triangle intersection algorithm [175]. The resulting cone gave an approximate representation of the acoustic window. The treatment angle with the largest cone was chosen as the starting treatment angle. Details for steps 1 and 4 are contained in the sections below. The code was run in MATLAB R2018b on a 14 core cluster node (2.6 GHz and 12.8 GB of RAM per core) alongside 1 NVIDIA Tesla V100 16GB GPU.






**Figure 8.1:** Workflow for calculating the starting treatment angle. Step 1: Tumour, organs at risk and bone (an acoustic obstruction) are segmented and meshed. Step 2: A set of vector rays, each representing patient treatment angles, are generated within 3D space. All rays originated from the tumour centroid, are in an axial plane and are angled at most  $40^\circ$  away from the vertical. Step 3: rays that intersect any organs at risk and acoustic obstructions, and have an origin-to-skin intersection distance greater than the MRgHIFU focal length (140 mm), are removed to improve computational performance. Step 4: For all remaining rays, the acoustic window is estimated by fitting a cone (originating at the tumour centroid, with the cone axis coincident with the ray). The treatment angle associated with the ray that has the largest cone is used as the starting treatment angle.

### 8.2.1.1 Meshing

As described previously (Section 7.2.2.1), segments were obtained as three-dimensional binary voxel arrays. Segments were transformed into triangulated surface meshes in order to allow use of the ray-triangle and cone-triangle intersection algorithms involved in Figure 8.1 Step 3 and 4, and later, to allow the calculation of tumour coverage. OARs and bones were meshed using the 'device-first' meshing method described in Section 7.2.2.1. This meshing method was used for meshing the OARs and acoustic obstructions when the target-first method was compared with the device-first method, with results that demonstrated consistency between the two methods under the same initial conditions (see Section 7.3.1). The tumour had already been meshed as a result of the segmentation process, using the marching cubes isosurface method [151, 152] (see Section 5.2.6).

The body outline was meshed using the marching cubes isosurface method as well, but this time with the isovalue set to  $10^{-6}$ . The marching cubes method was used for the body outline because the device-first meshing method produced staircased mesh surfaces (red line in Figure 8.2). This is thought to make identification of the skin points (where the vector rays intersect the surface of the body outline) more difficult, and hence introduced error into the identification of the MRgHIFU transducer home position and the calculation of the tumour coverage. This parameter was chosen in order for the mesh to delineate the outermost surface of the skin. The difference in the surface mesh produced by these two meshing schemes is shown for an example segment array in Figure 8.2.



**Figure 8.2:** A visual comparison of the mesh surfaces of two meshing methods applied to the same example segment, represented here by a two dimensional binary pixel array. Each pixel in the segment array is either '1' or '0', representing the segment tissue/material, or everything else, respectively. The mesh surface generated by the device-first meshing method (red, described in Chapter 7), which was used for OARs and bones, is compared to that generated by the marching cubes isosurface meshing method (cyan), which was used for the body outline.

### 8.2.1.2 Generation of Treatment Angles

Rays, each representing a different treatment angle, were generated. All treatment angles were angles of rotation around the inferior-superior axis, because even though the patient meshes were three dimensional, the patient was expected to be able to be rotated around that axis only. In later processes, one of these treatment angles will be identified as the starting treatment angle. Each ray represents the isocentre line (the line that would connect the magnetic isocentre and the transducer home position if the patient were actually oriented into that angle on the MRgHIFU bed) of a particular treatment position, with the associated treatment angle being the acute angle between the ray and the vertical isocentre line. In essence, if the patient were to be rotated into that particular treatment

position and placed on the MRgHIFU bed, the ray would ideally go through the magnetic isocentre as well as the tumour centroid (see Figure 1.1 Step 2). Rays up to 40° away from the anterior-posterior axis, where the 0° ray represented the supine position, were generated in 1° increments. The choice of 40° represents, in light of clinical experience, the expected maximum angle that a patient can tilt and still fit within the MR scanner bore during treatment.

### 8.2.1.3 Acoustic Window Estimation

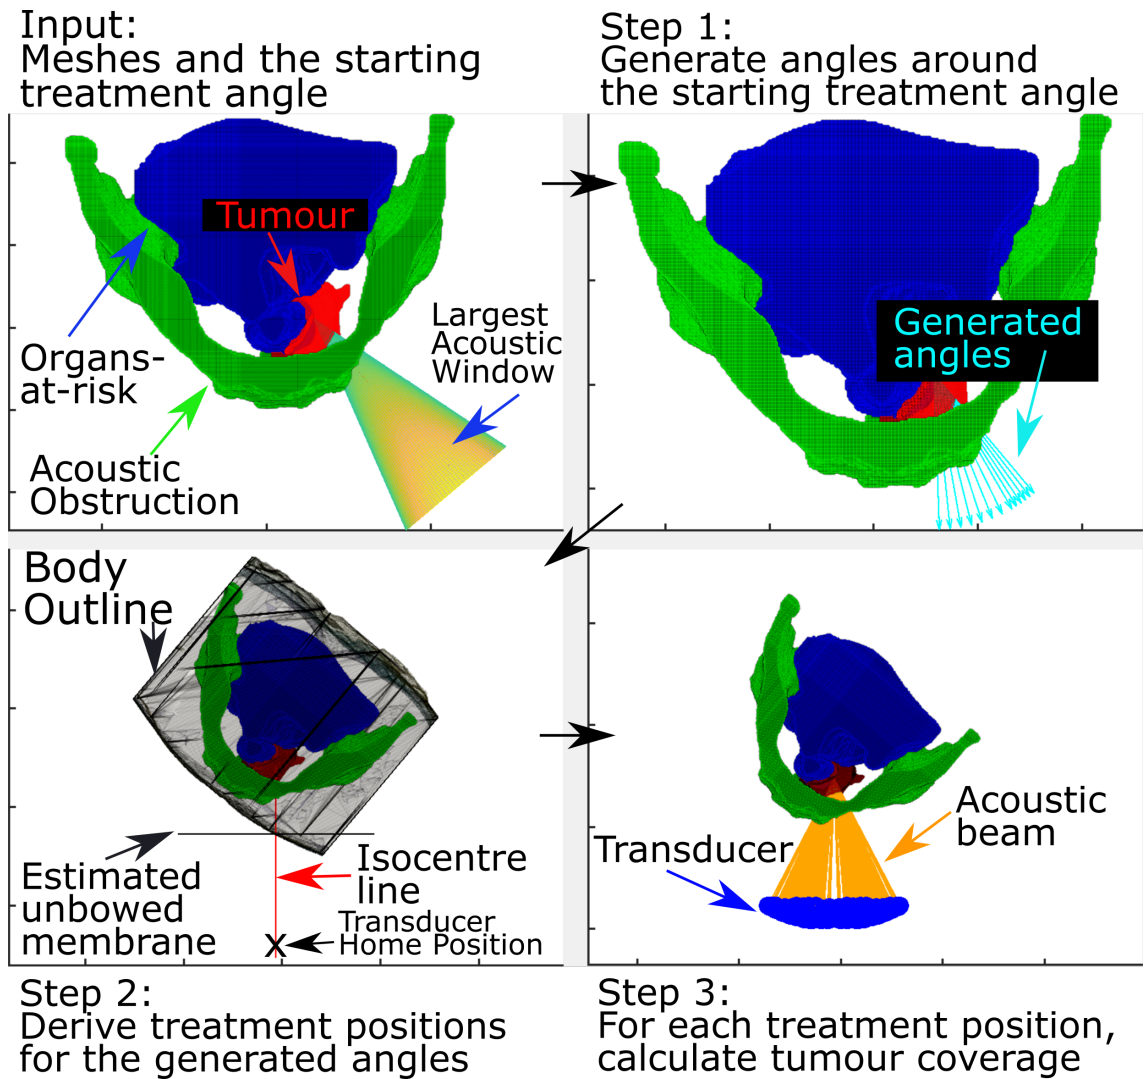
As stated above, the starting treatment angle was defined as that with the largest acoustic window. To identify it, acoustic windows were approximated for all remaining treatment angles by creating a right circular cone for each, originating at the tumour centroid (same as the ray origin), with the cone axis being the ray associated with that angle (see Figure 8.1 Step 4). The cones were widened by increasing their opening angle (the angle at the cone origin between the cone axis and the cone's curved sides), since the opening angle controlled the width of the cone base relative to the cone height. This was done using a custom-built optimiser. The cone opening angle was initialised at 0°, and was increased in initial steps of 45 degrees. For each increment, the cone was checked for intersection with any OAR or acoustic obstruction meshes using a cone-triangle intersection algorithm [175]. If intersection occurred, the cone opening angle was reset to that of the previous step, and the step size was halved. This was repeated until the opening angle increment was less than  $6 \times 10^{-5}^\circ$ . This limit was chosen because it was the minimum difference between cone opening angles observed for different treatment angles in preliminary experiments. An upper bound of 90° (where the cone becomes a plane) was enforced on the cone opening angle, as the cone-triangle intersection algorithm only functioned for cones with acute opening angles. The treatment angle associated with the cone that had the largest opening angle was chosen as the starting treatment angle.

## 8.2.2 Identifying the Ideal Treatment Angle

The ideal treatment angle is defined as that which results in the maximum tumour coverage. Identification of the ideal treatment angle was critical for the assess-

ment of patient suitability, and a major aspect of the prospective patient workflow which underpins this PhD project. The methodology used to identify the ideal treatment angle was also used to test whether the tumour coverage was maximal at the starting treatment angle, for which the approximate acoustic window was maximal. Furthermore, the process of identifying the ideal treatment angle allowed assessment of the variation in tumour coverage with angle and thus the effect of imprecise patient positioning.

An overview of the method developed for identifying the ideal treatment angle is shown in Figure 8.3 and described below. The inputs are the OAR, bone, tumour and body outline meshes, and the starting treatment angle. These were calculated as described in Section 8.2.1. In Step 1, a set of treatment angles was generated around the starting treatment angle, by generating rays originating from the tumour centroid within the same axial plane. As before, each ray represents the isocentre line of a particular treatment position, and was associated with a treatment angle. Details were provided in Section 8.2.2.1 below. In Step 2, for a particular treatment angle, the position of the MRgHIFU transducer relative to the tissue volume was found. First, the position of the associated ray's intersection with the skin, represented by the body mask mesh surface, was found. Secondly, because the ray is assumed to be the isocentre line for a particular treatment position, volunteer-derived estimates for the compressed gel-pad thickness and membrane bowing distance were used to estimate the position of the unbowed membrane, and from that the MRgHIFU transducer home position was estimated as described in Section 6.2.2.5 and Figure 6.4. In Step 3, the patient's tumour coverage for each treatment angle was calculated using the methodology described in Chapter 7. The tumour coverage was plotted against the treatment angle, in order to identify the ideal treatment angle, to identify how the tumour coverage would decrease as the patient is tilted away from the ideal treatment angle, and to compare between the starting, ideal and clinical treatment angles. Computation was performed on a 14-core cluster node (each core clocked at 2.6 GHz and with 12.8 GB of RAM) with a single NVIDIA Tesla V100 GPU.



**Figure 8.3:** Workflow for identifying the ideal treatment angle. Inputs are the body outline, tumour, organ-at-risk (OAR) and acoustic obstruction (i.e. bone) meshes, as well as the starting treatment angle. For ease of visual interpretation, the body outline is not shown except in Step 2, where it is used. Step 1: The ray representing the isocentre line of the starting treatment angle ('starting ray') is generated. Rays are generated perpendicular to the inferior-superior axis around the starting ray, each ray representing the isocentre line of a particular treatment angle. Step 2: for each ray, the transducer home position is identified using the techniques developed in Section 6.2.2.5. Each ray's intersection with the body mask mesh is found and the intersection point defined as the skin point. From the skin point, estimates of the compressed gel-pad thickness (9.8 mm) and the membrane bowing distance (10.0 mm) are used to obtain the estimated unbowed membrane, from which the transducer's home position is known to be a fixed distance (67.5 mm for the Sonalleve V2) posterior. The patient's meshes are positioned and rotated such that the tumour centroid, the skin point, the magnetic isocentre and the transducer's home position all lie on a single 'isocentre line' (red). Step 3: for each treatment angle, the tumour coverage is calculated. Ultimately, tumour coverage is plotted against treatment angle.

### 8.2.2.1 Generation of Potential Treatment Angles

In a similar fashion to the initial stages of the calculation of the starting treatment angle, a set of treatment angles was generated by representing them as vector rays originating from the tumour centroid, from which one would be identified as the ideal angle. These treatment angles were generated around the starting treatment angle in  $2^\circ$  increments up to  $10^\circ$  from the starting treatment angle, then in  $4^\circ$  increments beyond that. The smaller increments around the starting treatment angle were chosen because the ideal treatment angle is expected to be close to the starting treatment angle. The treatment angles were generated until they reached a lower limit of being supine (defined as  $0^\circ$ ), or an upper limit of a  $40^\circ$  tilt away from supine. The rationale for the  $40^\circ$  upper limit is the same as that for the starting treatment angle, namely that the patient was expected to be unable to fit within the MR scanner bore in that position. For the patient data used in this study,  $33^\circ$  was the greatest tilt observed (see Table 3.3).

### 8.2.2.2 Potential Treatment Positions

For each treatment angle, the MRgHIFU transducer home position was identified. This is required in order to identify the tumour coverage that can be achieved at that treatment angle. The MRgHIFU transducer was modelled on the Sonalleve<sup>®</sup> V2 transducer (Section 3.2), and its home position was known to be 67.5 mm posterior to the undeformed membrane (see Figure 3.1). The assumption was made that, in clinical practice, the patient would be angled at the treatment angle and positioned such that the tumour centroid, the magnetic isocentre and the transducer home position would be on the same anterior-to-posterior line ('isocentre line', as seen in Section 1.5). Given that assumption, the left-right and inferior-superior position of the transducer home position is known to be the same as that of the magnetic isocentre, and it is only the anterior-posterior placement of the transducer home position that is in question. This was identified using the method described in Section 6.2.2.5 and Figure 6.4. The intersection point between the isocentre line and the skin, represented by the body outline mesh, was found. The anterior-posterior position of the unbowed membrane was iden-

tified from the skin intersection point by using volunteer-derived estimates of the compressed gel-pad thickness and membrane bowing distance (9.8 mm and 10.0 mm, respectively, as in Section 6.3.1), with the exception of patient G24, who was treated on a customised gel-pad. For patient G24, the actual gel-pad thickness of 5.3 mm was used along with the volunteer-derived estimated membrane-bowing distance of 10.0 mm. The transducer home position was known to be 67.5 mm posterior to the estimated unbowed membrane (Section 3.2).

### 8.2.2.3 Calculation of Tumour Coverage

For each treatment angle, the tumour coverage was calculated using the method described in Section 7.2.2. The MRgHIFU transducer was modelled such that each element was represented by a point at its centre, and the acoustic beam was modelled as discrete lines connecting the transducer element points to the transducer focus. To be consistent with the previous work, the transducer was restricted to angulation only in the left-right direction, in  $2.5^\circ$  increments up to  $10^\circ$  off the anterior-posterior axis.

After tumour coverage had been calculated for all treatment angles, the ideal treatment angle, defined as that which results in the greatest tumour coverage, was identified. The starting treatment angle was compared to the ideal treatment angle to see whether the maximisation of the acoustic window size, as was done to derive the starting treatment angle, results in maximisation of the tumour coverage. Comparison between the ideal and the clinical treatment angles quantifies the quality of patient positioning provided by the computational method versus current clinical practice. An assessment of the effect of imprecise patient positioning, which is expected to be  $5^\circ$  in clinical practice, was performed for each patient by examining the tumour coverage at treatment angles  $5^\circ$  steeper and shallower than the ideal. The two resulting tumour coverage values are compared to the maximum tumour coverage at the ideal treatment angle, and the greatest difference between them is used to quantify the effect of imprecise patient positioning for that patient.



## 8.3 Results

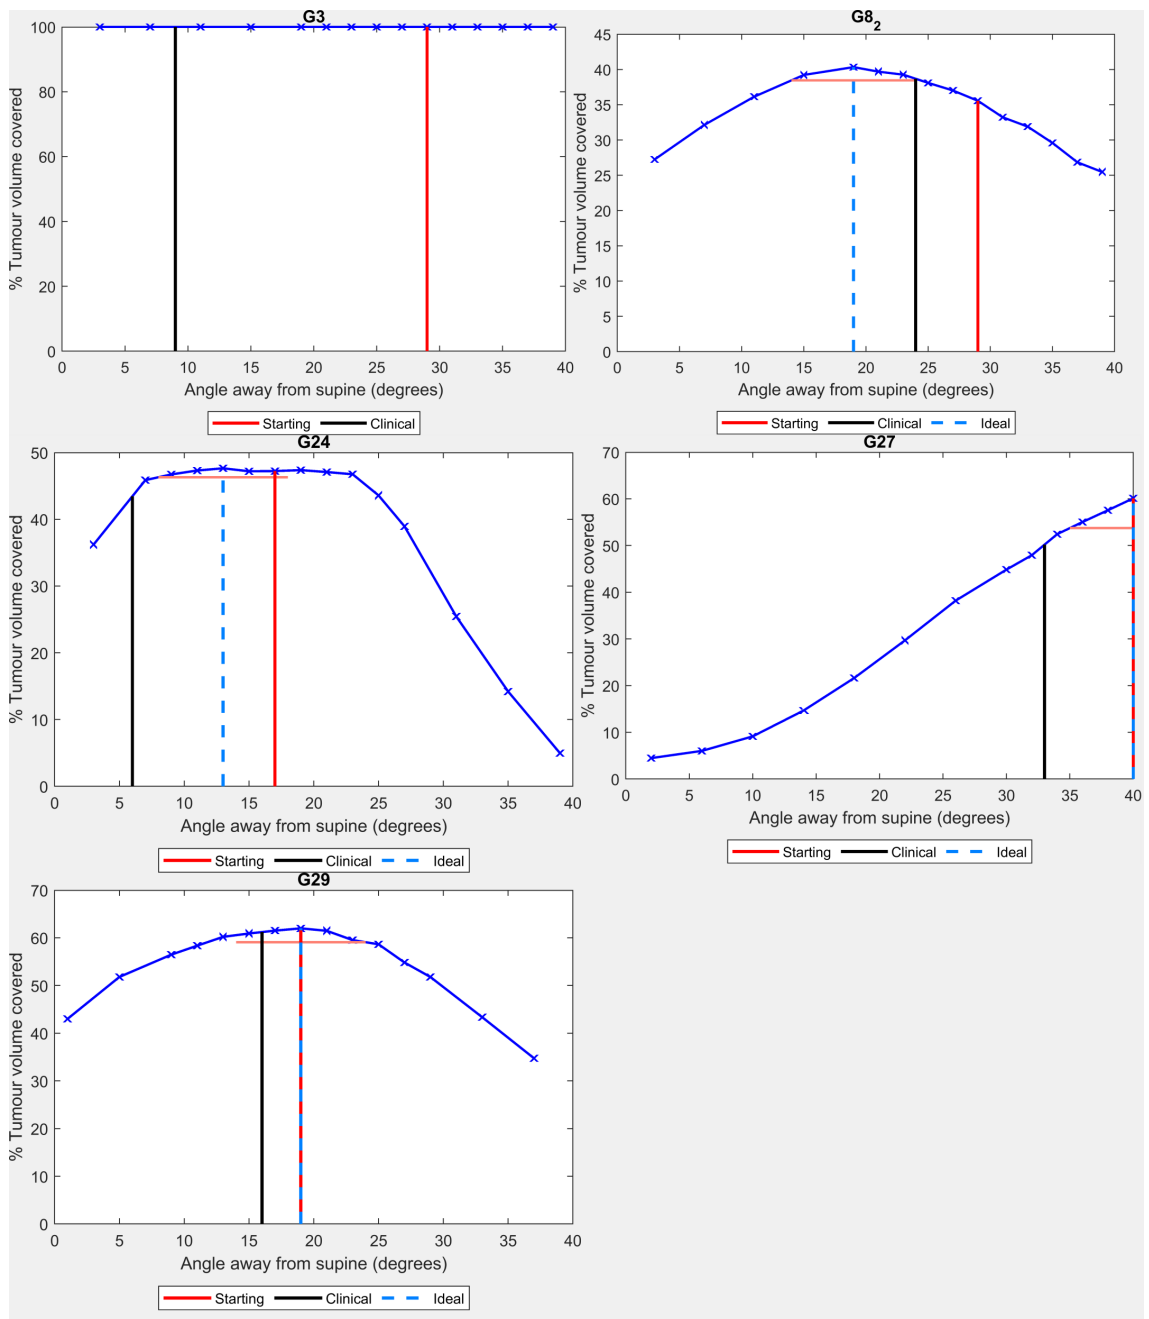
### 8.3.1 Starting Treatment Angle

A comparison of the starting treatment angle for each patient against the clinically used treatment angles and the ideal treatment angles, as shown in Figure 8.4, showed that the starting treatment angles differed from the ideal treatment angles by  $4\pm 4^\circ$  (range:  $0-10^\circ$ , 4 patients) and from the clinical treatment angles by  $9\pm 6^\circ$  (range:  $3-20^\circ$ , 5 patients). For patient G3, all treatment angles produced the same tumour coverage (100%). The computational time required to identify the starting treatment angle was  $3\pm 1$  hours (range: 2 hours (G27) to 5 hours (G3)).

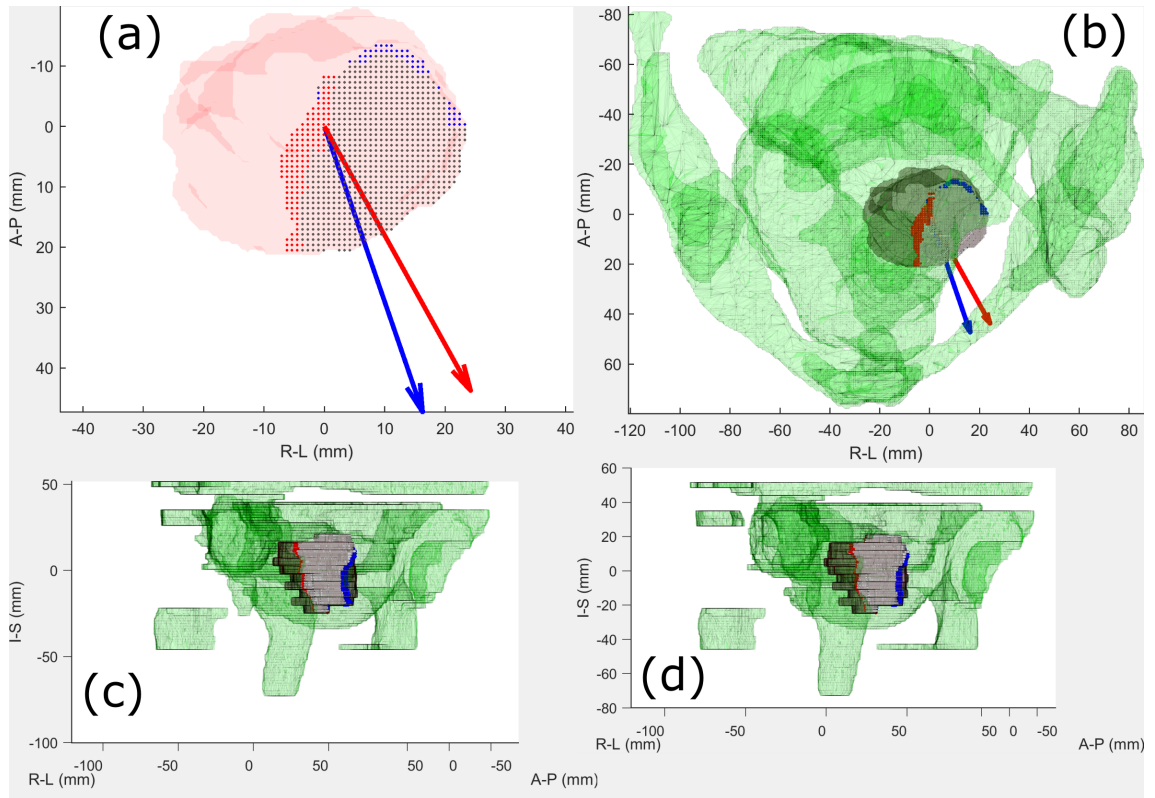
### 8.3.2 Ideal Treatment Angle

The variation in tumour coverage with patient tilt angle is shown in Figure 8.4. The ideal treatment angle differed from the clinical treatment angle by  $5\pm 2^\circ$  (range:  $2-7^\circ$ , 4 patients). Tumour coverage at the starting treatment angle was at, or close to (within 1%), maximum except for patient G8<sub>2</sub>, whose tumour coverage was 5% below the maximum. The tumour coverage for patient G3 was the same (100%) at all angles.

The effect of imprecise patient positioning was a decrease in tumour coverage of  $3\pm 2\%$  (range:  $0-7\%$ ), with the maximum decrease for patient G27. Patient G8<sub>2</sub> was examined in closer detail (Figure 8.5). Computational time required to identify the ideal angle, including the time required to identify the starting angle which initialises the search for the ideal angle, was  $39\pm 28$  hours (range: 5 (G3) hours to 78 hours (G8<sub>2</sub>)).



**Figure 8.4:** The variation of tumour coverage with treatment angle, for all patients. The starting treatment angle (red vertical), in which the approximate acoustic window was maximal, the ideal treatment angle (blue dashed vertical), in which the tumour coverage is maximal, and the clinical treatment angle (black vertical), which was used in treatment, are shown. Tumour volume coverage at angles 5° steeper and shallower than the ideal treatment angle were calculated, and the lower of the two was depicted by the horizontal pink line. For patients G27 and G29, the starting and ideal treatment angles are the same. For patient G3, tumour coverage was identical for at all angles tested.



**Figure 8.5:** Difference in estimated tumour volume covered between the starting treatment angle (red arrow), and the ideal treatment angle (blue arrow), for patient G8<sub>2</sub>. (a): Tumour volume is represented as discrete points; those covered by both the starting and ideal angles are grey, by the starting angle only were red, and by the ideal angle only were blue. The tumour itself is represented as the red mesh in (a) and is grey in the other subplots. (b): An axial view of the tumour within the pelvis showing covered volumes relative to the acoustic obstructions and organs-at-risk (green meshes). Views of the covered tumour volumes from (c) the ideal angle (19° from supine) and (d) the starting angle (29° from supine).

## 8.4 Discussion

The aim of this study was to develop a novel method of identifying the ideal treatment angle, in which the tumour coverage is the greatest, and to assess the acceptable margin of error in patient positioning, as part of a wider study into automating the evaluation of patient suitability for MRgHIFU therapy.

### 8.4.1 Comparison between the Starting and Ideal Treatment Angles

One of the objectives of this study was to investigate whether the treatment angle that maximised the acoustic window, i.e. the starting treatment angle, would be the same as that which maximises the tumour coverage, i.e. the ideal treatment angle. This was of interest because identifying the starting treatment angle requires less computing time ( $3\pm 1$  hours, range: 2-5 hours) than calculation of the ideal treatment angle (including time required to calculate the starting angle,  $39\pm 28$  hours, range: 5-78 hours). For all patients except G3, the ideal treatment angle differed from the starting treatment angle by  $4\pm 4^\circ$  (range: 0-10°). For patient G3, the tumour coverage was the same (100%) at all treatment angles tested, so every angle was "ideal". In comparison, the estimated precision with which a clinician could place a patient was expected to be approximately  $\pm 5^\circ$ . Although the starting angle for patient G27 reached the upper bound of 40°, clinical experience suggests that steeper angles could only be achieved in smaller patients.

The greatest difference between the starting and ideal treatment angles was 10° (patient G8<sub>2</sub>). The hypothesis that maximising the acoustic window size would result in the maximisation of tumour coverage was shown to be valid for 3/5 patients (G3, G27 and G29). The reason for the angular difference for patients G8<sub>2</sub> and G24 was that the starting treatment angle's acoustic window originated from the tumour centroid, so maximisation of the acoustic window just maximised the transducer's ability to reach this one position from a variety of angles, whilst access to the surrounding tissue was not considered. Thus, for some OAR, bone and tumour geometries, a treatment angle that provides less access to the centroid could result in a greater tumour volume coverage overall. Nonetheless, tumour coverage at the starting treatment angles were within 5% of the maximum tumour volume covered for all patients, so the starting treatment angle was shown to be a good initial angle from which to begin the search for the ideal treatment angle. The development of an optimisation-based method to identify the ideal

angle could be pursued, in order to reduce the computational time required for the current method.

#### **8.4.2 Comparison between the Ideal and Clinical Treatment Angles**

The difference between the ideal and the clinical treatment angle was  $5 \pm 2^\circ$  (range:  $0-7^\circ$ , 4 patients). For comparison, the estimated imprecision in patient positioning is  $\pm 5^\circ$ . Since the ideal and clinical treatment angles are so close, this suggests that the current clinical procedures involved in deciding patient placement are reasonably accurate. The clinical treatment angle represents a compromise between the clinically determined 'ideal' treatment angle and the practical considerations that prevent the patient from being positioned at that angle, for example whether the patient could fit in the MR scanner bore.

#### **8.4.3 Comparison between the Starting and Clinical Treatment Angles**

The starting treatment angle differed from the clinical treatment angle by  $9 \pm 6^\circ$ , with a maximum discrepancy at  $20^\circ$  for patient G3. The starting treatment angle ( $27 \pm 8^\circ$ ) appeared to be biased towards recommending a steeper treatment angle than that eventually used clinically ( $18 \pm 10^\circ$ ). This could be because patients could not physically fit within the MR bore at steeper treatment angles, and because patients may have settled stably into shallower angles prior to treatment commencing. Furthermore, whilst the starting angle is decided from fixed anatomy in referral images, the clinical treatment position is decided by clinicians on treatment day after accounting for soft tissue deformation and OAR displacement resulting from patient reorientation into the treatment position, clinical preparations such as bowel-filling and pre-treatment dieting, and biological variation. Clinicians may also have judged that a shallower angle would expose more of the tumour to ablation.

#### **8.4.4 Precision Required in Patient Positioning**

From clinical experience and practice, because patient reorientation into the treatment angle is done manually, it is expected that the patient could not be positioned exactly at the ideal treatment angle. The expected imprecision involved in patient placement is  $\pm 5^\circ$ . The decrease in tumour coverage resulting from this imprecision was  $3 \pm 2\%$  (range: 0-7%), with the maximum decrease for patient G27. By comparison, the average difference between the tumour coverage predicted from referral images and that calculated from treatment images (Section 6.3.3) was  $\pm 12\%$ , assuming perfect acoustic coupling between patient and the gel-pad. The uncertainty involved in calculating tumour coverage due to uncertainties in image segmentation and tumour discretisation was approximately 4.3% (Section 6.3.4). Therefore, imprecise patient positioning does not appear to present a major limitation to the clinical adoption of this method.

#### **8.4.5 Limitations of this Study**

There are several limitations to this study. One is the small patient sample size involved, due to the limited number of patients that underwent treatment. Volunteers could not be used to supplement the limited patient data without arbitrarily deciding on shapes, sizes and positions of dummy tumours. Another is that a sufficiently wide acoustic window does not mean that sufficient energy can be transmitted to the depth of the tumour to achieve tissue ablation – this is the subject of the next chapter. Furthermore, the ability to fit the patient within the MR bore at the ideal angle was only taken into account by setting an upper bound for the treatment angle at  $40^\circ$ . In practice, the larger the patient, the lower the maximum treatment angle achievable; within the available set of patients in this study, the steepest treatment angle achieved clinically was  $33^\circ$  (G27).

A limitation of the identification of the ideal angle was that perfect acoustic coupling was assumed. As a result of this, the tumour coverage for patient G3 was estimated to be 100% at all treatment angles. In clinical practice, extracorporeal air could be a major factor determining the coverage of the tumour as demonstrated in Section 6.3.3. When positioning the transducer home position

relative to the referral dataset, the same assumptions were made here as those made in Section 6.2.2.5, which were that: i) the gel-pad would be most compressed and the gel-pad most bowed at the isocentre line, and ii) after soft-tissue deformation resulting from the patient being placed on top of the gel-pad in the treatment position, the distance between the skin-intersection point (where the isocentre line met the skin) and the tumour centroid stayed constant. Because the same procedure was performed for all treatment angles, any relevant error in the tumour coverage was thought to apply at all angles.

## 8.5 Conclusions

A method for identifying the ideal treatment angle from referral imaging has been developed and tested on patient data for future use in the prospective patient workflow. The method first determines a 'starting' treatment angle, for which the acoustic window was maximised. The ideal treatment angle was then found, for each patient, by exhaustively examining steeper (up to 40°) and shallower (down to 0° i.e. supine) angles than their starting angle. For the five patients studied, the angular difference between the starting and ideal treatment angles was  $4 \pm 4^\circ$  (range: 0-10°), and that between the ideal and clinical treatment angles was  $5 \pm 2^\circ$  (range: 2-7°). The maximum effect of patient positioning up to 5° away from the ideal angle was estimated to be a 7% decrease in tumour coverage, suggesting that imprecise patient positioning does not present a major limitation to the clinical adoption of this method. Tumour coverage at the starting treatment angles was within 5% of the maximum tumour volume covered, which occurred at the ideal treatment angle, suggesting that the starting treatment angle could be used as a faster-to-calculate substitute for the ideal treatment angle. Future work related to identifying ideal patient positioning could be focused on reducing the computational time required to calculate the ideal treatment angle.

## Chapter 9

# Evaluating Patient Treatability

### 9.1 Introduction

Quantitative evaluation of the patient treatability, the topic of this chapter, is the final step in the prospective patient suitability workflow which was outlined in Section 1.5. A patient's treatability is defined as the percentage of their tumour(s) that could experience HIFU-induced coagulative necrosis. In the literature, this is interpreted as regions reaching a temperature threshold of 55°C for at least 1 second [1], or a thermal dose greater than 240 cumulative equivalent minutes at 43°C (CEM43) [24].

Currently, patient treatability is estimated by clinicians, who, as part of treatment planning, place notional treatment cells within and around the target tumour(s) on the screening image dataset, which may not be acquired at the ideal orientation for treatment. Clinicians concentrate on the placement of treatment cells, and assume that enough energy can be delivered from the transducer to each [18]. In the literature, acousto-thermal simulations designed to investigate how the different tissue layers between the transducer and target affect the ultrasound beam intensity and shape have been described [64, 83]. Acoustic simulation has been used in order to predict the acoustic energy that would be deposited within patient tissue. Thermal simulation have been used to predict the resultant temperature changes. Published methodologies for performing acoustic and thermal simulations are described in Sections 2.6 and 2.7, respectively. However, due to the high computational requirements of acoustic simulation (approximately



2 days per simulation [64, 83]), an exhaustive assessment of treatability has not been possible here.

The aim of the work described in this chapter, therefore, was the development of a method to estimate the percentage of a tumour that should be treatable. As discussed in Sections 2.6.4 and 2.7.4, acoustic and thermal simulations were performed using the open-source pseudo-spectral simulation package k-Wave [80]. This software package is publicly available, has been validated [84, 98] and has been used for simulating acoustic propagation in human tissue [64, 83]. This choice of simulation method affects several steps of the method, as detailed in the sections below. A broad overview of the method is provided below in Section 9.2.1.

## **9.2 Development of Simulation Methods**

### **9.2.1 Overview**

Treatability was assessed for the five patients described in Section 3.3.3. The workflow for this study is shown in Figure 9.1, with details for each step in the relevant subsections below. The inputs into the workflow were five patient treatment image datasets, the placements of treatment cells that contributed to their tumour coverage ('reachable' treatment cells), and the region of tumour volume covered (the latter two were derived from tumour coverage assessment, using the target-first method described in Chapter 7). For each patient dataset, the deepest (most anterior) and shallowest (most posterior) reachable treatment cells centered within the tumour were identified, with anterior and posterior defined as shown in Figures 3.1 and 9.2. In Step 1, the region of expected acoustic propagation was identified from the transducer position and angle required to reach the deepest or shallowest reachable treatment cells, and extracted from the image dataset in order to reduce the grid size for the upcoming simulations and therefore improve computational speed. In Step 2, each voxel of the simulation grid was assigned values for relevant mechanical (density), acoustic (speed of sound, attenuation coefficient) and thermal (thermal conductivity, specific heat capacity)

properties. In Step 3, acoustic simulation was performed using a pseudo-spectral method (source acoustic power 300 W, source frequency 1.22 MHz) to calculate the acoustic pressure fields resulting from sonication of a 4 mm treatment cell. The 4 mm cell was chosen because it was the treatment cell size most commonly used for treating patients in the recurrent gynaecological tumour trial whose patient data is used here [18]. In Step 4, a thermal simulation was performed using a pseudo-spectral method to identify the temperature history, and thermal dose received, in the tissue as a result of acoustic energy absorption over the duration of cell heating and cooling. The ablated tissue volume, defined as tissue which had received a cytotoxic thermal dose  $\geq 240$  CEM43, was identified for both the deepest and shallowest reachable treatment cells. In Step 5, the volume of ablated tissue along the line between the deepest and shallowest ablated tissue centroids was estimated by linear interpolation. The 'cell-equivalent' position on this line, i.e. that at which the volume of ablated tissue matches the treatment cell volume, was identified. The maximum treatable depth was defined, arbitrarily, as half a treatment cell length deeper into the patient than this cell-equivalent position. The treatable tumour volume was identified as the covered tumour volume that lay closer to the skin than this depth.

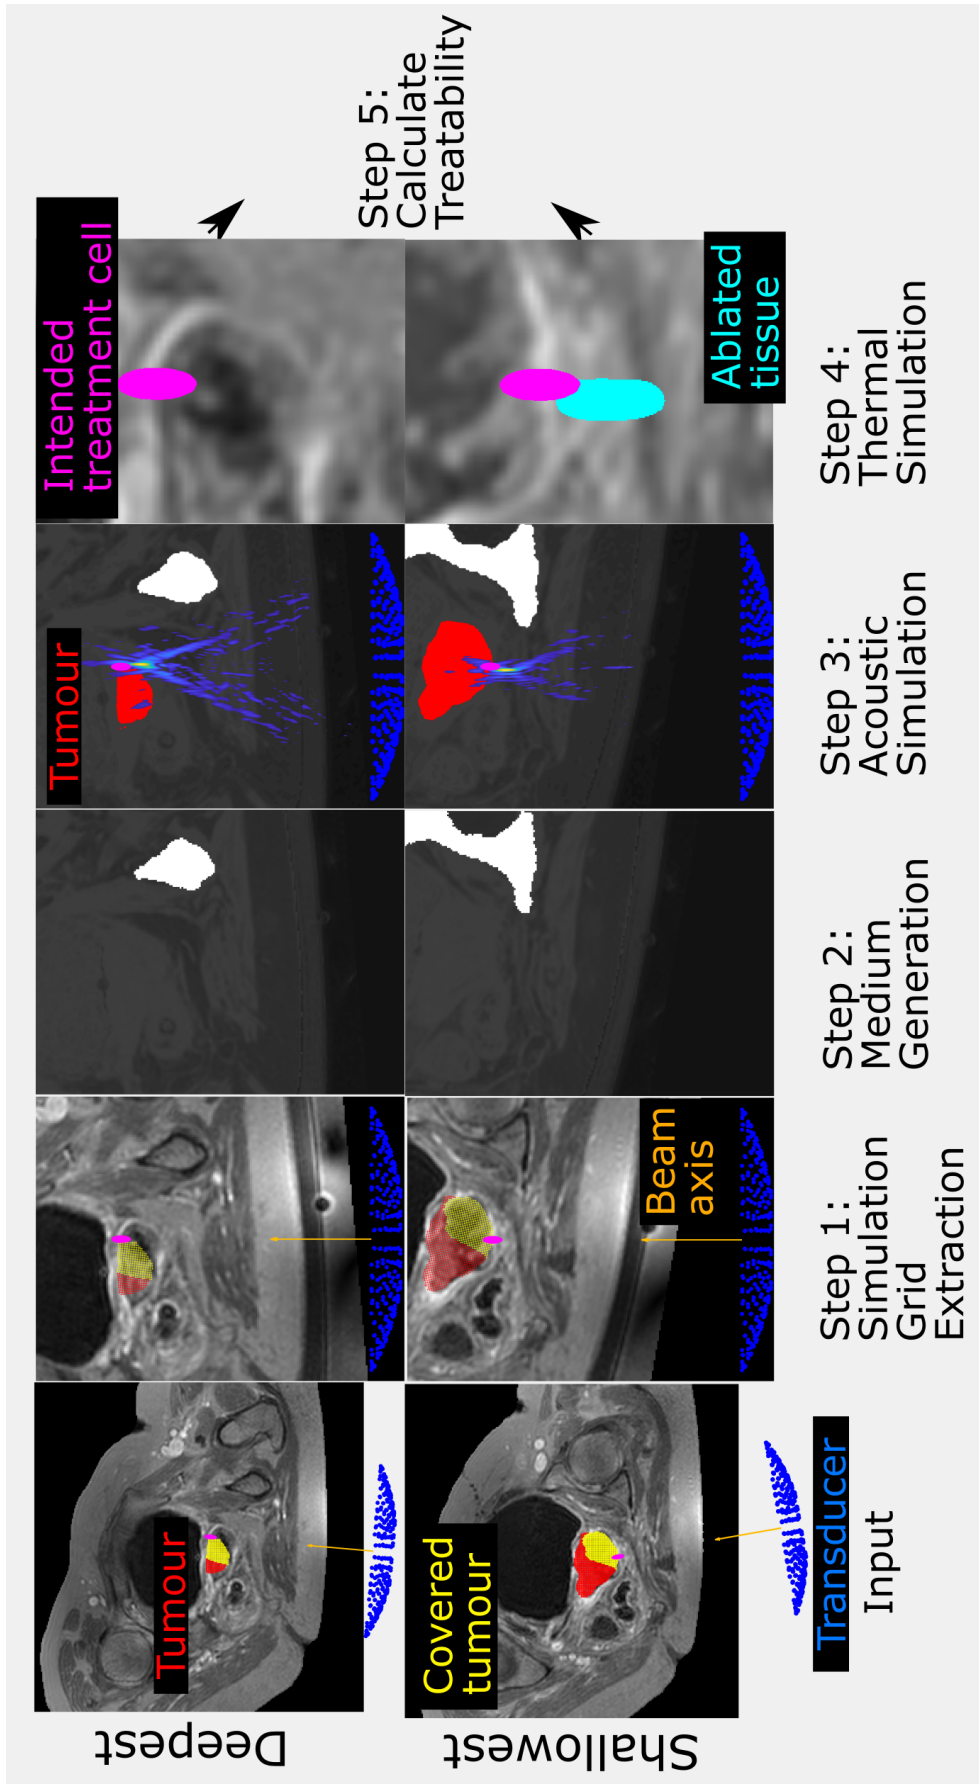


Figure 9.1: Schematic for the methodology used to quantify patient treatability. Inputs are the patient images (representative treatment image slices shown here), the deepest and shallowest reachable treatment cells (magenta), the covered tumour volume (yellow) and the entire tumour volume (red). The transducer (blue) was positioned and angled such that the geometric focus was placed at the reachable treatment cell centre. In step 1, the simulation grid was extracted from the larger dataset. In step 2, each voxel of the extracted region was assigned acoustic and thermal properties based on voxel intensity. A density map is displayed, with white denoting the highest density and black denoting the lowest density in the image. In step 3, an acoustic simulation is performed in order to estimate the acoustic pressure field in the tissue. The pressure field, overlaid on the density map and tumour, shows regions of high (yellow) and low (blue) pressure. In step 4, a thermal simulation is performed to identify the ablated tissue resulting from acoustic energy absorption and heat transfer. In step 5, the ablated tissue volumes (cyan) from sonicating the deepest and shallowest reachable treatment cells were used to estimate the maximum treatable depth and the proportion of tumour treatable.

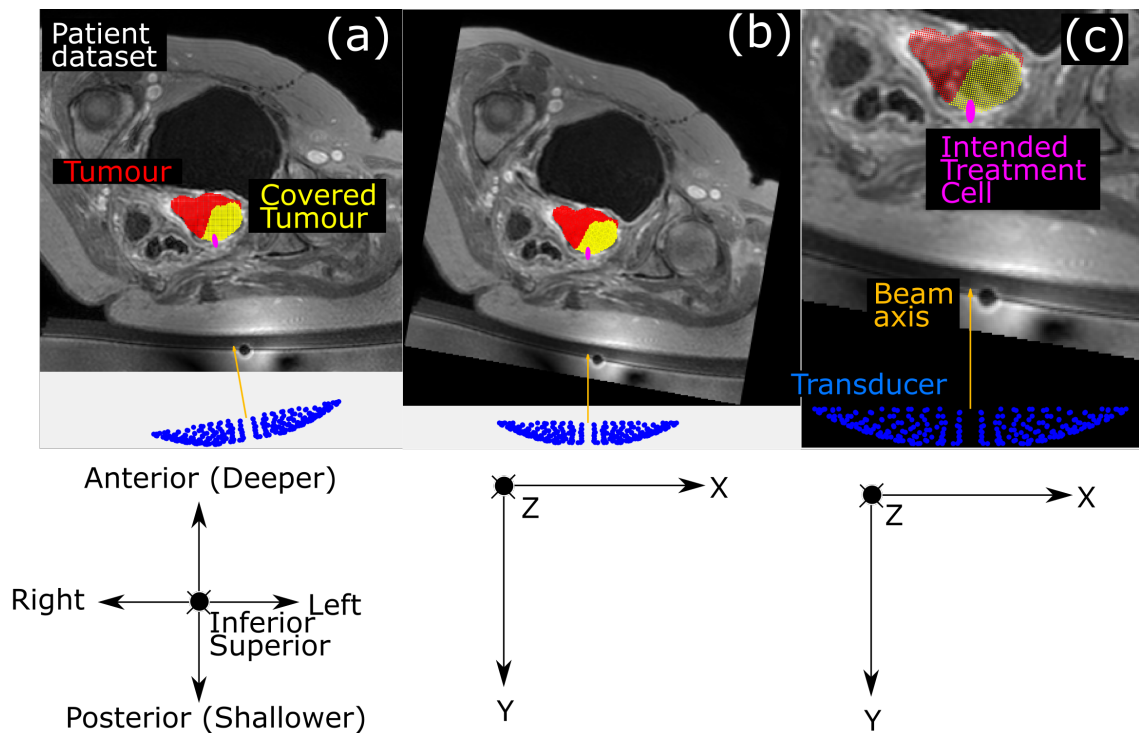
### 9.2.2 Inputs

Five patient treatment images were acquired as described in Section 3.3.3. Treatment images were assessed because the shape deformation of the gel-pad and human body was unknown for the registered-referral images. It is known that geometry of the medium is important to determining the acoustic energy deposited at the target, and hence, the ablated tissue volume [64, 83]. In the literature it has been reported that nonlinear acoustic simulation in three patients with and without refraction resulted in the spatial-peak time-average intensity decreasing by 92% and 77% on average compared to a simulation in water, respectively [83]. Further work could be done to simulate human body deformation. From the Fat and Water Dixon image datasets, the simulation grid was extracted. The deepest and shallowest reachable treatment cells, which lay within the tumour surface, were identified from the set of 'reachable' treatment cells which had been identified in the tumour coverage assessment (Section 7.2.2.3). Where multiple treatment cell centres had the same depth, the one closest to the isocentre line (a line going through the positions of the magnetic isocentre and transducer home position) was used as the deepest and shallowest reachable treatment cell centres. The transducer positions and angulations required to reach the deepest and shallowest treatment cells, obtained as a result of tumour coverage assessment, were used to extract the simulation grids from the dataset. The tumour coverage, determined previously, was used to provide an upper bound to the estimated proportion of tumour treatable. This is because, if a region of the tumour could not be reached and covered because it was blocked by an organ at risk or acoustic obstruction, it was assumed to be untreatable.

### 9.2.3 Simulation Grid Extraction

To reduce the computational time requirements for acousto-thermal simulation, a simulation grid (140 x 180 x 140 mm in X, Y and Z, axes defined in Figure 9.2) was identified from the larger Dixon image datasets (dimensions of 250.0 x 250.0 x 199.5 mm in Left-Right, Anterior-Posterior, and Inferior-Superior, axes defined in Figure 9.2). The simulation grid was extracted such that the transducer was

situated at one end of it, with the beam axis centred in the X-Y plane and parallel to the 180 mm-long Y-axis. The Z axis was perpendicular to the X and Y axes, as shown in Figure 9.2, and positive towards the patient head. If, in order to reach the deepest or shallowest treatment cell, the transducer position was outside of the field-of-view of the original image dataset (i.e. as seen in Figure 9.2(a)), then the region outside the dataset field-of-view around the transducer was assumed to be the oil bath within which the transducer sits in the MRgHIFU system. The simulation grid was selected to ensure that the entire Sonalleve<sup>®</sup> transducer was included, so the acoustic source could be accurately described during simulations in Step 3.



**Figure 9.2:** Extraction of the simulation grid from a patient image dataset, using the known transducer position and angulation required to target the intended treatment cell. (a): An example cross-section of the original image dataset is shown, with the tumour (red), covered tumour (yellow), shallowest reachable treatment cell (magenta) and associated transducer position and angulation (blue beam, orange beam axis) overlaid. (b): The original image dataset is rotated such that the transducer beam axis is made vertical, and the Y-axis is defined as parallel to that. (c): The simulation grid is extracted from the rotated dataset. The coordinate systems used to refer to directions within (a) the original image dataset (patient orientation-based directions), (b) the rotated and (c) the extracted grid are shown below the images.

## 9.2.4 Medium Generation

Each voxel within the extracted simulation grid had to be associated with appropriate material properties (density, sound-speed, attenuation coefficient, thermal conductivity, specific heat capacity) for muscle, fat, bone, gel-pad, oil and extracorporeal air [69, 176] for modelling acoustic propagation and the thermal bioeffects. Because of the way k-Wave models the power-law frequency-dependent attenuation within its governing equations [81, 96], the power-law frequency-dependent attenuation exponent  $\gamma$  (see Section 2.2 equation 2.3) was fixed for all medium voxels, and here, it was set to 1.1 [83].

A simple automatic method was used to assign different voxels their material properties. The patient Dixon Water (W) and Fat (F) images were combined into water-fraction (WF) and fat-fraction (FF) images:  $WF = \frac{W}{W+F}$  and  $FF = \frac{F}{W+F}$ . Bone voxels had already been identified by a manual segmentation (see Section 5.2.3), and were assigned appropriate material properties. A non-bone voxel within the body mask that had the value of 0.75 in the FF image necessarily had the value of 0.25 in the WF image. This was interpreted as a voxel having 75% adipose tissue and 25% aqueous tissue. This appeared plausible due to the partial volume effect in digitised images. Under this assumption, the sound speed, density and attenuation coefficients for each voxel were derived as the weighted average of the material properties of fat and muscle, respectively. The fluid within filled bladders was therefore given the material properties of muscle, but because the bladder was in the post-focal region of all patients tested, this was not expected to influence the acoustic energy absorption at the focus and pre-focal region. For simplification, tumour voxels were also modelled as muscle. In previous studies, the acoustic sound-speed of tumours has been measured to range from 1532 m/s at 22°C to 1581 m/s at in-vivo temperatures [69], and the attenuation coefficient for metastatic livers has been measured to be 0.26-0.70 dB/(MHz $^\gamma$ ) [69], similar to that of muscle.

A voxel outside the body mask represents either extracorporeal air, gel-pad or oil bath. Extracorporeal air had already been segmented for patients (see Section

5.2.3 and Figure 5.5). Discussions with Profound Medical revealed that the gel-pad consisted almost entirely of water. The weighted average voxel labelling scheme described above was used to assign material properties to the gel-pad and oil-bath voxels that came from the original image dataset. As mentioned previously, for some patients, the transducer was positioned outside the image dataset field-of-view, and hence the voxels in the region beyond the image dataset field-of-view were assigned the acousto-thermal properties of oil.

Since the transducer was manipulated in order that its focus reached a geometric focal point that is known to be reachable, organs at risk and acoustic obstructions, such as bone and extracorporeal air, were expected to be at or beyond the periphery of the acoustic beam. However, extracorporeal air voxels had to be handled carefully, because for some methods, the simulation can become numerically unstable if the medium contains interfaces between materials with large impedance mismatches [177], such as that between gel-pad and extracorporeal air (on the order of  $10^4$ ). In one study, an impedance ratio of 16-to-1 resulted in up to 6% error in transmission [178]. The acoustic impedance ratio between bone and soft tissue was approximately 5-to-1, suggestive of an error of approximately 2% [178]. Therefore, all extracorporeal air voxels were assigned the same sound-speed and density as water, and in order to model the acoustic opacity of air, an artificial attenuation coefficient of  $20 \text{ dB}/(\text{MHz}^\gamma \text{ cm})$  was assigned to air voxels where diffracted acoustic waves intersected an air boundary. The reflection of diffracted waves from the air interface was expected to be negligible. In order to reduce reflection from the grid boundaries, perfectly matched layers [93] surrounded the computational grid.

The material properties of the materials that were used in the simulation are shown in Table 9.1. The gel-pad is expected to warm to patient skin temperature ( $33.5^\circ\text{C}$ [179]) over the course of treatment. At this temperature, the gel-pad sound speed was calculated from the Bilaniuk and Wong equation for the sound speed of water [180] to be  $1517 \text{ m/s}$ , whilst the density was  $995.1 \text{ kg/m}^3$  [38]. The acoustic properties of tissue are known to vary with temperature [37]. The effect of thermo-acoustic lensing was expected to be minor ( $1 \text{ mm}$  displacement

in peak focus position after heating to 80°C at the focus [73]). For simplicity, and because k-Wave does not support coupled acousto-thermal simulations, acoustic and thermal properties for materials were assumed to remain constant. The acoustic properties were measured from tissue samples at 35-37°C [69]. The sound-speed and density of the oil used in the Sonalleve<sup>®</sup> V1 had been measured at room temperature [181], and it was assumed that the same oil was used in the Sonalleve<sup>®</sup> V2. The thermal properties of the oil were assumed to be those of mineral oil, measured at 40°C [182]. This temperature was chosen, because the mineral oil was assumed to be heated by ultrasound absorption over the course of the treatment. The attenuation coefficient of oil was obtained from measurements performed by John Civalo in the ICR therapeutic ultrasound group using a buoyancy method [183]. Acoustic energy deposition into extracorporeal air and thermal conduction from sonicated tissue deep within the patient was expected to be negligible. Furthermore, because of the form of the bioheat equation (Section 2.2 equation 2.4), the pseudo-spectral method used for thermal simulation is affected by discontinuities in thermal conductivity. Therefore, for simplicity, extracorporeal air voxels were assigned the thermal properties of water [83].



**Table 9.1:** Acoustic and thermal properties of different tissues and materials. Sources for these properties are referenced.

	Water (Gel-Pad)	Oil	Muscle	Adipose Tissue	Bone
Sound-speed (m/s)	1517	1380 [181]	1575 [83]	1478 [83]	4080 [83]
Density (kg/m <sup>3</sup> )	995.1	840 [181]	1055 [83]	950 [83]	1908 [83]
Pressure Attenuation Coefficient (dB/(MHz <sup>γ</sup> cm))	0.00217 [83]	0.1	0.6 [83]	0.48 [83]	20 [83]
Thermal Conductivity (W/(m K))	0.6045 [176]	0.130 [182]	0.49497 [176]	0.21145 [176]	0.32 [176]
Specific Heat Capacity (J/(kg K))	4178 [176]	1974 [182]	3421 [176]	2348 [176]	1312 [176]

### 9.2.5 Acoustic Simulation

Simulations were performed to obtain acoustic pressure fields, which are used to calculate acoustic intensity and energy deposition. Acoustic simulations were performed using the open-source k-Wave v1.3 [80, 81] package, which uses pseudo-spectral methods to model acoustic wave behaviour from coupled first-order partial differential equations. Simulations were performed on a regular Cartesian spatial grid.

A linear simulation to analyse patient treatability was developed, in preference to more extensive nonlinear simulations. Linear simulations have been used to assess how tissue geometry affected kidney ablation treatments [64]. The computational resources required to perform nonlinear simulations were prohibitive. In preliminary work, nonlinear simulations performed at a spatial resolution of 13.5 points per minimum wavelength (the resolution at which simulations had been validated against experimental data [84]) were estimated to require more

than 1 week of computation on a commercially available cloud node (128 CPUs, 2 TB RAM).

Three-dimensional acoustic simulations were performed using k-Wave, with the simulation code implemented in C++, on a 20 CPU-core node (each core clocked at 2.60 GHz and with 12.8 GB RAM per core). The k-Wave GPU code could not be used because each simulation required more video memory than was available.

For all acoustic simulations performed, the acoustic pressure field emerging from the voxelised Sonalleve<sup>®</sup> transducer described in Section 3.2 was simulated, with source pressure corresponding to 300 W of acoustic power across its entire surface (the maximum output allowed in the clinical uterine configuration) and frequency of 1.22 MHz. The total number of time steps simulated was that required to travel from one corner of the simulation grid to the opposite corner at the slowest sound-speed within the medium. Further details are provided below.

### 9.2.5.1 Modelling Bone with Reduced Sound Speed

#### 9.2.5.1.1 Method

The speed of sound in bone is much higher than that in other materials of interest (Table 9.1). This means that use of the correct value would require finer time resolution than if bone were not present in order to maintain the same numerical accuracy in acoustic simulation [184]. This would result in a longer computational time. However, the acoustic impedance (the product of the medium's speed of sound and density) determines the reflection and transmission at an interface between two materials (see Section 2.2, [35]). Therefore, changing the sound speed of bone in the simulation to that of water, and then compensating by increasing the density accordingly left the acoustic impedance unchanged. This could provide a speedup in computational time without sacrificing accuracy with which reflections at bone interfaces were simulated. In addition, because the transducer was positioned and angled such that bone was only at the periphery of, or beyond, the acoustic beam, the acoustic energy propagating towards bone was expected to be negligible. Any erroneous refraction of transmitted acoustic

waves in bone was ignored.

The effect of this change was tested using patient G8<sub>2</sub> (chosen arbitrarily). Two simulations were run, one with the correct bone sound-speed and density (the 'unmodified' simulation) and the other with the sound-speed of water and elevated density (5131.6 kg/m<sup>3</sup>) assigned to bone (the 'modified' simulation). In both cases, the acoustic impedance in bone was 7.78 MPa.s/m<sup>3</sup>. The chosen spatial resolution of 2.2 grid points per minimum wavelength (1.1 mm), corresponding to a spatial resolution of 514 µm, which is just above the Nyquist sampling limit of 2 points per shortest wavelength, was chosen to minimise the time required for the validation study. For both simulations, the Courant–Friedrichs–Lewy (CFL) number, which defines the time resolution as a function of the spatial resolution and the maximum medium sound-speed [184], was set to approximately 0.1 such that each wave period would be represented by an integer number of time steps. This corresponded to time steps of 13 ns (simulated for approximately 14900 steps) and 33 ns (simulated for approximately 5900 steps) for the unmodified and modified simulations, respectively. This choice of CFL number ensured that the time step in the simulation would be small enough to yield stable simulations. The transducer's geometric focus was centred on the tumour centroid with the beam axis perpendicular to the MR bed surface.

In order to investigate the potential for further increase in the time step, the CFL numbers 0.2, 0.25 and 0.3 were studied, and the peak intensity was examined. The computational times required for the simulations were compared. Differences in the peak intensity compared to the unmodified simulation were used to quantify the effect this modification had on simulation accuracy.

#### **9.2.5.1.2 Results**

The results of the simulations are shown in Table 9.2. At CFL = 0.2, the unmodified simulation became unstable and returned meaningless results (denoted by N/A in the table). The same was also observed at CFL = 0.3 for the modified simulation. At CFL = 0.1, the difference in the peak intensity when modifying

the bone sound speed was 1%, resulting in a speedup of 2.5 times. A greater speedup of 6.7 times was achieved by increasing the CFL number to 0.25 for the modified simulation without adverse effects on accuracy (effect on peak intensity was 1.2%).

**Table 9.2:** Results for the simulations involving the modification, or lack thereof, of bone sound speed.

<b>CFL Number</b>	<b>Simulation Type</b>	<b>Peak Intensity (W/cm<sup>2</sup>)</b>	<b>% Difference from Unmodified CFL=0.1</b>	<b>Simulation Time Elapsed (mins)</b>
0.1	Unmodified	2399	0	240
0.1	Modified	2422	1.0	96
0.2	Unmodified	N/A	N/A	N/A
0.2	Modified	2429	1.3	48
0.25	Unmodified	N/A	N/A	N/A
0.25	Modified	2427	1.2	36
0.3	Unmodified	N/A	N/A	N/A
0.3	Modified	N/A	N/A	N/A

### 9.2.5.1.3 Discussion

Changing the bone sound speed to that of water and increasing the density such that the acoustic impedance of bone remains the same was done with the goal of reducing computational time without affecting simulation accuracy. Results showed an acceptably small (1%) increase in peak acoustic intensity and a speedup of 2.5 times (from 240 minutes to 96 minutes). The speedup was improved to 6.7x (from 240 minutes to 36 minutes) by additionally increasing the CFL number from 0.1 to 0.25, with a total increase in peak acoustic intensity of 1.2% compared to the unmodified simulation. This modification was therefore used for all subsequent acoustic simulations, including the linear model validation and the simulation of treatment cell sonication.

## 9.2.5.2 Linear Acoustic Model Validation

### 9.2.5.2.1 Method

In order to quantify the simulation accuracy, k-Wave's linear model was compared with the O'Neil solution [185] built into the k-Wave package, as had previously been done by others [98]. In general, the higher the grid resolution, the greater the maximum frequency that could be simulated and the better the simulation accuracy. However, increasing spatial resolution increases computational resource requirements. Thus, a compromise was sought. The relationship between the spatial resolution, simulation accuracy and computational time was examined as described below. The spatial resolution with the least computational time, that resulted in the simulated pressure being within an acceptable margin of error ( $\pm 10\%$ ) of that produced by an analytical model, was chosen for subsequent simulations. This accuracy threshold was chosen as it corresponds to the accepted experimental uncertainty for pressure measurements using hydrophones [84].

Acoustic propagation from a bowl transducer with similar specifications to the Sonalleve<sup>®</sup> transducer (130 mm diameter, 140 mm radius of curvature, source frequency 1.22 MHz, source pressure 350 kPa which corresponds to approximately 300 W acoustic power in water) through a homogeneous, linear, non-attenuating medium (sound-speed 1480 m/s, density 998.2 kg/m<sup>3</sup>) was modelled. The simulation grid (134 x 134 x 190 mm) was aligned with the transducer beam axis, such that the longest grid axis was parallel to the beam axis. The spatial resolutions tested corresponded to 2.2, 4.2, 6.2, 8.2 and 10.2 grid points per minimum wavelength (1.2 mm), corresponding to spatial resolutions of 551, 289, 196, 148, 119  $\mu\text{m}$  respectively. The CFL number was set to approximately 0.25 for each choice of spatial resolution, such that each wave period would be represented by an integer number of time steps. This choice was made in light of the results of Section 9.2.5.1.2.

k-Wave and the O'Neil model, an integral equation that can describe the pressure at every point in the field [185], were each used to generate an estimate of the pressure amplitude along the beam axis and along a line perpendicular to


the beam axis through the transducer focal point. The average errors between the k-Wave and O’Neil pressures along the beam axis (‘axial’) and along the line perpendicular to the beam axis (‘lateral’) were quantified using the L2 relative error, which has been used in the literature for this purpose [98], using equation 9.1:

$$L2 = 100\% \times \sqrt{\frac{\sum_x (p_{sim}(x) - p_{ref}(x))^2}{\sum_x (p_{ref}(x))^2}} \quad (9.1)$$


where  $p_{sim}(x)$  is the pressure amplitude at position  $x$  simulated by k-Wave, and  $p_{ref}(x)$  is the reference pressure amplitude calculated by the O’Neil solution. To identify the shape error between the k-Wave and O’Neil pressures, they were normalised such that their peak pressures were identical, and the error between them was then quantified by the L2 relative error. The non-normalised error between them was calculated to quantify the overall error.

#### 9.2.5.2.2 Results

The pressure amplitudes generated along the beam axis and perpendicular to the beam axis at the transducer focus by k-Wave were plotted with that generated by the O’Neil solution, for the spatial resolutions of 2.2 and 10.2 grid points per minimum wavelength in Figures 9.3 and 9.4 respectively. At 2.2 points per minimum wavelength, the near-field pressure generated by k-Wave disagrees with that predicted by the O’Neil model. Even at the highest spatial resolution tested (10.2 points per minimum wavelength), there is a discrepancy between the simulated and theoretical peak pressures of approximately 2 MPa.



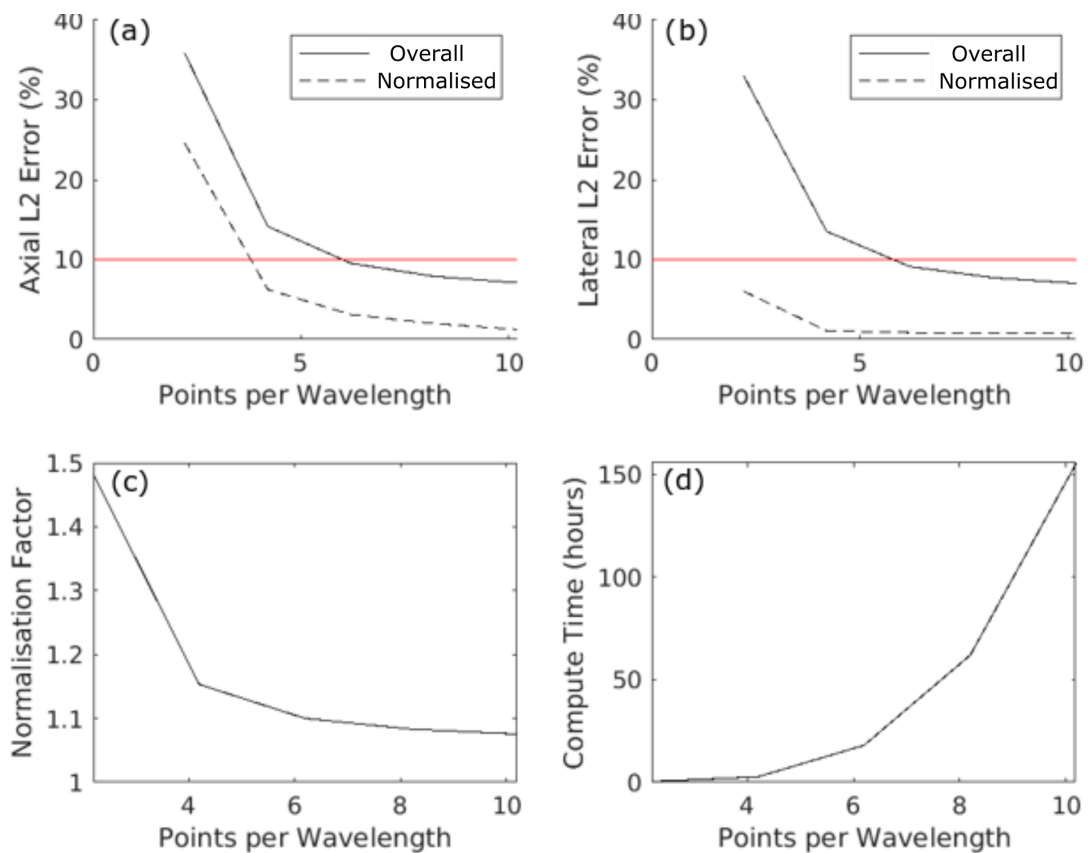
**Figure 9.3:** Comparison of the pressure amplitude simulated by k-Wave and that calculated by O'Neil's solution at the spatial resolution of  $551 \mu\text{m}$  (2.2 grid points per minimum wavelength) along the beam axis (top) and across the beam axis through the transducer focus (bottom). k-Wave underestimated the maximum focal peak by 9 MPa (33%).



**Figure 9.4:** Comparison of the pressure amplitude simulated by k-Wave and that calculated by O'Neil's solution at the spatial resolution of  $119\ \mu\text{m}$  (10.2 grid points per minimum wavelength) along the beam axis (top) and across the beam axis through the transducer focus (bottom). k-Wave underestimated the maximum focal peak by 2 MPa (7%).



The results of the linear acoustic model validation are summarised in Figure 9.5. Panels (a) and (b) show the overall and normalised L2 relative error. The latter indicates the error in the shape of the pressure amplitude envelope (as seen in the near field in Figure 9.3). The red line, at 10%, is the accepted uncertainty for pressure measured by hydrophones [84]. The spatial resolution with the fastest computational time that falls under this threshold of error (both laterally and axially) is 6.2 grid points per wavelength, which corresponds to a spatial resolution of 196  $\mu\text{m}$ .



**Figure 9.5:** Validation of the linear simulation against the O'Neil analytical solution as a function of spatial resolution. (a): the overall (solid black) and normalised (dashed) L2 relative error between the k-Wave and O'Neil pressures along the beam axis are shown. The red line, at 10%, is the accepted uncertainty for pressure measured by hydrophones [84]. (b): the overall (solid black) and normalised (dashed) L2 relative error between the k-Wave and O'Neil pressures across the beam axis are shown. The red line, at 10%, is the accepted uncertainty for pressure measured by hydrophones [84]. (c): the normalisation factor that the k-Wave pressures are multiplied by in order to match the O'Neil pressure at peak pressures, at each spatial resolution. (d): the computational time required for each simulation at each spatial resolution.

### 9.2.5.2.3 Discussion

The linear model was validated at several spatial resolutions. This was done to identify the spatial resolution at which acoustic and thermal simulation of treatment cell sonication should be performed in order to balance sufficient simulation accuracy with reasonable computation time. The chosen spatial resolution of 6.2 grid points per wavelength, corresponding to approximately 196  $\mu\text{m}$  grid spacing, required the least amount of computational time whilst satisfying the threshold of acceptable simulation accuracy (10%, which is the accepted hydrophone error for pressure measurements in the literature [84]). That the normalised L2 relative error is lower than the overall L2 relative error (by approximately 7% at 6.2 points per wavelength) suggests that the main contribution to the error is the amplitude and not the shape of the simulated pressure, as seen in Figure 9.4. Underestimated pressure amplitudes at the focal peak suggest underestimated acoustic energy deposition into the tissue, and therefore underestimated temperature changes at the focus and underestimated tissue ablation volume [39]. A proposed reason for this underestimation in amplitude by Martin, Ling and Treeby in 2016 [98] was the discretisation of the transducer bowl in k-Wave compared to the perfectly spherical bowl assumed in the O'Neil solution. In k-Wave, the transducer bowl is discretised into non-zero voxels within a binary array, each of which represents a pressure source. Martin et al. claim that due to this discretisation, the density of source voxels in the grid directions (i.e. X, Y and Z) and in the diagonal directions are different – for example, a horizontal line would consist of  $\sqrt{2}$  more source points than a diagonal line of the same length – and hence, cause the underestimation in focus pressure. The discretised transducer is also posited to focus less effectively. As the spatial resolution becomes finer, the discretisation converges closer to the ideal spherical bowl, resulting in a decrease in overall error.

The choice to perform a linear simulation was made in light of the limited computational resources. However, nonlinear modelling would have been preferable. In nonlinear acoustics, as acoustic waves travel through a nonlinear medium, the

amplitude of higher-frequency harmonic components is increased. These higher harmonics tighten the beam focus. Furthermore, because attenuation in tissue is modelled as power-law frequency-dependent, the higher harmonics transfer acoustic energy to tissue more effectively. Published studies suggest that nonlinear simulation results in greater focal temperatures and a greater volume of ablated tissue compared to linear simulation [39]. Preliminary investigations of the convergence of peak focal pressure as spatial resolution increased were performed for nonlinear simulations to validate a nonlinear acoustic model. Unfortunately, convergence was not observed even at 10.2 grid points per minimum wavelength, which required approximately 60 hours of computational processing on a commercially-available computational platform, suggesting that an even finer spatial resolution was required for accurate nonlinear simulation. In the literature, 13.5 grid points per minimum wavelength has been used to validate k-Wave simulation against experimental data on a supercomputer cluster, using custom-built code that is not publicly available [84]. This was untenable given the available computational resources for this PhD project.

### 9.2.5.3 Acoustic Simulation of a Patient Sonication

In the Sonalleve<sup>®</sup> MRgHIFU system, treatment cells are created by electronically steering the transducer focus to 'trajectory points' arranged in concentric circles around the treatment cell centroid [42, 65]. The acoustic power at each trajectory point within each concentric circle is constant, but varies between concentric circles in order to compensate for the loss in focal intensity as the focus is steered further away from the beam axis. For a 4 mm-diameter treatment cell, the transducer focus is steered between 8 trajectory points, which are situated on a 2 mm-radius circle. In preliminary investigations, it was found that each simulation of acoustic propagation to one of these trajectory points, at a spatial resolution of 6.2 points per minimum wavelength (chosen as a result of the validation) and a CFL number of 0.25, required more than 24 hours to complete, suggesting that simulations for all 8 trajectory points could be untenable.

To reduce computation time, an alternative strategy was pursued. Firstly, an

on-axis pressure field, with no phase difference between the transducer elements, was simulated. Secondly, to approximate the steered pressure field when sonicating a trajectory point, the on-axis pressure field was translated to the steered beam focus position, such that the intended beam focus in the on-axis pressure field is translated to the trajectory point. This was done for all eight trajectory points. Thirdly, the translated acoustic pressure field was transformed into acoustic intensity using equation 2.1. Acoustic intensity was then transformed into volumetric heat deposition for thermal simulation, as described below. All patient simulations were conducted with the source acoustic power of 300 W.

The focal peak pressure was quantified for each patient, for the deepest and shallowest reachable treatment cells. The relationship between the focal peak pressure and the beam path length, the distance along the beam axis between the centroid of the intended treatment cell (at the transducer geometric focus) and the skin, was examined.

## **9.2.6 Thermal Simulation**

The temperature change in tissue resulting from acoustic energy deposition was modelled using the Pennes bioheat equation (Section 2.2 [40]). The thermal dose received by tissue was calculated from the tissue's temperature history (equation 2.6, [23]). Software within the k-Wave package was available for simulating thermal diffusion, which solved the Pennes bioheat equation using a pseudo-spectral and k-space method [127]. All thermal simulations were performed on a 14 CPU-core node (each core clocked at 2.60 GHz and with 12.8 GB RAM per core). Thermal properties for all tissues are listed in Table 9.1. For simplicity, the blood perfusion rate in all tissues was assumed to be zero, as is the case in some of the literature [64]. This was done for simplicity, as the tumour blood flow is known to be heterogeneous [186], with the tumour core generally being underperfused [187]. This presented a best-case scenario in terms of thermal dose accumulation.

Clinical ablation using treatment cells involves electronically steering the transducer focus to each trajectory point in a cell and sonicating it for 50 ms,

with the time required to switch between trajectory points being <10 ms [42]. This continues for a prescribed length of time. After this prescribed time is reached, the Sonalleve<sup>®</sup> is forbidden from sonicating another treatment cell for a period of time [2] in order to prevent burns in healthy tissue. This was simulated as described. The heating part of the simulation was run for a total of 16 seconds, the Sonalleve<sup>®</sup> configuration exposure time for a 4 mm treatment cell [2]. When sonicating each trajectory point, the acoustic intensity at that point was transformed into volumetric heat deposition (using equation 2.5 in Section 2.2, with the absorption coefficient being defined as 75% that of the attenuation coefficient [36]) and used as input into the Pennes bioheat equation. The simulation was propagated in time by 50 ms. Then, to represent switching, the heat source was set to zero and the simulation was propagated in time by 10 ms. After that, the heat source was changed to that associated with the acoustic intensity field at another trajectory point. When the heating part of the simulation was complete, the cooling part of the simulation began. The heat source was set to zero. The total cooling time to be simulated was 30 seconds with a time step of 0.1 seconds [121]. The cooling duration of 30 seconds was chosen from clinical experience. The thermal dose [23] resulting from this was calculated using equation 2.6.

The centroid, and volume, of the ablated tissue, defined as the tissue that received  $\geq 240$  CEM43, was identified and recorded. The percentage overlap between the treatment cell and the ablated tissue was quantified. Differences between the centroid of the intended treatment cell (i.e. the geometric focus) and that of the ablated volume were also recorded.

## 9.2.6.1 Software Validation

### 9.2.6.1.1 Methods

The thermal simulation software was validated by comparison with an analytical solution derived from Green's functions [129], which had been built into k-Wave and which only worked for homogeneous data. The medium density, thermal conductivity and specific heat capacity were set to be those of muscle (see Table 9.1). The thermal simulation grid had a field-of-view of 88 x 88 x 88 grid points

and spatial resolution was 182  $\mu\text{m}$ , which was the same as the grid spacing for acoustic simulation. The time step was set to be 50 ms and a total of 320 time steps, corresponding to 16 seconds, was simulated, to match the thermal simulations in patients. A 8 x 8 grid point square, centered in the simulation grid, was set as a constant heat source of 50  $\text{W}/\text{cm}^3$ . This was the approximate maximum heat deposition observed for patient G29 (arbitrarily chosen) when sonicating the shallowest treatment cell. The difference between the k-Wave simulated and the analytically-derived temperature fields was quantified using the overall L2 relative error (equation 9.1).

#### **9.2.6.1.2 Results**

The overall L2 relative error between the analytically-derived and k-Wave simulated temperature fields was on the order of  $10^{-14}\%$ .

#### **9.2.6.1.3 Discussion**

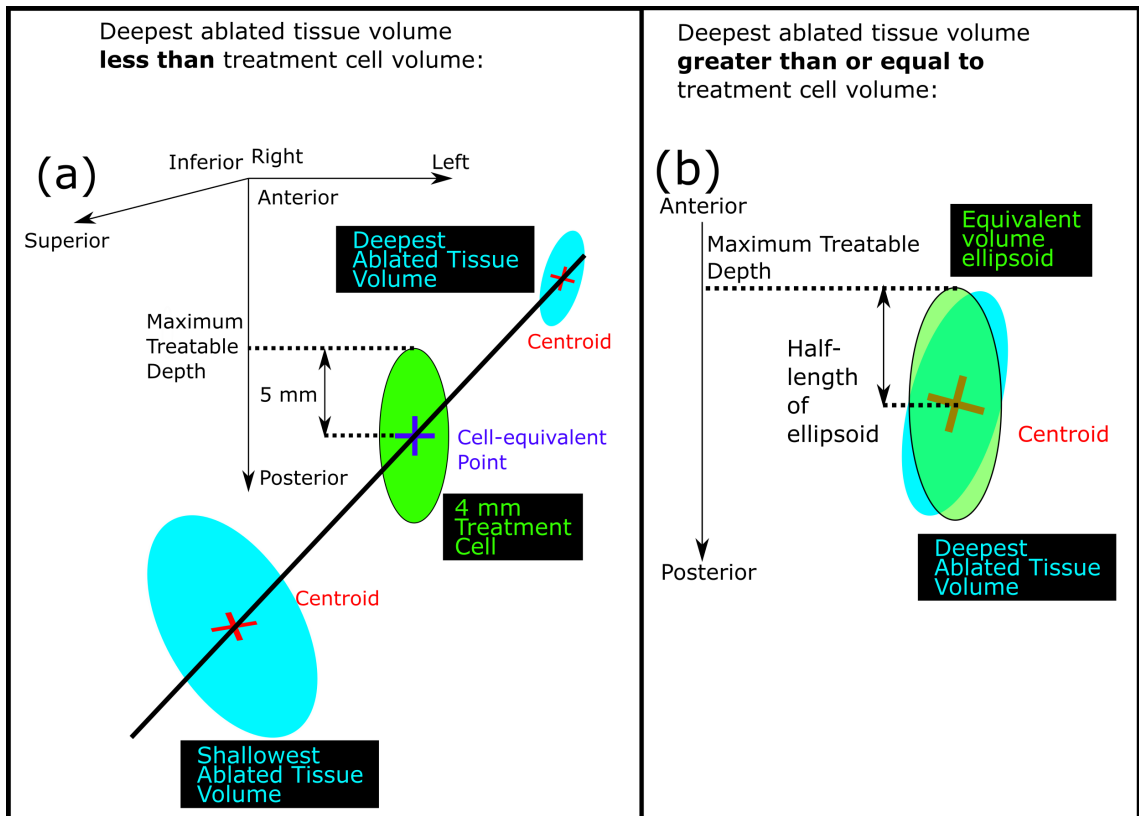
For thermal simulation validation, the difference between the analytical solution and the simulation was on the order of  $10^{-14}\%$ . This was believed to arise from finite-precision representation of real numbers in computers. These errors were small in comparison to other errors involved in treatability analysis, e.g. the error between simulated and analytically-derived acoustic pressure was approximately 9%. Hence the thermal simulation method was assumed to be fit for purpose. In the future, a more robust validation could be performed in heterogeneous media by comparing the k-Wave simulation to experimental data, or to results generated by a finite-difference method.

### **9.2.7 Calculation of Tumour Treatability**

Patient treatability was defined as the percentage of their tumour predicted to be able to receive a cytotoxic thermal dose. Because calculation of ablated tissue volumes for each treatment cell position required at least 20 hours of computational time, a simple method was developed to estimate treatability, as described below.

After the ablated tissue volume was calculated for both the deepest and shallowest reachable treatment cells, their position was found in the coordinate space of the original image dataset (see Figure 9.2 a). The maximum treatable depth was estimated in one of two ways, depending on whether the deepest ablated tissue volume was smaller than or greater than the 4 mm treatment cell volume, as seen in Figure 9.6 (a) and (b) respectively. If the deepest ablated tissue volume was smaller than the treatment cell volume, a line was drawn through the centroids of the ablated tissue. If, for a particular patient, the deepest ablated tissue volume was zero, the ablated tissue volume centroid was estimated by identifying the average offset between the geometric focus and the ablated tissue volume centroids for other patients and target positions, and then applying this average offset to the deepest geometric focus. The ablated tissue volume was linearly interpolated along this line. Using this patient-specific linear interpolation approach, the 'cell-equivalent' point, where the ablated tissue volume was equal to the treatment cell volume ( $84 \text{ mm}^3$  for a 4 mm cell), is identified on this line. The maximum treatable depth was taken to be an additional half-treatment cell length anterior to the cell-equivalent point, which would be approximately the deepest point of the ablated tissue centered at the cell-equivalent point. If the deepest ablated tissue volume was larger than, or equal to, the treatment cell volume, an ellipsoid with the same volume as the deepest ablated tissue volume, and the same ratio of diameter to length as the 4 mm treatment cell, is created and centred on the deepest ablated tissue centroid. The maximum treatable depth was chosen to be an additional half-length of the ellipsoid deeper than the centroid.

All covered tumour volume more superficial to the transducer than the maximum treatable depth was interpreted as being treatable, and the percentage of tumour that was treatable was calculated. The estimated treatable tumour volume was compared to the covered tumour volume (calculated as described in Chapter 7), in order to assess how well the tumour coverage prediction agreed with predicted tumour treatability. In order to investigate how tumour treatability depends on the tumour's depth, the vertical distance from the tumour centroid to the skin ('tumour-skin distance') was recorded for each patient. To get an idea



**Figure 9.6:** Schematic diagram illustrating the calculation of the maximum treatable depth. (a): if the deepest ablated tissue volume was less than the treatment cell volume, a line was drawn between the centroids of the deepest and shallowest ablated tissue volumes. The cell-equivalent point was found by interpolating the ablated tissue volume along this line until the cell-equivalent point, where it was the same volume as the treatment cell ( $84 \text{ mm}^3$  for a 4 mm cell), is found. The maximum treatable depth is defined to be half a treatment cell length (5 mm for a 4 mm cell) anterior of the cell-equivalent point. (b): if the deepest ablated tissue volume was greater than the treatment cell volume, an equivalent volume ellipsoid, with the same ratio of diameter to length as the treatment cell was created and with the same volume as the deepest ablated tissue volume, is centred at the deepest ablated tissue centroid. The maximum treatable depth was defined to be half the ellipsoid length anterior of the deepest ablated tissue centroid.

of how far the HIFU focus could reach into the tumour, the anterior-posterior distance between the deepest and shallowest ablated tissue centroids (the 'ablated tissue separation distance') was measured for each patient.

### 9.2.7.1 Clinical Treatability

With the exception of patient G8<sub>2</sub>, all patients had been treated using 4 mm and 8 mm diameter treatment cells, with the source acoustic power for sonicating each



cell being determined by the clinician. Test sonications, used for calibrating the Sonalleve<sup>®</sup> control software, had been conducted at acoustic powers of 80 W for patient G3 and 110 W otherwise for 16 seconds for all cell sizes. For non-test exposures, 4 mm cells had been sonicated for 16 seconds, and 8 mm cells for 20 seconds (see Table 3.1). Patient G8<sub>2</sub> had been treated with feedback cells [65, 106], in which the source acoustic power was selected by the treatment team but the heating duration was determined by the Sonalleve<sup>®</sup> control software. For each patient, the number of completed sonications, the range of acoustic powers used per sonication and the total energy delivered were recorded. The clinicians had estimated the clinically treated tissue volume for each patient by first estimating the ablated tissue volume for each sonication, using proton resonance frequency shift MR thermometry data to identify the 240 CEM43 dose contours. Then, the clinically treated tissue volume was estimated by summing the ablated tissue volumes.

## 9.3 Results

### 9.3.1 Acoustic Simulation of a Patient Sonication

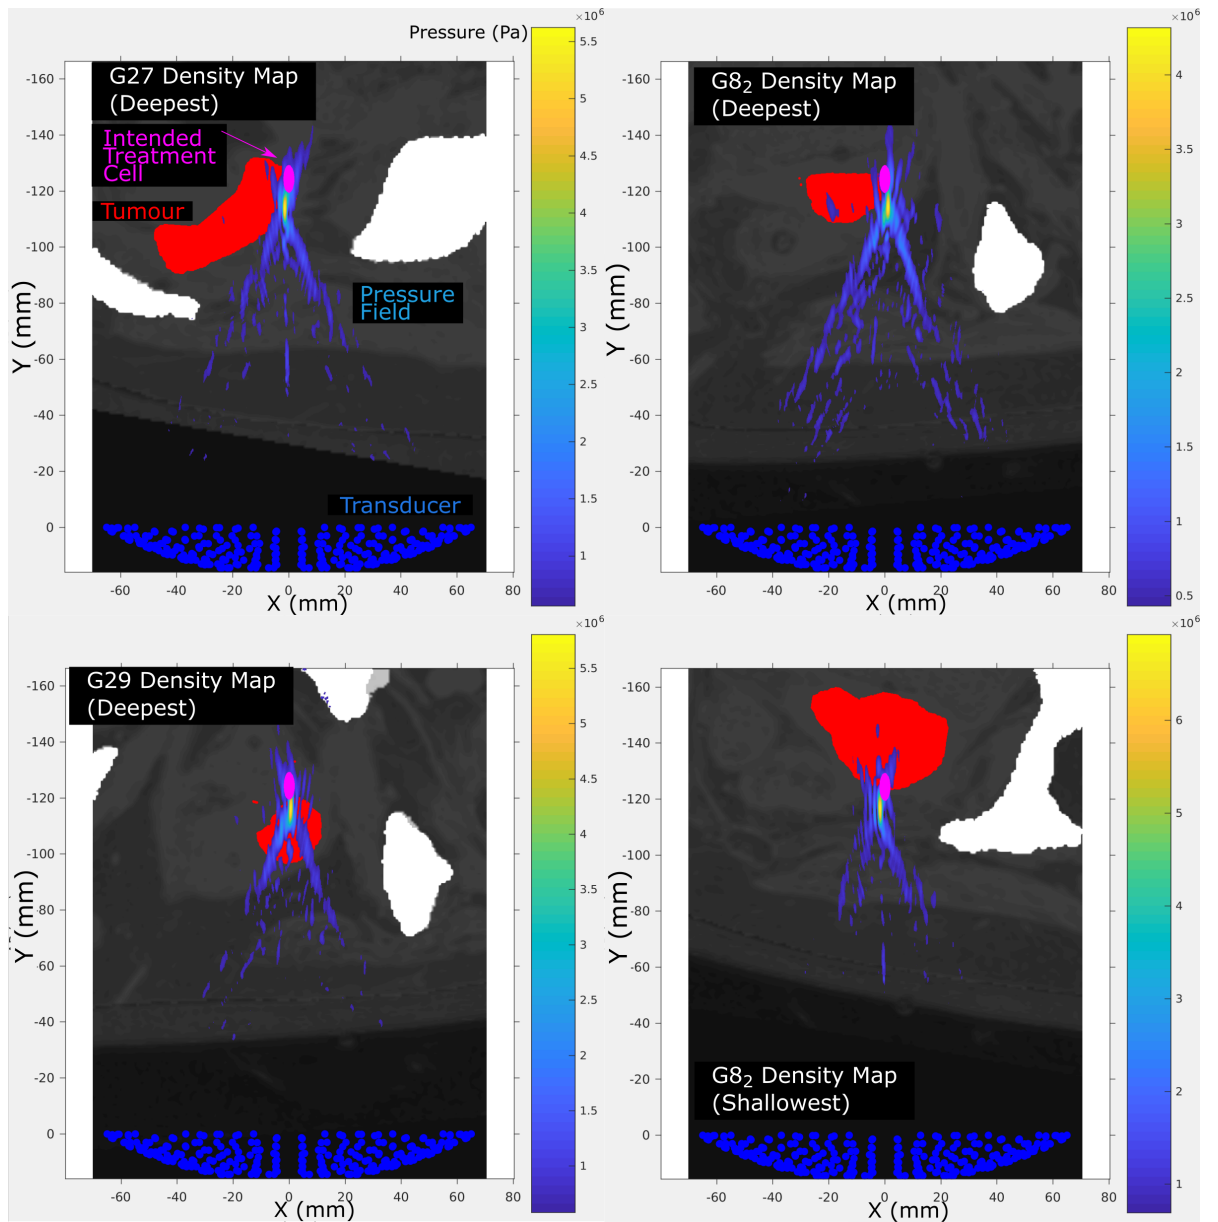
Acoustic simulation was performed without electronic steering. Four representative examples of the simulated on-axis pressure fields (i.e. no electronic steering) are shown in Figure 9.7. The acoustic pressure focus (yellow) was offset from the intended treatment cell (magenta) both along the beam axis, and perpendicular to it. Images are overlaid on a density map of the image, where lighter voxels represent higher densities). Visually, absorption in the acoustic beam appears greater in the pre-focal region than the post-focal region. It is noted that even with a steeply angled patient (G27, angled at 33°), the lateral offset in focal pressure is on the order of several millimetres from the intended geometric focus.

The peak pressures were calculated for all patients and for both target points (the deepest and the shallowest reachable treatment cells), as shown in Table 9.3. The peak focal pressure appears to decrease linearly with path length (R-squared of 0.69), as seen in Figure 9.8.

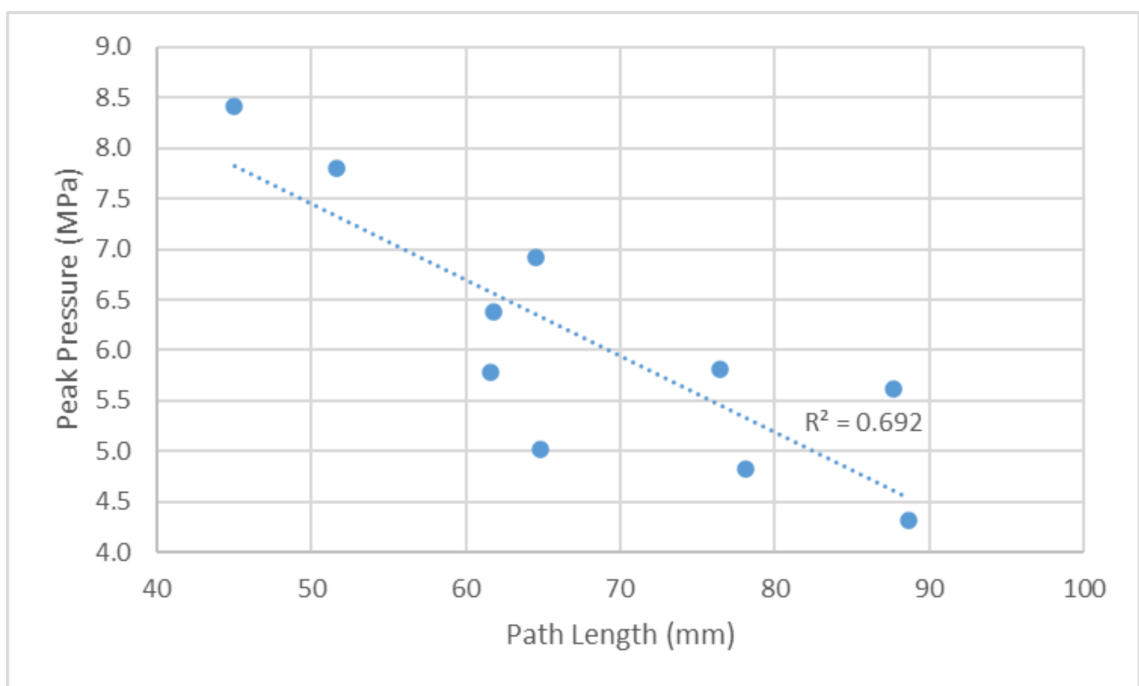
The time required for an acoustic simulation, averaged over all five patients and for both the deepest and shallowest reachable treatment cells, was  $40 \pm 6$  hours (range: 24-44 hours). From the linear model validation, the pressure is estimated to be underestimated by 10%. Uncertainty in path length was estimated to be 0.9 mm, the in-plane voxel dimension for the image dataset.

**Table 9.3:** Peak focal pressure and path length in tissue travelled by the acoustic beam for each patient and target treatment cell. All acoustic simulations were performed using a source acoustic power of 300 W.

Patient	Target Point	Maximum Pressure (MPa)	Path Length (mm)
G3	Shallowest	6.4	61.8
G8 <sub>2</sub>	Shallowest	6.9	64.5
G24	Shallowest	5.8	61.6
G27	Shallowest	8.4	45.0
G29	Shallowest	7.8	51.6
G3	Deepest	5.0	64.8
G8 <sub>2</sub>	Deepest	4.3	88.7
G24	Deepest	4.8	78.1
G27	Deepest	5.6	87.7
G29	Deepest	5.8	76.5



**Figure 9.7:** Four representative examples of the acoustic pressure field (see colour bar, only pressure values >10% of the focal peak pressure are shown for clarity), generated with no electronic steering applied, targeting both deepest and shallowest reachable treatment cells (magenta). All acoustic simulations were performed at source acoustic power of 300 W. The acoustic field is overlaid on the patients' density maps (lighter means material is denser, bone is white and oil is black). Tumour is shown in red. The cross-section is the X-Y slice in which contains the peak focal pressure. The focal peak is offset both axially and laterally from the intended treatment cell.

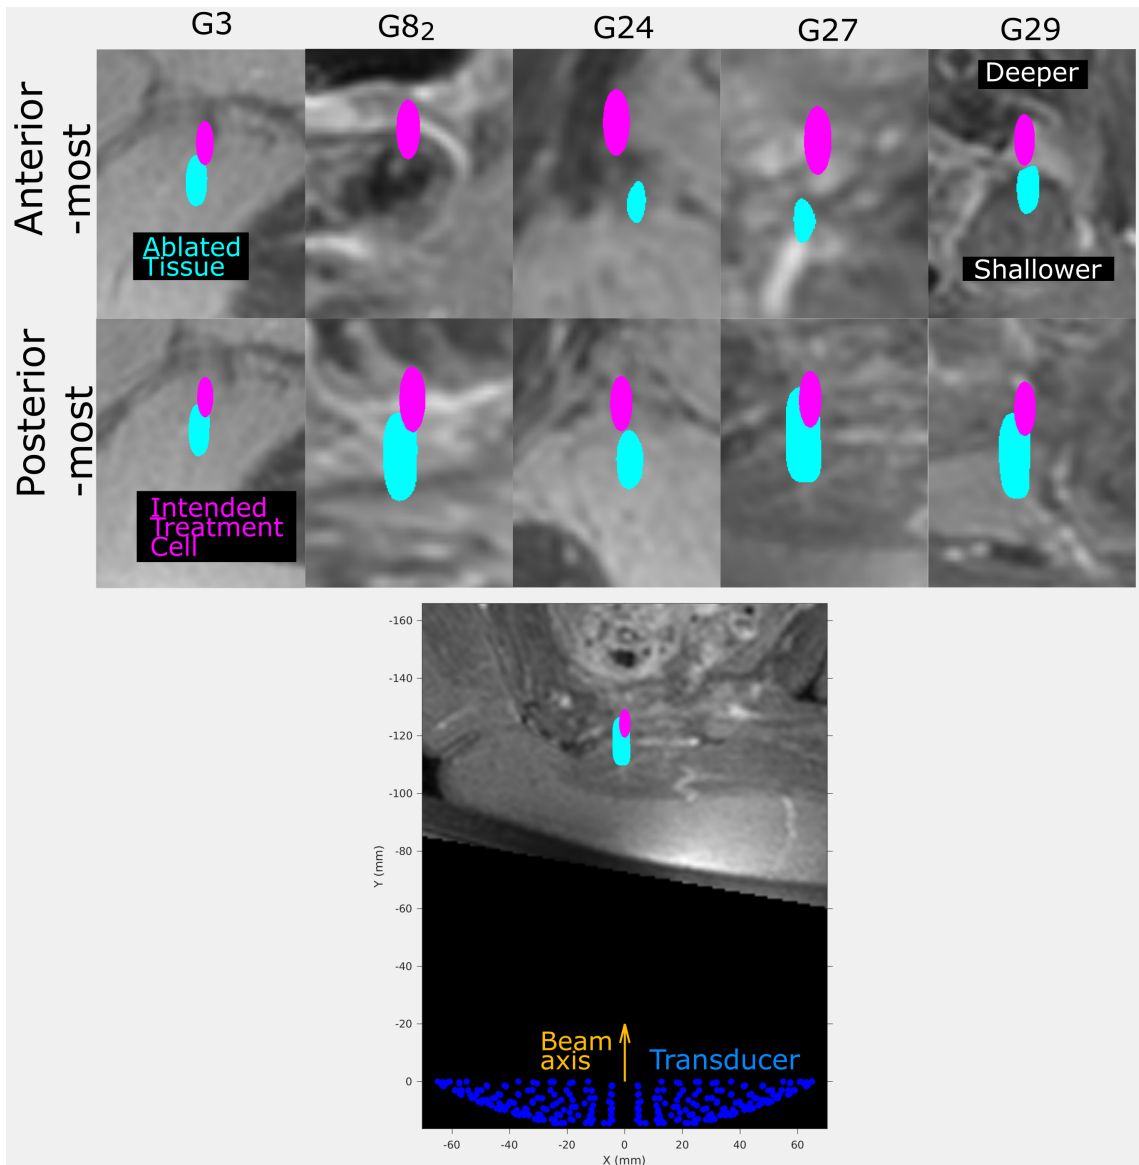


**Figure 9.8:** Focal peak pressure is plotted against path length in tissue. The dotted line is a linear regression.

### 9.3.2 Thermal Simulation

The differences between the centroids of the ablated tissue and the treatment cell are shown in Table 9.4 and Figure 9.9. The differences are defined such that Y is the direction along the HIFU beam axis, while X and Z are perpendicular to the beam axis, as defined in Figure 9.2. A positive X, Y and Z centroid position difference means that the centroid of the ablated tissue is more to the patient's left-hand side, closer to the transducer surface, and more towards the head than that of the treatment cell, respectively. On average, the ablated tissue was offset from the treatment cell by  $-0.6 \pm 1.7$  mm in X,  $9.3 \pm 1.6$  mm in Y and  $-0.5 \pm 0.4$  mm in Z, i.e. the main offset was that the ablation occurred closer to the transducer than the geometric focus. When sonicating the deepest and shallowest reachable treatment cells, the mean volume of the resulting ablated tissue ( $60 \pm 52$  and  $250 \pm 100$  mm<sup>3</sup>, respectively) was approximate to or exceeded the treatment cell volume (84 mm<sup>3</sup>). As seen in Figure 9.10, the path length appears to be linearly related to ablated tissue volume (R-squared of 0.88) and the Y-offset (R-squared of 0.72). An error in the path length of 12 mm, as observed for the registered-referral dataset of Volunteer 2 due to an inaccurate estimate of the soft tissue deformation resulting from reorientation into the treatment position (see Section 6.4.2), would have resulted in an error of approximately 100 mm<sup>3</sup> in ablated tissue volume. Thermal simulations took, on average,  $22 \pm 2$  hours (range: 18-24 hours) per patient per target point.


Uncertainty in centroid position difference was estimated to be 0.2 mm, the simulation spatial resolution, and uncertainty in ablated tissue volume due to the discretised nature of the simulation grid was estimated to be 1.7% (see Section 6.3.4). Uncertainty in path length was estimated to be 0.9 mm, the in-plane voxel dimension for the image dataset.



**Figure 9.9:** Cross-sections of the ablated tissue (cyan) when sonicating the deepest and shallowest reachable treatment cell (magenta, centered at the geometric focus of the transducer but shown in images with full cross-section for scale, 4 mm at the widest diameter and 10 mm maximum length) are shown overlaid on MR images of patient anatomy. Thermal diffusion was simulated for a total heating time of 16 seconds (standard for 4 mm cells) and a cooling time of 30 seconds. The X-Y cross-sections shown are those with the largest ablated tissue cross-sectional area. For all patients and both target positions, the ablated tissue is offset from the treatment cell as quantified in Table 9.4. The treatment cells and ablated tissue are positioned relative to the transducer (blue) in the X-Y plane as shown in the large subfigure.

**Table 9.4:** Differences between the intended treatment cell and the ablated tissue. For context, a 4 mm treatment cell is 4 mm in diameter (across X and Z), 10 mm in length (along Y), and 84 mm<sup>3</sup> in volume. SD is standard deviation. Thermal diffusion was simulated for a total heating time of 16 seconds (standard for 4 mm cells) and a cooling time of 30 seconds.

Subject	Target Point	Centroid Position Difference (mm)			Ablated Tissue Volume (mm <sup>3</sup> )	% of Treatment Cell Ablated	Path Length (mm)
		X	Y	Z			
G3	Shallowest	-1.5	8.0	-0.4	201	15	61.8
G8 <sub>2</sub>	Shallowest	-1.9	8.7	-0.3	213	8	64.5
G24	Shallowest	1.4	10.1	-0.7	126	0	61.6
G27	Shallowest	-1.2	6.6	-0.0	414	62	45.0
G29	Shallowest	-1.8	8.5	-0.5	318	21	51.6
Mean ± SD	Shallowest	-1.0 ± 1.2	8.4 ± 1.1	-0.4 ± 0.2	250 ± 100	21 ± 22	56.9 ± 7.4
G3	Deepest	-1.9	8.6	-0.7	143	1	64.8
G8 <sub>2</sub>	Deepest	N/A	N/A	N/A	0	0	88.7
G24	Deepest	3.0	11.9	-0.8	27	0	78.1
G27	Deepest	-2.0	11.4	-0.0	35	0	87.7
G29	Deepest	0.6	9.4	-1.3	93	0	76.5
Mean ± SD	Deepest	-0.1 ± 2.0	10.3 ± 1.4	-0.7 ± 0.4	60 ± 52	0.3 ± 0.5	79.2 ± 8.7
Mean ± SD	Overall	-0.6 ± 1.7	9.3 ± 1.6	-0.5 ± 0.4	160 ± 130	11 ± 9	68 ± 14



**Figure 9.10:** The path length in tissue is plotted against (top) ablated tissue volume and (bottom) the offset between the transducer geometric focus and the ablated tissue centroid in the Y-axis, with positive offset meaning the centroid is closer to the transducer than the geometric focus. Dotted lines are linear regression. The Y-offset associated with G8<sub>2</sub> shallowest target point, which was estimated from the Y-offsets of other patients and target points, is not included in the Y-offset plot.




### 9.3.3 Tumour Treatability

Details of the patients' tumours are shown in Table 9.5. Estimated treatable tumour volumes are shown qualitatively in Figure 9.11 and quantitatively in Table 9.6. On average, tumour treatability was less than tumour coverage by  $32\pm 14\%$  (range: 15-50%) of the tumour volume. Since the treatment cells were discovered from thermal simulation results to be, on average, centred 9.3 mm closer to the transducer surface than the geometric focus, tumour coverage and treatability (which depends on tumour coverage) was recalculated with the offset included. Those results are also shown in Table 9.6. On average, with the offset included, tumour treatability was less than tumour coverage by  $12\pm 10\%$  (range: 3-31%) of the tumour volume. The total time required to calculate tumour treatability from inputs was estimated to be  $62\pm 7$  hours (range: 42-67 hours). Without the offset, patients G27 and G29 have a closer agreement between the tumour treatability and the tumour coverage ( $<19\%$  difference), than patients G3, G8<sub>2</sub> and G24 ( $>33\%$  difference), as seen in Figure 9.12. With the offset, the agreement between tumour treatability and tumour coverage is  $<15\%$  except for patient G24, where the difference is  $>30\%$ . Tumour-skin distance (the anterior-posterior distance from tumour centroid to skin) has an uncertainty of 0.9 mm. The ablated tissue separation distance (the anterior-posterior component of the distance between the ablated tissue centroids) has an uncertainty of 0.2 mm.

**Table 9.5:** Details of patient tumours.


Subject	Tumour Volume (cm <sup>3</sup> )	Tumour-Skin Distance (mm)	Ablated Tissue Separation Distance (mm)
G3	1.20	64.6	2.4
G8 <sub>2</sub>	43.1	81.5	25.9
G24	13.2	85.0	12.9
G27	29.1	69.5	37.1
G29	26.3	67.5	27.3



**Figure 9.11:** Cross-sections of the treatable volumes (cyan) for each patient are shown overlaid on MR images of patient anatomy with the tumour coverage (yellow) and the tumour itself (red). The treatable tumour volume is a subset of the covered tumour volume, which is a subset of the tumour. The cross-sections shown here are from slices in the same plane as the cell-equivalent point (which had been interpolated or extrapolated from the centroids of the deepest and shallowest ablated tissue volumes). The transducer (130 mm aperture width) is shown for scale.

**Table 9.6:** Tumour treatability versus tumour coverage.

Subject	% of Tumour Estimated Treatable	% of Tumour Covered	% of Tumour Estimated Treatable with offset	% of Tumour Covered with offset
G3	2.5	47	0.5	5
G8 <sub>2</sub>	19	53	15	29
G24	0.3	50	0.3	31
G27	38	56	26	34
G29	67	82	59	62



**Figure 9.12:** The tumour-skin distance is plotted against the percentage difference between the tumour treatability and the tumour coverage.

### 9.3.3.1 Clinical Treatability

Data from the clinical treatment is shown in Table 9.7. If the offset between the ablated tissue and the geometric offset was not accounted for, the estimated tumour volume treatable (Table 9.6) differed from the clinically treated tissue volume by  $23\pm 25\%$  of the tumour volume. If the offset was accounted for, the difference between estimated and clinically treated tissue volume reduced to  $19\pm 21\%$  of the tumour volume.

**Table 9.7:** Clinical data on patient treatments.

Subject	Clinically Treated Tissue Volume (ml)	(as % of tumour volume)	Energy Delivered (kJ)	Acoustic Power Range (W)	No. Completed Sonications with Cell Size:	
					4 mm	8 mm
G3	0.08	7	36.4	140-250	4	5
G8 <sub>2</sub>	4.52	10	23.3	270-300	3	4
G24	0.08	1	61.7	200-270	4	11
G27	0.00	0	52.8	220-270	9	5
G29	0.63	2	97.1	250-290	22	1

## 9.4 Discussion

### 9.4.1 Acoustic Simulation of a Patient Sonication

Focal peak intensity was found to decrease linearly with increased path length in tissue, as seen in Figure 9.8 and Table 9.3. This was probably caused by an accumulation of the effects of acoustic refraction, diffraction and attenuation as the acoustic waves propagate deeper into the human body. The clinical implication of this is that the deeper treatment cells would require either greater input acoustic power or more time to ablate than those that are shallower.

In this study, the time required for each acoustic simulation was long ( $40\pm 6$  hours, range: 24-44 hours), so the acoustic pressure fields generated during sonication of the trajectory points within a 4 mm treatment cell were approximated by translation of the on-axis pressure field. The accuracy of the pressure fields for

each trajectory point may therefore be limited. In an experiment designed to assess the simulated acoustic intensity variation with lateral beam steering, linear acoustic propagation was simulated within a homogeneous water medium, with the transducer focus i) on-axis without electronic steering and ii) electronically steered laterally off-axis by 2 mm (the radius of the 4 mm treatment cell, where the trajectory points are). The peak intensity of the steered beam was 98.5% that of the on-axis beam, suggesting that the acoustic beam field differences should be negligible. Within the Sonalleve system, this energy loss due to steering is compensated for by increasing acoustic power [42, 65], with the exact compensation mechanism being proprietary. Furthermore, the heating was expected to be localised to a region within and around the intensity peak. Therefore, for this initial study, it appears reasonable to ignore the effect of electronic steering for such a small lateral deflection.

The effects of not using nonlinear simulations are that the focal peak temperature and the volume of each ablation will be underestimated [39], because the peak focal pressure is underestimated and the frequency-dependent absorption of higher frequency harmonics is ignored. Studies have been performed on the thresholds at which nonlinear effects begin to show [70]. For the settings involved in the acoustic simulations in this study, acoustic intensity was estimated as varying by approximately 6%. Hence, the linear model provides a worst-case scenario for acoustic energy deposition (see equation 2.3 and 2.5) and therefore, thermal bioeffects.

### **9.4.2 Thermal Simulation**

Thermal simulations were performed to identify whether tissue within the treatment cell would receive a cytotoxic thermal dose. A parallel to this is performed in clinical practice, where clinicians fire low acoustic power 'test shots' at a particular depth within the patient tumour and observe the resulting temperature rise induced. The purpose of this is i) to ensure that the position of the focus coincides with the heated tissue volume, and ii) to identify the acoustic power setting required to achieve clinically relevant heating (thermal ablation) at depth whilst

reducing harm to healthy tissue.

On average, the ablated tissue was found to be  $9.3 \pm 1.6$  mm (range: 6.6 to 11.9 mm) closer to the transducer along the HIFU beam axis than the transducer's geometric focus. There appears to be a relationship between the path length through tissue and the offset magnitude (see Figure 9.10). This suggests that the tumour coverage assessment methodology could be improved by centering the intended 4 mm treatment cells closer to the transducer than the geometric focus. The negative correlation between ablated tissue volume and path length (see Figure 9.10) reflects the negative correlation between peak focal pressure and path length (Figure 9.8), which arises from refraction and diffraction blurring the focus and attenuation reducing the energy that reaches the focus. The difference in the ablated tissue volume between two patients with similar path lengths ( $201 \text{ mm}^3$  and  $126 \text{ mm}^3$  when targeting the shallowest reachable treatment cell in patients G3 and G24, respectively) was probably due to the different patient geometries. For G3, the acoustic beam path mainly travels through adipose tissue before reaching the tumour (see Figure 9.11), whereas with patient G24, a section of the acoustic beam travels through adipose tissue, muscle, and then re-enters adipose tissue (same figure). This highlights the importance of using patient-specific simulations to determine patient treatability.

In this study, acoustic energy deposition calculation was based on linear acoustics, representing the worst-case scenario for tissue heating. All simulations were performed with the source acoustic power at 300 W, which is the maximum acoustic power output of the Sonalleve<sup>®</sup> system in the clinical uterine configuration. The ablated tissue volume when targeting the shallowest reachable treatment cell ( $250 \pm 100 \text{ mm}^3$ ) exceeded the treatment cell volume ( $84 \text{ mm}^3$ ), suggesting that the acoustic source power of 300 W was excessive for tissue ablation at this depth. However, when targeting deeper treatment cells, the ablated tissue volume ( $60 \pm 52 \text{ mm}^3$ ) was less than the treatment cell volume, suggesting that more acoustic power than 300 W is required. There are several caveats to be noted. Thermal simulations were conducted with zero blood perfusion, which would also imply an overestimation in thermal dose. For patient G27, when target-

ing the shallowest reachable treatment cell, a thermal simulation including constant blood perfusion (blood density  $1050 \text{ kg m}^{-3}$ , blood specific heat  $3617 \text{ J kg}^{-1} \text{ K}^{-1}$ ) and blood perfusion rate  $0.01 \text{ s}^{-1}$ ) was performed. The ablated tissue volume was reduced to  $261 \text{ mm}^3$  and computational time was increased to 25 hours. Future work could be done on developing a more accurate model of HIFU-induced heating.

### 9.4.3 Tumour Treatability

Due to prohibitive computation time and resource requirements for simulations, a method which relied on identifying the maximum treatable depth was used to estimate the tumour treatability (see Figure 9.6) and then marking all covered tumour tissues shallower than that as treatable. For each patient, the maximum treatable depth was estimated from a patient-specific linear relationship generated by ablated tissue estimates at two treatment cell positions. This was done because tissue distributions are unique to each patient, and linear interpolation required only two acoustic simulations per patient, which is the extent of data available at that time. Given that ablated tissue volume was discovered to have a negatively correlated relationship with the path length of the acoustic beam through tissue (Figure 9.10), this interpolation seems reasonable.

If the offset between the ablated tissue volume and the geometric focus was accounted for, the difference between the tumour treatability and coverage decreased from an average of  $32 \pm 14\%$  to an average of  $12 \pm 10\%$  of the tumour volume. Generally, with the offset included, the difference between tumour coverage and treatability (see Figure 9.12 and Tables 9.5 and 9.6) increased with tumour depth. The difference between tumour coverage and treatability probably arose from the variation in ablated tissue volume with depth, and from the ablated tissue shape being non-ellipsoidal due to acoustic wave propagation through the intervening tissues. Nonetheless, the reduction in the difference between tumour coverage and estimated tumour treatability suggests that incorporating the results of the simulations into the tumour coverage methodology could result in better quality predictions of patient suitability for MRgHIFU therapy of pelvic tu-

mours. However, a generalised method for calculating tumour coverage, which incorporates the effects of the intervening tissue on the position, shape and size of the ablation region, has not yet been developed. Further development and refinement of the tumour coverage methodology could be pursued prior to clinical deployment, in order to deploy .

For patient G8<sub>2</sub>, the estimated treatability may have been overestimated. This is because, even though the shallowest target point for G8<sub>2</sub> has zero ablated tissue volume, it may not be the shallowest target point that has no ablation. With the current method, a shallower zero-ablation target point implies a shallower maximum treatable depth. Hence, the treatability for G8<sub>2</sub> could be overestimated. In future work, more target points within the tumour could be assessed so the maximum treatable depth could be determined more accurately.

#### 9.4.3.1 Clinical Treatability

The clinical trial in which the patients had been enrolled had set the threshold for proceeding from screening to treatment at 50% tumour coverage [18]. All patients mentioned in this thesis had been assessed visually for this threshold, and all had proceeded to treatment. From quantitative tumour coverage analysis, 4/5 patients (all except G3) would have passed screening if the offset was not accounted for, and 1/5 patients (G29) would have passed otherwise. From the estimated treatability, only 1/5 patients (G29) would have been predicted to receive a cytotoxic thermal dose to >50% of the tumour.

In practice, the volume of tissue ablated in all patients was <10% that of their tumour volumes. The discrepancy between the estimated and clinically treated tissue volume was  $23\pm 25\%$  of the tumour volume without accounting for offset, and  $19\pm 21\%$  of the tumour volume otherwise. An important point to note is that, according to the clinical protocols, the volume of tissue clinically ablated may have included tissue outside the tumour. Therefore, the difference between estimated treatability and the clinical treated tumour volume could be greater than the numbers shown. The discrepancy between estimated and clinical treatability could have arisen from several factors. Firstly, the treatment could have stopped



early due to skin burns and subcutaneous fat damage. Secondly, because tumours are clinically treated by first targeting the deepest portions of the tumour and then progressively shallower regions [28], it is possible the clinicians could have stopped the treatment once they realised that less than 50% of the tumour would be treatable. Thirdly, blood perfusion into tissue had not been accounted for in the thermal simulations performed here. According to the bioheat equation (eq 2.4), non-zero perfusion would reduce the maximum temperatures reached within the tissue, thus reducing the volume of ablated tissue and making the maximum treatable depth shallower. Fourthly, treatability had been estimated using simulations where the source acoustic power was 300 W, whereas the maximum acoustic powers of all patients (except G8<sub>2</sub>) was less than that. Fifthly, the MR thermometry could have been inaccurate, perhaps due to the presence of fat surrounding or within the tumour. Lastly, the methodology for estimating tumour treatability from linear interpolation of two points could be inaccurate.

In clinical practice, if tissue does not reach ablation temperatures or does not receive a cytotoxic thermal dose, clinicians could select a smaller treatment cell and pack treatment cells more closely in order to cumulatively deliver the required treatment dose. However, the 4 mm treatment cell was the smallest cell available clinically [2]. In clinical experience, the largest ablated tumour volume (4.5 cm<sup>3</sup>) was observed for patient G8<sub>2</sub>, with the next greatest ablated volume (0.63 cm<sup>3</sup>) observed for G29. The predicted tumour treatability for G8<sub>2</sub> was 8.2 cm<sup>3</sup> and for G29, was 17.6 cm<sup>3</sup>. Clinical measurements of ablated volume were made from MR thermometry, in which spatial resolution, temporal resolution and temperature resolution were balanced against each other, leading to uncertainties. Differences between predicted and clinical treatability suggest that the treatability assessment method needs to be improved. The development of treatment protocols for gynaecological cancer ablation, such as increasing the sonication time for treatment cells, should be pursued as part of future work.

#### 9.4.4 Limitations of this Study

There are several limitations to this study, many of which stem from the resource-intensive acoustic and thermal simulations. A linear acoustic simulation was used, because a nonlinear simulation required such a fine spatial resolution (13.5 grid points per wavelength [84]) to generate accurate results that simulation of each exposure was estimated to require more than 1 week of computational time on our current computational system, or a supercomputer with several hundred cores, which was not available to us. This resulted in an underestimate of the acoustic pressure at the focus [39] and thus of ablated tissue volume. Furthermore, because of the long acoustic simulation time, the tumour treatability had to be estimated using the minimal number of simulations, instead of taking months to assess patient treatability by exhaustive simulation of every reachable treatment cell.

Acoustic simulations could be sped up by using a faster computational approach, such as the angular-spectrum methods discussed in Section 2.6.3. Simplification to a two-dimensional system would speed up computation, but would also generate inaccurate results because the source is three-dimensional and the geometry of the body is heterogeneous and has no underlying symmetries. Fast acoustic simulations could allow the incorporation of nonlinearity and the simulation of steered acoustic beams within tissue when generating acoustic pressure fields for the sonication of treatment cells without increasing the time requirements unreasonably. Thermal simulation could be sped up with improved software engineering, such as implementation of the code in a compiled language such as C++ or GPU implementation. The k-Wave acoustic code has already been optimised with these methods [84]. Faster thermal simulation could allow incorporation of the switching time, when the transducer is switched off and re-focused to another trajectory point, into the thermal simulation without drastically increasing computation time. Faster acousto-thermal simulations may also allow simulation of sonicating more regions within the tumour, thus improving the accuracy of tumour treatability estimation without increasing computational time

unreasonably.

Several simplifications had been made over the course of the project. First, blood perfusion was not incorporated into thermal simulations. Blood perfusion into tissue (healthy and malignant) is known to vary with temperature [187]. Secondly, the tumour has been assigned material properties that are interpolated from the material properties of fat and muscle. Inclusion of more realistic information could result in more accurate simulations.

As mentioned in previous Chapters, the limited size of the patient cohort means that conclusions about the relationship between patient treatability and tumour depth, or between ablated tissue volume and path length in tissue, are not definitive. With a larger patient data set, instances of outliers, such as patient G3, could be clearly distinguished.

## 9.5 Conclusions

A method for estimating patient treatability has been developed using treatment image data and acousto-thermal simulations. Focal peak pressure, ablated tissue volume, and the offset between the ablated tissue centroid and the geometric focus, varied linearly with acoustic path length. In the future, simple models could be built using this information to assist in the ad-hoc treatment planning process currently employed clinically. Tumour treatability, estimated from interpolation of ablation volumes, was less than tumour coverage as calculated in Chapter 7 by  $32\pm 14\%$  (range: 15-50%), which reduced to  $12\pm 10\%$  (range: 3-31%) if the offset between the ablated tissue centroid and the geometric focus was accounted for. Generally, the magnitude of disagreement between tumour coverage and treatability increased with tumour depth. The reduced difference between tumour coverage and treatability after accounting for the offset suggests that, after incorporating the results of simulations into tumour coverage analysis, tumour coverage prediction could be developed as a rapid screening tool for clinical use. To fully assess whether the assumptions made in estimating tumour treatability are correct, more work is required, particularly in the development of a method to predict the soft tissue deformation resulting from reorientation of the patient from

supine to treatment position.

The estimated tumour treatability differs from the clinically treated tissue volume by  $19\pm 21\%$  even after accounting for offset, suggesting that the methodology for estimating treatability requires more development. Several hypotheses for explaining the difference were posited, and research into the reasons behind the difference is the subject of future work.

## Chapter 10

# Overall Conclusions and Future Work

### 10.1 Overall Conclusions

The overarching aim of this PhD project has been to develop aspects of a quantitative screening workflow for assessing whether patients with pelvic tumours are suitable for MRgHIFU therapy. Significant progress has been made in developing various aspects of the workflow outlined in Chapter 1 (Figure 1.1).

Novel methods for predicting the MRgHIFU target coverage from supine MR referral images have been developed using 10 volunteer datasets and were then applied retrospectively to 5 patient datasets. For patients, the difference between the target coverage computed using referral images and that using treatment images was within  $12\pm 7\%$  on average (range: 4-21%), when one patient, in which inadequate acoustic coupling was achieved during treatment, was excluded from analysis. Tumour coverage assessment was refined by the development of the target-first method, which derived possible transducer positions and angulations from consideration of the target, and GPU acceleration, resulting in a sevenfold increase in computational speed (from  $19.7\pm 8.8$  to  $2.8\pm 2.0$  hours) with minimal ( $0.31\pm 0.08\%$  difference, range: 0.18-0.38%) effect on tumour coverage results. This allowed tumour coverage to be assessed in a clinically relevant timeframe.

A method for identifying the ideal treatment angle in which to position the patient, in order to maximise tumour coverage, from referral images has been

developed and tested on patient data for future prospective use in the workflow. The method first determines a 'starting' treatment angle, for which the acoustic window was maximised. The ideal treatment angle was then found, for each patient, by exhaustively examining steeper (up to 40°) and shallower (down to 0° i.e. supine) angles than their starting angle. For the five patients studied, the angular difference between the starting and ideal treatment angles was  $4\pm 4^\circ$  (range: 0-10°), and that between the ideal and clinical treatment angles was  $5\pm 2^\circ$  (range: 2-7°). The maximum effect of patient positioning up to 5° away from the ideal angle was estimated to be a 7% decrease in tumour coverage, suggesting that imprecise patient positioning does not present a major limitation to the clinical adoption of this method. Tumour coverage at the starting treatment angles was within 5% of the maximum tumour volume covered, which occurred at the ideal treatment angle, for all patients, suggesting that the starting treatment angle could be used as a faster-to-calculate substitute for the ideal treatment angle.

A method for estimating patient treatability has been developed using treatment image data and acousto-thermal simulations. Focal peak pressure, ablated tissue volume, and the offset between the ablated tissue centroid and the geometric focus, varied linearly with acoustic path length. Simple models could be built using this information to assist in the ad-hoc treatment planning process currently employed clinically. Tumour treatability, estimated from interpolation of ablation volumes, was less than tumour coverage as calculated in Chapter 7 by  $32\pm 14\%$  (range: 15-50%) of the tumour volume, and the difference was reduced to  $12\pm 10\%$  (range: 3-31%) if the offset mentioned previously was accounted for when assessing tumour coverage. Generally, the magnitude of disagreement between tumour coverage and treatability increased with tumour depth. The reduced difference between tumour coverage and treatability after accounting for the offset suggests that, after incorporating the results of simulations into tumour coverage analysis, tumour coverage prediction could be developed as a rapid screening tool for clinical use. To fully assess whether the assumptions made in estimating tumour treatability are correct, more work is required, particularly in the development of a method to predict the soft tissue deformation resulting from

reorientation of the patient from supine to treatment position.

As patients were treated by MRgHIFU, clinicians monitored tissue temperatures within patients using MR thermometry. The estimated tumour treatability differs from the clinically treated tissue volume by  $19\pm 21\%$  even after accounting for offset, suggesting that the methodology for estimating treatability requires more development. Several hypotheses for explaining the difference were posited, and research into the reasons behind the difference is the subject of future work.

Despite the relatively small patient cohort size, the focus on pelvic tumours, and the limited range of patient positions and MRgHIFU equipment on which the methodology was devised and tested, results are encouraging. The development of a quantitative screening workflow, which should obviate the need for patient suitability to be assessed using clinical judgement based on operator experience, should be feasible with further work.

## **10.2 Future Work**

There is considerable opportunity for future work designed to bring the quantitative screening workflow to a clinically useful state. Several areas, described in subsections below, were identified as possible paths for future research to take.

### **10.2.1 Treatability Estimation**

The highest priority problem for future work to tackle would be to explain why the estimated tumour treatability differed greatly from the clinically treated tissue volume. As discussed previously, this could have resulted from clinical factors, as well as inaccuracies in tumour treatability estimation. Here, several strategies for improving the treatability estimation method are proposed.

In the current methods described in the thesis, blood perfusion had not been accounted for. Nonzero blood perfusion in tissue would reduce the temperature rise resulting from ultrasound absorption. This could explain why the estimated treatability is greater than the clinical treated tissue volume. Afterwards, if treatability was less than the clinically treated tissue volume, nonlinear acoustic simulations could be performed in order to assess whether the increased amplitudes

and absorption resulting from nonlinearity could explain the difference.

Due to the high computational requirements of acousto-thermal simulation, the tumour treatability assessment methodology presented here relied on some simplifying assumptions. Future work could focus on testing whether the assumptions made were valid, and developing a more sophisticated tumour treatability assessment methodology. For example, instead of the current method involving the calculation a maximum treatable depth in the anterior-posterior direction only, a surface representing the edge of the treatable region could be interpolated from multiple acousto-thermal simulations and used to determine tumour treatability. This may require the development of rapid methods for acoustic and thermal simulations.

Research performed here could also be used in applications beyond patient screening. One benefit of MRgHIFU is the ability to perform MR thermometry, which allows clinicians to use low-power sonications ('test shots') in order to calibrate the control software, so that treatment cells in the software align with the observed region of temperature increase. Furthermore, clinicians use the thermal response of the test shots to subjectively determine the suitable source acoustic power for each treatment cell. Currently, MR thermometry is used to control the sonication of feedback cells [65, 106]. In the future, software could be developed to quantitatively predict patient treatability based on test shots. For example, a machine-learning model could be built that uses thermal response information from test shots to predict the expected temperature increases at higher acoustic powers. The advantage of such software would be that, if a patient was untreatable, the clinicians would be informed of this without sonicating at maximum power and damaging the patient's healthy tissues.

Tumour coverage assessment could be further improved by taking into account the variation in offset distance between the ablated tissue volume and the geometric focus, and in the ablated tissue volume, with acoustic path length in tissue. Through improvements in the accuracy of calculating tumour coverage, identification of the ideal treatment position could also be improved. Furthermore, incorporating the results of simulations into tumour coverage assessment could



assist in the development of a rapid alternative to full simulation in determining patient suitability for MRgHIFU.

### **10.2.2 Structural Deformation**

After an accurate treatability estimation method is developed, the next priority would be to develop a method to simulate the geometric deformation of the patient after reorientation from referral to treatment position. By estimating the human body shape and tissue distribution upon reorientation from referral to treatment position, tumour treatability estimation from referral imaging could be made more plausible because patient geometry is known to be important in determining the acoustic intensity at the focus [64, 83], the ablated tissue volume and hence, tumour treatability. An accurate soft tissue deformation simulation would allow patient treatability assessment from referral imaging, as aspired to in the prospective patient workflow (Figure 1.1). Soft tissue deformation could be simulated using finite element modelling [188–196]. Simulating the patient's geometric deformation could also improve tumour coverage assessment by improving estimation of the MRgHIFU transducer position relative to referral images that have been rotated into treatment orientations. Furthermore, soft tissue deformation simulations could model how organs at risk deform and move within the body as the patient is rotated into treatment positions.

### **10.2.3 Robustness**

If more patient data could be obtained and accessed, a better understanding of the robustness and quality of the methodology could be obtained. A wider variety of patient data could also allow screening techniques to be developed for tumour target sites outside the pelvic region. Furthermore, development of automatic image segmentation methods, such as atlas-based or statistical shape methods, which require many manually segmented datasets to construct automatic segmentation models, could be achieved with a larger patient cohort.

## 10.2.4 Clinical Translation

Clinical translation is easier if useful results could be obtained within a reasonable timeframe. Currently, long computation times present obstacles to the adoption of these quantitative screening techniques. The time required for each component of the prospective patient workflow is summarised here: manual image segmentation (estimated to be approximately 8-12 hours of continuous work), identification of the ideal patient position ( $39\pm 28$  hours) and treatability assessment ( $62\pm 7$  hours). In general, computational time could be reduced with better software engineering (such as writing the algorithms in a compiled language such as C++) and taking advantage of advances in computational hardware, such as GPUs and FPGAs.

In this thesis, image segmentation was performed manually. For patients, organs at risk and acoustic obstructions were only segmented in the region of the expected beam paths. Recent advances in machine learning and deep learning could enable the automation of image segmentation [197, 198], and reducing the time required for manual contouring and hence making the quantitative screening workflow more practical for clinical use.

Software could be developed to allow real-time manipulation of the 3D graphical models of organs at risk, bone and the tumour generated from meshing image segments (see Section 7.2.2.1). This visualisation of the patient anatomy could quickly and effectively inform clinicians on any overlap between the acoustic obstructions and organs at risk, and the proposed acoustic beam. Clinicians could then use this anatomical information to quickly determine suitable placement of treatment cells, and hence interactively derive tumour coverage and patient suitability within the current clinical workflow. In the proposed quantitative workflow, clinician decisions on suitable placement of treatment cells could be used to inform the placement of other treatment cells. Furthermore, clinicians could use this anatomical information to propose an initial estimate for the ideal treatment angle.

In this thesis, the methodology for identifying the ideal treatment angle was


exhaustive, with it being identified by calculating tumour coverage for a range of possible treatment angles between the patient lying supine and the estimated maximum tilt achievable within the MR scanner bore. An optimisation-based method, which would select the next treatment angle to be assessed for tumour coverage from the previously obtained tumour coverage results, could reduce the number of calculations, and therefore speed up the process into a more clinically reasonable timeframe.

Improvements could be made to the acoustic simulation methodology, by using a faster method, such as hybrid angular spectrum, to reduce the time required to calculate treatability with the methodology introduced in this thesis. This would require development and validation of a software package that uses this methodology. With improvements in computational speed, nonlinearity and temperature-dependent acoustic properties could be incorporated into acoustic simulations without increasing computational time unreasonably. In the literature, deep learning has been developed as a method of solving similarly complex problems (structural deformation [199], fluid dynamics [200]) that compares favourably with simulation from first principles, suggesting that development of a fast, deep learning-based method for acoustic and thermal simulation model is feasible. A fast, accurate method of quantifying patient treatability could allow more suitable patients access to MRgHIFU therapy, and reduce hospital resources spent on patients who turn out to be unsuitable.


## **Appendix A**

# **Volunteer and Patient Data**

Representative axial slices of volunteers and patients are shown in Figures A.1 and A.2 respectively. Volunteer slices are the axial slices closest to the magnetic isocentre. Patient slices are the axial slices closest to the tumour centroid.



**Figure A.1:** Representative axial slices of volunteers, obtained at the magnetic isocentre. Only the referral, steep and shallow treatment images were used in this project.



**Figure A.2:** Representative axial slices of patients, obtained at the tumour centroid. The tumour is highlighted in red.

## Appendix B

# Standard Operating Procedure for Image Registration

The below SOP is taken from an appendix in my paper [161] (see Appendix D).

1. Open Horos on Mac OS X. Make sure the pyOsiriX plugin [144] is installed.
2. Import the in-phase MRI datasets that are to be registered. Double click them to bring them up together.
3. Select a dataset. Then, at the top menu, select 2D Viewer → Sort By... → Slice Location Ascending.
4. Below the menu bar, in a section titled “Mouse button function”, select the point function. Use the point function to mark an anatomical feature on one dataset and the same anatomical feature on the other. The same point names, e.g. “Point 1”, must correspond to the same anatomical features in both datasets. Repeat this for the list of anatomical features mentioned below. If the same anatomical feature cannot be found in one or both of the datasets, ignore that anatomical feature and continue down the list. At least 10 features should be marked by the end.
  - (a) Femur/pelvis landmark marks where the two bones meet in the head-most direction (Right and Left)
  - (b) Ischial spine (Right and Left)

- (c) Superior-most or inferior-most of ischial tuberosity (Right and Left)
  - (d) Pubic arch
  - (e) Anterior-facing spur in axial plane where pelvis first encloses femur head (Right and Left)
  - (f) Sacral nerve bundle (S1 and S2) when just-enclosed by bone (Right and Left)
  - (g) Spinal nerves splitting from spinal cord (Right and Left)
  - (h) Sacrum/L5 disc
  - (i) Coccyx
5. Open the pyOsiriX console within Horos. A Python script can be used to extract point data from a dataset in Horos and save it as a pickle, which can then be processed in an external Python environment. Do this for both datasets.



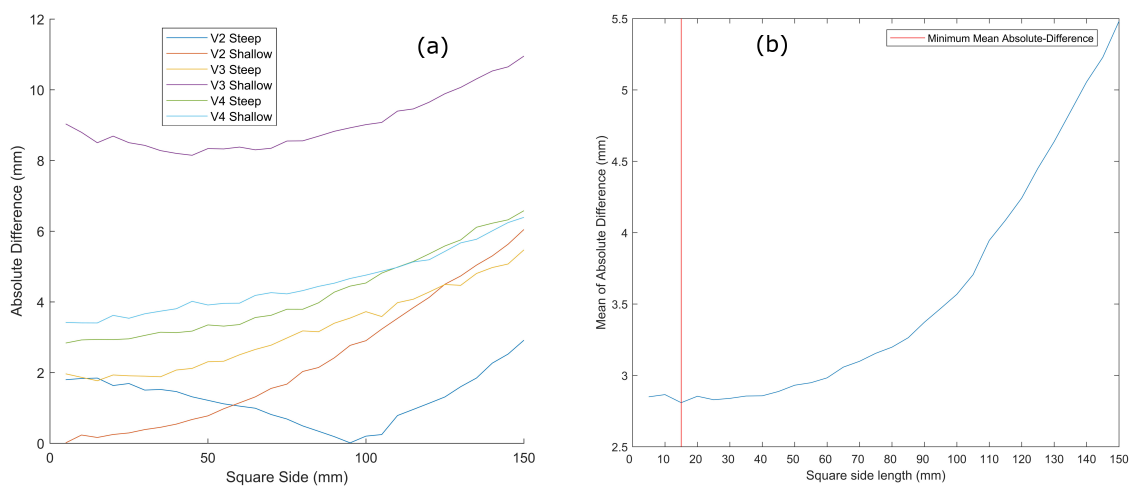
## Appendix C

# Investigation of the Ideal Sample Square Size

An investigation was performed to identify the ideal size of the square within which to begin sampling the skin position, for the purposes of accounting for expected body deformation (see Section 6.2.2).

Squares of different side lengths ranging from 5 mm to 150 mm in 5mm increments, all with the centre at the Left-Right Inferior-Superior zero-position and with the face normal to the Left-Right and Inferior-Superior axes. The number of samples taken within each square was the side length in millimetres squared. Three volunteers' (Volunteers 1, 2, 3) data was used, in both Steep and Shallow configurations. Rather than a time-consuming tumour coverage calculation, the home position predicted from the registered-referral image dataset (using various square sizes) was compared to that calculated from the treatment image dataset. The square size with the minimum distance between these positions is the ideal size of the square.

Results are shown in Figure C.1. The square side length with the minimum mean absolute difference is 15 mm.



**Figure C.1:** Graphs showing (a): the absolute difference between treatment and registered-Referral virtual transducer zero-position for each volunteer and treatment position tested and (b): the mean absolute difference in zero-position over all tested volunteers and treatment positions. The red line denotes the square side length with the minimum mean absolute difference, which is 15 mm.

## **Appendix D**

**Paper: Prediction of pelvic tumour coverage by magnetic resonance-guided high-intensity focused ultrasound (MRgHIFU) from referral imaging**



## Prediction of pelvic tumour coverage by magnetic resonance-guided high-intensity focused ultrasound (MRgHIFU) from referral imaging

Ngo Fung Daniel Lam , Ian Rivens , Sharon L. Giles , Emma Harris , Nandita M. deSouza & Gail ter Haar

To cite this article: Ngo Fung Daniel Lam , Ian Rivens , Sharon L. Giles , Emma Harris , Nandita M. deSouza & Gail ter Haar (2020) Prediction of pelvic tumour coverage by magnetic resonance-guided high-intensity focused ultrasound (MRgHIFU) from referral imaging, International Journal of Hyperthermia, 37:1, 1033-1045, DOI: [10.1080/02656736.2020.1812736](https://doi.org/10.1080/02656736.2020.1812736)

To link to this article: <https://doi.org/10.1080/02656736.2020.1812736>



© 2020 The Author(s). Published with license by Taylor & Francis Group, LLC



Published online: 01 Sep 2020.



Submit your article to this journal [↗](#)



Article views: 254




View related articles [↗](#)



View Crossmark data [↗](#)

## Prediction of pelvic tumour coverage by magnetic resonance-guided high-intensity focused ultrasound (MRgHIFU) from referral imaging

Ngo Fung Daniel Lam<sup>a</sup> , Ian Rivens<sup>a</sup> , Sharon L. Giles<sup>b</sup> , Emma Harris<sup>a</sup> , Nandita M. deSouza<sup>b</sup>  and Gail ter Haar<sup>a</sup> 

<sup>a</sup>Joint Department of Physics, The Institute of Cancer Research, London, UK; <sup>b</sup>The CRUK Cancer Imaging Centre, The Institute of Cancer Research and The Royal Marsden NHS Foundation Trust, London, UK

### ABSTRACT

**Background:** Patient suitability for magnetic resonance-guided high intensity focused ultrasound (MRgHIFU) ablation of pelvic tumors is initially evaluated clinically for treatment feasibility using referral images, acquired using standard supine diagnostic imaging, followed by MR screening of potential patients lying on the MRgHIFU couch in a ‘best-guess’ treatment position. Existing evaluation methods result in  $\geq 40\%$  of referred patients being screened out because of tumor non-targetability. We hypothesize that this process could be improved by development of a novel algorithm for predicting tumor coverage from referral imaging.

**Methods:** The algorithm was developed from volunteer images and tested with patient data. MR images were acquired for five healthy volunteers and five patients with recurrent gynaecological cancer. Subjects were MR imaged supine and in oblique-supine-decubitus MRgHIFU treatment positions. Body outline and bones were segmented for all subjects, with organs-at-risk and tumors also segmented for patients. Supine images were aligned with treatment images to simulate a treatment dataset. Target coverage (of patient tumors and volunteer intra-pelvic soft tissue), i.e. the volume reachable by the MRgHIFU focus, was quantified. Target coverage predicted from supine imaging was compared to that from treatment imaging.

**Results:** Mean ( $\pm$ standard deviation) absolute difference between supine-predicted and treatment-predicted coverage for 5 volunteers was  $9 \pm 6\%$  (range: 2–22%) and for 4 patients, was  $12 \pm 7\%$  (range: 4–21%), excluding a patient with poor acoustic coupling (coverage difference was 53%).

**Conclusion:** Prediction of MRgHIFU target coverage from referral imaging appears feasible, facilitating further development of automated evaluation of patient suitability for MRgHIFU.

### ARTICLE HISTORY

Received 1 August 2019  
Revised 13 August 2020  
Accepted 16 August 2020

### KEYWORDS

Treatment planning; magnetic resonance imaging guidance; high intensity focused ultrasound; human body deformation; pelvis; referral imaging



## 1. Introduction

Magnetic resonance guided high-intensity focused ultrasound (MRgHIFU) is a noninvasive, non-ionizing treatment modality which has a number of established clinical applications including the ablation of uterine fibroids and bone nerves (for pain palliation) [1], and the treatment of essential tremor [2]. In addition, MRgHIFU is being trialed in the UK for the thermal ablation of recurrent gynaecological tumors (NCT02714621) [3].

MRgHIFU therapy of pelvic tumors is particularly challenging because of the depth of the tumors within the body. MRgHIFU systems can only treat targets within the focal length constraints of their transducers, and identifying acoustic access which is free from obstruction by acoustically opaque tissues, such as gas and bone, and from organs at risk is challenging [3]. Failure to correctly identify suitable patients for MRgHIFU therapy could deprive them of their only treatment option, while failure to identify patients who cannot be treated wastes patient time and hospital resources on

screening sessions. Patients must therefore be carefully assessed prior to being accepted for treatment. We hypothesize that an algorithm could be developed, that could accurately predict target tumor coverage by HIFU from referral imaging.

Currently, the clinical evaluation process relies heavily on experience and opinion. The process is as follows: patients are referred to the MRgHIFU clinic on the basis of supine diagnostic imaging, often follow-up imaging after unsuccessful prior treatment [3,4] and referred to here as the ‘referral image dataset’. If at this point treatment appears qualitatively feasible, patients progress to the screening stage. At screening, patients are imaged with treatment conditions being mimicked as closely as possible. Patients are asked to lie in one or two ‘best guess’ treatment positions on the MRgHIFU couch. The ‘best guess’ positions are identified by the treatment team using prior clinical experience and subjective judgment. Suitable patients, those for whom a majority of the tumor can be reached or who fulfill clinical trial eligibility criteria, are invited back for treatment. The current

**CONTACT** Ngo Fung Daniel Lam  [Daniel.Lam@icr.ac.uk](mailto:Daniel.Lam@icr.ac.uk)  Institute of Cancer Research, 15 Cotswold Road, Sutton, Surrey, SM2 5NG, UK

© 2020 The Author(s). Published with license by Taylor & Francis Group, LLC

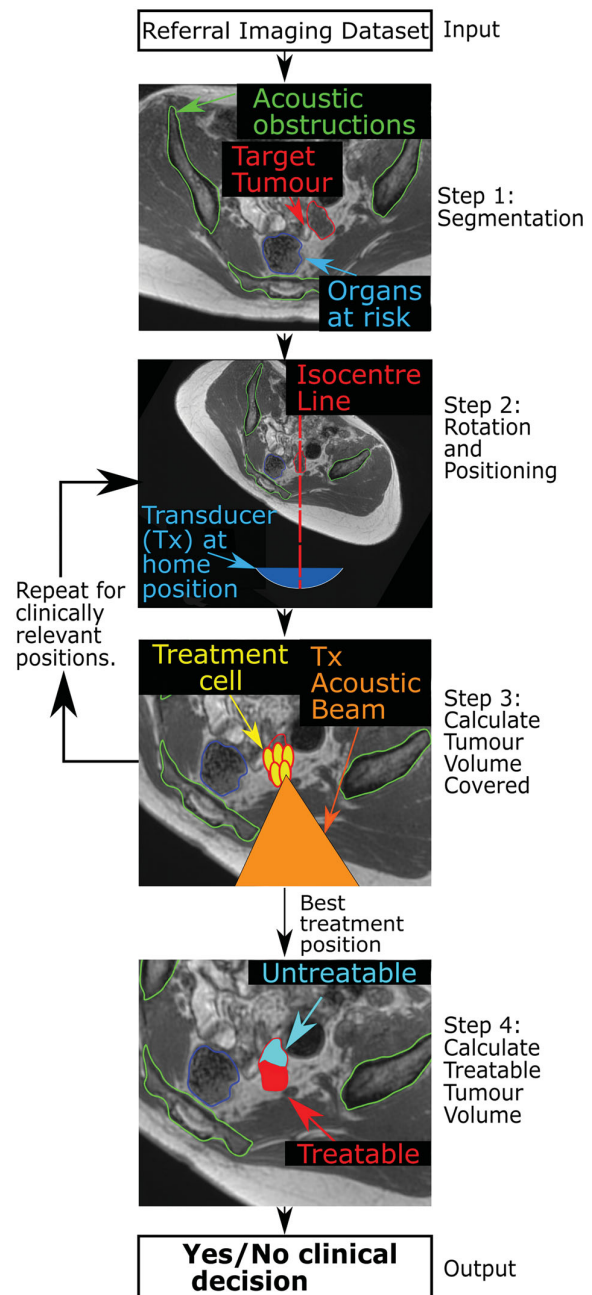
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

process is challenging. In a previous metastatic bone pain palliation trial, 16 of 37 patients (43%) initially considered for treatment were found at screening not to satisfy eligibility criteria because of disease that could not be targeted, for reasons that include tumor accessibility and size [4]. In a pilot planning study which assessed MRgHIFU for the treatment of recurrent gynaecological tumors, 9 of 20 eligible patients (45%) who underwent screening imaging were subsequently assessed as untreatable because of an eligibility criterion, namely, that >50% tumor coverage could be achieved without risk of damage to surrounding structures [3]. These two studies suggest that, for abdominal pelvic tumors, the current evaluation process may overestimate the number of patients that are suitable for MRgHIFU by more than 40%.

Given the relatively poor results of the current subjective method, we propose a workflow that would ultimately be suitable for the quantitative assessment of patient suitability for MRgHIFU therapy (Figure 1). In this paper, we focus on a core aspect of that workflow, as explained below. If the workflow were to be successfully implemented, the number of patients incorrectly denied treatment could be minimized, and the number who would benefit from a screening scan could be maximized. In the long-term, it may even be possible to avoid the need for a screening visit, which could mean that a sick patient will no longer need to travel to the magnetic resonance (MR) imaging unit and undergo what may be a lengthy session in which optimal treatment positions are investigated, only to return days to weeks later for a treatment session. This may also reduce the load on the resources of a busy clinical MR department.

The proposed patient workflow (Figure 1) comprises three steps. In Step 1, key anatomical components that could prevent access to targets, such as acoustic obstructions and organs at risk, are segmented from the referral images. In Step 2, the referral imaging dataset is orientated into plausible potential treatment positions. In Step 3, the percentage of tumor volume that can be reached by the HIFU focus (% target volume covered) is calculated at each orientation. In Step 4, acoustic and thermal modeling are used to calculate the treatable target volume, in order to facilitate a quantitative clinical decision as to whether a patient should proceed to screening.

The focus of this paper is Step 3, the calculation of tumor coverage. As far as the authors are aware, no previous work has been done on predicting target tumor coverage from referral images. A novel method has been developed to identify the tumor coverage that could be achieved in the presence of acoustic obstructions and organs at risk, and using this methodology, a feasibility study has been performed to determine whether it is possible to accurately predict tumor coverage from referral imaging by comparison with predictions made using subjects lying in treatment orientations. For this purpose, volunteer imaging data were obtained, and used to develop novel data processing and analysis techniques for the calculation of tumor coverage. Subsequently, the method was tested using patient data obtained in a concurrently started clinical trial.



**Figure 1.** Schematic of proposed patient workflow. Workflow designed to assess the potentially MRgHIFU-treatable percentage of a patient's target tumor. Using a supine referral image dataset, step 1 involves segmentation of important structures: organs at risk, acoustic obstructions, and the target tumor. Step 2 rotates the referral imaging dataset into possible treatment positions, with the tumor centroid lying, by idealized design, along a vertical line through the magnetic isocentre and, by system design, the transducer's home position. In step 3, target coverage (i.e. percentage of target volume coverable by an 8 mm treatment cell) is calculated. Cycling through steps 2 & 3 identifies the patient orientation with the maximum target volume coverage. In step 4, the treatable percentage of the target volume is quantified, using acoustic and thermal modeling of MRgHIFU treatment. This allows a clinical decision of whether to progress to treatment to be made.

## 2. Methods

### 2.1. Overview

In order to evaluate the developed methodology for the calculation of tumor coverage, estimations of target (tumor)

coverage from referral and treatment images obtained for each volunteer (patient) were compared. Here, the referral imaging dataset is the expected input into the prospective patient workflow and is used to predict target volume coverage. We assume the treatment images depict the subject positioned in a plausible (volunteer) or actual (patient) treatment position, respectively, on the MRgHIFU bed. The treatment imaging dataset is used to calculate the target coverage. The workflow used in this study is shown in Figure 2. As the treatment position is known from the treatment images, the referral imaging dataset was oriented into the known treatment position to compare the predicted target coverage with the actual target coverage. This was achieved by an affine registration of the referral imaging dataset to the treatment imaging dataset (Step 1 in Figure 2). Segmentation of the acoustic obstructions and organs at risk (Step 2 in Figure 2) from both datasets was performed to identify tissues that could prevent target coverage. This was followed by calculation of the target (tumor) coverage (Step 3 in Figure 2) and comparison of the results for predictions from referral imaging datasets with those from treatment imaging datasets.

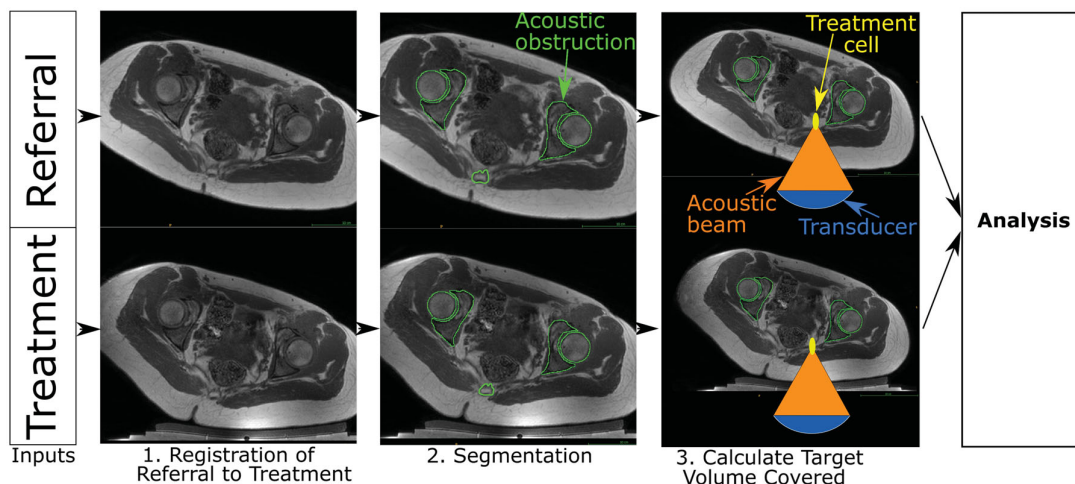
At the start of the project, clinical trial data were not available. The method was therefore developed using volunteer imaging data, with the goal of testing it on anticipated clinical datasets. As a result of significant anatomical differences between volunteers and patients, some adaptation was necessary. Firstly, volunteers lacked target tumors. This could have been addressed by the creation of dummy tumors, but in the absence of an obvious method for defining the size, shape and position of dummy tumors in an unbiased and clinically relevant way, all the soft tissue in the pelvis was defined as 'target tissue'. Secondly, while patients undergo dietary and physical bowel preparation prior to treatment in order to minimize the risk of bowel and rectal damage, volunteers were

not required to do so. These tissues were therefore not considered to be organs-at-risk when processing volunteer data. While these two limitations present challenges, they do not prevent like-for-like comparison between target coverage predictions from referral and treatment imaging datasets. Datasets from 5 volunteers, comprising pseudo-referral and pseudo-treatment imaging datasets were available for the development of the method. The methodology was subsequently tested on 5 patients who had undergone ablative MRgHIFU treatment for recurrent gynaecological tumors.

## 2.2. Input images

All subjects were scanned on a 3.0T Philips Achieva<sup>®</sup> MR scanner (Amsterdam, Netherlands), using a multi-point Dixon sequence [5] (TE1/TE2 = 1.186 (out-of-phase)/2.372 (in-phase) ms, TR = 3.62 ms, number of echoes = 2, flip angle = 10°). This produced four 3D image sets for each referral and treatment imaging dataset: in-phase (IP'), out-of-phase (OP'), water-only (Water') and fat-only (Fat') image sets. Patients were further imaged using, amongst others, a T2w Large Field-of-View (T2wLFOV) sequence.

All referral imaging datasets were acquired with subjects lying supine on the standard MR bed using SENSE XL torso coils (Philips, Netherlands) wrapped around the pelvis. Treatment imaging datasets were acquired with subjects lying oblique supine decubitus on a gel-pad, which was placed on top of an acoustically transparent membrane on the top surface of the Sonallevé<sup>®</sup> V2 MRgHIFU couch (Profound Medical, Mississauga, Canada), using two Sonallevé<sup>®</sup> coils – one integrated into the acoustic window, and an external pelvic coil. The subject's body weight caused the gel-pads to compress and the membrane to bow. Subjects were positioned by a radiographer experienced in



**Figure 2.** Schematic of developmental methodology used in this study. The accuracy of the methodology to calculate target coverage from referral imaging was assessed using this workflow. The target volume coverage by MRgHIFU was calculated from a subject's treatment image dataset, acquired with the subject placed in a plausible or actual treatment position (bottom row) for volunteers or patients, respectively. Comparison with the target volume coverage predicted from a supine referral image dataset allowed assessment of the methodology. Step 1: the referral imaging dataset is rotated into the same orientation as the treatment imaging dataset using affine registration both to allow comparison with the treatment imaging dataset. Step 2: segmentation of acoustic obstructions (e.g. bones, shown), organs at risk (patients only) and the target tumors (patients only) was performed to identify tissues that impede target coverage. Step 3: Target volume coverage was calculated for the registered-referral imaging dataset and the treatment imaging dataset, and finally, the two quantities were compared to assess the predictive capacity of the methodology.

MRgHIFU. Cohort-specific imaging information for volunteers is given in Section 2.2.1, and for patients, in Section 2.2.2.

Treatment angles were measured using ITK-Snap 3.6.0 software [6] (University of Pennsylvania, USA), by manually drawing a line between the axial-plane positions of the left and right ischial spines, and finding the angle between this and a horizontal line.

### 2.2.1. Volunteers

Five female volunteers (age: 28–44 years, weight: 55–72 kg, body mass index: 20.2–26.4 kg/m<sup>2</sup>), were scanned (with ethics approval from The Royal Marsden and ICR Committee for Clinical Research (internal protocol CCR1406)). In addition to the supine referral imaging dataset described above, each volunteer was scanned in two ‘treatment’ positions deemed to be plausible from experience of treating patients with pelvic bone pain with MRgHIFU [3,4]. These positions were nominally ‘steep’ and ‘shallow’, but were dependent on a subject’s size and shape, which affected how they fitted into the bore of the MR scanner. This generated two treatment imaging datasets per volunteer. The volunteers, wearing thin trousers, were placed with their left buttock roughly centered over the acoustic window and with their right side elevated using angled foam pads. They were scanned from the L5-sacrum disk to the inferior-most point of the ischial tuberosity in the axial direction. Fields-of-view were chosen to include the full body outline in the axial slices. 15 mm-thick gel-pads were used to provide acoustic coupling between the skin and the Sonallevé® acoustic window for all volunteers. The voxel size for referral imaging and treatment imaging datasets was approximately 0.78 × 0.78 × 1.50 mm<sup>3</sup>. Volunteer details are recorded in Table 1.

### 2.2.2. Patients

Five patient datasets were acquired after volunteer image acquisition began, as part of a recurrent gynaecological

tumor clinical trial (NCT02714621, REC: 15/WM/0470) [3]. For treatment imaging datasets, patients were oriented into a clinically judged treatment position, with the tumor as close to the magnetic isocentre as possible. Because pretreatment diagnostic referral imaging was not available, the earliest (Day-7) follow-up supine images were used as ‘referral’ imaging datasets. These were chosen to minimize anatomical changes between the two imaging datasets. 15 mm-thick gel-pads were used for patients P2 to P5. For patient P1, a 40 mm-thick gel-pad was manually cut out to provide a degassed-water-filled recess, into which the patient was lowered. Patient details are recorded in Table 2. Weight data had been collected from patients as part of the trial data, but height data (and therefore BMI data) had not.

Patient referral and treatment imaging datasets were acquired after gadolinium contrast injection for improved contrast, and were acquired with a Field-of-View (FoV) of 288 × 288 × 133 voxels and voxel size 0.87 × 0.87 × 1.50 mm<sup>3</sup>. As part of a separate study, patient’s tumors were segmented from patient T2wLFOV datasets (TE = 90 ms, TR = 3620.4 ms, number of echoes = 16, flip angle = 90°, FoV 672 × 672 × 40 voxels, voxel size 0.45 × 0.45 × 4.5 mm<sup>3</sup>) obtained immediately pretreatment. These segments were used to define the target tumor volume for each patient.

### 2.3. Image registration

Registration of referral imaging datasets to treatment imaging datasets rotated the referral imaging dataset into the same treatment orientation as used in the treatment imaging dataset, which allowed the target coverage predicted from the registered-referral imaging dataset to be compared to that calculated from the treatment imaging dataset. Each subject’s referral imaging dataset was registered to their treatment imaging dataset(s) by aligning 10 or more manually placed bony landmark points, distributed throughout the pelvis, using Horos v2.4.0 (Horos Project) [7]. Registration was

Table 1. Details of volunteers participating in this study.

Volunteer	1	2	3	4	5	Mean ± Standard Deviation
Age (years)	28	44	29	27	36	33 ± 6
Body Mass Index (kg/m <sup>2</sup> )	20.2	26.4	23.5	23.8	20.9	23 ± 2
Height (cm)	165	165	170	160	168	166 ± 3
Weight (kg)	55	72	68	61	59	63 ± 6
Pelvic tilt from supine (°)						
Steep,	23,	19,	17,	24,	29,	22 ± 4,
Shallow	17	12	8	13	16	13 ± 3
Gel-pad Thickness (mm)						
Steep,	10.2,	N/A,	9.8,	9.7,	N/A,	9.8 ± 0.3
Shallow	9.8	N/A	9.8	9.3	10.0	
Membrane Bowing (mm)						
Steep,	10.4,	N/A,	8.6,	10.9,	N/A,	10.0 ± 1.3
Shallow	11.7	N/A	9.4	10.9	7.8	

Table 2: Details of patients participating in this study.

Patient	P1	P2	P3	P4	P5	Mean ± Standard Deviation
Age (years)	64	53	72	74	59	64 ± 8
Weight (kg)	42	76	57	61	61	59 ± 11
Treatment Angle (°)	6	33	16	9	24	18 ± 10
Gel Pad Thickness (mm, mean ± SD)	5.3 ± 0.5	10.9 ± 0.6	8.6 ± 0.4	12.3 ± 0.4	8.0 ± 0.4	10 ± 2
(Nominal)	(40)	(15)	(15)	(15)	(15)	(15)
Membrane Bowing (mm, mean ± SD)	4.1 ± 0.2	10.0 ± 0.5	9.0 ± 0.5	5.0 ± 0.2	10.0 ± 0.1	7.6 ± 2.8



performed following the standard operating procedure described in the [Appendix](#). The software calculated the required affine transformation and applied it to the referral imaging dataset [8] to generate the registered-referral imaging dataset.

To quantify the quality of this registration, the intra-observer (3 volunteer datasets) error and inter-observer (3 observers, 1 volunteer dataset) error associated with the referral-to-treatment registration was calculated. The errors were quantified as the mean Euclidean distance between corresponding points.

## 2.4. Image segmentation

The presence of acoustic obstructions and organs at risk in the beam path prevents safe sonication of the target, and hence they were segmented in order to identify acoustic access to the target. The tumor defined the target volume for patients, and hence was segmented. The body outline was segmented to assist with the other segmentation processes, and to assist in positioning the MRgHIFU system relative to the registered-referral imaging dataset. Organs at risk, bone (an acoustic obstruction) and the tumor were manually segmented from the MR datasets (as shown in [Figure 2](#), Step 2). The body outline and extracorporeal air (an acoustic obstruction) were segmented automatically, as described below.

### 2.4.1. Body outline

The body outline delineates the skin surface, and, particularly for treatment imaging datasets, needs to be separated from the gel-pad the subject lies on. An automatic process involving Otsu thresholding [9] was developed to separate the body from surrounding extracorporeal air and the gel-pad. Connected-components labeling [10] was used to collate segments of the body, and morphological operations [11] and flood-filling [12] were employed to link disparate segments and fill holes within segments.

### 2.4.2. Acoustic obstructions

Internal acoustic obstructions, primarily bone, were segmented by manual contouring of axial slices using OsiriX Lite v10.0.4 [13] (Pixmeo, Geneva, Switzerland) and Horos. For volunteers, pelvic bones were manually segmented from referral imaging datasets. The registered-referral imaging dataset pelvic bone segments were applied to the corresponding treatment imaging dataset in order to reduce the burden of manual contouring. Femora were manually segmented separately from referral and treatment imaging datasets, because of the likelihood of different articulation between datasets (unlike the more rigid pelvis). For patients, the treatment region was considerably smaller and therefore pelvic bones as well as femora close to the target (tumor) could be manually segmented in a realistic time. However, contouring was restricted to  $\pm 10$  axial slices from the edges of the tumor to reduce the time burden of manual segmentation. The pelvic bones at the greater sciatic notch were always segmented, because the notch defines the superior

edge of the sciatic foramen through which the acoustic beam is expected to sonicate the tumor.

Air gaps between the patient and the gel-pad act as acoustic obstructions. Extracorporeal air in volunteer treatment imaging datasets was not segmented, because the trousers worn by volunteers during image acquisition prevented skin-to-gel-pad acoustic coupling. Instead, volunteer acoustic coupling limits in the left-right direction were manually identified, as shown in [Figure 4](#). For volunteers, it was assumed that the intergluteal cleft would be filled with acoustic-coupling gel as part of clinical preparations, and hence, they were not treated as acoustic obstructions. Extracorporeal air in the patient treatment imaging datasets was segmented to define the limits of acoustic coupling, using an automatic segmentation algorithm inspired by Kullberg *et al.* [14]. In some cases, the intergluteal cleft was seen to contain air, and was therefore manually contoured and included as part of the extracorporeal air segment.

### 2.4.3. Target volume

As part of a separate study, patient tumors had been contoured by an experienced radiographer (SG) using in-house software (Adept v0.2, The Institute of Cancer Research, UK) [3] on referral and treatment imaging T2wLFOV images, where the slice thickness was 10 times that of the in-plane voxel dimensions. Segmented tumors were registered to align with the Dixon imaging datasets using the same procedures described above in order to obtain tumor outlines in the Dixon images. Since healthy volunteers had no tumors, all soft tissue within the pelvic region was designated as the target.

### 2.4.4. Organs at risk

Organs at risk, namely the uterus, rectum, bladder, and intestines were manually segmented for patients. Some patients had previously undergone pelvic exenteration surgery resulting in the removal of most pelvic organs.

### 2.4.5. Evaluation of automated segmentation quality

Automatic segmentation quality for the body outline and for extracorporeal air was assessed by comparing randomly selected image slices with corresponding manually segmented slices (body: five slices per dataset, from three 'steep' treatment imaging datasets and two 'steep' registered-referral imaging datasets originating from three volunteers; air: five slices per dataset from three patient treatment datasets). In order to determine the ability of the segmentation to determine acoustic coupling between patient and transducer, only the extracorporeal air segments around the body/gel-pad interface were assessed.

The assumption that the manually-segmented pelvic bone in volunteer registered-referral datasets could be used to automatically segment the pelvic bones in the treatment imaging dataset was similarly tested against manual contouring performed on the treatment imaging dataset (five slices per treatment dataset, four treatment datasets originating

from three volunteers). The segmentation quality of the volunteer bony pelvis and femora was taken to be indicative of the segmentation quality for all manually segmented tissues. Quality metrics were Dice Similarity Coefficient (DSC) and mean contour-to-contour distance [15,16].

## 2.5. Prediction of target volume coverage

### 2.5.1. Overview

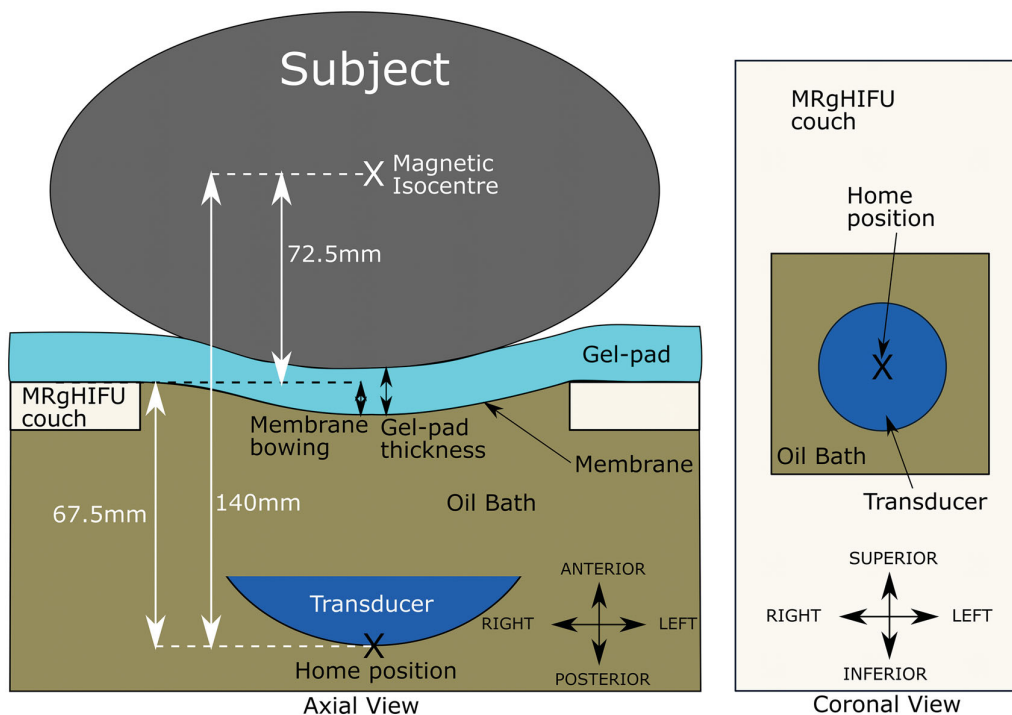
To calculate the target volume that can be covered, an MRgHIFU transducer was simulated. Positioning of the MRgHIFU transducer was known for the treatment imaging datasets, but had to be derived for the registered-referral imaging datasets. In the process of positioning the virtual transducer/referral imaging dataset, patient-induced compression of the gel-pad and bowing of the oil-bath membrane had to be taken into account. To reduce the computational time required, additional practical and clinically-relevant restrictions were placed on transducer translation, as described in greater detail below. The target volume covered by treatment cells was calculated for corresponding pairs of registered-referral and treatment datasets, and then, for each subject, the two volumes were compared. The details of these procedures are presented below.

### 2.5.2. MRgHIFU system characteristics

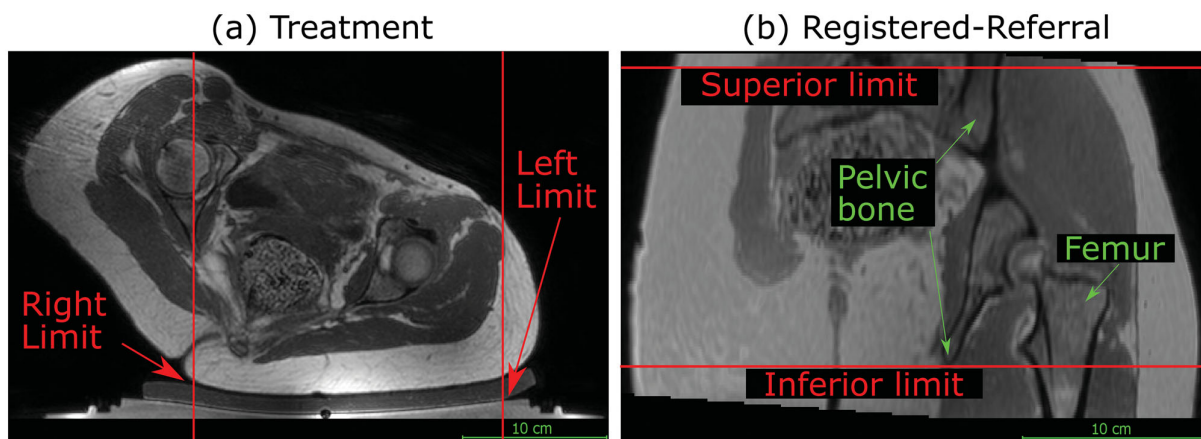
The simulated transducer was modeled on The Royal Marsden Hospital's MRgHIFU system, the Sonalleve<sup>®</sup> V2. The system replaces the imaging couch in the bore of the MR scanner for treatment. The 256-element phased-array

transducer (130 mm diameter, focal length 140 mm, source frequency 1.22 MHz) is mounted on a robotic positioner with 3 linear and 2 rotational motion capabilities in an oil bath, and faces the patient through a thin (50  $\mu\text{m}$  thick) acoustically transparent membrane. The transducer's home position (black cross in Figure 3) always lies 140 mm below the magnetic isocentre, and the undeformed membrane-to-isocentre distance is 72.5 mm. Acoustic coupling is achieved using a degassed-water wetted gel-pad (either 15 or 40 mm thick). When a subject is in place, the gel-pad is compressed and the acoustic membrane bowed under their weight. From its home position, the transducer can translate in 50  $\mu\text{m}$  steps up to: 72.5 mm left or right and inferior or superior, and 34 mm toward the patient (anterior) and 33 mm away (posterior). The transducer can be angled up to 10° away from the perpendicular in the left-right and inferior-superior directions.

The transducer was simulated in MATLAB R2018b. It consisted of 256 points that represented the center of each transducer element. Ultrasound rays traced from each element on the transducer surface to the transducer focal point were used to represent the acoustic beam. The transducer was restricted to being able to tilt  $\pm 10^\circ$  in 2.5° steps in the left-right direction only, in order to avoid incomplete registered-referral dataset image slices resulting from registration, but otherwise possessed the translational extents of the clinical device as described above. The transducer is assumed to produce a perfect 8 mm treatment cell, i.e. an 8 mm x 21.84 mm ellipsoid [17,18] centered at the focal point with its long-axis aligned to the beam axis.



**Figure 3.** Schematic of the Sonalleve<sup>®</sup> V2 MRgHIFU system: LEFT - a subject lying on the MR bed will compress the acoustic-coupling gel-pad and bow the acoustic membrane, which seals the oil bath. Ideally, target tissue would be centered directly above the transducer's home position and the center of the membrane/gel pad and below the magnetic isocentre. RIGHT - a coronal view of the MRgHIFU couch showing the transducer's home position below the center of the membrane.



**Figure 4.** Transducer translation restrictions for volunteer data. Practical restrictions applied to the transducer's translation capabilities (solid red lines) for volunteer datasets only. (a) For a treatment imaging dataset, the left-right translation was limited by the extent of acoustic coupling between the volunteer's skin and the gel pad. The corresponding registered-referral imaging dataset shared these left-right restrictions. (b) For a registered-referral imaging dataset, the transducer's inferiorsuperior translation was restricted by the extent of pelvic bone and the requirement for a full body outline within the image. The corresponding treatment imaging dataset shared these inferior-superior restrictions.

### 2.5.3. Practical and clinically-relevant restrictions on transducer translation

In order to improve computational efficiency of target coverage prediction, transducer translation in the left-right and inferior-superior axes was restricted to the left-right and inferior-superior extents of the targets. For patients, practical restrictions on left-right and inferior-superior translation were calculated from the left-right and inferior-superior extents of the tumor. For volunteers, the target is all soft tissue within the pelvic region. Hence, practical and clinically-relevant limits were manually identified (see Figure 4) and implemented. The left-right limits represent the extents of acoustic coupling. The inferior-superior limits represent the inferior-superior extents of the registered-referral imaging dataset containing complete body outlines and pelvic bone.

### 2.5.4. Estimated patient deformation resulting from reorientation into the treatment position

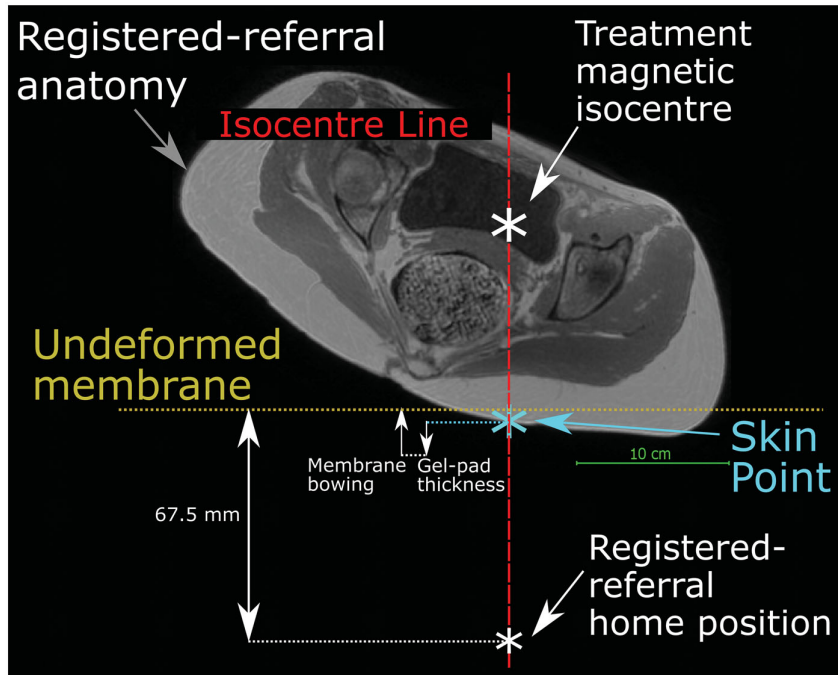
In this study, the treatment position was known from the treatment imaging dataset. In treatment imaging datasets, the isocentre, and hence the transducer's home position (Section 2.5.2), was known. In the registered-referral imaging dataset, because the treatment position is the same, the transducer's home position left-right and inferior-superior coordinates were taken from the treatment imaging dataset. However, to mimic the prospective workflow, the anterior-posterior coordinate had to be estimated from data within the registered-referral imaging dataset. The method of doing so is shown in Figure 5. Briefly, it was assumed that: i) the gel-pad would be most compressed and the membrane most bowed at the isocentre line, and ii) after soft tissue deformation resulting from the reorientation into the treatment position, the isocentre-to-skin point distance would remain the same. The membrane bowing distance and gel-pad thickness for patients was assumed to be that calculated for volunteers. These quantities were obtained by determining the average gel-pad thickness and

membrane bowing distance close to the isocentre line, using ITK-Snap, in the 7/10 volunteer treatment imaging datasets in which measurement was possible. From this, the position of the undeformed membrane, and hence the transducer anterior-posterior home position, was estimated (see Figure 3). Patient P1 had been treated on a customized gel-pad, the thickness of which was independently measured and used for positioning. For comparison, the actual patient gel-pad thicknesses and membrane bowing distances were measured and compared to the volunteer-derived averages.


### 2.5.5. Calculation of target coverage

For volunteers, a regular grid of target points, one per image voxel, was created in the soft tissue (see Figure 6); for patients, this grid was created solely within the tumor [19]. The transducer acoustic beam had been discretized into 256 rays, linking the center of a transducer element to the focus. Each ray was discretized into regularly spaced (0.2 mm) points along its length, and each was tested for intersection with acoustic obstructions or organs at risk. If no point intersected these, an 8-mm treatment cell was drawn around the focal point, and all grid points within this were marked as covered (Figure 6). This was repeated as the transducer was exhaustively translated and tilted. The number of grid points covered, multiplied by the image voxel volume, was used to quantify the target volume covered. For volunteers, the transducer was translated in 4 mm steps, whereas for patients, 2 mm steps were used in order to ensure coverage of the smaller tumor volume.

For volunteers, the accuracy of the methodology was quantified by calculating how much of the soft tissue volume coverage calculated from the treatment imaging dataset was predicted to be covered from the registered-referral imaging dataset, as described in Equation (1). In effect, the treatment imaging dataset covered soft tissue volume becomes the target volume for the registered-referral imaging dataset,



**Figure 5.** Method used to predict the transducer’s anterior-posterior home position in a registered-referral imaging dataset. The treatment dataset magnetic isocentre is known because the registered-referral imaging dataset had been registered to the treatment imaging dataset. A line was drawn downwards from the treatment dataset isocentre and intersected the skin at the skin point. From this skin point, the home position was calculated using the average compressed gel-pad thickness, the average membrane bowing distance, and the calibrated distance between undeformed membrane and home position of 67.5 mm (see Figure 3).



**Figure 6.** Method for quantifying target volume covered within a dataset (volunteer treatment imaging dataset in this example). A regular 3D grid of potentially accessible points was created (blue crosses) within the target: soft tissue (volunteers) or tumor (patients). For each transducer position and tilt identified in Section 2.5.3, the acoustic beam was checked for intersection with any acoustic obstructions (green contours) or organs at risk. If no obstruction exists, an 8 mm treatment cell was created around the focus (yellow ellipse). Grid points within a treatment cell were marked as ‘covered’ (red crosses).

allowing calculation of the percentage target volume covered ( $TVC_{vol}$ ).

$$TVC_{vol} = 100\% \times \frac{CV_{RegisteredReferral} \cap CV_{Treatment}}{CV_{Treatment}} \quad (1)$$

where CV is the covered target volume.

For patients, the accuracy of the methodology was quantified using the difference between the percentage tumor volumes covered ( $TVC_{pat}$ ), calculated from treatment imaging dataset and that calculated from registered-referral imaging dataset.  $TVC_{pat}$  is given by:

$$TVC_{pat} = 100\% \times \frac{CV}{TV} \quad (2)$$

where CV is the covered tumor volume and TV is the total tumor volume.

### 3. Results

#### 3.1. Subjects

Details for the volunteers involved in the study are recorded in Table 1, and those for patients in Table 2, as are the (pseudo)-treatment angle(s), compressed gel-pad thickness and membrane bowing distance for each subject. For volunteers, 15 mm gel-pads were compressed to an average of  $9.8 \pm 0.3$  (mean  $\pm$  standard deviation, with range: 9.3 to 10.2) mm, and the average membrane bowing distance close to the isocentre line was  $10.0 \pm 1.3$  (range: 7.8 to 11.7) mm. The weight ranges of volunteers and patients (patients:  $59 \pm 11$  kg vs volunteers:  $63 \pm 6$  kg) were similar. The range of patient treatment angles ( $6$ - $33^\circ$ ) slightly exceeded the range of volunteer angles ( $8$ - $29^\circ$ ).

#### 3.2. Image registration quality

Between three observers, the mean distance between corresponding points for the referral imaging dataset for one volunteer, registered to one of their treatment imaging datasets, was on average  $1.2 \pm 0.2$  mm. For one observer, the mean distance between corresponding points for the referral imaging datasets for three volunteers, each registered to one of their corresponding treatment imaging datasets, was on average  $1.3 \pm 0.2$  mm. These distances are less than the axial slice thickness of the Dixon image datasets and less than double the in-plane image resolution.

#### 3.3. Segmentation quality

##### 3.3.1. Automatic segmentation quality

Automatically segmented body outlines agreed with validation slices with a mean DSC of  $0.991 \pm 0.003$  and an average mean contour-to-contour distance of  $0.9 \pm 0.4$  mm. Automatic extracorporeal air segmentation of patient data agreed with validation slices with a mean DSC of  $0.89 \pm 0.06$  and an average mean contour-to-contour distance of  $0.25 \pm 0.16$  mm.

##### 3.3.2. Manual segmentation quality

Volunteer treatment image pelvic bone segmentation agreed with the validation slices, with mean DSC of  $0.93 \pm 0.01$  and an average mean contour-to-contour distance of  $0.76 \pm 0.10$  mm. Volunteer femur segmentation agreed with the validation slices with mean DSC of  $0.96 \pm 0.01$  and an average mean contour-to-contour distance of  $0.53 \pm 0.11$  mm.

#### 3.4. Prediction of target volume coverage

The  $TVC_{vol}$  for each volunteer in each of their two treatment positions is shown in Figure 7(a). For volunteers, the registered-referral imaging dataset predicted target volume coverage of  $91 \pm 6\%$  (range: 78 to 98%) of that calculated from the corresponding treatment imaging dataset. The  $TVC_{pat}$  for each patient's treatment imaging and referral imaging are shown in Figure 7(b). Patient P4 appears to be an outlier. Excluding their data, for patients, registered-referral  $TVC_{pat}$  predicted the treatment  $TVC_{pat}$  to within an average of  $12 \pm 7\%$  (range: 4 to 21%). Representative images of the target (tumor) volumes covered for volunteers and patients are shown in Figure 7(c) and (d), respectively.

### 4. Discussion

The aim of this study was to develop a novel method to calculate tumor coverage and assess the feasibility of predicting tumor coverage from (supine) referral imaging, as part of a wider study into automating the evaluation of patient suitability for MRgHIFU therapy.

#### 4.1. Subjects

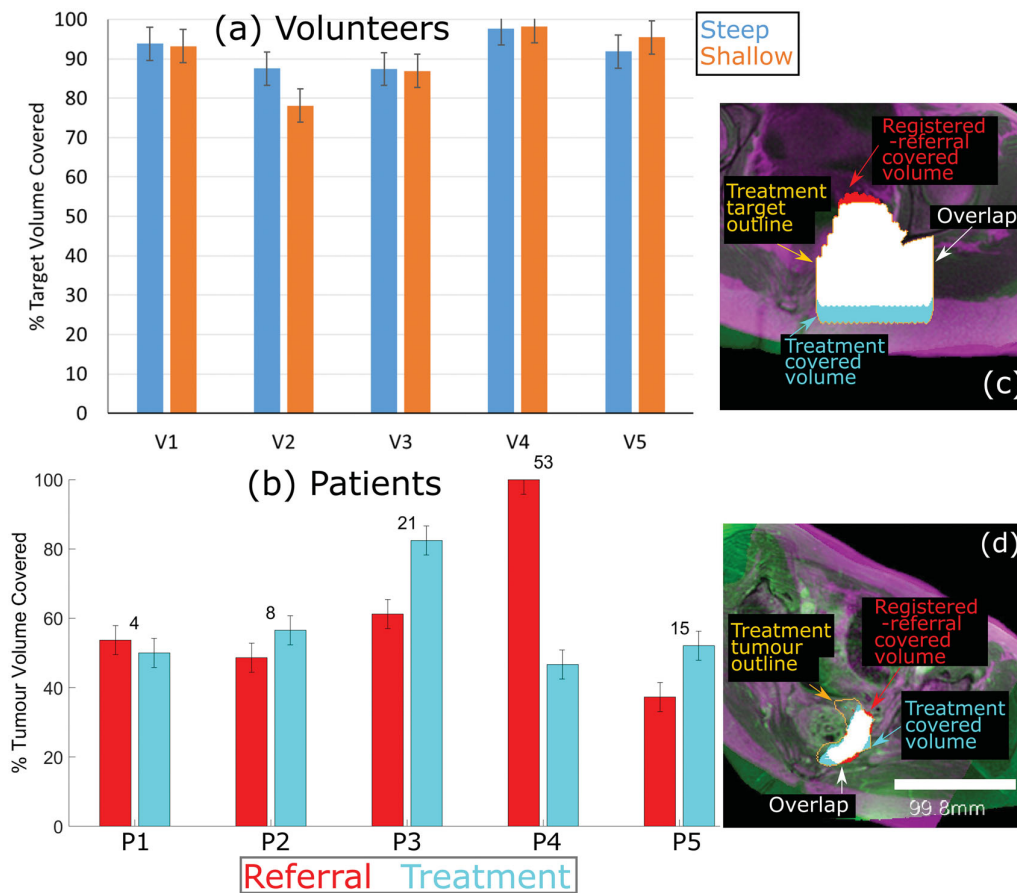
Although patient mean age was nearly double that of the volunteers, their weights were similar. Compressed gel-pad thickness and membrane bowing for volunteers varied minimally (mean  $\pm$  standard deviation being  $9.8 \pm 0.3$  mm and  $10.0 \pm 1.3$  mm respectively), suggesting that use of mean values for the prediction of patient tumor coverage should be acceptable. Minimum and maximum patient tilt angles exceeded those of volunteers by at most  $4^\circ$  despite acquiring the volunteer imaging before the patient data was available.

#### 4.2. Image registration

Mean post-registration misalignment between referral and treatment images was found to be less than the axial slice thickness of the Dixon MR imaging, in line with results from literature [20].

#### 4.3. Image segmentation

Automatic and manual segmentation of acoustic obstructions, organs at risk and the body outline resulted in mean DSCs  $\geq 0.89$  and mean contour-to-contour distances that were less than the axial slice thickness (1.5 mm). A mean



**Figure 7.** Percentage of target volume covered. (a) For volunteers, the agreement between the referral and treatment covered volumes is shown, where the treatment covered volume is the ground-truth. (b) For patients, the percentage of the registered-referral tumor (red) and the treatment tumor (blue) that was covered is shown. The numbers on top of each set of bars represent the difference in % Tumor Volume Covered predicted from the registered-referral dataset, and that calculated from treatment dataset. Representative examples of target coverage for volunteers (c) and tumor coverage for patients (d) are shown, with a scale bar in (d). The anatomy is shaded purple in the registered-referral dataset, and green in the treatment dataset.

contour-to-contour distance of 2.81 mm has been deemed acceptable for breast-air boundary segmentation from MR imaging (voxel size: isotropic 2.5 mm) [21]. The DSC for extracorporeal air segmentation (patient treatment imaging datasets only) was less than that for body outline segmentation (volunteers and patients, treatment and registered-referral datasets) while the mean contour-to-contour distance was better than that for body outline segmentation. This was probably due to the smaller size of the air segments around the patient/gel-pad interface, causing a misidentified voxel to have a greater effect than for the larger body outline. From the DSC (0.96) and mean contour-to-contour distance (0.53 mm) values, the assumption that pelvic bone segments identified on treatment images were identical to post-registration, manually outlined referral image segments appears to be valid (Section 3.3.2).

Since tumors were manually segmented by an expert, any segmentation imprecision or inaccuracy was ignored. Tumors were segmented on datasets with slice thickness (4.5 mm) 10 times the in-plane resolution (0.45 mm), and thus rotation during registration could introduce relatively large discrepancies between the interpolated and actual tumor outlines,

thus increasing uncertainty in the  $TVC_{pat}$  predicted from referral imaging datasets.

#### 4.4. Target coverage

##### 4.4.1. Volunteer study

The volunteer's results show an average target coverage agreement between treatment and registered-referral imaging datasets of 91% (range: 78–98%), corresponding to a mean difference of 9%. This suggests that the techniques used for positioning the transducer in the registered-referral imaging datasets were sufficient to proceed to testing with patient data. The worst agreement (78%, for Volunteer 2 tilted at a 12° treatment angle) was attributed to inaccurate placement of the transducer's home position, caused by the skin point directly below the isocentre (see Figure 5) not remaining at constant position between the registered-referral and treatment datasets, as had been assumed. Consequently, the HIFU focus was predicted to reach 12 mm deeper into the volunteer than it could. The next worst agreements, (88% for Volunteer 2 tilted at 19° and Volunteer 3 tilted at 17° and 8°)

were due to the same cause, resulting in overestimation of the focal depth by 6 mm.

For a single volunteer, the difference between target coverage predicted from registered-referral datasets and that from treatment datasets results from differing femur segments and differing transducer home positions. Since angulation was restricted to tilting left-right only, and the transducer was restricted to prevent translation beyond the inferior-superior extents of the pelvis, differences in femur segmentations were judged to have only a small effect. Refinement of the transducer positioning technique, by sampling within a  $15 \times 15$  mm region around the isocentre line (see Figure 5) instead of using a single skin position to predict the anterior-posterior position, provided no statistically significant improvement (data not presented).

#### 4.4.2. Patient study

The goal of this study was to develop and test a method for quantitatively assessing tumor coverage from referral imaging, as opposed to the current clinical practice of qualitative assessment, and to assess the feasibility of the new methodology. From the results, quantitative prediction of tumor coverage from referral imaging appears feasible. Despite the simplicity of the technique used to account for the expected body deformation resulting from reorientation from supine into a treatment position, the  $TVC_{pat}$  predicted from the registered referral and the treatment imaging datasets had a mean difference of 12% (range: 4–21%), excluding an outlier for whom the difference was 53% (see below). In the literature, a median difference of 21% in automatic segmentation had been judged as acceptable [22]. In the context of the current clinical practice, where  $\geq 40\%$  of referred patients fail screening, these results are encouraging [3,4]. The small cohort involved in this study (5 volunteers, 5 patients) represents lower than expected patient recruitment for the clinical trial. However, other published studies have also involved small patient cohorts, e.g. a transcranial simulation study involved 5 patients [23], a simulation study for kidney ablation examined 4 patients [24], and in various therapeutic feasibility studies, between 10 and 13 patients were considered [25–27]. In addition, an automatic geometric optimization technique for the packing of HIFU treatment cells demonstrated its capabilities using test objects and the publicly available dataset of a single volunteer [28]. Results from these small-cohort feasibility studies also demonstrate high variance in results. For example, in the transcranial simulation study, simulation results differed from measured data by up to  $40 \pm 13\%$  [23]. The results here indicate a step toward the long-term objective of widespread quantitative analysis of patient suitability for MRgHIFU therapy, with the aim of improving clinical decision-making and minimizing the impact on patient and hospital time and resources.

The outlier referred to above was patient P4, whose poor results were due to the assumption of perfect acoustic coupling between patient and gel-pad when calculating  $TVC_{pat}$  for the registered-referral imaging dataset. In practice, treatment imaging showed that the tumor periphery was obstructed by air between the patient and gel-pad. This

highlights a possible advantage of the proposed workflow. Having established that a greater tumor coverage could have been achieved at the referral stage, clinicians may have been able to improve the clinical preparations, and increase tumor coverage.

In general, the marginally poorer results for patients compared to volunteers (excluding the outlier patient) may be partially due to volunteer target volumes being over 10 times larger ( $\sim 300,000 \pm 100,000$  mm<sup>3</sup>) than patient targets ( $\sim 20,000 \pm 10,000$  mm<sup>3</sup>). A missed voxel has a larger proportional effect for smaller target volumes.

A source of error for the patient cohort may arise from the differences in the actual gel-pad thickness and membrane bowing (Table 2) compared to the mean values determined from the volunteer cohort which were used in the predictive calculations. Membrane bowing differences from the average of 10.0 mm ranged from 0.9 mm to 4.7 mm for patients, and from 0.4 mm to 2.2 mm for volunteers. Gel-pad thickness differences from the average of 9.8 mm ranged from 1.1 to 2.5 mm for patients who were treated on 15 mm gel-pads, and from 0.0 to 0.7 mm for volunteers. To evaluate the effect of this, the TVC was recalculated with the actual gel-pad thickness and membrane bowing distance for all patients. The maximum difference in  $TVC_{pat}$  Registered-Referral that resulted from using the average membrane bowing and gel-pad thickness, rather than the actual measured values, was 0.3% (patient P1). As more data from clinical studies becomes available, modeling the relationship between membrane bowing distance, or compressed gel-pad thickness, and patient weight and orientation may generate more accurate predictions of the transducer home position from referral imaging.

Deformation and translation of organs at risk, due to reorientation from referral to treatment position, clinical preparation such as pretreatment dieting and bowel-preparation, and the time between referral and treatment (1 week), may explain why the patient results show worse agreement overall than the volunteer results. In clinical experience, organs at risk such as the rectum are known to vary substantially and unpredictably in shape, position and volume [29,30]. The overall accuracy of the proposed patient workflow is expected to be limited by the patient-specific soft tissue deformation and coupling to the gel-pad. At the very least, the methodology presented here allows quantitative assessment of tumor coverage prior to the screening stage, reducing the need for clinical experience, and the influence of subjective opinion, on the assessment of patient suitability for progression through the treatment pathway.

#### 4.5. Limitations of the study

One of the major limitations is the small volunteer and patient cohort, which restricts the statistical certainty of the results. This study is also limited to predicting pelvic tumor coverage. However, the proposed patient workflow may be adaptable for other tumor sites. Assessment of the tumor volume that can be successfully ablated will require acoustic propagation and thermal bioeffects modeling. This is the

subject of extensive ongoing work. Patient deformation resulting from orientation into the treatment position was only accounted for using the simple assumption that the iso-centre-to-skin point distance would remain constant. This produced acceptable results for tumor coverage. However, accurate acousto-thermal modeling requires an accurate description of the medium of propagation, which may require simulation of soft tissues deformation between the gel-pad and the target.

Only reorientations from supine to oblique supine decubitus positions were tested in this study. While the results of this study are only applicable to the specific diagnostic MR bed and MRgHIFU couch used, the core principles are expected to be applicable to other HIFU devices, and referral datasets obtained from X-ray tomographic imaging. Furthermore, since the patient mean age was almost twice that of the volunteers, patient soft tissue could have different elastic properties than that of volunteers and therefore exhibit different deformation behavior. This could have affected the developed methodology.

## 5. Conclusion

Novel methodology for predicting the MRgHIFU target coverage from supine (MR) referral imaging was developed using 10 volunteer datasets and was retrospectively applied to 5 patient datasets. The difference between the target coverage computed using referral and treatment image datasets was within 12% on average (range: 4–21%), after one patient, with inadequate acoustic coupling during treatment, was excluded from analysis. Despite the relatively small cohort size, the focus on pelvic tumors, and the limited range of patient positions and MRgHIFU equipment on which the methodology was devised and tested, these results suggest quantitative, automated screening and treatment planning should be feasible, eventually obviating the need for patient suitability to be assessed using qualitative clinical judgment based on operator experience.

## Geolocation information

The study was conducted at the Institute of Cancer Research and the Royal Marsden Hospital, Sutton, Surrey, United Kingdom.

## Acknowledgements

The authors would like to thank Matthew Blackledge, Simon Doran and Matthew Orton from the Institute of Cancer Research (ICR) for their technical support, and Ari Partanen and others from Profound Medical for their support. We are grateful to Philips for their loan of the Sonalleve system to The Royal Marsden Hospital (RMH), and we acknowledge the support of the RMH MRI team, volunteers and patients, the Focused Ultrasound Foundation, CRUK and EPSRC in association with MRC & Department of Health (C1060/A10334, C1060/A16464), the NHS, the NIHR Biomedical Research Centre, the Clinical Research Facility in Imaging, and the Cancer Research Network. This study resulted from research performed as part of a studentship supported by Philips. The views expressed are those of the authors, and not necessarily those of

the National Health Service (NHS), the Department of Health, the ICR, the RMH, Profound Medical, Philips or the NIHR.

## Disclosure statement

The lead author is the recipient of a studentship supported by Philips.

## Funding

The first author (NFDL) produced this work whilst on a studentship supported by Philips. Funding from the NIHR Research for Patient Benefit programme (PB-PG-0815-20001) enabled the acquisition of data for the study.

## ORCID

Ngo Fung Daniel Lam  <http://orcid.org/0000-0002-8575-7162>  
 Ian Rivens  <http://orcid.org/0000-0001-5316-8884>  
 Sharon L. Giles  <http://orcid.org/0000-0002-8528-3745>  
 Emma Harris  <http://orcid.org/0000-0001-8297-0382>  
 Nandita M. deSouza  <http://orcid.org/0000-0003-4232-476X>  
 Gail ter Haar  <http://orcid.org/0000-0001-8909-0775>

## Data availability statement

The data that support the findings of this study are not available due to limitations in the ethical review.

## References

- [1] Hurwitz MD, Ghanouni P, Kanaev SV, et al. Magnetic resonance-guided focused ultrasound for patients with painful bone metastases: phase III trial results. *JNCI*. 2014;106(5):dju082.
- [2] Siedek F, Yeo SY, Heijman E, et al. Magnetic resonance-guided high-intensity focused ultrasound (MR-HIFU): technical background and overview of current clinical applications (Part 1). *Rofo*. 2019;191(6):522–530.
- [3] Giles SL, Imseeh G, Rivens I, et al. MR guided high intensity focused ultrasound (MRgHIFU) for treating recurrent gynaecological tumours: a pilot feasibility study. *Br J Radiol*. 2019; 92(1098):20181037.
- [4] Giles SL, Brown MRD, Rivens I, et al. Comparison of imaging changes and pain responses in patients with intra- or extraosseous bone metastases treated palliatively with magnetic resonance-guided high-intensity-focused ultrasound. *J Vasc Interv Radiol*. 2019; 30(9):1351–1360.e1.
- [5] Eggers H, Brendel B, Duijndam A, et al. Dual-echo Dixon imaging with flexible choice of echo times. *Magn Reson Med*. 2011;65(1): 96–107.
- [6] Yushkevich PA, Piven J, Hazlett HC, et al. User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. *Neuroimage*. 2006;31(3): 1116–1128.
- [7] HOROS Dicom Viewer. Version 2.4.0. HOROS Project; 2015. Available from: <https://horosproject.org/>
- [8] Späth H. Fitting affine and orthogonal transformations between two sets of points. *Mathematical Communications*. 2004;9(1): 27–34.
- [9] Otsu N. A threshold selection method from gray-level histograms. *IEEE Trans Syst Man Cybern*. 1979;9(1):62–66.
- [10] Shapiro LG. Connected component labeling and adjacency graph construction. In: Kong TY, Rosenfeld A, editors. *Machine intelligence and pattern recognition*. Vol. 19. North-Holland; 1996. p. 1–30.



- [11] Gonzalez RC, Woods RE. Digital image processing. 4th ed. New York, USA: Pearson; 2018.
- [12] Soille P. Morphological Image Analysis: Principles and Applications. Berlin (Germany): Springer-Verlag; 1999.
- [13] Rosset A, Spadola L, Ratib O. OsiriX: an open-source software for navigating in multidimensional DICOM images. *J Digit Imaging*. 2004;17(3):205–216.
- [14] Kullberg J, Ahlström H, Johansson L, et al. Automated and reproducible segmentation of visceral and subcutaneous adipose tissue from abdominal MRI. *Int J Obes (Lond)*. 2007;31(12):1806–1817.
- [15] Xia Y, Fripp J, Chandra SS, et al. Automated bone segmentation from large field of view 3D MR images of the hip joint. *Phys Med Biol*. 2013;58(20):7375–7390.
- [16] Fripp J, Crozier S, Warfield SK, et al. Automatic segmentation of the bone and extraction of the bone-cartilage interface from magnetic resonance images of the knee. *Phys Med Biol*. 2007;52(6):1617–1631.
- [17] Köhler MO, Mougnot C, Quesson B, et al. Volumetric HIFU ablation under 3D guidance of rapid MRI thermometry. *Med Phys*. 2009;36(8):3521–3535.
- [18] Enholm JK, Köhler MO, Quesson B, et al. Improved volumetric MR-HIFU ablation by robust binary feedback control. *IEEE Trans Biomed Eng*. 2010;57(1):103–113.
- [19] Sannholm F. Automated Treatment Planning in Magnetic Resonance guided High Intensity Focused Ultrasound [Master's Thesis]. Espoo, Finland: Aalto University School of Electrical Engineering; 2011.
- [20] Ostergaard NK, Denis de Senneville B, Elstrom UV, et al. Acceleration and validation of optical flow based deformable registration for image-guided radiotherapy. *Acta Oncol*. 2008;47(7):1286–1293.
- [21] Wang L, Platel B, Ivanovskaya T, et al., editors. Fully automatic breast segmentation in 3D breast MRI. 2012 9th IEEE International Symposium on Biomedical Imaging (ISBI); 2012 2–5 May 2012; Barcelona, Spain.
- [22] Thomson D, Boylan C, Liptrot T, et al. Evaluation of an automatic segmentation algorithm for definition of head and neck organs at risk. *Radiat Oncol*. 2014;9(1):173.
- [23] Pulkkinen A, Werner B, Martin E, et al. Numerical simulations of clinical focused ultrasound functional neurosurgery. *Phys Med Biol*. 2014;59(7):1679–1700.
- [24] Abbas MA, Coussios C-C, Cleveland RO, editors. Patient Specific Simulation of HIFU Kidney Tumour Ablation. 2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC); 2018 18–21 July 2018; Honolulu, HI, USA.
- [25] Marinova M, Rauch M, Mücke M, et al. High-intensity focused ultrasound (HIFU) for pancreatic carcinoma: evaluation of feasibility, reduction of tumour volume and pain intensity. *Eur Radiol*. 2016;26(11):4047–4056.
- [26] Welch BT, Schmitz JJ, Kurup AN, et al. Feasibility and outcomes of percutaneous thermal ablation of hepatocellular carcinoma in a transplanted allograft. *Abdom Radiol (NY)*. 2018;43(6):1478–1481.
- [27] Miyamoto K, Kapa S, Mulpuru SK, et al. Safety and efficacy of cryoablation in patients with ventricular arrhythmias originating from the Para-Hisian region. *JACC Clin Electrophysiol*. 2018; 4(3):366–373.
- [28] Williamson T, Everitt S, Chauhan S. Automated geometric optimization for robotic HIFU treatment of liver tumors. *Comput Biol Med*. 2018;96:1–7.
- [29] Roeske JC, Forman JD, Mesina CF, et al. Evaluation of changes in the size and location of the prostate, seminal vesicles, bladder, and rectum during a course of external beam radiation therapy. *Int J Radiat Oncol Biol Phys*. 1995;33(5):1321–1329.
- [30] Scaife J, Harrison K, Romanchikova M, et al. Random variation in rectal position during radiotherapy for prostate cancer is two to three times greater than that predicted from prostate motion. *Br J Radiol*. 2014;87(1042):20140343.

### Appendix: Registration standard operating procedure

1. Open Horos™ on Mac OS X. Make sure the pyOsiriX plugin is installed.
2. Import the in-phase MRI datasets that are to be registered. Double click them to bring them up together.
3. Select a dataset. Then, at the top menu 2D Viewer → Sort By... → Slice Location Ascending.
4. Below the menu bar, in a section titled 'Mouse button function', select the point function. Use the point function to mark an anatomical feature on one dataset and the same anatomical feature on the other. The same point names, e.g. 'Point 1', must correspond to the same anatomical features in both datasets. Repeat this for the list of anatomical features mentioned below. If the same anatomical feature cannot be found in one or both of the datasets, ignore that anatomical feature and continue down the list. At least 10 features should be marked by the end.  
Anatomical Features:
  - a. Femur/pelvis landmark marks where the two bones meet in the head-most direction (Right and Left)
  - b. ischial spine (Right and Left)
  - c. Superior-most or inferior-most of ischial tuberosity (Right and Left)
  - d. Pubic arch/top of pubic arch connection
  - e. Anterior-facing spur in axial plane where pelvis first encloses femur head (Right and Left)
  - f. Sacral nerve bundle (S1 and S2) when just-enclosed by bone (Right and Left)
  - g. Spinal nerves splitting from spinal cord (Right and Left)
  - h. Sacrum/L5 disk border
  - i. Coccyx
5. Open the pyOsiriX console within Horos. A Python script can be used to extract point data from a dataset in Horos and save it in a format that can be processed in an external Python environment. Do this for both datasets.

# Bibliography

- [1] G. R. ter Haar and C. Coussios. “High intensity focused ultrasound: Physical principles and devices”. In: *Int. J. Hyperthermia* 23.2 (2007), pp. 89–104. ISSN: 0265-6736 1464-5157. DOI: 10.1080/02656730601186138.
- [2] Philips. *Sonalleve MR-HIFU: Instructions for Use: Uterine Application and Bone Application*. Release 3. Vantaa, Finland: Philips Oy, Apr. 2012.
- [3] Elizabeth A. Stewart et al. “Clinical outcomes of focused ultrasound surgery for the treatment of uterine fibroids”. In: *Fertility and Sterility* 85.1 (2006), pp. 22–29. ISSN: 0015-0282. DOI: 10.1016/j.fertnstert.2005.04.072.
- [4] Elizabeth A. Stewart et al. “Sustained Relief of Leiomyoma Symptoms by Using Focused Ultrasound Surgery”. In: *Obstetrics & Gynecology* 110.2 (2007), pp. 279–287. ISSN: 0029-7844. DOI: 10.1097/01.AOG.0000275283.39475.f6.
- [5] Jonathan Hindley et al. “MRI Guidance of Focused Ultrasound Therapy of Uterine Fibroids: Early Results”. In: *American Journal of Roentgenology* 183.6 (2004), pp. 1713–1719. ISSN: 0361-803X. DOI: 10.2214/ajr.183.6.01831713.
- [6] Mark D. Hurwitz et al. “Magnetic Resonance–Guided Focused Ultrasound for Patients With Painful Bone Metastases: Phase III Trial Results”. In: *JNCI: Journal of the National Cancer Institute* 106.5 (2014), dju082. ISSN: 0027-8874 1460-2105. DOI: 10.1093/jnci/dju082. URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4112926/>.

- [7] B. Liberman et al. "Pain Palliation in Patients with Bone Metastases Using MR-Guided Focused Ultrasound Surgery: A Multicenter Study". In: *Ann. Surg. Oncol.* 16.1 (2008), pp. 140–146. ISSN: 1534-4681. DOI: 10.1245/s10434-008-0011-2. URL: <http://dx.doi.org/10.1245/s10434-008-0011-2> <http://link.springer.com/article/10.1245%5C%2Fs10434-008-0011-2>.
- [8] Mohammed Alkhorayef et al. "High-Intensity Focused Ultrasound (HIFU) in Localized Prostate Cancer Treatment". In: *Polish Journal of Radiology* 80 (2015), pp. 131–141. ISSN: 1733-134X 1899-0967. DOI: 10.12659/PJR.892341. URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4360749/>.
- [9] Sebastien Crouzet et al. "High intensity focused ultrasound (HIFU) for prostate cancer: Current clinical status, outcomes and future perspectives". In: *Int. J. Hyperthermia* 26.8 (2010), pp. 796–803. ISSN: 0265-6736. DOI: 10.3109/02656736.2010.498803. URL: <http://dx.doi.org/10.3109/02656736.2010.498803>.
- [10] A. Blana et al. "High-intensity focused ultrasound for the treatment of localized prostate cancer: 5-year experience". In: *Urology* 63.2 (2004), pp. 297–300. ISSN: 1527-9995 (Electronic) 0090-4295 (Linking). DOI: 10.1016/j.urology.2003.09.020. URL: <http://www.ncbi.nlm.nih.gov/pubmed/14972475>.
- [11] T. Uchida et al. "Treatment of localized prostate cancer using high-intensity focused ultrasound". In: *BJU Int.* 97.1 (2006), pp. 56–61. ISSN: 1464-4096 (Print) 1464-4096 (Linking). DOI: 10.1111/j.1464-410X.2006.05864.x. URL: <http://www.ncbi.nlm.nih.gov/pubmed/16336329>.
- [12] W. Jeffrey Elias et al. "A Randomized Trial of Focused Ultrasound Thalamotomy for Essential Tremor". In: *New England Journal of Medicine* 375.8 (2016), pp. 730–739. DOI: 10.1056/NEJMoa1600159. URL: <http://www.nejm.org/doi/full/10.1056/NEJMoa1600159>.

- [13] F. Wu et al. “A randomised clinical trial of high-intensity focused ultrasound ablation for the treatment of patients with localised breast cancer”. In: *Br. J. Cancer*. 89.12 (2003), pp. 2227–33. ISSN: 0007-0920 (Print) 0007-0920 (Linking). DOI: 10.1038/sj.bjc.6601411. URL: <http://www.ncbi.nlm.nih.gov/pubmed/14676799>.
- [14] F. Wu et al. “Extracorporeal high intensity focused ultrasound ablation in the treatment of 1038 patients with solid carcinomas in China: an overview”. In: *Ultrason. Sonochem.* 11.3-4 (2004), pp. 149–54. ISSN: 1350-4177 (Print) 1350-4177 (Linking). DOI: 10.1016/j.ultsonch.2004.01.011. URL: <http://www.ncbi.nlm.nih.gov/pubmed/15081972>.
- [15] Brian H. H. Lang, Yu-Cho Woo, and Keith Wan-Hang Chiu. “Identifying predictive factors for efficacy in high intensity focused ultrasound (HIFU) ablation of benign thyroid nodules – a retrospective analysis”. In: *International Journal of Hyperthermia* 37.1 (2020), pp. 324–331. DOI: 10.1080/02656736.2020.1747646.
- [16] Michel Azizi et al. “Endovascular ultrasound renal denervation to treat hypertension (RADIANCE-HTN SOLO): a multicentre, international, single-blind, randomised, sham-controlled trial”. In: *The Lancet* 391.10137 (2018), pp. 2335–2345. ISSN: 0140-6736. DOI: [https://doi.org/10.1016/S0140-6736\(18\)31082-1](https://doi.org/10.1016/S0140-6736(18)31082-1). URL: <http://www.sciencedirect.com/science/article/pii/S0140673618310821>.
- [17] Pawel Piotr Dobrakowski et al. “MR-Guided Focused Ultrasound: A New Generation Treatment of Parkinson’s Disease, Essential Tremor and Neuropathic Pain”. In: *Interventional Neuroradiology* 20.3 (2014), pp. 275–282. DOI: 10.15274/INR-2014-10033.
- [18] Sharon L Giles et al. “MR guided high intensity focused ultrasound (MRgHIFU) for treating recurrent gynaecological tumours: a pilot feasibility study”. In: *The British Journal of Radiology* 92.1098 (2019),

- p. 20181037. DOI: 10.1259/bjr.20181037. URL: <https://www.birpublications.org/doi/abs/10.1259/bjr.20181037>.
- [19] Sharon L. Giles et al. “Comparison of Imaging Changes and Pain Responses in Patients with Intra- or Extraosseous Bone Metastases Treated Palliatively with Magnetic Resonance-Guided High-Intensity–Focused Ultrasound”. In: *Journal of Vascular and Interventional Radiology* 30.9 (2019), 1351–1360.e1. ISSN: 1051-0443. DOI: <https://doi.org/10.1016/j.jvir.2019.02.019>. URL: <http://www.sciencedirect.com/science/article/pii/S1051044319302593>.
- [20] D. Harding et al. “Evaluation of Quality of Life Outcomes Following Palliative Treatment of Bone Metastases with Magnetic Resonance-guided High Intensity Focused Ultrasound: An International Multicentre Study”. In: *Clinical Oncology* (2017). ISSN: 0936-6555. DOI: <https://doi.org/10.1016/j.clon.2017.12.023>. URL: <http://www.sciencedirect.com/science/article/pii/S0936655517305575>.
- [21] Yasutoshi Ishihara et al. “A precise and fast temperature mapping using water proton chemical shift”. In: *Magnetic Resonance in Medicine* 34.6 (1995), pp. 814–823. ISSN: 1522-2594. DOI: 10.1002/mrm.1910340606. URL: <http://dx.doi.org/10.1002/mrm.1910340606>.
- [22] Elizabeth Ramsay et al. “Temperature-dependent MR signals in cortical bone: Potential for monitoring temperature changes during high-intensity focused ultrasound treatment in bone”. In: *Magnetic resonance in medicine* 74.4 (2015), pp. 1095–1102. ISSN: 0740-3194 1522-2594. DOI: 10.1002/mrm.25492. URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4431954/>.
- [23] S. A. Sapareto and W. C. Dewey. “Thermal Dose Determination in Cancer Therapy”. In: *Int. J. Radiat. Oncol. Biol. Phys.* 10 (1984), pp. 787–800.

- [24] Pavel S. Yarmolenko et al. “Thresholds for thermal damage to normal tissues: An update”. In: *International journal of hyperthermia : the official journal of European Society for Hyperthermic Oncology, North American Hyperthermia Group* 27.4 (2011), pp. 320–343. ISSN: 0265-6736 1464-5157. DOI: 10.3109/02656736.2010.534527. URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3609720/>.
- [25] Barbara Gerbert et al. “Decision Support Software to Help Primary Care Physicians Triage Skin Cancer: A Pilot Study”. In: *Archives of Dermatology* 136.2 (2000), pp. 187–192. ISSN: 0003-987X. DOI: 10.1001/archderm.136.2.187. URL: <https://dx.doi.org/10.1001/archderm.136.2.187>.
- [26] Michael Schwenke et al. “An integrated model-based software for FUS in moving abdominal organs”. In: *International Journal Of Hyperthermia: The Official Journal Of European Society For Hyperthermic Oncology, North American Hyperthermia Group* 31.3 (2015), pp. 240–250. ISSN: 1464-5157. DOI: 10.3109/02656736.2014.1002817.
- [27] A. Scherrer, S. Jakobsson, and K.-H. Kufer. *On the advancement and software support of decision making in focused ultrasound therapy*. Report. Fraunhofer ITWM, 2015.
- [28] A. Scherrer et al. *A hybrid optimization method for focused ultrasound plan computation*. Report. Fraunhofer ITWM, 2016.
- [29] R. J. Fedewa et al. “Automated treatment planning for prostate cancer HIFU therapy”. In: *IEEE Ultrasonics Symposium, 2005*. Vol. 2, pp. 1135–1138. ISBN: 1051-0117. DOI: 10.1109/ULTSYM.2005.1603050.
- [30] Stephanie Tung et al. “Magnetic Resonance–guided High-intensity Focused Ultrasound (MRgHIFU) Virtual Treatment Planning for Abdominal Neuroblastoma Utilizing Retrospective Diagnostic 3D CT Images”. In: *Journal of Pediatric Hematology/Oncology* Epub ahead of print (). ISSN: 1077-4114. DOI: 10.1097/MPH.0000000000001563.

- [31] A. W. Dutton. "A three dimensional geometrical patient treatment planning program for scanned focussed ultrasound hyperthermia". Thesis. 1990.
- [32] David Kinnaird et al. "Treatment planning and patient positioning for MR-guided high intensity focused ultrasound treatment: a systematic approach". In: *Journal of Therapeutic Ultrasound* 3.1 (2015), P65. ISSN: 2050-5736. DOI: 10.1186/2050-5736-3-s1-p65. URL: <http://dx.doi.org/10.1186/2050-5736-3-s1-p65>.
- [33] Young-sun Kim et al. "Screening Magnetic Resonance Imaging–Based Prediction Model for Assessing Immediate Therapeutic Response to Magnetic Resonance Imaging–Guided High-Intensity Focused Ultrasound Ablation of Uterine Fibroids". In: *Investigative Radiology* 51 (1 Jan. 2016), pp. 15–24. DOI: 10.1097/RLI.000000000000199.
- [34] Zhongqiong Liu et al. "Establishment of a scoring system for predicting the difficulty level of high-intensity focussed ultrasound ablation of uterine fibroids". In: *International Journal of Hyperthermia* 34 (1 2018), pp. 77–86. DOI: 10.1080/02656736.2017.1325015.
- [35] C. R. Hill, J. C. Bamber, and G. R. Ter Haar, eds. *Physical Principles of Medical Ultrasonics*. 2nd ed. West Sussex, England: John Wiley & Sons, 2004.
- [36] H. Pauly and H. P. Schwan. "Mechanism of Absorption of Ultrasound in Liver Tissue". In: *The Journal of the Acoustical Society of America* 50.2B (1971), pp. 692–699. DOI: 10.1121/1.1912685.
- [37] J. C. Bamber and C. R. Hill. "Ultrasonic Attenuation and Propagation Speed in Mammalian Tissues as a Function of Temperature". In: *Ultrasound Med. Biol.* 5 (1979), pp. 149–157.
- [38] U.S. Department of the Interior, Bureau of Reclamation. "Ground Water Manual". In: *The Water Encyclopedia*. Ed. by Pedro Fierro Jr and Evan K. Nyer. Third Edition, Hydrologic Data and Internet Resources. 1977. DOI: 10.1201/9781420012583.

- [39] E. A. Filonenko and V. A. Khokhlova. “Effect of Acoustic Nonlinearity on Heating of Biological Tissue by High-Intensity Focused Ultrasound”. In: *Acoustical Physics* 47.4 (2001), pp. 541–549.
- [40] H. H. Pennes. “Analysis of Tissue and Arterial Blood Temperatures in the Resting Human Forearm”. In: *J. Appl. Physiol.* 1.2 (1948), pp. 93–122.
- [41] Matthew A. Lewis, Robert M. Staruch, and Rajiv Chopra. “Thermometry and Ablation Monitoring with Ultrasound”. In: *International journal of hyperthermia : the official journal of European Society for Hyperthermic Oncology, North American Hyperthermia Group* 31.2 (2015), pp. 163–181. ISSN: 0265-6736 1464-5157. DOI: 10.3109/02656736.2015.1009180. URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4467963/>.
- [42] Max O. Köhler et al. “Volumetric HIFU ablation under 3D guidance of rapid MRI thermometry”. In: *Medical Physics* 36.8 (2009), pp. 3521–3535. DOI: doi:10.1118/1.3152112. URL: <https://aapm.onlinelibrary.wiley.com/doi/abs/10.1118/1.3152112>.
- [43] P Wust et al. “Hyperthermia in combined treatment of cancer”. In: *The Lancet Oncology* 3 (8 2002), pp. 487–497. DOI: 10.1016/S1470-2045(02)00818-5.
- [44] Jian-Jun Li et al. “Complications of High Intensity Focused Ultrasound for Patients with Hepatocellular Carcinoma”. In: *Technol. Cancer Res. T.* 8.3 (2009), pp. 217–224. DOI: 10.1177/153303460900800306. URL: <http://tct.sagepub.com/content/8/3/217.abstract>.
- [45] Stefan Klein et al. “Automatic segmentation of the prostate in 3D MR images by atlas matching using localized mutual information”. In: *Medical Physics* 35.4 (2008), pp. 1407–1417. ISSN: 2473-4209. DOI: 10.1118/1.2842076. URL: <http://dx.doi.org/10.1118/1.2842076>.
- [46] Ying Xia et al. “Automated bone segmentation from large field of view 3D MR images of the hip joint”. In: *Physics in Medicine & Biology* 58.20



- (2013), p. 7375. ISSN: 0031-9155. URL: <http://stacks.iop.org/0031-9155/58/i=20/a=7375>.
- [47] Jurgen Fripp et al. "Automatic segmentation of the bone and extraction of the bone–cartilage interface from magnetic resonance images of the knee". In: *Physics in Medicine & Biology* 52.6 (2007), p. 1617. ISSN: 0031-9155. URL: <http://stacks.iop.org/0031-9155/52/i=6/a=005>.
- [48] Zhengyi Yang et al. "Automatic bone segmentation and bone-cartilage interface extraction for the shoulder joint from magnetic resonance images". In: *Physics in Medicine & Biology* 60.4 (2015), pp. 1441–1459. ISSN: 0031-9155. URL: <http://stacks.iop.org/0031-9155/60/i=4/a=1441>.
- [49] Barry J. Broderick et al. "Technique for the computation of lower leg muscle bulk from magnetic resonance images". In: *Medical Engineering & Physics* 32.8 (2010), pp. 926–933. ISSN: 1350-4533. DOI: <http://dx.doi.org/10.1016/j.medengphy.2010.06.008>. URL: <http://www.sciencedirect.com/science/article/pii/S1350453310001359>.
- [50] S. Orgiu et al. "Automatic muscle and fat segmentation in the thigh from T1-Weighted MRI". In: *J Magn Reson Imaging* 43.3 (2016), pp. 601–610. ISSN: 1522-2586 (Electronic) 1053-1807 (Linking). DOI: [10.1002/jmri.25031](https://doi.org/10.1002/jmri.25031). URL: <https://www.ncbi.nlm.nih.gov/pubmed/26268693>.
- [51] Zhen Ma et al. "Segmentation of female pelvic organs in axial magnetic resonance images using coupled geometric deformable models". In: *Computers in Biology and Medicine* 43.4 (2013), pp. 248–258. ISSN: 0010-4825. DOI: <https://doi.org/10.1016/j.combiomed.2012.12.012>. URL: <http://www.sciencedirect.com/science/article/pii/S0010482513000152>.
- [52] Nobuyuki Otsu. "A Threshold Selection Method from Gray-Level Histograms". In: *IEEE Transactions on Systems, Man, and Cybernetics* 9.1

- (1979), pp. 62–66. ISSN: 0018-9472. DOI: 10 . 1109 / TSMC . 1979 . 4310076.
- [53] J. Kullberg et al. “Automated and reproducible segmentation of visceral and subcutaneous adipose tissue from abdominal MRI”. In: *International Journal Of Obesity* 31 (2007), p. 1806. DOI: 10 . 1038 / sj . ijo . 0803671. URL: <http://dx.doi.org/10.1038/sj.ijo.0803671>.
- [54] Joel Kullberg et al. “Automated assessment of whole-body adipose tissue depots from continuously moving bed MRI: A feasibility study”. In: *Journal of Magnetic Resonance Imaging* 30.1 (2009), pp. 185–193. ISSN: 1522-2586. DOI: 10 . 1002 / jmri . 21820. URL: <http://https://doi.org/10.1002/jmri.21820>.
- [55] Neeraj Sharma and Lalit M. Aggarwal. “Automated medical image segmentation techniques”. In: *Journal of Medical Physics* 35.1 (2010). PMID: 11701515, pp. 3–14. DOI: 10 . 4103 / 0971 - 6203 . 58777.
- [56] Zhaobin Wang et al. “Review of pulse-coupled neural networks”. In: *Image and Vision Computing* 28.1 (2010), pp. 5–13. ISSN: 02628856. DOI: 10 . 1016 / j . imavis . 2009 . 06 . 007.
- [57] R. C. Gonzalez and R. E. Woods. *Digital Image Processing*. Fourth. New York, USA: Pearson, 2018. ISBN: 9780133356724.
- [58] P. Soille. *Morphological Image Analysis: Principles and Applications*. Berlin, Germany: Springer-Verlag, 1999.
- [59] David Thomson et al. “Evaluation of an automatic segmentation algorithm for definition of head and neck organs at risk”. In: *Radiation Oncology (London, England)* 9 (2014), pp. 173–173. ISSN: 1748-717X. DOI: 10 . 1186 / 1748 - 717X - 9 - 173. URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4123306/>.
- [60] John C. Roeske et al. “Evaluation of changes in the size and location of the prostate, seminal vesicles, bladder, and rectum during a course of external beam radiation therapy”. In: *International Journal of Radia-*

*tion Oncology\*Biology\*Physics* 33.5 (1995), pp. 1321–1329. ISSN: 0360-3016. DOI: [https://doi.org/10.1016/0360-3016\(95\)00225-1](https://doi.org/10.1016/0360-3016(95)00225-1). URL: <http://www.sciencedirect.com/science/article/pii/S0360301695002251>.

- [61] J Scaife et al. “Random variation in rectal position during radiotherapy for prostate cancer is two to three times greater than that predicted from prostate motion”. In: *The British Journal of Radiology* 87.1042 (2014), p. 20140343. DOI: 10.1259/bjr.20140343. URL: <https://www.birpublications.org/doi/abs/10.1259/bjr.20140343>.
- [62] Wicger K H Wong, Lucillus H T Leung, and Dora L W Kwong. “Evaluation and optimization of the parameters used in multiple-atlas-based segmentation of prostate cancers in radiation therapy”. In: *The British Journal of Radiology* 89.1057 (2016), p. 20140732. DOI: 10.1259/bjr.20140732.
- [63] Andrea Schenk, Guido Prause, and Heinz-Otto Peitgen. “Efficient Semi-automatic Segmentation of 3D Objects in Medical Images”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2000* (Pittsburgh, PA, USA). Ed. by Scott L. Delp, Anthony M. DiGoia, and Branislav Jaramaz. Berlin, Heidelberg, Germany: Springer Berlin Heidelberg, 2000, pp. 186–195. ISBN: 978-3-540-40899-4.
- [64] Magda A. Abbas, Constatin-C. Coussios, and Robin O. Cleveland. “Patient Specific Simulation of HIFU Kidney Tumour Ablation”. In: *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (Honolulu, HI, USA). 2018, pp. 5709–5712.
- [65] J. K. Enholm et al. “Improved Volumetric MR-HIFU Ablation by Robust Binary Feedback Control”. In: *IEEE Transactions on Biomedical Engineering* 57.1 (2010), pp. 103–113. ISSN: 0018-9294. DOI: 10.1109/TBME.2009.2034636.
- [66] Fredrik Sannholm. “Automated Treatment Planning in Magnetic Resonance guided High Intensity Focused Ultrasound”. Thesis. 2011.

- [67] Tom Williamson, Scott Everitt, and Sunita Chauhan. “Automated geometric optimization for robotic HIFU treatment of liver tumors”. In: *Computers in Biology and Medicine* 96 (2018), pp. 1–7. ISSN: 0010-4825. DOI: <https://doi.org/10.1016/j.compbiomed.2018.02.014>. URL: <http://www.sciencedirect.com/science/article/pii/S0010482518300441>.
- [68] Kenji Shimada and David C. Gossard. “Bubble Mesh: Automated Triangular Meshing of Non-Manifold Geometry by Sphere Packing”. In: *Proceedings of the Third ACM Symposium on Solid Modeling and Applications. SMA '95*. Salt Lake City, Utah, USA: Association for Computing Machinery, 1995, pp. 409–419. ISBN: 0897916727. DOI: 10.1145/218013.218095. URL: <https://doi.org/10.1145/218013.218095>.
- [69] F. A. Duck. *Physical Properties of Tissue: A Comprehensive Reference Book*. London: Academic Press, 1990. ISBN: 978 1 903613 50 4.
- [70] J. E. Soneson and M. R. Myers. “Thresholds for nonlinear effects in high-intensity focused ultrasound propagation and tissue heating”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 57.11 (2010), pp. 2450–2459. ISSN: 0885-3010. DOI: 10.1109/TUFFC.2010.1711.
- [71] Gianmarco Pinton et al. “Effects of nonlinear ultrasound propagation on high intensity brain therapy”. In: *Medical Physics* 38.3 (2011), pp. 1207–1216. DOI: 10.1118/1.3531553. eprint: <https://aapm.onlinelibrary.wiley.com/doi/pdf/10.1118/1.3531553>. URL: <https://aapm.onlinelibrary.wiley.com/doi/abs/10.1118/1.3531553>.
- [72] Ibrahim M. Hallaj and Robin O. Cleveland. “FDTD simulation of finite-amplitude pressure and temperature fields for biomedical ultrasound”. In: *The Journal of the Acoustical Society of America* 105.5 (1999), pp. L7–L12. DOI: 10.1121/1.426776. URL: <https://asa.scitation.org/doi/abs/10.1121/1.426776>.

- [73] Ibrahim M. Hallaj, Robin O. Cleveland, and Kullervo Hynynen. “Simulations of the thermo-acoustic lens effect during focused ultrasound surgery”. In: *The Journal of the Acoustical Society of America* 109.5 (2001), pp. 2245–2253. DOI: 10.1121/1.1360239. URL: <https://asa.scitation.org/doi/abs/10.1121/1.1360239>.
- [74] Aki Pulkkinen et al. “Numerical simulations of clinical focused ultrasound functional neurosurgery”. In: *Physics in medicine and biology* 59.7 (2014), pp. 1679–1700. ISSN: 0031-9155 1361-6560. DOI: 10.1088/0031-9155/59/7/1679. URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4083098/>.
- [75] Christopher W. Connor and Kullervo Hynynen. “Bio-acoustic thermal lensing and nonlinear propagation in focused ultrasound surgery using large focal spots: a parametric study”. In: *Physics in Medicine & Biology* 47.11 (2002), p. 1911. ISSN: 0031-9155. URL: <http://stacks.iop.org/0031-9155/47/i=11/a=306>.
- [76] G. F. Pinton et al. “A heterogeneous nonlinear attenuating full-wave model of ultrasound”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 56.3 (Mar. 2009), pp. 474–488. ISSN: 1525-8955. DOI: 10.1109/TUFFC.2009.1066.
- [77] U. Vyas and D. Christensen. “Ultrasound beam simulations in inhomogeneous tissue geometries using the hybrid angular spectrum method”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 59.6 (2012), pp. 1093–1100. ISSN: 0885-3010. DOI: 10.1109/TUFFC.2012.2300.
- [78] Mingjun Wang and Yufeng Zhou. “Simulation of non-linear acoustic field and thermal pattern of phased-array high-intensity focused ultrasound (HIFU)”. In: *International Journal of Hyperthermia* 32.5 (2016), pp. 569–582. ISSN: 0265-6736. DOI: 10.3109/02656736.2016.1160154. URL: <http://dx.doi.org/10.3109/02656736.2016.1160154>.

- [79] Roger J. Zemp, Jahangir Tavakkoli, and Richard S. C. Cobbold. “Modeling of nonlinear ultrasound propagation in tissue from array transducers”. In: *The Journal of the Acoustical Society of America* 113.1 (2003), pp. 139–152. ISSN: 0001-4966. DOI: 10.1121/1.1528926.
- [80] B. E. Treeby and B. T. Cox. “k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields”. In: *J. Biomed. Opt.* 15.2 (2010), p. 021314. ISSN: 1560-2281 (Electronic) 1083-3668 (Linking). DOI: 10.1117/1.3360308. URL: <http://www.ncbi.nlm.nih.gov/pubmed/20459236>.
- [81] B. E. Treeby et al. “Modeling nonlinear ultrasound propagation in heterogeneous media with power law absorption using a k-space pseudospectral method”. In: *J Acoust Soc Am* 131.6 (2012), pp. 4324–36. ISSN: 1520-8524 (Electronic) 0001-4966 (Linking). DOI: 10.1121/1.4712021. URL: <http://www.ncbi.nlm.nih.gov/pubmed/22712907>.
- [82] B. T. Cox et al. “k-space propagation models for acoustically heterogeneous media: Application to biomedical photoacoustics”. In: *The Journal of the Acoustical Society of America* 121.6 (2007), pp. 3453–3464. DOI: 10.1121/1.2717409.
- [83] V. Suomi et al. “Full Modeling of High-Intensity Focused Ultrasound and Thermal Heating in the Kidney Using Realistic Patient Models”. In: *IEEE Trans Bio-Med Eng* 65.11 (2018), pp. 2660–2670. DOI: 10.1109/TBME.2017.2732684.
- [84] E. Martin, J. Jaros, and B. Treeby. “Experimental validation of k-Wave: Nonlinear wave propagation in layered, absorbing fluid media”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 67 (1 2019), pp. 81–91. DOI: 10.1109/TUFFC.2019.2941795.
- [85] P. M. Meaney et al. “A 3-D Finite-element model for computation of temperature profiles and regions of thermal damage during focused ultrasound surgery exposures.pdf>”. In: *Ultrasound Med. Biol.* 24.9 (1998), pp. 1489–1499.

- [86] Yang-Sub Lee. “Numerical solution of the KZK equation for pulsed finite amplitude sound beams in thermoviscous fluids”. PhD thesis. Austin, TX, USA: The University of Texas at Austin, Dec. 1993.
- [87] Yang-Sub Lee and Mark F. Hamilton. “Time-domain modeling of pulsed finite-amplitude sound beams”. In: *The Journal of the Acoustical Society of America* 97.2 (1995), pp. 906–917. DOI: 10.1121/1.412135.
- [88] E. A. Zabolotskaya and R. V. Khokhlov. “Quasi-Plane Waves in the Non-linear Acoustics of Confined Beams”. In: *Sov. Phys. Acoust.* 15.1 (1969), pp. 35–40.
- [89] V. P. Kuznetsov. “Equations of Nonlinear Acoustics”. In: *Sov. Phys. Acoust.* 16.4 (1971), pp. 467–470.
- [90] Robin O. Cleveland, Mark F. Hamilton, and David T. Blackstock. “Time-domain modeling of finite-amplitude sound in relaxing fluids”. In: *The Journal of the Acoustical Society of America* 99.6 (1996), pp. 3312–3318. DOI: 10.1121/1.414983.
- [91] P. J. Westervelt. “Parametric Acoustic Array”. In: *J. Acoust. Soc. Am.* 35.4 (1963), pp. 535–537.
- [92] M. F. Hamilton and C. L. Morfey. “3. Model Equations”. In: *Nonlinear Acoustics*. Ed. by M. F. Hamilton and D. T. Blackstock. California, USA: Academic Press, 1998.
- [93] Jean-Pierre Berenger. “A perfectly matched layer for the absorption of electromagnetic waves”. In: *J. Comput. Phys.* 114.2 (1994), pp. 185–200. ISSN: 0021-9991. DOI: 10.1006/jcph.1994.1159.
- [94] Tabei M., Mast T. D., and Waag R. C. “A k-space method for coupled first-order acoustic propagation equations”. In: *J. Acoust. Soc. Am.* 111 (2002), p. 53.
- [95] T. D. Mast et al. “A k-space method for large-scale models of wave propagation in tissue”. In: *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* 48 (2001), p. 341.

- [96] B. E. Treeby and B. T. Cox. "Modeling power law absorption and dispersion for acoustic propagation using the fractional Laplacian". In: *J. Acoust. Soc. Am.* 127.5 (2010), pp. 2741–48. ISSN: 1520-8524 (Electronic) 0001-4966 (Linking). DOI: 10.1121/1.3377056. URL: <http://www.ncbi.nlm.nih.gov/pubmed/21117722>.
- [97] Maarten Hornikx, Thomas Krijnen, and Louis van Harten. "openPSTD: The open source pseudospectral time-domain method for acoustic propagation". In: *Computer Physics Communications* 203 (2016), pp. 298–308. ISSN: 0010-4655. DOI: <https://doi.org/10.1016/j.cpc.2016.02.029>. URL: <http://www.sciencedirect.com/science/article/pii/S0010465516300443>.
- [98] E. Martin, Y. T. Ling, and B. E. Treeby. "Simulating Focused Ultrasound Transducers Using Discrete Sources on Regular Cartesian Grids". In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 63.10 (Oct. 2016), pp. 1535–1542. DOI: 10.1109/TUFFC.2016.2600862.
- [99] Anthony Grisey et al. "Simulation of high-intensity focused ultrasound lesions in presence of boiling". In: *Journal of Therapeutic Ultrasound* 4.11 (2016). DOI: 10.1186/s40349-016-0056-9.
- [100] Anthony Grisey et al. "Influence of Skin and Subcutaneous Tissue on High-Intensity Focused Ultrasound Beam: Experimental Quantification and Numerical Modeling". In: *Ultrasound in Medicine & Biology* 42.10 (2016), pp. 2457–2465. ISSN: 0301-5629. DOI: <https://doi.org/10.1016/j.ultrasmedbio.2016.06.013>. URL: <http://www.sciencedirect.com/science/article/pii/S0301562916301442>.
- [101] Joseph W. Goodman. *Introduction to Fourier Optics. 2nd Ed.* New York: McGraw-Hill, 1996.



- [102] Robert J. McGough, Thaddeus V. Samulski, and James F. Kelly. “An efficient grid sectoring method for calculations of the near-field pressure generated by a circular piston”. In: *The Journal of the Acoustical Society of America* 115.5 (2004), pp. 1942–1954. DOI: 10.1121/1.1687835. eprint: <https://doi.org/10.1121/1.1687835>. URL: <https://doi.org/10.1121/1.1687835>.
- [103] Robert J. McGough. “Rapid calculations of time-harmonic nearfield pressures produced by rectangular pistons”. In: *The Journal of the Acoustical Society of America* 115.5 (2004), pp. 1934–1941. DOI: 10.1121/1.1694991. eprint: <https://doi.org/10.1121/1.1694991>. URL: <https://doi.org/10.1121/1.1694991>.
- [104] X. Zeng and R. J. McGough. “Optimal simulations of ultrasonic fields produced by large thermal therapy arrays using the angular spectrum approach”. In: *J Acoust Soc Am* 125.5 (2009), pp. 2967–77. ISSN: 1520-8524 (Electronic) 0001-4966 (Linking). DOI: 10.1121/1.3097499. URL: <https://www.ncbi.nlm.nih.gov/pubmed/19425640>.
- [105] Sara L. Johnson et al. “Validation of hybrid angular spectrum acoustic and thermal modelling in phantoms”. In: *International Journal of Hyperthermia* 35.1 (2018), pp. 578–590. ISSN: 0265-6736. DOI: 10.1080/02656736.2018.1513168. URL: <https://doi.org/10.1080/02656736.2018.1513168>.
- [106] M. J. Voogt et al. “Volumetric feedback ablation of uterine fibroids using magnetic resonance-guided high intensity focused ultrasound therapy”. In: *European Radiology* 22.2 (2012), pp. 411–417. ISSN: 0938-7994 1432-1084. DOI: 10.1007/s00330-011-2262-8. URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3249029/>.
- [107] Scott Almquist, Dennis L. Parker, and Douglas A. Christensen. “Rapid full-wave phase aberration correction method for transcranial high-intensity focused ultrasound therapies”. In: *Journal of Therapeutic Ultrasound* 4 (30 2016). DOI: 10.1186/s40349-016-0074-7.

- [108] Joseph Blackmore, Robin O. Cleveland, and Joel Mobley. “Spatial filters suppress ripple artifacts in the computation of acoustic fields with the angular spectrum method”. In: *The Journal of the Acoustical Society of America* 144.5 (2018), pp. 2947–2951. DOI: 10.1121/1.5079637.
- [109] Yun Jing. “On the use of an absorption layer for the angular spectrum approach (L)”. In: *The Journal of the Acoustical Society of America* 131.2 (2012), pp. 999–1002. DOI: 10.1121/1.3675967.
- [110] M. E. Frijlink et al. “Abersim: A simulation program for 3D nonlinear acoustic wave propagation for arbitrary pulses and arbitrary transducer geometries”. In: *2008 IEEE Ultrasonics Symposium*. 2008, pp. 1282–1285. DOI: <https://doi.org/10.1109/ULTSYM.2008.0310>.
- [111] Xiaozheng Zeng and Robert J. McGough. “Evaluation of the angular spectrum approach for simulations of near-field pressures”. In: *The Journal of the Acoustical Society of America* 123.1 (2008), pp. 68–76. DOI: 10.1121/1.2812579. eprint: <https://doi.org/10.1121/1.2812579>. URL: <https://doi.org/10.1121/1.2812579>.
- [112] François Varray et al. “CREANUIS: A Non-linear Radiofrequency Ultrasound Image Simulator”. In: *Ultrasound in Medicine & Biology* 39.10 (2013), pp. 1915–1924. ISSN: 0301-5629. DOI: 10.1016/j.ultrasmedbio.2013.04.005.
- [113] F. Varray et al. “Fundamental and second-harmonic ultrasound field computation of inhomogeneous nonlinear medium with a generalized angular spectrum method”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 58.7 (2011), pp. 1366–1376.
- [114] J. Gu and Y. Jing. “Modeling of wave propagation for medical ultrasound: a review”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 62.11 (2015), pp. 1979–1992. ISSN: 0885-3010. DOI: 10.1109/TUFFC.2015.007034.

- [115] Ping Wu, Rymantas Kazys, and Tadeusz Stepinski. “Optimal selection of parameters for the angular spectrum approach to numerically evaluate acoustic fields”. In: *The Journal of the Acoustical Society of America* 101.1 (1997), pp. 125–134. DOI: 10.1121/1.418013. eprint: <https://doi.org/10.1121/1.418013>. URL: <https://doi.org/10.1121/1.418013>.
- [116] Joshua E. Soneson. “A User-Friendly Software Package for HIFU Simulation”. In: *AIP Conference Proceedings* 1113.1 (2009), pp. 165–169. DOI: 10.1063/1.3131405. URL: <https://aip.scitation.org/doi/abs/10.1063/1.3131405>.
- [117] C. Lafon et al. “Cylindrical thermal coagulation necrosis using an interstitial applicator with a plane ultrasonic transducer: in vitro and in vivo experiments versus computer simulations”. In: *International Journal of Hyperthermia* 16 (6 2000), pp. 508–522. DOI: 10.1080/02656730050199359.
- [118] C. Garnier, C. Lafon, and J. Dillenseger. “3-D Modeling of the Thermal Coagulation Necrosis Induced by an Interstitial Ultrasonic Transducer”. In: *IEEE Transactions on Biomedical Engineering* 55 (2 2008), pp. 833–837.
- [119] M. Necati Ozisik et al. *Finite Difference Methods in Heat Transfer*. 2nd. Boca Raton, USA: CRC Press, 2017. ISBN: 9781482243468.
- [120] R. J. McGough et al. “Treatment Planning for Hyperthermia with Ultrasound Phased Arrays”. In: *IEEE Trans Ultrasonics Ferroelectrics Frequency Control* 43.6 (1996), pp. 1074–1084.
- [121] Jean-Louis Dillenseger and Simon Esneault. “Fast FFT-based bio-heat transfer equation computation”. In: *Computers in Biology and Medicine* 40.2 (2010), pp. 119–123. ISSN: 0010-4825. DOI: <http://dx.doi.org/10.1016/j.compbiomed.2009.11.008>. URL: <http://www.sciencedirect.com/science/article/pii/S0010482509002030>.

- [122] Z. S. Deng and J. Liu. “Mathematical modeling of temperature mapping over skin surface and its implementation in thermal disease diagnostics”. In: *Comput. Biol. Med.* 34.6 (2004), pp. 495–521. ISSN: 00104825. DOI: 10.1016/s0010-4825(03)00086-6.
- [123] Z. S. Deng and J. Liu. “Monte Carlo Method to Solve Multidimensional Bio-heat Transfer Problem”. In: *Numer. Heat Tr. B-Fund.* 42.6 (2002), pp. 543–567. ISSN: 1040-7790 1521-0626. DOI: 10.1080/10407790260444813.
- [124] Kai-Yang Li et al. “The noninvasive reconstruction of 3D temperature field in a biological body with Monte Carlo method”. In: *Neurocomputing* 72.1 (2008). Machine Learning for Signal Processing (MLSP 2006) / Life System Modelling, Simulation, and Bio-inspired Computing (LSMS 2007), pp. 128–133. ISSN: 0925-2312. DOI: <https://doi.org/10.1016/j.neucom.2008.03.016>. URL: <http://www.sciencedirect.com/science/article/pii/S0925231208003615>.
- [125] H. M. M. ten Eikelder et al. “Modelling the temperature evolution of bone under high intensity focused ultrasound”. In: *Physics in Medicine & Biology* 61.4 (2016), p. 1810. ISSN: 0031-9155. URL: <http://stacks.iop.org/0031-9155/61/i=4/a=1810>.
- [126] Nick Todd, Allison Payne, and Dennis L. Parker. “Model predictive filtering for improved temporal resolution in MRI temperature imaging”. In: *Magnetic Resonance in Medicine* 63 (5 2010), pp. 1269–1279. DOI: 10.1002/mrm.22321.
- [127] Bradley E. Treeby, Elliott S. Wise, and B. T. Cox. “Nonstandard Fourier Pseudospectral Time Domain (PSTD) Schemes for Partial Differential Equations”. In: *Communications in Computational Physics* 24 (3 2018), pp. 623–634. DOI: 10.4208/cicp.OA-2017-0192.
- [128] *Heat Diffusion In A Homogeneous Medium*. Online. Accessed: 2020-09-18. URL: [http://www.k-wave.org/documentation/example%5C\\_diff%5C\\_homogeneous%5C\\_medium%5C\\_diffusion.php](http://www.k-wave.org/documentation/example%5C_diff%5C_homogeneous%5C_medium%5C_diffusion.php).

- [129] B. Gao, S. Langer, and P. M. Corry. "Application of the time-dependent Green's function and Fourier transforms to the solution of the bioheat equation". In: *International Journal of Hyperthermia* 11.2 (1995), pp. 267–285. ISSN: 0265-6736. DOI: 10 . 3109 / 02656739509022462. URL: <https://doi.org/10.3109/02656739509022462>.
- [130] Zhong-Shan Deng and Jing Liu. "Monte Carlo Simulation of the Effects of Large Blood Vessels During Hyperthermia". In: *Computational and Information Science*. Ed. by Jun Zhang, Ji-Huan He, and Yuxi Fu. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 437–442. ISBN: 978-3-540-30497-5.
- [131] Ruihua Ding and Robert McGough. *FDTD*. Version No version number. Feb. 22, 2009. URL: <https://www.egr.msu.edu/~fultras-web/download.php>.
- [132] Philips. *Sonalleve MR-HIFU: Release 3: Application Guide: Uterine Therapy*. Philips, 2012.
- [133] Merel Huisman et al. "Feasibility of volumetric MRI-guided high intensity focused ultrasound (MR-HIFU) for painful bone metastases". In: *Journal of Therapeutic Ultrasound* 2.1 (2014), p. 16. ISSN: 2050-5736. DOI: 10 . 1186 / 2050 - 5736 - 2 - 16. URL: <http://dx.doi.org/10.1186/2050-5736-2-16>.
- [134] Holger Eggers et al. "Dual-echo Dixon imaging with flexible choice of echo times". In: *Magnetic Resonance in Medicine* 65.1 (2011), pp. 96–107. ISSN: 0740-3194. DOI: 10 . 1002 / mrm . 22578. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.22578>.
- [135] Paul A. Yushkevich et al. "User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability". In: *NeuroImage* 31.3 (2006), pp. 1116–1128. ISSN: 1053-8119. DOI: <https://doi.org/10.1016/j.neuroimage.2006.01.015>. URL: <http://www.sciencedirect.com/science/article/pii/S1053811906000632%20http://www.itksnap.org/>.

- [136] Horos Project. *Horos*. Computer Program. 2015. URL: <https://horosproject.org/>.
- [137] Antoine Rosset, Luca Spadola, and Osman Ratib. “OsiriX: An Open-Source Software for Navigating in Multidimensional DICOM Images”. In: *Journal of Digital Imaging* 17.3 (2004), pp. 205–216. ISSN: 0897-1889 1618-727X. DOI: 10.1007/s10278-004-1014-6. URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3046608/>.
- [138] Helmuth Späth. “Fitting affine and orthogonal transformations between two sets of points”. In: *Mathematical Communications* 9.1 (2004), pp. 27–34. ISSN: 1331-0623.
- [139] Berthold K. P. Horn. “Closed-form solution of absolute orientation using unit quaternions”. In: *Journal of the Optical Society of America A* 4.4 (1987), pp. 629–642. DOI: 10.1364/JOSAA.4.000629. URL: <http://josaa.osa.org/abstract.cfm?URI=josaa-4-4-629>.
- [140] Travis E. Oliphant. *NumPy: A guide to NumPy*. [Online; accessed <today>]. 2006. URL: <http://www.numpy.org/>.
- [141] Stéfan van der Walt, S. Chris Colbert, and Gaël Varoquaux. “The NumPy Array: A Structure for Efficient Numerical Computation”. In: *Computing in Science & Engineering* 13.2 (2011), pp. 22–30. DOI: 10.1109/MCSE.2011.37. eprint: <https://aip.scitation.org/doi/pdf/10.1109/MCSE.2011.37>. URL: <https://aip.scitation.org/doi/abs/10.1109/MCSE.2011.37>.
- [142] Eric Jones, Travis Oliphant, Pearu Peterson, et al. *SciPy: Open source scientific tools for Python*. [Online; accessed 29/08/2019]. 2001. URL: <http://www.scipy.org/>.
- [143] Terry S. Yoo et al. *Engineering and Algorithm Design for an Image Processing API: A Technical Report on ITK - the Insight Toolkit*. Report. 2002. DOI: 10.3233/978-1-60750-929-5-586.

- [144] Matthew D. Blackledge et al. “Rapid development of image analysis research tools: Bridging the gap between researcher and clinician with pyOsiriX”. In: *Computers in Biology and Medicine* 69 (2016), pp. 203–212. ISSN: 0010-4825. DOI: <https://doi.org/10.1016/j.compbiomed.2015.12.002>. URL: <http://www.sciencedirect.com/science/article/pii/S0010482515003868>.
- [145] R. Castillo et al. “A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets”. In: *Phys Med Biol* 54 (2009). DOI: 10.1088/0031-9155/54/7/001. URL: <https://doi.org/10.1088/0031-9155/54/7/001>.
- [146] N. K. Ostergaard et al. “Acceleration and validation of optical flow based deformable registration for image-guided radiotherapy”. In: *Acta Oncol* 47 (7 2008), pp. 1286–1293. DOI: 10.1080/02841860802258760. URL: <https://doi.org/10.1080/02841860802258760>.
- [147] Dzung L. Pham, Chenyang Xu, and Jerry L. Prince. “Current Methods in Medical Image Segmentation”. In: *Annual Review of Biomedical Engineering* 2.1 (2000). PMID: 11701515, pp. 315–337. DOI: 10.1146/annurev.bioeng.2.1.315. eprint: <https://doi.org/10.1146/annurev.bioeng.2.1.315>. URL: <https://doi.org/10.1146/annurev.bioeng.2.1.315>.
- [148] Andriy Fedorov et al. “3D Slicer as an image computing platform for the Quantitative Imaging Network”. In: *Magnetic Resonance Imaging* 30 (9 2012), pp. 1323–1341. DOI: 10.1016/j.mri.2012.05.001.
- [149] Ron Kikinis, Steve D. Pieper, and Kirby G. Vosburgh. “3D Slicer: A Platform for Subject-Specific Image Analysis, Visualization, and Clinical Support”. In: *Intraoperative Imaging and Image-Guided Therapy*. Ed. by Ferenc A. Jolesz. New York, NY: Springer New York, 2014, pp. 277–289. ISBN: 978-1-4614-7657-3. DOI: 10.1007/978-1-4614-7657-3\_19. URL: [https://doi.org/10.1007/978-1-4614-7657-3\\_19](https://doi.org/10.1007/978-1-4614-7657-3_19).

- [150] C. Pinter et al. “SlicerRT – Radiation therapy research toolkit for 3D Slicer”. In: *Med. Phys.* 39.10 (2012), pp. 6332–6338.
- [151] William E. Lorensen and Harvey E. Cline. “Marching cubes: A high resolution 3D surface construction algorithm”. In: *SIGGRAPH Comput. Graph.* 21.4 (1987), pp. 163–169. ISSN: 0097-8930. DOI: 10.1145/37402.37422.
- [152] Martin Helm and Peter Hammer. *Marching Cubes*. <https://uk.mathworks.com/matlabcentral/fileexchange/32506-marching-cubes>. Version 1.3.0.0. 2013.
- [153] Sven Holcombe. *inpolyhedron - are points inside a triangulated volume?* <https://uk.mathworks.com/matlabcentral/fileexchange/37856-inpolyhedron-are-points-inside-a-triangulated-volume>. Version 3.3.0.0. 2015.
- [154] Guido Gerig, Matthieu Jomier, and Miranda Chakos. “Valmet: A New Validation Tool for Assessing and Improving 3D Object Segmentation”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2001*. Springer Berlin Heidelberg, 2001, pp. 516–523. ISBN: 978-3-540-45468-7.
- [155] D. P. Huttenlocher, G. A. Klanderman, and W. J. Rucklidge. “Comparing images using the Hausdorff distance”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15.9 (1993), pp. 850–863. DOI: 10.1109/34.232073.
- [156] Sarah A. Mason et al. “Towards ultrasound-guided adaptive radiotherapy for cervical cancer: Evaluation of Elekta’s semiautomated uterine segmentation method on 3D ultrasound images”. In: *Medical Physics* 44.7 (2017), pp. 3630–3638. ISSN: 2473-4209 0094-2405. DOI: 10.1002/mp.12325. URL: <https://www.ncbi.nlm.nih.gov/pubmed/28493295> % 20<https://www.ncbi.nlm.nih.gov/pmc/PMC5575494/>.



- [157] Viviana Carillo et al. “Contouring Variability of the Penile Bulb on CT Images: Quantitative Assessment Using a Generalized Concordance Index”. In: *International Journal of Radiation Oncology • Biology • Physics* 84.3 (2012), pp. 841–846. ISSN: 0360-3016. DOI: 10.1016/j.ijrobp.2011.12.057. URL: <http://dx.doi.org/10.1016/j.ijrobp.2011.12.057>.
- [158] Shalini K Vinod et al. “High-risk CTV delineation for cervix brachytherapy: Application of GEC-ESTRO guidelines in Australia and New Zealand”. In: *Journal of Medical Imaging and Radiation Oncology* 61.1 (2017), pp. 133–140. DOI: 10.1111/1754-9485.12509. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/1754-9485.12509>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/1754-9485.12509>.
- [159] William F. Ames. *Numerical Methods for Partial Differential Equations (Third Edition)*. Computer Science and Scientific Computing. Academic Press, 1992.
- [160] R. Seip et al. “Automated HIFU treatment planning and execution based on 3D modeling of the prostate, urethra, and rectal wall”. In: *Ultrasonics Symposium, 2004 IEEE*. Vol. 3, 1781–1784 Vol.3. ISBN: 1051-0117. DOI: 10.1109/ULTSYM.2004.1418172.
- [161] Ngo Fung Daniel Lam et al. “Prediction of pelvic tumour coverage by magnetic resonance-guided high-intensity focused ultrasound (MRgHIFU) from referral imaging”. In: *International Journal of Hyperthermia* 37.1 (2020), pp. 1033–1045. DOI: 10.1080/02656736.2020.1812736.
- [162] M. Köhler et al. “MR-only simulation for radiotherapy planning”. In: *Philips White Paper* (2015). URL: <http://incenter.medical.philips.com/>.
- [163] Koji Miyamoto et al. “Safety and Efficacy of Cryoablation in Patients With Ventricular Arrhythmias Originating From the Para-Hisian Region”. In: *JACC: Clinical Electrophysiology* 4.3 (2018), pp. 366–373. ISSN: 2405-

500X. DOI: <https://doi.org/10.1016/j.jacep.2017.12.013>.  
URL: <http://www.sciencedirect.com/science/article/pii/S2405500X17312586>.

- [164] Milka Marinova et al. "High-intensity focused ultrasound (HIFU) for pancreatic carcinoma: evaluation of feasibility, reduction of tumour volume and pain intensity". In: *European Radiology* 26 (11), pp. 4047–4056. DOI: 10.1007/s00330-016-4239-0.
- [165] Brian T. Welch et al. "Feasibility and outcomes of percutaneous thermal ablation of hepatocellular carcinoma in a transplanted allograft". In: *Abdominal Radiology* 43 (6), pp. 1478–1481. DOI: 10.1007/s00261-017-1323-0.
- [166] Chuanxing Li et al. "Superficial malignant tumors: Non-invasive treatment with ultrasonographically guided high-intensity focused ultrasound". In: *Cancer Biology & Therapy* 8.24 (2009), pp. 2398–2405. DOI: 10.4161/cbt.8.24.10277.
- [167] Department of Health. *Review of Cancer Waiting Times Standards*. Tech. rep. 15283. London, UK: United Kingdom Government, Jan. 2011. URL: <https://www.gov.uk/government/publications/the-national-cancer-strategy>.
- [168] P. Cosper et al. "Cervical Tumor Volume Doubling Time: A Pilot Study". In: *International Journal of Radiation Oncology\*Biological\*Physics* 93.3, Supplement (2015). Proceedings of the American Society for Radiation Oncology 57th Annual Meeting, E258–E259. ISSN: 0360-3016. DOI: <https://doi.org/10.1016/j.ijrobp.2015.07.1198>. URL: <http://www.sciencedirect.com/science/article/pii/S036030161501929X>.
- [169] GM Zharinov and VA Gushchin. "[The rate of tumor growth and cell loss in cervical cancer]". In: *Voprosy Onkologii* 35.1 (1989), pp. 21–25. ISSN: 0507-3758. URL: <http://europepmc.org/abstract/MED/2919502>.

- [170] Tomas Möller and Ben Trumbore. “Fast, minimum storage ray-triangle intersection”. In: *Journal of Graphics Tools* 2.1 (1997), pp. 21–28.
- [171] J. J. Jiménez et al. “Performance analysis for GPU-based ray-triangle algorithms”. In: *2014 International Conference on Computer Graphics Theory and Applications (GRAPP)*. 2014, pp. 1–8.
- [172] Herbert Edelsbrunner, David G. Kirkpatrick, and Raimund Seidel. “On the shape of a set of points in the plane”. In: *IEEE Transactions on Information Theory* 29.4 (1983), pp. 551–559.
- [173] Thomas. *Hardware accelerated ray-triangle intersection*. Version 1.2.0.0. Apr. 8, 2020. URL: <https://www.mathworks.com/matlabcentral/fileexchange/49670-hardware-accelerated-ray-triangle-intersection>.
- [174] Jesús P. Mena-Chalco. *Ray/Triangle Intersection*. <https://www.mathworks.com/central/fileexchange/25058-ray-triangle-intersection>. Computer Program. 2009.
- [175] Philip Schneider and David H. Eberly. *Geometric Tools for Computer Graphics*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2002. ISBN: 9780080478029.
- [176] Online Database. 2018. DOI: 10.13099/VIP21000-04-0. URL: [itis.swiss/virtual-population/tissue-properties/downloads/database-v4-0/](https://itis.swiss/virtual-population/tissue-properties/downloads/database-v4-0/).
- [177] Gerard T. Schuster. “Chapter 8: Numerical Solutions to the Wave Equation”. In: *Seismic Inversion*. 2017, pp. 85–100. DOI: 10.1190/1.9781560803423.ch8.
- [178] J. L. Robertson, B. T. Cox, and B. E. Treeby. “Quantifying numerical errors in the simulation of transcranial ultrasound using pseudospectral methods”. In: *2014 IEEE International Ultrasonics Symposium*. 2014, pp. 2000–2003.

- [179] William Bierman. "THE TEMPERATURE OF THE SKIN SURFACE". In: *Journal of the American Medical Association* 106.14 (Apr. 1936), pp. 1158–1162. ISSN: 0002-9955. DOI: 10.1001/jama.1936.02770140020007.
- [180] Nykolai Bilaniuk and George S. K. Wong. "Speed of sound in pure water as a function of temperature". In: *The Journal of the Acoustical Society of America* 93.3 (1993), pp. 1609–1612. DOI: 10.1121/1.406819.
- [181] W. Kreider et al. "Characterization of a multi-element clinical HIFU system using acoustic holography and nonlinear modeling". In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 60.8 (2013), pp. 1683–1698.
- [182] Nadolny, Zbigniew and Dombek, Grzegorz. "Thermal properties of mixtures of mineral oil and natural ester in terms of their application in the transformer". In: *E3S Web Conf.* 19 (2017), p. 01040. DOI: 10.1051/e3sconf/20171901040. URL: <https://doi.org/10.1051/e3sconf/20171901040>.
- [183] Adam Shaw. "A Buoyancy Method for the Measurement of Total Ultrasound Power Generated by HIFU Transducers". In: *Ultrasound in Medicine & Biology* 34.8 (2008), pp. 1327–1342. ISSN: 0301-5629. DOI: <https://doi.org/10.1016/j.ultrasmedbio.2008.01.008>. URL: <http://www.sciencedirect.com/science/article/pii/S0301562908000136>.
- [184] R. Courant, K. Friedrichs, and H. Lewy. "Über die partiellen Differenzgleichungen der mathematischen Physik". In: *Mathematische Annalen* 100 (1 1928), pp. 32–74. DOI: 10.1007/BF01448839.
- [185] H. T. O'Neil. "Theory of Focusing Radiators". In: *The Journal of the Acoustical Society of America* 21.5 (1949), pp. 516–526. DOI: 10.1121/1.1906542. URL: <https://asa.scitation.org/doi/abs/10.1121/1.1906542>.

- [186] Robert J. Gillies et al. "Causes and Effects of Heterogeneous Perfusion in Tumors". In: *Neoplasia* 1.3 (1999), pp. 197–207. ISSN: 1476-5586. DOI: 10.1038/sj.neo.7900037.
- [187] Chang W. Song. "Effect of Local Hyperthermia on Blood Flow and Microenvironment: A Review". In: *Cancer Research* 44.10 Supplement (1984), 4721s–4730s. ISSN: 0008-5472.
- [188] M. Bro-Nielsen. "Finite element modeling in surgery simulation". In: *Proceedings of the IEEE* 86.3 (1998), pp. 490–503. ISSN: 0018-9219. DOI: 10.1109/5.662874.
- [189] Morten Bro-Nielsen and Stephane Cotin. "Real-time Volumetric Deformable Models for Surgery Simulation using Finite Elements and Condensation". In: *Computer Graphics Forum* 15.3 (1996), pp. 57–66. DOI: doi : 10 . 1111 / 1467 - 8659 . 1530057. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-8659.1530057>.
- [190] M. Grujicic et al. "Seat-cushion and soft-tissue material modeling and a finite element investigation of the seating comfort for passenger-vehicle occupants". In: *Materials & Design* 30.10 (2009), pp. 4273–4285. ISSN: 0261-3069. DOI: <https://doi.org/10.1016/j.matdes.2009.04.028>. URL: <http://www.sciencedirect.com/science/article/pii/S0261306909001770>.
- [191] Wen-Ming Chen et al. "Effects of internal stress concentrations in plantar soft-tissue—A preliminary three-dimensional finite element analysis". In: *Medical Engineering & Physics* 32.4 (2010), pp. 324–331. ISSN: 1350-4533. DOI: <https://doi.org/10.1016/j.medengphy.2010.01.001>. URL: <http://www.sciencedirect.com/science/article/pii/S1350453310000020>.
- [192] Wookjin Lee, Byeong Hee Won, and Seong Wook Cho. "Finite element modeling for predicting the contact pressure between a foam mattress and the human body in a supine position". In: *Computer Methods in Biome-*

- chanics and Biomedical Engineering* 20.1 (2016), pp. 104–117. ISSN: 1025-5842. DOI: 10.1080/10255842.2016.1203421. URL: <https://doi.org/10.1080/10255842.2016.1203421>.
- [193] Rolf M. Koch et al. “Simulating facial surgery using finite element models”. In: *SIGGRAPH '96 Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM, pp. 421–428. DOI: 10.1145/237170.237281. URL: <http://doi.acm.org/10.1145/237170.237281>.
- [194] Ridha Hambli, Abdessalam Chamekh, and Hédi Bel Hadj Salah. “Real-time deformation of structure using finite element and neural networks in virtual reality applications”. In: *Finite Elements in Analysis and Design* 42.11 (2006), pp. 985–991. ISSN: 0168-874X. DOI: <https://doi.org/10.1016/j.finel.2006.03.008>. URL: <http://www.sciencedirect.com/science/article/pii/S0168874X06000461>.
- [195] Karol Miller et al. “Total Lagrangian explicit dynamics finite element algorithm for computing soft tissue deformation”. In: *Communications in Numerical Methods in Engineering* 23.2 (2007), pp. 121–134. DOI: [doi:10.1002/cnm.887](https://doi.org/10.1002/cnm.887). URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cnm.887>.
- [196] Stian F. Johnsen et al. “NiftySim: A GPU-based nonlinear finite element package for simulation of soft tissue biomechanics”. In: *International Journal of Computer Assisted Radiology and Surgery* 10.7 (2015), pp. 1077–1095. ISSN: 1861-6429. DOI: 10.1007/s11548-014-1118-5. URL: <https://doi.org/10.1007/s11548-014-1118-5>.
- [197] Zeynettin Akkus et al. “Deep Learning for Brain MRI Segmentation: State of the Art and Future Directions”. In: *Journal of Digital Imaging* 30.4 (2017), pp. 449–459. ISSN: 0897-1889 1618-727X. DOI: 10.1007/s10278-017-9983-4. URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC5537095/>.

- [198] F. Milletari, N. Navab, and S. Ahmadi. “V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation”. In: *2016 Fourth International Conference on 3D Vision (3DV)* (Stanford, CA, USA). IEEE, Oct. 2016, pp. 565–571. DOI: 10.1109/3DV.2016.79.
- [199] Liang Liang et al. “A deep learning approach to estimate stress distribution: a fast and accurate surrogate of finite-element analysis”. In: *Journal of The Royal Society Interface* 15.138 (2018), p. 20170844. DOI: 10.1098/rsif.2017.0844. eprint: <https://royalsocietypublishing.org/doi/pdf/10.1098/rsif.2017.0844>. URL: <https://royalsocietypublishing.org/doi/abs/10.1098/rsif.2017.0844>.
- [200] Luning Sun et al. “Surrogate modeling for fluid flows based on physics-constrained deep learning without simulation data”. In: *Computer Methods in Applied Mechanics and Engineering* 361 (2020), p. 112732. ISSN: 0045-7825. DOI: 10.1016/j.cma.2019.112732. URL: <http://www.sciencedirect.com/science/article/pii/S004578251930622X>.