

The Effects of Vertebral Variation on the Mechanical Outcomes of Vertebroplasty

Gavin Antony Day

Submitted in accordance with the requirements for the degree of
Doctor of Philosophy

University of Leeds

School of Mechanical Engineering

December 2018

The candidate confirms that the work submitted is his own, except where work which has formed part of jointly authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

Chapter 2 of this thesis contains material published in a jointly authored paper:

Zapata-Cornelio, F.Y., Day, G.A., Coe, R.H., Sikora, S.N., Wijayathunga, V.N., Tarsuslugil, S.M., Mengoni, M. and Wilcox, R.K., 2017. Methodology to Produce Specimen-Specific Models of Vertebrae: Application to Different Species. *Annals of biomedical engineering*, 45(10), pp.2451-2460.

Specimens dissected, tested and modelled from Chapter 2 formed a larger set of vertebral specimens in the publication. Other work carried out in the publication included analysis of the trabecular bone properties of the bovine tail vertebrae.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

©2018 The University of Leeds and Gavin Antony Day

The right of Gavin Antony Day to be identified as Author of this work has been asserted by Gavin Antony Day in accordance with the Copyright, Designs and Patents Act 1988.

Acknowledgements

I would like to thank all of the people who helped me throughout this PhD project. Of these, I would especially like to thank my primary supervisor, Professor Ruth Wilcox, for her tireless optimism, guidance and ability to understand my rambling during meetings. Another special thank you goes to my co-supervisor, Dr Alison Jones, without whose continuous stream of ideas my project would likely be less interesting and my thesis an even dryer read. I would like to thank Dr Ruth Coe and Dr Sebastian Sikora for being so welcoming at the start of my PhD and for their continued valuable input.

All of the members of IMBE have made the past four years both a pleasure and an incredible learning opportunity. The wealth of knowledge and help I have been given from the people around me has been tremendously valuable and has certainly helped me write this thesis. Members of both the main PhD office and “writing up” office have made some of the longer days seem much shorter, for which I am immensely thankful.

On a personal note, I would like to thank my Mum and Dad for their continuous support and resolute belief in me. Without their care for my education at a young age I may not have found science as fascinating as I did and still do. I would like to thank my girlfriend, Rachel, for all of the love, compassion and patience she has shown me. For sticking with me throughout my undergraduate degree and PhD I will be forever grateful.

This work was funded by the ERC and EPSRC.

Abstract

Osteoporotic vertebral compression fractures are a commonly encountered clinical problem that causes a reduced quality of life for a large proportion of those affected. One of the treatments for this type of fracture is vertebroplasty, where the injection of bone cement into the vertebral body aims to stabilise the vertebra and relieve pain. Despite being a frequently used treatment a number of studies and randomised clinical trials have questioned the efficacy of the procedure. These clinical trials and studies have suggested that the procedure is no more effective than a placebo in terms of pain relief.

Finite Element (FE) models allow an investigation into the structural and geometric variation that affect the response to augmentation. However, current specimen specific FE models are limited due to the poor reproduction of cement augmentation behaviour.

The aims of this thesis were to develop new methods of modelling the vertebral body in an augmented state, using these models as an input to a statistical shape and appearance model (SSAM). Methods were developed for experimental testing, cement augmentation and modelling through a specimen specific modelling approach to create and solve FE models. These methods were initially used with bovine tail vertebrae and then refined for the use of human lumbar vertebrae. These latter models formed the input set for the creation of a SSAM, where vertebral and augmentation variables were examined.

Models of augmentation in human lumbar vertebrae achieved a good agreement with their experimental counterparts through the development of novel modelling techniques. A new SSAM method has been developed for human lumbar vertebrae and applied to evaluate the mechanical performance of vertebroplasty. The tools developed can now be applied to examine wider patient cohorts and other clinical therapies.

Contents

1 Literature Review	1
1.1 Introduction	1
1.2 The Spine	2
1.2.1 Vertebrae	3
1.3 Vertebral Fractures & Vertebroplasty	5
1.3.1 Vertebral Fracture Types	5
1.3.2 Risk Factors	6
1.3.3 Diagnosis	7
1.3.4 Vertebroplasty	7
1.4 Experimental Studies	11
1.4.1 Methods of Characterising Bone Quality and Structure	11
1.4.2 Studies on Vertebral Fractures and Vertebroplasty	14
1.4.3 Discussion of Experimental Studies	20
1.5 Finite Element Modelling (FE)	20
1.5.1 Geometry & Meshing	21
1.5.2 Material Properties	25
1.5.3 Validation of Models	26
1.5.4 Modelling Vertebroplasty	26
1.5.5 Discussion of FE studies	30
1.6 Capture of Population Variation	31
1.6.1 Statistical Shape and Appearance Models	32
1.6.2 Discussion	36
1.7 Conclusion	37
1.8 Aims & Objectives	37
2 Bovine Tail Vertebrae Study	39
2.1 Introduction	39
2.2 Experimental Methods	39

2.2.1	Introduction	39
2.2.2	Specimen Preparation	40
2.2.3	Axial Compression	42
2.2.4	Vertebroplasty	45
2.2.5	MicroCT Scanning	48
2.2.6	Results	48
2.3	Finite Element Modelling	51
2.3.1	Introduction	51
2.3.2	Software and Imaging Resources	52
2.3.3	Model Creation	53
2.3.4	Material Properties	57
2.3.5	Sensitivity Tests	60
2.3.6	Statistical Analysis	65
2.3.7	Results	66
2.4	Discussion	70
2.4.1	Experimental Discussion	70
2.4.2	Computational Results Discussion	72
3	Non-augmented Human Lumbar Vertebra Study	75
3.1	Introduction	75
3.2	Experimental Methods	75
3.2.1	Specimens	76
3.2.2	Dissection & Potting	76
3.2.3	Imaging	77
3.2.4	Loading	77
3.2.5	Mechanical Data Processing	78
3.2.6	Image Processing	79
3.3	Computational Methods	81
3.3.1	Vertebral Modelling Using the BV/TV Method	82
3.3.2	Sensitivity Studies	84
3.3.3	Material Properties	86
3.3.4	Vertebral Density Visualisation	87
3.4	Results	87
3.4.1	Experimental Results	88
3.4.2	Computational Results	89
3.4.3	Sensitivity Study Results	92

3.5	Discussion	97
3.5.1	Experimental Results	97
3.5.2	Limitations Regarding Specimen Numbers & Vertebral Features . . .	97
3.5.3	Computational Results	98
4	Augmented Human Lumbar Vertebrae Study	104
4.1	Introduction	104
4.2	Experimental Methods	104
4.2.1	Vertebroplasty	104
4.2.2	Loading and Repeated Loading Tests	106
4.3	Computational Methods	108
4.3.1	Modelling Augmentation	109
4.3.2	Measuring the Vertebral Body & Augmentation Volume	112
4.3.3	Cement Distribution Visualisation	113
4.4	Results	113
4.4.1	Experimental Results	113
4.4.2	Computational Modelling	121
4.5	Discussion	122
4.5.1	Experimental Results	122
4.5.2	Computational Results	124
4.5.3	Quantification of Error	126
4.6	Conclusion	127
5	Statistical Shape & Appearance Modelling	129
5.1	Introduction	129
5.1.1	General Methodology	130
5.1.2	Chapter Overview	131
5.2	Methods	132
5.2.1	Statistical Shape and Appearance Model Creation	133
5.2.2	The Generation of Vertebral Models to Understand Variation	134
5.2.3	Measurement of Vertebral Geometry	136
5.2.4	Validation of Vertebral Variation	138
5.2.5	Load Position Sensitivity Tests	139
5.3	Results from the Generated Models	139
5.3.1	General Results of the Statistical Shape and Appearance Model . . .	139
5.3.2	Validation of Generated Vertebrae	140
5.3.3	Measuring & Interrogating Vertebral Variation	144

5.4	Discussion of Non-Augmented Generated Vertebrae	156
5.4.1	Validation	156
5.4.2	Variation	158
5.4.3	Load Position Effect	161
5.4.4	Summary of the Non-Augmented Generated Vertebra	162
5.5	Methods of Characterising Augmentation	162
5.5.1	Artificial Augmentation of Generated Models	162
5.5.2	Augmented Region Volume Sensitivity	163
5.5.3	Augmented Region Position Sensitivity	163
5.6	Results of Augmentation in Generated Models	164
5.6.1	Influence of Augmentation Volume on Vertebral Stiffness	164
5.6.2	Influence of Augmented Region Position on Vertebral Stiffness . . .	167
5.6.3	Loading Position Effect on Model Stiffness	170
5.7	Discussion of Augmented Generated Vertebrae	172
5.7.1	Response to Augmentation	172
5.7.2	Load Position Effect	175
5.8	Conclusion	176
6	Discussion and Conclusion	178
6.1	Introduction	178
6.2	Discussion	178
6.2.1	Experimental Methodologies and Results	178
6.2.2	Computational Methodologies and Results	180
6.2.3	Application to Clinical Practice	182
6.3	Future Recommendations	183
6.3.1	Further Computational Modelling	183
6.3.2	Continuation of Statistical Shape and Appearance Modelling	184
6.4	Conclusion	185

List of Figures

1.1	Curvature of the vertebral column with the four regions labelled. Adapted from [5].	3
1.2	Superior view of a lumbar vertebrae, with key features labelled. Adapted from [5].	4
1.3	Types of VCF. Left: Intact FSU. Left-Mid: Anterior wedge fracture. Right-Mid: Posterior wedge fracture. Right: Vertebral body collapse.	5
1.4	Three approaches to vertebroplasty. A, transpedicular approach, B, parapedicular approach, C, oblique approach. Adapted from [5].	8
1.5	A, the mean BMD from cylindrical cores taken from the lumbar spine of five different species. B, the mean fracture stress for the same five samples. Bars indicate the range of values. Adapted from [41].	14
1.6	Voxel finite element mesh of L2 vertebrae used by Buckley et al. with voxel size of 2 mm [61].	24
1.7	A: Top view μ CT scans of human vertebrae augmented <i>in vitro</i> , showing gradual reduction of cement opacity to the edges of the internal cement volume, adapted from Belkoff et al. [24]. B: The augmented model generation for the study carried out by Liebschner et al. [72].	27
1.8	Femoral shape variations for the first three modes from the PCA. From 115 bones the maximum eigenvalue is shown with the minimum eigenvalue shown in wireframe. Taken from [82].	35
2.1	Flow-chart detailing the experimental process from initial dissection to final load test.	40
2.2	Photograph and diagram depicting the method of creating end-caps for the specimens. Red arrows indicate the location of the removed pedicles.	41
2.3	The experimental setup for axial loading the vertebral specimens.	43
2.4	The difference between failure (A) and non-failure (B) for bovine tail vertebra compressed to a maximum load of 9500 N or until a peak was observed.	43

2.5	A typical load displacement curve showing how the gradient was taken from 0.3 mm long sections incremented at 0.1 mm across the length of the curve.	44
2.6	The difference seen when measuring the greatest gradient (stiffness) using different portions of the load displacement curve. From 0 to 1500 N, 0 to 5000N and 0 to 9500N.	45
2.7	The 9 gauge vertebroplasty needles used for augmentation, showing the hollow needle and the pointed insert.	46
2.8	A: μ CT scan of an augmented vertebrae, with some visible PMMA residing in the needle channel. B: μ CT scan of an augmented vertebrae using PMMA mixed with barium sulphate.	47
2.9	A: μ CT scan of an augmented vertebrae showing the cement leaking from vascular channels on the anterior side. B: Photograph of an augmented vertebrae cut into four quarters showing a vascular channel leading into the spinal canal.	48
2.10	μ CT scans of two vertebra, showing the cement leaking into the spinal canal and out of the vascular channels and the vertebral surface.	48
2.11	The maximum stiffness of 12 bovine tail vertebrae between 0 and 5000 N taken from the load - displacement data. Stiffness values for the intact vertebrae, following fracture and following subsequent vertebroplasty are shown. * Indicates those specimens that achieved a clear failure below 9500 N.	49
2.12	Axial μ CT slices of T2-CC2 (A) and T8-CC2 (B), with cement masked in red, showing the extend of cement fill at the point where the cement was most anterior.	50
2.13	μ CT scans of four augmented vertebra using a steel rod to fill the spinal canal and blu-tac to cover the external vascular channels. Shows greatly reduced cement content within the spinal canal with less cement at the surface of vascular channels. (A) T2 CC1, (B) T6 CC1, (C) T9 CC3, (D) T9 CC2.	51
2.14	A: The difference between the post augmentation and fractured stiffness against the intact stiffness. B: The difference between the post augmentation and intact stiffness against the intact stiffness.	51
2.15	Side and top view of a vertebral μ CT scan showing the effect of the down-sample from 82 μ m to 1mm cubed.	54
2.16	Side view of a vertebral model showing segmented vertebra, including the internal void that is filled.	54

2.17	Side & top down view of a vertebral model showing the alignment of the analytical rigid plane.	56
2.18	A lateral slice through an augmented bovine tail vertebra, showing the cement mask in red at the full 82 μm resolution.	57
2.19	An illustration of the gradient descent method used to optimise the greyscale conversion factor for the models.	59
2.20	Depiction of the yielding material interface creation	60
2.21	Mid-slice through an augmented vertebra, cyan: vertebral body, red: cement. A, element size of 1 x 1 x 1 mm. B, element size of 2 x 2 x 2 mm.	61
2.22	The effect of moving the cement volumes from the most anterior position to the most posterior position on the recorded stiffness of the vertebra, when using uniform 12 % fill volume of cement.	63
2.23	The effect of moving the cement volumes from the most superior position to the most inferior position on the recorded stiffness of the vertebra, when using uniform 12 % fill volume of cement.	63
2.24	A density map of the T12 CC2 vertebra. Yellow/orange elements are the most dense, while blue elements are the least.	64
2.25	The <i>in silico</i> compared with <i>in vitro</i> stiffness for intact specimens (triangles) and augmented specimens (circles). The dotted line shows a one-to-one correspondence.	66
2.26	The percentage decrease in the vertebral stiffness after reducing the elastic modulus of the cement volume within 12 augmented vertebrae.	68
2.27	The effect of reducing the elastic modulus of the cement volume within 12 augmented vertebrae. Shows the <i>in silico</i> stiffness for the six elastic moduli tested against their <i>in vitro</i> stiffness.	68
2.28	A comparison of the different approaches to modelling augmentation in the bovine tail vertebrae.	69
3.1	The removal of the posterior elements, through cutting at the narrowest part of the pedicle.	77
3.2	The holding of the vertebrae during potting in PMMA endcaps.	77
3.3	The effect of reducing the segment size on the maximum stiffness reported from four human vertebrae loaded to 2000 N. Using an increment size of 1 data point (0.0017 mm) and segment sizes of 100 to 1 data point (0.17 mm to 0.0017 mm).	79

3.4	The normalised (with respect to the volume of the ROI) histogram data for the 14 lumbar vertebrae.	80
3.5	A comparison of the direct-greyscale method (A, B and C) and the bone volume fraction method (D, E and F). A shows the full (82 μm) resolution scan, B shows the same image downsampled to 1 mm resolution and C shows the segmented scan. D shows the segmented bone at 82 μm , E shows this image downsampled to 1 mm and F shows this image after segmentation.	84
3.6	The effect of having tied contacts between the endcaps and the bone, with the increases in Von Mises stress indicated.	85
3.7	The loading positions used for the load variation tests. Position X is the original central loading position for ap and lr loading, other loading positions are 1 mm, 2 mm, 10 mm and 20 mm away from the centre.	86
3.8	The relationship between the bone volume fraction and the degree of anisotropy with a correlation coefficient of 0.74.	89
3.9	The results of using the direct-greyscale method and the bone volume fraction method to model non-augmented human lumbar vertebrae. Red shows the agreement using the direct-greyscale method and blue shows the agreement using the bone volume fraction method, with the orange dotted line showing perfect agreement.	90
3.10	A and B: The resultant 1 mm ³ greyscale background from the direct-greyscale and bone volume fraction method respectively, where the brightness in A has been increased for visibility. C and D: a heat map of the greyscale values through the mid slice of models from the direct-greyscale and bone volume fraction method respectively.	90
3.11	The variation in the greyscale distribution across the mid-slice of the vertebrae from Spine 1, 2, 3 and 4. The general density changes as well as shifting density distributions are visible.	91
3.12	The effect of the loading position for the human lumbar vertebrae, shown as a percentage change compared to a central loading position, for load positions from the posterior to anterior according to Figure 3.7	93
3.13	The effect of the loading position for the human lumbar vertebrae, shown as a percentage change compared to a central loading position, for load positions from the left to the right according to Figure 3.7.	94

3.14	The change in stiffness following a change to the thickness of (a) the inferior, (b) the superior and (c) both cement endcaps. In each case, the effect of $\pm 1, 2$ and 3 mm change in thickness is shown. The endcaps had a tied contact to the vertebrae.	95
3.15	The change in stiffness following a change to the thickness of (a) the inferior, (b) the superior and (c) both cement endcaps. In each case, the effect of $\pm 1, 2$ and 3 mm change in thickness is shown. The endcaps have a frictionless contact to the vertebrae.	96
3.16	An example showing the ability to create a greyscale background similar to that achieved using the BV/TV method, but instead adjusting the contrast of the background from the direct-greyscale method.	101
3.17	An illustration of the slipping that occurs in Spine 4 L5 when the endcap depth is reduced by 3 mm on the superior and inferior endcaps	102
4.1	An illustration showing the approach to needle insertion and cement position.	105
4.2	A μ CT scan showing the injected volume of cement at the anterior of the vertebral body with the cement track from the exiting needle to the right. For this specimen cement leaked through the anterior wall, limiting the quantity of cement injected.	105
4.3	The stiffness of the augmented vertebral specimens over the course of an initial load and three repeated loads.	107
4.4	The load - displacement results for the initial load for all used vertebra . . .	108
4.5	The load - displacement results for the Spine 3 L2 vertebra. Showing results of the intact load and post augmentation load as well as the three repeats. .	108
4.6	An illustration of the registration process, registering the non-augmented scan (A) with the augmented scan (B) and showing the combined image in C. The images are at 1 mm resolution, showing the artifacts created by the barium sulphate when viewed at this resolution, while to increase accuracy the registration carried out in Slicer 3D was at the full $82 \mu\text{m}$ resolution. . .	110
4.7	An illustration showing the origin of the masks from the non-augmented scan (A), the augmented scan (B) and the combination in C.	111
4.8	Adding the needle track to registered models.	111
4.9	A diagram showing the fitting of the cylinders to the vertebral body, where the point cloud describing the vertebra comes from an STL file, the red ring shows the points included in the mid-slice and the central red dot indicates the centre of the mid-slice.	112

4.10	The experimental stiffness of the vertebrae pre and post augmentation. Showing that 5 of the 14 vertebrae showed an increase in stiffness following augmentation.	115
4.11	The relationship between the percentage cement fill of the total vertebral volume achieved through the augmentation procedure and the bone volume fraction value for each vertebra.	115
4.12	The effect of using the fitted cylinder as a measure of vertebral body volume, on the relationship between augmentation fill volume and change in stiffness following augmentation. The red points show the relationship using the total vertebral volume, with the relationship using the fitted cylinder shown in blue.	116
4.13	Transverse and sagittal 3D projection views of the Spine 1 spine following augmentation, with the cement region being the brighter material in the vertebral body. The L1 vertebra is defined as having a dispersed volume of cement while the remaining L2, L3, L4, L5 vertebrae have concentrated volumes of cement.	118
4.14	Transverse and sagittal 3D projection views of the Spine 2 spine following augmentation, with the cement region being the brighter material in the vertebral body. The vertebra is defined as having a dispersed volume of cement.	119
4.15	Transverse and sagittal 3D projection views of the Spine 3 spine following augmentation, with the cement region being the brighter material in the vertebral body. The L1 vertebra is defined as having a dispersed volume of cement while the remaining L2 and L3 vertebrae have concentrated volumes of cement.	119
4.16	Transverse and sagittal 3D projection views of the Spine 4 spine following augmentation, with the cement region being the brighter material in the vertebral body. The L1, L2 and L5 vertebrae are defined as having a dispersed volume of cement while the remaining L3 and L4 vertebrae have concentrated volumes of cement.	120
4.17	The relationship between the percentage fill of the total vertebral volume and the percentage change in the vertebral stiffness following augmentation. The line and r value are for the red points where the cement volume was characterised as concentrated. The remaining blue points indicated the vertebra where the cement volume was characterised as dispersed.	121

4.18	The results of using the initial method (red circles), the registration method (blue triangles) and using the defined needle tracks (green triangles), showing the agreement to the perfect $x=y$ line in orange.	122
5.1	A flow chart describing the study design and usage of different generated models.	132
5.2	The three cuboids fitted to the vertebra point cloud in the axial plane. . . .	137
5.3	The three cuboids fitted to the vertebra point cloud in the coronal plane. . .	137
5.4	The three cuboids fitted to the vertebra point cloud in the sagittal plane. .	138
5.5	Geometric variation in the generated vertebrae.	140
5.6	Histogram of the models generated from principal component 1, with $\pm 3, 2$ & 1 standard deviations away from the mean. The values are normalised with respect to the total volume of the vertebrae, allowing comparisons with different sized model. The grey background represents the range of histograms seen in the input set.	142
5.7	Histogram of the models generated from principal component 2, with $\pm 3, 2$ & 1 standard deviations away from the mean. The values are normalised with respect to the total volume of the vertebrae, allowing comparisons with different sized model. The grey background represents the range of histograms seen in the input set.	142
5.8	Histogram of the models generated from principal component 3, with $\pm 3, 2$ & 1 standard deviations away from the mean. The values are normalised with respect to the total volume of the vertebrae, allowing comparisons with different sized model. The grey background represents the range of histograms seen in the input set.	143
5.9	Greyscale backgrounds for the axial mid-slice from, A, the Spine 1 L1 input vertebra and B, the mean generated model.	143
5.10	The vertebral body volume variation for the first three principal components, including the variation at $\pm 1, 2$ & 3 standard deviations away from the mean and including the mean.	144
5.11	Three dimensional views of the surface point clouds of the vertebral models from the first three principal components, showing the mean, +3 and -3 standard deviations away from the mean. Showing how the geometry is captured in the first three principal components.	146

5.12	Axial views of the mid slice of the vertebral models from the first three principal components, showing the mean, +3 and -3 standard deviations away from the mean. Showing how the geometry is captured in the first three principal components.	147
5.13	Sagittal views of the mid slice of the vertebral models from the first three principal components, showing the mean, +3 and -3 standard deviations away from the mean. Showing how the geometry is captured in the first three principal components.	148
5.14	Coronal views of the mid slice of the vertebral models from the first three principal components, showing the mean, +3 and -3 standard deviations away from the mean. Showing how the geometry is captured in the first three principal components.	149
5.15	The mean greyscale variation for the first three principal components, including the variation at ± 1 , 2 & 3 standard deviations away from the mean and including the mean.	152
5.16	The variation in the greyscale distribution across the mid-slice of the vertebrae generated from the first three principal components.	153
5.17	The stiffness variation for the first three principal components, including the variation at ± 1 , 2 & 3 standard deviations away from the mean and including the mean.	154
5.18	Effect of loading position on the vertebral stiffness.	155
5.19	Sagittal views of the L1 and L5 vertebrae from spine 1 and spine 4 of the input set in A and B respectively, showing the opposing wedge shapes of the two.	160
5.20	An illustration of the loading positions in the mean model.	163
5.21	Effect of augmentation in principal component 1.	165
5.22	Effect of augmentation in principal component 2.	165
5.23	Effect of augmentation in principal component 3.	166
5.24	The stiffness response to augmented region movements on the X axis, from left to right when viewing from the coronal plane.	168
5.25	The stiffness response to augmented region movements on the Y axis, from the posterior to the anterior of the vertebral body.	168
5.26	The stiffness response to augmented region movements on the Z axis from inferior to superior positions.	169
5.27	Effect of loading position on the augmented vertebral stiffness.	171

List of Tables

1.1	Two randomised clinical trials, with their similar inclusion and exclusion criteria and results for the study.	9
1.2	Comparison of the methods used in studies carrying out vertebroplasty experimentally on cadaveric specimens.	19
1.3	The geometry generation, meshing and material property assignment methods for 7 studies modelling single vertebrae to acquire stiffness and strength data.	22
1.4	Method of geometry generation, cement position, location and materials used for five finite element studies of vertebroplasty.	28
2.1	The volume of cement and the vertebra volume for the 12 specimens used, along with the percentage cement fill and an indication as to whether the stiffness values of the augmented vertebrae were greater than the fractured stiffness. This information was measured from the down-sampled models generated from μ CT scans of the vertebrae.	50
2.2	The difference between interaction properties, tied and not tied for four augmented vertebrae specimens.	62
2.3	The variability the modelling approaches of four users with 8 models, each users models undergoing separate greyscale optimisation and FE model solving	65
2.4	The mean, standard deviation and concordance correlation coefficient (CCC) of the intact and augmented vertebrae (simple tied interaction) for <i>in vitro</i> and <i>in silico results</i>	67
3.1	Details of the lumbar sections used from four cadaveric spines.	76
3.2	Settings used for the ImageJ plugin, BoneJ: Optimise Threshold.	81
3.3	The suggested threshold from the optimise threshold BoneJ tool for the 14 lumbar vertebrae in the set.	81

3.4	The vertebrae used for the calibration set and the validation set in the optimisation process.	87
3.5	BV/TV and degree of anisotropy (DA) values found using the ImageJ plugin BoneJ tools Volume Fraction and Anisotropy respectively using a threshold of 18.	88
3.6	The experimental stiffness results.	88
4.1	The material properties for the injected cement volumes and its surrounding interface region for the three main modelling methods used. P.P., indicates perfectly plastic	112
4.2	The experimental stiffness results for pre and post augmentation.	114
4.3	Vertebral body volume compared to the volume of the fitted cylinder.	116
4.4	The characterisation of the augmented vertebrae into vertebrae with dispersed and concentrated volumes of cement.	117
5.1	Descriptions of the models used for characterisation of generated models and of the variation described by different principal components. These models were also used for subsequent tests on the effect of augmentation.	134
5.2	Standard deviations of the models generated to validate the geometric, material property and stiffness variation. The vertebrae form a “cube” of variation that should describe all possible variation within 1 S.D of the mean for the first three principal components. 0 describes the mean position for the principle component and the 0,0,0 vertebrae represents the overall mean model.	135
5.3	The measurements and abbreviations taken from the vertebrae, used to compare and validate the generated models to the input models.	138
5.4	Variation captured in the first five principal components.	139
5.5	The measurements of the 18 variables (abbreviations in Table 5.3), shown as a fraction of the mean model’s measurements. Colouration shows the reduced measurements in red and larger measurements in green showing the geometric variation quantitatively. The colouration in the table allows easy identification of the changes to the vertebral geometry compared to the mean generated model. For example, while the changes to the measurements are near uniform in PC 1 & 3, the non-uniform nature of PC 2 shows the different shapes that are present and described above. The acronyms for the measurements are given in Table 5.3	150

5.6 The stiffness of the generated models from PC1, 2 and 3, showing the fractional change to the mean model in each case. 152

List of Abbreviations

2D	Two-Dimensional
3D	Three-Dimensional
BCs	Boundary Conditions
CC	Coccyx Vertebrae
CCC	Concordance Correlation Coefficient
CT	Computed Tomography
DICOM	Digital Imaging and Communications in Medicine
STL	Stereo Lithography File
E	Young's modulus
FE	Finite Element
PCA	Principal Component Analysis
PC	Principal Component
SSAM	Statistical Shape and Appearance Model
SSM	Statistical Shape Model
RMS	Root Mean Squared
ν	Poisson's ratio

Chapter 1

Literature Review

1.1 Introduction

It is estimated that approximately 2.5 million people in England and Wales have osteoporosis [1]. The prevalence of osteoporotic vertebral compression fractures is difficult to estimate because not all fractures come to the attention of clinicians and they are not always recognised on X-ray images. However, clinically evident osteoporotic vertebral compression fractures are associated with an increase in mortality [1]. Vertebral compression fractures occur when the anterior of the vertebral body collapses as a result of trauma, cancer or osteoporosis. Osteoporotic vertebral compression fractures usually result in pain for the patient and the loss of anterior vertebral height; the pain can lead to difficulty breathing, sleeping and many other effects including a general worsening of the patient's quality of life. Most treatments aim to relieve pain, restore mobility and attempt to reduce the occurrence of further fracture. Conservative treatments for the fracture and associated pain have many significant side effects. An alternative treatment is vertebroplasty, which is a minimally invasive percutaneous injection of cement into the vertebral body. The cement used is usually polymethylmethacrylate (PMMA). It is injected into the vertebra (alongside anaesthetics and analgesics), with aims to relieve the pain and stabilise the bone to limit the occurrence of further fractures.

The results following cement augmentation show a reported decrease in post-augmentation pain in many short term clinical follow-ups [2], however longer term studies and clinical trials have questioned the efficacy of the procedure [3, 4]. A weakness of these studies is a lack of investigation into the response of different patient subsets and therefore different types of vertebrae to augment. Hence, there is a need to investigate how these different vertebrae respond to augmentation from a mechanical perspective.

The aim of the work in this thesis was to identify features and modes of variation that lie within different vertebrae and to understand how the variation affects the outcomes of vertebroplasty from a mechanical perspective. A greater understanding of the effects of vertebral variation will aid clinicians in identifying patients that are suited to the procedure and those that are best suited to alternative, more conservative treatments. To achieve this aim, a finite element approach was used, with models of different vertebrae constructed and compared to corresponding experimental specimens. These models were then used with a statistical shape and appearance tool to describe the type of variation found within them and provide an understanding of how different types of variation affect the mechanical response to vertebroplasty.

In this chapter, the background literature is reviewed. The detailed aims and objectives of the study are then presented.

1.2 The Spine

The human spine is a complex structure that has a host of biomechanical functions. At each level, there are three joints (a disc and two facet joints) which, along with two vertebrae, muscle and ligamentous tissue form a functional unit. The primary functions of the spine are: to protect the spinal cord, to provide stability and mobility, and to allow the transmittance of movement of the upper and lower extremities.

The spinal column (Figure 1.1) consists of 24 articulating vertebrae and nine fused vertebrae, divided into five regions. These regions consist of seven cervical, 12 thoracic, five lumbar, five fused sacral and three to five fused coccygeal vertebrae, which all vary in size, shape and curvature.

The curvature of the spinal column features lordosis (concave curvature, when viewed from the posterior) in the cervical and lumbar regions and kyphosis (convex curvature) in the thoracic and sacro-coccygeal regions. It has been postulated that the curvature exists to increase rotational stability – moving mass away from the centre line increases the centre of inertia about the skull-pelvis axis. The change in curvature during gait cycle reduces the loads on the skull due to absorption of energy in the tendons and musculature of the surrounding areas. The curvature is created by vertebral geometry in the thoracic and sacral regions, while in the cervical and lumbar regions the curvature is created by the intervertebral discs being wedge shaped.




Figure 1.1: Curvature of the vertebral column with the four regions labelled. Adapted from [5].

1.2.1 Vertebrae

Vertebrae vary greatly between each region and gradually within them. All vertebrae consist of the vertebral body on the anterior portion and the neural arch on the posterior portion along with pedicles and bony processes (Figure 1.2).

The vertebral body has a similar structure to the femoral head, with a low-density trabecular structure internally, with a more dense cortical shell. Due to approximately 80 percent of the compressive load of the spine being carried by the vertebral body [6], the trabeculae are orientated vertically to aid load transfer, with horizontal trabeculae providing resistance to compressive buckling. Especially in older patients, often with osteoporosis, the vertebral body can be subject to vertebral fracture, more commonly on the anterior portion, causing anterior collapse and a wedge shape when viewed from the side [6].

Vertebral bodies vary between level in terms of cross sectional area, height and strength, with the axial strength of vertebrae increasing from approximately 1300 N at the third cervical level to over 8000 N at the fourth lumbar level [7]. Similar variations have been found in the bone mineral density (BMD) [8], with the BMD varying with the loads they are expected to carry. While a reduction in BMD is most often associated with osteoporosis,

similar increases in the risk of fracture are attributed to metastatic bone disease. The shape of the vertebral body also changes between each level. Focussing on the lumbar spine, the L1 vertebra have an anterior wedge shape, where the anterior height is smaller than the posterior height. Following an inflection point at the middle (L3) lumbar vertebrae, the L5 vertebra has a posterior wedge shape with a smaller posterior height compared to the anterior height [9].




Figure 1.2: Superior view of a lumbar vertebrae, with key features labelled. Adapted from [5].

The posterior region of the vertebrae consists of two pedicles, followed by the neural arch and processes. The spinous and transverse processes allow for greater leveraging by the muscles and ligaments while the superior and inferior articular facets of neighbouring vertebrae form the facet (zygapophysial) joints.

Vertebrae at different levels have differing structures to permit their required role. Cervical vertebrae are considerably smaller and lighter than vertebrae at other levels due to the reduced weight they are required to carry (the head and neck, and forces from stabilising muscles). One identifying feature is the foramen transversarium lateral to the vertebral body which allows for the passage of the vertebral artery and vein [10].

The thoracic vertebrae vary substantially through T1 to T12 with the size of the vertebral body gradually increasing and the pedicles changing in orientation and shape. The addition of adjoining ribs through the thoracic region greatly increase the stiffness of the section due to the ligaments and costovertebral joints.

The lumbar vertebrae carry the greatest load of the spinal regions [7] and consequently have the lowest height to width ratio [11].

The sacral vertebrae are fused to form the sacrum which are followed by three to four elements that form the coccyx and fuse in adulthood.

1.3 Vertebral Fractures & Vertebroplasty

1.3.1 Vertebral Fracture Types

The two most frequent types of vertebral fracture are compression and burst fractures. These types originate from different conditions of both the intact vertebrae and the loads applied. There are different ways of characterising and classifying vertebral fractures, here they are described using the three columns defined by Denis [12].

1.3.1.1 Compression Fractures

A vertebral compression fracture (VCF) is defined as a failure of the vertebral body under the anterior column with the middle column remaining intact, a feature unique to this type of fracture [12]. Such fractures can also be identified through compression of the cancellous bone and lack of fragmentation, especially compared to burst fractures. The two main types of VCF are anterior and lateral, with the mechanism being anterior flexion and lateral flexion respectively. These types can be further divided based on whether the superior or inferior end-plate experienced failure, however failure of the superior end-plate is more common [12]. Radiographically, the anterior height of the vertebral body is reduced with no visible change or damage to the posterior region of the vertebral body [12], although occasionally in juvenile and osteoporotic vertebrae, the damage is limited to end-plate impaction [13]. A more extreme case is a vertebral body collapse, found in osteoporotic spines, where occasionally fragments of bone are produced which can violate the spinal canal. This is more common when both end-plates are impacted and such cases are treated as burst fractures [13].




Figure 1.3: Types of VCF. Left: Intact FSU. Left-Mid: Anterior wedge fracture. Right-Mid: Posterior wedge fracture. Right: Vertebral body collapse.

The upper lumbar and mid thoracic regions have been found to be the most common sites for VCFs, with the lateral type generally occurring only in the lumbar spine [12].

1.3.1.2 Demographic

Patients suffering from vertebral compression fractures (VCFs) can experience severe pain for extended amounts of time, dramatically changing daily activities in addition to the pain, through reducing the lower vital capacity and forced expiratory volume when compared to patients without such fractures [14]. Mortality rates were also reported to increase with VCFs in women, with the rate increasing further with the addition of more fractures [14].

Due to VCFs occurring most commonly through the loads applied to the spine exceeding the axial strength of the vertebral body, the demographic of sufferers is dominated by those with osteoporosis or tumours within the vertebrae [14]. The occurrence of osteoporosis is usually age related with the frequency of VCF increasing from 25 percent in postmenopausal women to 40 percent in women aged 80 years old in the United States [15]. However, as discussed by Melton & Kallmes [16], there is difficulty in defining VCFs due to the number of systems developed to define them and the lack of a clear gold standard for classification. This study found that the prevalence varied from 7% to 19% for women in the age range of 50-80 years depending on the method by which the fracture was defined and 4% to 17% for men in the same age range. The prevalence of osteoporotic VCFs in women was found to be twice that of men in an age-adjusted study [17].

VCFs due to tumour infiltration into the vertebrae is another subset of patients and was the motivation for the first image guided percutaneous vertebroplasty procedure in 1984 [18], however it is difficult to assess the rate of such occurrences. The increasing ability to treat osteolytic metastases and myeloma, leave more patients open to vertebral collapse. This is further increased by possible secondary osteoporosis induced by the treatment of malignant lesions [14]. However, specific pathologies and trauma account for a mere 3% and 14% respectively of all clinically relevant fractures [16].

1.3.2 Risk Factors

Many of the risk factors of VCFs are the same as osteoporosis due to their linked nature and can be categorised into potentially modifiable and non-modifiable risk factors [15]. Non-modifiable factors include age, gender, Caucasian race and history of existing fractures; modifiable factors include insufficient physical activity, calcium and vitamin D deficiency and alcohol and tobacco use [15].

1.3.3 Diagnosis

Approximately two thirds of VCFs are undiagnosed [19] due to back pain often being regarded as a consequence of aging and not reported by patients. Care is often required to ensure that pain is directly related to a VCF and not another spinal entity such as facet arthropathy, herniated disc and spinal stenosis [14]. Indicators for vertebroplasty are often described as pain localised to the area, which lacks suggestions of nerve or cord compression and includes an increase in pain under weight bearing [14]. Vertebral pain for 1-6 weeks that fails to reduce after oral analgesics has also been defined as an indicator [20].

Due to the strong relationship between BMD and bone strength, with stronger vertebrae exhibiting a higher BMD [21], dual energy X-ray absorptiometry (DEXA) or quantitative CT can be used to predict vertebral fractures. Increasingly, the ability to assess the bone structure at the clinical level gives a greater ability to clinicians to predict fractures more accurately.

1.3.4 Vertebroplasty

The main group of patients that receive vertebroplasty are those with either osteoporotic VCFs or those with tumour infiltration. Early use of vertebroplasty was limited to those patients who responded poorly to conservative treatments, including analgesics, bed rest, physical therapy and in some cases bracing [14]. However due to the low complication rate of vertebroplasty, the indication of an osteoporotic VCF is often enough and helps to reduce further compression of the vertebral body [20]. Vertebroplasty consists of the injection of bone cement into the vertebral body from the posterior side of the body. The addition of cement into an osteoporotic, fractured vertebral body allows stabilisation of the fracture with the aim of reducing pain for the patient.

1.3.4.1 Vertebroplasty Procedure

Performing vertebroplasty requires careful monitoring during cement injection and cannula placement; this is carried out with the use of fluoroscopic guidance, which helps to limit the possibility of extravasation (cement leakage from the vertebral body). Biplane fluoroscopy is often used, allowing the procedure to be carried out more rapidly, while single plane fluoroscopy requires checking both lateral and anteroposterior projections [14].

There are three common approaches to the vertebral body: transpedicular, parapedicular, and oblique, shown in Figure 1.4. Transpedicular vertebroplasty provides a route into the vertebral body through the pedicle, which acts as a tunnel reducing the risk of dural

puncture. Depending on patient vertebrae size and level, the pedicle may be insufficient to allow passage of the vertebroplasty needle and hence other approaches may be required. The transpedicular approach also often requires injection from both sides of the vertebrae into the vertebral body to prevent build-up of cement on one side.

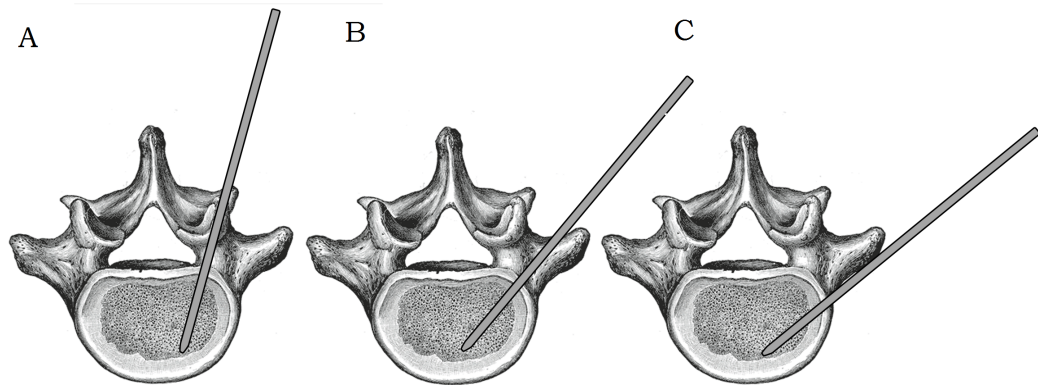


Figure 1.4: Three approaches to vertebroplasty. A, transpedicular approach, B, parapedicular approach, C, oblique approach. Adapted from [5].

The parapedicular approach involves the needle being inserted transversely across the pedicle until the vertebral body wall is reached. It has the advantage of allowing the needle to be positioned ideally for injection, removing the need for multiple injections to evenly distribute the cement. Although allowing ideal positioning of the needle tip, it requires the needle to pass close to the basilar vein, therefore increasing risk of puncturing the vein.

Finally, the oblique approach avoids the pedicle entirely, entering the vertebral body in the posterolateral corner and is often used in the thoracic region where the needle can pass over the top of the rib into the vertebral body. Disadvantages of this approach include difficulties in positioning due to the longer needle required and the increased risk of cement extravasation at the needle entry point around the exiting nerve root [22].

Table 1.1: Two randomised clinical trials, with their similar inclusion and exclusion criteria and results for the study.

Study Author	Number of patients	Type of pathologies	Inclusion Criteria	Exclusion Criteria	Results for osteoporotic patients
Buchbinder et al. [23]	71 patients, 35 in VP group, 36 in placebo group	Vertebral compression fractures	Pain < 12 months duration and presence of one or two vertebral fractures	> 90 % collapse, presence of spinal cancer, retropulsion of fragments, hip fracture or infection	No beneficial effect of vertebroplasty over a sham procedure at 1 week or at 1, 3, or 6 months among patients
Kallmes et al. [4]	131 patients, 68 VP group and 63 placebo group	Vertebral compression fractures	> 50 years of age 1-3 painful, osteoporotic vertebral compression fractures between vertebral levels T4 and L5. Fracture < 1 year old	Evidence of neoplasm, retropulsion of fragments, hip fracture or infection	No significant difference between groups one month after the procedure on measures of back pain intensity, functional disability, and quality of life

1.3.4.2 Results of Vertebroplasty and Clinical Trials

The pain relief reported following vertebroplasty is currently not fully understood; theories include effects on the nerve endings, both thermal and/or chemical, and the general stabilisation of the vertebrae due to the material properties of the cement which can restore or improve the vertebral material properties [24]. The identification of factors that affect pain relief are often difficult to spot, especially with the relatively small patient populations reported in studies, for example the study by Barr et al. [25], with 47 patients was able to find trends showing patients with single level fractures responded better to the procedure. However, the importance of other factors such as degree of kyphosis and compression, age, gender and fracture location require a much larger population in order to obtain any statistical significance.

A large systematic review by Hulme et al. [26] of 69 clinical studies achieved contrasting results to the study described above. A large proportion of patients, 87 %, had pain relief of some degree out of 1552 patients from 32 studies. However, the review also found higher than expected leakage rates, with leakage occurring in 41 % of vertebroplasty procedures and frequent new fractures were found above and below the augmented level. These fractures give an example of the need for larger, comparative, blinded and randomised clinical trials, which could determine whether these fractures are a feature of altered loading, increased patient activity or whether the new fractures would have occurred regardless of the vertebral augmentation.

The two studies summarised in Table 1.1 detail blinded randomised and controlled studies by Buchbinder et al. [23] and Kallmes et al. [4]. These studies raised questions over the risks and evidence for vertebroplasty due to their conclusions that there is no difference between the vertebroplasty and placebo groups. The near simultaneous publication of such results caused many to disregard much of the positive evidence [26, 27] for vertebroplasty. However there are some considerations regarding both of the trials: the inclusion and exclusion criteria detailed in Table 1.1 were neither clear nor well defined, failing to take results of MRI scans (Kallmes et al. [4]) and physical examinations (both) into consideration. Such results would link to accepted indications of bone marrow oedema and pain on palpation respectively [28]. In addition to a population bias towards fractures less than six weeks old, there was no statistical significance between chronic and sub-chronic patients (due the small numbers), prohibiting any insight into which subgroups of patients respond best to vertebroplasty. Additionally, the National Institute For Health And Clinical Excellence guidance document suggests that consideration for receiving the treatment should be after six weeks, due to the severity of pain reducing over that time and many patients being

pain free at six weeks post fracture [1]. The standard deviation for both pain intensity and Roland–Morris Disability Questionnaire scores were generally high, especially at one month post-procedure. For example the mean (\pm SD) RDQ score in the vertebroplasty group was 12.0 ± 6.3 , as compared with 13.0 ± 6.4 in the control group. This highlights the need to understand where this variation originates from and hence which subset of patients the treatment is better suited to. Despite the lack of significance between the sham/treatment groups, both studies reported a reduction in pain following augmentation. Suggesting that if the correct subset of suitable patients can be identified, then the treatment is a viable alternative to conservative treatments.

1.4 Experimental Studies

A range of experimental studies have been undertaken to characterise vertebrae in terms of material and bone properties, response to loading and the *in vitro* response to augmentation. The findings of such studies are outlined in this section.

1.4.1 Methods of Characterising Bone Quality and Structure

The biomechanical properties of the vertebrae are known to rely heavily on the trabecular structure, especially the compressive strength and stiffness which relates to the failure behaviour and elastic behaviour respectively. The compressive strength originates from the architecture of the load-bearing trabeculae, which is characterised by thick trabeculae columns or plates oriented vertically and held in place by much thinner horizontal trabeculae. This internal structure changes with age; the vertical plates are reduced to columns through bone remodelling and horizontal supports are often removed [29, 30]. Osteoporosis is often defined by a reduction in the Bone Mineral Density (BMD) of 2.5 standard deviations below that of a young, healthy member of the population of the same gender. However, reports of poor correlations between the BMD and vertebral fracture rates suggest that a measure of BMD is not sensitive enough to solely determine fracture risks [31], hence the trabecular architecture must be studied using more sensitive and specialised tests.

Methods for studying the relationship between structure and mechanical function of the vertebral body usually involve mechanical testing of the vertebrae or trabecular bone samples in conjunction with a study of the bone density. The trabecular structure and bone architecture in general can be identified through calculation of the ash density and comparisons of the trabecular structure through histological and μ CT examination. The following measurable parameters are usually measured using μ CT and/or histological images of the trabecular bone: the bone volume fraction (BV/TV), connectivity density

(Conn.D), Structural Model Index (SMI), degree of anisotropy (DA) and the trabecular separation, number and plate thickness (Tb.Sp, Tb.N and Tb.Th respectively) and are discussed in detail in studies by Hulme et al. [32] and Mosekilde et al. [33].

The ash density allows comparison of the densities of bone samples through the removal of water and soft tissue in addition to a calculation of the mineral content using an additional measurement of the dry weight. Bone marrow and other remaining soft tissue (fat) is often removed prior to incineration using high-pressure water jets and acetone washes [34]. The dry weight is measured following drying using a recommended 100°C furnace for an hour [34] and the dry density (gcm^{-3}) is the weight divided by the specimen volume. To ash the specimens, they are usually placed in a muffle furnace at 650°C for 18-24 hours [33, 34], which, following cooling, can be reweighed. The mineral content can be calculated through dividing the ash weight by the dry weight and ash density by dividing the ash weight by the initial specimen volume. However, with the adoption of μCT in most studies regarding trabecular structure, ash density calculations are rarely used in more recent studies.

Other parameters of bone architecture (BV/TV, Conn.D, SMI, DA, Tb.Th, Tb.N and Tb.Sp) are usually defined through μCT scans of the bone and analysis using accompanying software to the μCT scanner. The bone volume fraction, BV/TV, usually quoted as a ratio or percentage, is a measurement of the proportion of the total volume of interest that is bone tissue. Conn.D indicates the number of trabecular connections per volume of interest. The SMI provides a quantification of numbers of different types of trabecular element, usually on a scale from 0 to 3 (from rods to parallel plates)[35]. Tb.Sp and Tb.Th are measured in length and identify the average separation between trabeculae and the thickness of trabeculae respectively, while Tb.N quantifies the number of trabeculae per unit length. Finally, the degree of anisotropy measures the average alignment of the trabeculae along a specific axis, where a value of one usually specifies isotropic behaviour and less than one equates to various degrees of anisotropy [32].

The region of interest for taking specimens from the vertebral body depends on the nature of the study. If the study requires the capture of an average value for the vertebrae, then the largest possible volume of cancellous bone is required, while avoiding the cortical bone and areas where the basivertebral vein intersects [8]. Other studies have included all cancellous bone between the cranial and caudal endplates of vertebrae and have used subsections of the vertebral body to identify differences and changes with age to specific regions [32]. From such measurements bone morphology has been shown to vary greatly between different regions of the vertebrae [32, 36] and between different ages of vertebrae [36, 37].

Animal models are commonly used for the *in vitro* and *in vivo* biomechanical models of the spine while testing the performance of various treatments. However, despite providing a basic understanding of spinal function, the differences in vertebrae at different levels and the changes that the vertebrae experience with ageing mean a single species cannot be used to model the entire spine. Hence many different large mammals have been used, including bovine, ovine, porcine and cervine vertebrae, all of which have already been characterised in the literature [38–40]. These studies compare the anatomical variation in terms of vertebral body width, height and depth, spinal canal size and pedicle height and width. Studies detailing similar anatomical properties for human vertebrae also exist, allowing comparison and assessment of whether a particular animal model is appropriate [10, 11]. The anatomical differences found between the animal study papers listed above and the human studies show many differences between all of the parameters, furthering the idea of multiple animals being used for different studies. The studies suggested that sheep spines were much larger than humans, with the vertebral body height being particularly greater [39, 40], while the mean vertebral width and depth of the human spine are greater than that of all the animals in the studies (with the exception of the upper thoracic segments in the deer [40]). Much of the geometric and structural variation between human and animal vertebrae may be due to the orientation during walking and therefore load differences. Despite this disparity in posture, muscle forces add additional loads to the vertebrae, reducing the total loading differences.

A study regarding the trabecular composition of bone samples from the lumbar spine compared human, dog, pig, cow and sheep and concluded that care was required when choosing a suitable animal model for a particular study [41]. This was due to the large interspecies differences in terms of the bone quality (both density and structure), height and area, as well as fracture stress. They found that human Bone Mineral Content (BMC) and volumetric BMD (vBMD) were significantly lower when compared to the other animals in the study, aligning with the greatly reduced fracture stress also reported; the details of this can be seen in Figure 1.5. However, the study also reports a higher fracture stress in sheep compared to cows, yet similar values for the BMC and vBMD, with a similar trend also found between pigs (higher BMD) and dogs (lower BMD). Having similar density properties yet higher fracture stress, aligns with reports of poor correlations between the BMD and vertebral fracture in the study by Hordon et al [31].




Figure 1.5: A, the mean BMD from cylindrical cores taken from the lumbar spine of five different species. B, the mean fracture stress for the same five samples. Bars indicate the range of values. Adapted from [41].

1.4.2 Studies on Vertebral Fractures and Vertebroplasty

There have been a considerable number of experimental studies that have investigated the effects of vertebroplasty through mechanical testing. Often these studies are carried out in conjunction with computational studies, where properties are defined experimentally and used to validate computational models, for example, the study by Wijayathunga et al [42]. These studies are reviewed in the following section, with a focus on the preparation of specimens and the testing methods, with references made to notable findings.

1.4.2.1 Specimen Preparation

μCT scans can be used to characterise bone as described above, assess fractures and identify the level of cement penetration from vertebroplasty. Hence, pre and post-mechanical loading μCT scans (including scans before and after cement augmentation) are usually captured.

To simulate hydration of physiologic conditions, specimens are often wrapped in phosphate-buffered-solution (PBS, saline) soaked gauze following dissection [25]. Studies report storing specimens at -20°C and thawing for 24 hours before testing at a non-physiological room temperature [43, 44]. This difference between room and body temperature may affect the results, with constituents of the specimens, such as the bone marrow, having different properties at the two temperatures. However, there is a lack of research investigating the effects temperature has on mechanically testing specimens in the literature.

1.4.2.2 Mechanical Testing

The generation of vertebral fractures experimentally attempts to match the natural creation of such fractures; compression fractures are usually generated with an application of force at a low rate, mimicking the gradual creation of fractures in osteoporotic vertebrae over time. Natural burst fractures are usually the result of a traumatic event, with high energy impacts causing high rate axial loads through the vertebrae; experimentally burst fractures are usually generated through dropping a load of known mass onto the specimen from a calculated height [43, 45], generating comparable forces to a natural traumatic impact. However, other methods involving biaxial hydraulic testing machines, where high loads (50 – 100 % of the animal’s body weight) were applied over a short period of time, generating the burst fracture [46]. The remainder of this section will consider vertebral compression fractures as these are the main focus of the present study.

Mounting of specimens prior to loading in material testing machines allows the position of vertebrae to be maintained during testing, whilst not restricting or confining compression more than necessary. Methods of mounting allow parallel positioning of the exterior surfaces and allow perpendicular loading of the vertebrae along the vertebral axis. These requirements for the mounting are usually achieved through potting the specimens in PMMA [43, 44, 47], semi cured moulding material [48] or low viscosity resin [49].

Methods of loading both a pre-fracture and post fracture/augmentation rely on mimicking the natural loading of the spine, with similar methods being used throughout the literature. Loading is carried out with a materials testing machine, applying loads either under pure axial compression or allowing flexion of the upper endplate through various methods. Methods of applying the “natural” load include applying the load or displacement through a steel ball, requiring a specified loading point and allowing natural motion [43, 44]. Applying bilateral loads through pneumatic cylinders and allowing controlled flexion with simultaneous compression [47] and similarly with the application of two anterior and two posterior loads using pneumatic cylinders positioned so that more of the load was applied

to the anterior side causing flexion [49] were other methods used in the literature. The positioning of the applied load differs in the literature varying between the mid point of the superior endplate [50] and one quarter of the distance of the vertebral body in from the anterior margin of the superior endplate [24]. With some studies identifying the effect of loading position on the measured stiffness of the vertebrae [42].

The definition of the point of fracture varies in the literature, most commonly the fracture point was defined as the peak on the recorded load-displacement graph. However, Furtado et al. [44] defined the point of fracture creation as 75% of the original vertebral body height, with the failure strength defined as either the value at the end of the experiment (75% of VB height) or the peak on the load-displacement graph. Pneumáticos et al. [49] defined the point of fracture (or maximum fracture load) using the load-displacement graph, where catastrophic failure could be observed through a sudden jump in the displacement.

1.4.2.3 Results of Mechanical Loading

Mechanically loading vertebrae in the study by Furtado et al. [44] showed a significant correlation between failure strength and the product of BMD and endplate area, with the range of failure loads being 900 N to 2200 N for human thoracic vertebrae between T2 and T12. However, the failure load was defined as the load at 75 percent of the original vertebral height or the peak load before reaching this deformation, potentially suggesting that the maximum load was not achieved until after 25 percent strain. This may explain the discrepancy between these results and those of Pneumáticos et al. [49] where the average reported failure load was 6724 ± 3291 N for intact specimens using vertebrae between the thoracic levels T2 to T11 from four human cadaveric spines.

1.4.2.4 Experimental Vertebroplasty Compared with Clinical Vertebroplasty

Experimentally, transpedicular (uni or bi-pedicular) approaches are more common in the literature [44, 47, 49], although extrapedicular approaches have also been used [44].

Although bilateral transpedicular vertebroplasty is the more common clinical procedure [51], unipedicular vertebroplasty is used under some mitigating circumstances often requiring the patient to return for the second injection. A study comparing bilateral and unilateral vertebroplasty found no significant difference between the two procedures in terms of vertebral height and stiffness, possibly attributed to the central positioning of the cement despite the unilateral approach [51]. The authors reached the conclusion that a unipedicular approach is a valid alternative to bipedicular vertebroplasty and may be especially useful during multilevel vertebroplasty by reducing the number of injections and hence risk of

cement leaking outside the vertebral walls (extravasation).

1.4.2.5 Results following Experimental Vertebroplasty

The methods used to recreate the vertebroplasty procedure experimentally vary considerably between studies, especially given the sensitivity that certain studies suggest factors such as fill volume have on outcomes. Table 1.2 shows the methods and some of the more important variables used in a selection of studies carrying out vertebroplasty on cadaveric specimens. Below is discussion of some of the finer points of the studies, including comparisons to clinical studies.

Identifying the effect of vertebroplasty on intact (non-fractured) specimens in the thoracolumbar region, Higgins et al. [52] found the strength was increased by a statistically significant amount of $\sim 36\%$ using 20% fill volume and that when comparing the strength increases with BMD, it was found that those vertebrae with a lower BMD showed a more dramatic increase in the strength. Conversely Graham et al. [53], found that highly osteoporotic vertebrae showed the least improvement compared to intact vertebrae in terms of strength and stiffness.

Higgins found that the upper thoracic vertebrae failed to show any significant result following augmentation of both 10 % and 20 % compared to the intact controls. Belkoff et al. [54] suggested that 7.7 mL of cement was required to restore the original strength of fractured osteoporotic vertebrae, this corresponded to $\sim 24\%$ volume fill in the lumbar vertebrae tested.

When comparing different cements Belkoff [55] found that in order to obtain a significant increase in the strength 6 mL or approximately a fill of 18 % was required. Similarly, Lee et al. [56], found that between 25 % and 30 % volume fill was required to restore strength to the intact level for lower thoracic and lumbar vertebrae in a clinical study. Graham et al. [53], required a fill volume of 24 % (average 7 mL) to achieve statistical significance, however this did not return the stiffness to the intact level, and only returned the strength to the intact level in those specimens with the highest BMD of their group.

An increase in strength following fracture and augmentation is a desirable outcome especially for osteoporotic vertebrae where returning the strength to that of the intact vertebrae may not prevent fractures. This is often carried out as prophylactic percutaneous vertebroplasty, where the procedure is undertaken on vertebrae that are likely to fracture or on levels adjacent to fractured vertebrae already undergoing vertebroplasty. Restoring the stiffness however is believed to be responsible for pain relief due to the internal stabilisation

and prevention of micro motion, providing a more suitable environment for healing. A study examining the quantity of cement required to restore the strength and stiffness of osteoporotic vertebrae having undergone a compression fracture found that as little as 2 mL of cement is enough to restore the strength of the vertebrae to the intact value, with the quantity required to restore the stiffness being between 4 mL and 6 mL depending on the level (lumbar vertebrae requiring more cement) [24].

Regarding extravasated cement during the vertebroplasty procedure, Higgins et al. [52] found that an increase in the BMD greatly increased the tendency for cement to leak, however, this typically was from the anterior wall of the vertebra (more favourable than into the spinal canal) and was independent of vertebral level.

Despite previous studies suggesting BMD alone could not be used as an indicator for vertebral fracture [31], Higgins et al. [52] concluded that it was one of the most important factors, conceding however that the BMD measured ex-vivo is not directly comparable to that measured in a clinical setting. This was due to differences in the BMD calculation due to the surrounding soft tissue for the *in vivo* scans.

The results from Furtado et al. [44] using an extrapedicular vertebroplasty approach achieved approximately 25 % fill for specimens which had previously undergone loading to generate a compression fracture, equating to an average volume of cement of approximately 4.5 ml. This augmentation caused a significant increase in the failure load, a factor of 1.72 increasing the average failure load from $1.61 \text{ kN} \pm 0.49 \text{ kN}$ to $2.63 \text{ kN} \pm 0.85 \text{ kN}$. A similar result was achieved by Tohmeh et al. [51], showing that the post augmentation strength is significantly higher than the fractured and intact vertebrae. Pneumaticos et al. [49] found no significant difference between intact and post augmentation intact vertebrae, with an average failure load of $5.77 \pm 2.13 \text{ kN}$ for the augmented intact specimens with 6 mL of cement injected.

Table 1.2: Comparison of the methods used in studies carrying out vertebroplasty experimentally on cadaveric specimens.

Author	Type of specimen	Procedure type & fill volume	Cement Type	Key Finding
Belkoff et al. [55]	Five vertebral bodies (L1-L5) from four female cadaveric spines (age, 80 ± 5 years)	Transpedicular, 6 ml fill volume throughout	A bioactive cement, Orthocomp (Orthovita, Malvern, PA) & a PMMA based cement, Simplex P (Howmedica, Rutherford, NJ)	Significantly greater strength following injection of cement, compared to the intact vertebrae
Furtado et al. [44]	Twenty-six single vertebrae from 2 female cadavers (age, 88 and 89 years)	Extrapedicular into anterior third of vertebral body, 20 % volume fill, based on height x endplate surface area	PMMA with 20 % by dry weight of barium sulfate	Increased failure strength by a factor of 1.72 post vertebroplasty compared to the intact vertebrae
Higgins et al. [52]	Human cadaveric, 61 vertebrae from 5 cadavers with mean age of 81 years	Unipedicular, 10% and 20% cement fill by volume, with unfilled as controls	PMMA	A statistically significant 36% strength increase as compared with the unfilled controls regardless of density levels
Pneumatics et al. [49]	40 vertebrae from four human cadaveric thoracic spines (age range 65-69 years)	Transpedicular, 6 ml fill volume throughout	PMMA, with 5 to 1 ration of barium sulphate	A reduced failure load post vertebroplasty, although non-significant

1.4.3 Discussion of Experimental Studies

The results presented above describe a wide ranges of approaches to measuring vertebral variables and performing vertebroplasty. The experimental setup for the mechanical testing of vertebrae is broadly similar throughout the literature. It usually involves potting the vertebrae in a material through which loads are applied axially. These loads often attempt to mimic natural loading using either hydraulics or devices that allow flexion of the vertebral endplates. Failure loads in the literature spanned a wide range, likely originating from the quality / osteoporotic nature of the vertebrae.

Experimental vertebroplasty attempts used a range of approaches for the injection, with transpedicular being the most common. Clinically the approach would be chosen by the ease of access and the level in question. Fill volumes ranged between 10 and 30 %, with suggestions that between 25 to 30 % fill was required to restore the strength to that of an intact vertebrae.

1.5 Finite Element Modelling (FE)

The use of finite element models of the spine, spinal segments and vertebrae has been rising rapidly over the past decades. Finite element models are being used to examine a range of interventions and devices as well as to aid our understanding of spinal biomechanics. The main benefit of the finite element approach is the ability to test a range of variables and properties for the same model or specimen, with an additional benefit of assessing parameters that cannot easily be measured experimentally, such as stress and strain fields.

Important factors for the generation of biomechanical finite element models have been discussed previously [57, 58] and include verification, sensitivity testing and validation of models created. The first of these factors is an assessment of the numerical accuracy of the model; given that most studies use commercially available software for FEA this verification has usually already been carried out, with documentation available. Additional verification can be carried out regarding mesh sensitivity and convergence, where the level of detail of the model is investigated with considerations of accuracy and computational cost. Sensitivity testing determines the sensitivity of a model to various input variables and the errors that these input variables have on the system. Such tests may include the response of the system to various boundary conditions and material properties. Validation of models is a proof that the computational results agree with either *in vitro* or *in vivo* results, however, proof that the results agree with the *in vitro* model do not mean that the

model is a valid representation of the *in vivo* scenario.

Due to the importance of trabecular bone in many aspects of bone and spinal research, the methods used to model it are quite detailed throughout the literature. There are currently two dominant approaches to modelling trabecular bone using FEA, discussed in detail by Mengoni et al. [59]. These are μ FE models, where the micro structure of the bone is expressed explicitly, and continuum level FE models which represent the trabecular structure through a continuous model, often using an inhomogeneous material property to represent the variations in micro-structure. The increased computational cost of μ -FE models, especially for full sized vertebrae and larger functional spinal units, means that many studies use continuum level FEA, despite the reduced level of detail. However, the development of continuum level models allows the use of much larger CT scan data sets, given that the resolution of clinical μ CT scanners is upwards of 1.2 mm³.

The studies summarised in Table 1.3 use various methods to acquire vertebrae geometry, generate models and apply appropriate material properties, either from combinations of scan data and experimental data, or homogenous experimentally defined properties. The variations between methods and results which can be drawn from these studies is described below.

1.5.1 Geometry & Meshing

With the growing availability of μ CT scanners and the increasing resolution available with them, studies using *in vitro* measurements and handcrafted geometries such as the study by Higgins et al. [52] in 2007 are being replaced with geometries developed from μ CT data. One of the more common methods of generating specimen specific models of vertebrae is the conversion of voxels from down-sampled μ CT images into hexahedral elements. The direct conversion of voxels into elements has the advantage of increasing the simplicity of model generation, however, this requires real specimens (in the form of animal or human tissue) which can be difficult to acquire and scan.

Table 1.3: The geometry generation, meshing and material property assignment methods for 7 studies modelling single vertebrae to acquire stiffness and strength data.

Author	Geometry Generation	Meshing	Material Properties
Robson Brown et al. (2014) [60]	Used μ CT with resolution of 0.074 x 0.074 x 0.074 mm	Using ScanIP software, segmented CT images down sampled to 1 x 1 x 1 mm and meshed in ScanFE using hexahedral elements and tetrahedral elements for the smooth vertebral surface	Material properties assigned using density values from μ CT. No division between cortical and cancellous bone. Relationship between elastic modulus and BV/TV value investigated
Buckley et al. (2006) [61]	Used μ CT data with resolution of 1 x 1 x 1 mm	Conversion of voxels to hexahedral elements of size 1 x 1 x 1 mm	Material properties assigned using density values from μ CT. No division between cortical and cancellous bone
Chevalier et al. (2009) [62]	Used HR-pQCT data with resolution of 0.082 x 0.082 x 0.082 mm	Used μ CT data to find cortical wall thickness. Surface meshes defined through local triangulation for each cell for the surface and interior cortical wall boundary. Trabecular bone modelled with hexahedrons of size 1.312 x 1.312 x 1.312 mm using GMSH software	Elastic and yield properties assigned to each trabecular element from μ CT density data

Eswaren et al. (2007) [63]	Used μ CT data with resolution of 0.03 x 0.03 x 0.03 mm	Conversion of voxels to hexahedral elements of size 0.06 x 0.06 x 0.06 mm. 8 node hexahedral elements using in-house software. Endplates and cortical shell modelled with higher definition using thickness from μ CT Bone within 180 μ m of the outer structure was identified as belonging to the cortical shell.	Used a constant material property for entire model
Higgins et al. (2007) [52]	Taken from in vitro measurements of L1 vertebrae, using flat endplates and curved side walls	Composed of 12 node brick elements using in-house code. Used separate mesh for cortical walls and different thicknesses for the endplates and posterior and anterior cortical shell.	Anisotropic through model, using different material properties for cortical shell, endplates and vertebral body.
Kinzl et al. (2012) [64]	Used μ CT with resolution of 0.082 x 0.082 x 0.082 mm	Used μ CT data to find cortical wall thickness. Surface meshes defined through local triangulation for each cell for the outer and interior cortical wall boundary, meshed using pentahedrals. Trabecular bone modelled with tetrahedrals	Material Mapping based on porosity (density). Uses parameters for elastic, plastic and damage behaviour. Uses elasto-plastic material for the cement augmented region
Wijayathunga et al. (2008) [42]	Taken from μ CT data with resolution 0.074 x 0.074 x 0.074 mm	Hexahedral and tetrahedral elements using ScanFE software, smoothing is applied to vertebral surface to improve geometry	Material properties assigned using density values from μ CT. No division between cortical and cancellous bone

The cortical surfaces of these models are often rough when just using voxel to element meshing, for example the models created by Buckley et al. [61] in Figure 1.6. However, many studies introduce smoothing to the surface of a vertebra model. The error introduced from possible stress raisers when omitting smoothing from the model was questioned by Chevalier et al. [62]. This study suggested that smoothing, in the form of surface meshes representing the cortical walls, successfully removed the stress raisers and artificial damage zones from their models. The studies by the group that produced the Chevalier et al. and Kinzl et al. papers in Table 1.3 use methods for fitting the vertebrae surface of the model to that of the scanned specimen, rather than simply smoothing the voxel to element mesh. Kinzl et al. [64] used a “marching tetrahedral” method to extract the surface of the scanned vertebrae from Treece et al. [65], which uses an adaptation of the popular marching cubes method to obtain an iso-surface from discretised three-dimensional data (from μ CT data). Other more recent studies use entirely tetrahedral meshes [66], which although are often less accurate at portraying the propagation of stress, give uniform element types and smoothed model surfaces.




Figure 1.6: Voxel finite element mesh of L2 vertebrae used by Buckley et al. with voxel size of 2 mm [61].

There is a reduced difference between the structure of the cortical shell and the cancellous bone structure of the vertebra when compared to other human bones [67]. Despite this, the cortical walls have been shown to share a significant quantity of the load, especially in the axial midsection, where the cross section is narrowest and the load carried by the cortical wall is greatest [67]. An understanding and correct modelling of the load sharing between the cortical wall and trabecular bone becomes more important when considering the variation of the cortical wall with age, location, level and osteoporotic nature of the vertebrae in question [68]. Chevalier et al. [62] found that adding a surface faced cortical shell to their models unloaded the centre trabecular bone and increased the strength and stiffness values of their models to that of their experimental results. Approaches that rely on the μ CT greyscale values to assign material properties to elements, give additional stiffness to the cortical walls due to the denser and therefore brighter cortical regions in

the μ CT images, without the need for modelling the regions separately.

The increased resolution of the Eswaran et al. [63] study allowed for a uniform element size and material assignment for the cortical shell and trabecular bone, relying on the resolution to represent the higher density of bone within this region. In this case, separate meshes were used for the two bone types (cortical and trabecular) to extract information about load sharing between them. However, the required use of a supercomputer to analyse the FE models limits such models being used within research, especially where large datasets are required to investigate population variation.

Convergence studies on FE mesh size are reliant on a range of factors that include the inhomogeneous material properties and complex geometry of bones. Jones and Wilcox [69] showed that the models that contained inhomogeneous properties based on an underlying greyscale background presented inconsistent convergence behaviour. This inconsistent behaviour was highlighted again by Zhao [70] showing that as resolutions increased elements began to represent either trabecular bone or trabecular space, rather than the combination or average of the two. Because of these reasons, most continuum level studies used element sizes of between $1 \times 1 \times 1 \text{ mm}^3$ and $2 \times 2 \times 2 \text{ mm}^3$.

1.5.2 Material Properties

Material properties for continuum level FE models are in general derived from μ CT greyscale data for each voxel and assigned to the matching element. These greyscale values are used to derive elastic modulus values using a conversion equation. This conversion equation is often optimised such that the error between experimental and computational results are minimised. Robson Brown et al. [60] performed an investigation into the effect of a linear or non-linear relationship between the elastic modulus and BV/TV . However, due to lower order errors in the form of the tissue modulus and degree of anisotropy, no benefit was found when using the higher order relationship. A good agreement was found with the linear relationship in the Robson Brown et al. study when using a similar method in the study by Wijayathunga et al.[42].

Chevalier et al. [62] and Kinzl et al. [64] used other methods of representing the material properties. In these studies an enhanced continuum FE model was used where, rather than using the more common method of converting μ CT voxels directly into hexahedral elements, they used high-resolution CT images, which included the cortex and used fabric-elasticity relationships to describe bone morphology. This deviation away from the more common methods used by others [42, 60], potentially allows μ CT at lower resolutions to be used, which both reduces computational cost and introduces possibilities for the use of clinically

relevant resolutions. The use of lower resolutions was achieved by using $82\ \mu\text{m}$ scans and coarsening them to $1.312\ \text{mm}$ to mimic common clinical CT scans. With a focus on identifying modelling errors, a study by Pahr and Zysset [71] compared this enhanced continuum FE models with standard methods. This study looked at both the effect of smooth cortex modelling (also used by Chevalier and Kinzl) and the morphology-elasticity relationship. Pahr and Zysset [71] concluded that this method of modelling provided statistically equivalent results for the stiffness of two different models when compared to the more common method, whilst analysing the model at least 100 times faster. This reduction in time to analyse the models was due to the reduced sensitivity to mesh size, meaning mesh density and therefore number of nodes / elements could be reduced.

1.5.3 Validation of Models

Validation of the models through comparisons of the computationally derived results to the experimentally obtained results is important to give clinical relevance to results. Most studies provide validation of non-augmented models, often using two sets of models, one to derive the material properties and one to test or validate them against the experimental results [42].

Of the studies compared in Table 1.3, only the studies by Buckley et al. [61], Wijayathunga et al. [42] and Robson Brown et al. [60], performed validation against experimental results. Wijayathunga et al. [42] and Robson Brown et al. [60] used similar methods of validation using two sets for calibration and validation. These studies validated both strength and stiffness, while Buckley et al. [61] validated their models by comparing strength alone using an non-described method. All three studies that performed validation found excellent agreement between the model and experiment using their measured metric.

1.5.4 Modelling Vertebroplasty

The methods used to model vertebroplasty from a selection of studies are summarised in Table 1.4. The variation in geometry generation can be seen to vary similarly to studies modelling vertebrae alone, however, in addition to this, the generation of cement geometry and material properties is also presented due to the range of methods used.

Methods of modelling vertebroplasty usually involve using μCT scans of augmented vertebrae, masking the internal volume of cement though thresholding the greyscale background. However, the methods used in the studies by Liebschner et al., Polikeit et al. and Baroud et al. [72–74] used approximations of the cement distribution, with the former modelling the distribution according to images from *in vitro* experiments and clinical




Figure 1.7: A: Top view μ CT scans of human vertebrae augmented *in vitro*, showing gradual reduction of cement opacity to the edges of the internal cement volume, adapted from Belkoff et al. [24]. B: The augmented model generation for the study carried out by Liebschner et al. [72].

trial CT scans, while the other studies modelled the cement more approximately. Such approximations of the internal volume of cement can be seen in Fig. 1.7:B. The μ CT scans in Fig. 1.7:A show the much more loosely defined boundaries of the more random internal cement regions, which are created by the trabecular bone that the cement is injected into. The voxels within the defined cement region are often given linear orthotropic elastic properties described by a rule of mixture [75], or more basic material properties with constant elastic modulus and Poisson ratio [42, 72–74].

Table 1.4: Method of geometry generation, cement position, location and materials used for five finite element studies of vertebroplasty.

Author	Geometry	Cement position & shape	Cement material
Baroud et al. [73]	Hand modelled, described in Smit et al. 1997, two level model	Used 70% fill of cement	Used previously obtained values for cement infiltrated bone, 46 times stronger, 12 times stiffer than surrounding trabecular bone.
Chevalier et al. [75]	μ CT scanned pre and post augmentation, model of vertebroplasty combined cement region with non-augmented scan.	Cement position and structure taken directly from μ CT scans	Bone-cement mixture described by a rule of mixture using a tensor for the isotopic stiffness and for bone stiffness.
Liebschner et al. [72]	Taken from μ CT data	Cement capsule design, cylinder with rounded edges. Positioned to investigate uni/bipedicular vertebroplasty and centred cement positioning.	Unspecified, PMMA
Polikeit et al. [74]	Taken from CT scans at 1 mm resolution, manually constructed details not visible. Two level model	Cement modelled as barrels using radiographs as guides. Positioned to investigate uni/bipedicular vertebroplasty including one model with 100 % fill	Constant Young's modulus (3000 MPa) & Poisson ratio (0.41)
Wijayathunga et al. [42]	μ CT scanned post augmented vertebrae	Identified from μ CT scan using constant threshold value based on greyscale	Used constant properties for Young's modulus and yield stress. Examined effect of lowering cement modulus to align with that of cement impregnated bone

1.5.4.1 Validation of Augmentation

The studies by Chevalier et al. [75] and Wijayathunga et al. [42] were the only ones to use experimental cement distributions in their FE models. The other studies listed in Table 1.4 rely on previously validated models of non-augmented vertebrae, with approximately shaped and placed cement. While such studies can report changes in the overall stiffness, it is difficult to identify how the load is transferred through the internal cement and whether it has been modelled with the correct material properties and boundary conditions for the cement-bone interface. Wijayathunga et al. [42] found large errors between experimental and computational augmented vertebral behaviour, concluding that the representation of the internal cement region as a homogeneous material was possibly inadequate. In light of these results the lack of a comparison to in vitro results in the Chevalier et al. [75] study (in order to study prophylactic vertebroplasty) generates questions of the accuracy of cement modelling, in both material properties and boundary conditions.

Validation is rarely carried out for vertebral models having undergone vertebroplasty with one of the few examples being the study by Wijayathunga et al. [42] described above.

1.5.4.2 Bone-Cement Interface

Studies have shown that simple approaches to modelling augmentation do not accurately describe the results seen experimentally [42]. Hence, modelling the bone-cement interface is an important step in portraying the experimental results.

Zhao et al. and Tozzi et al. [76, 77], used a μ FE method to investigate the bone-cement interface. They assigned homogenous and orthotropic elastic-plastic properties respectively along with examining the effect of friction between the two surfaces (coefficient of friction 0.3 and 0.4 respectively). Both studies used stepwise compression testing of their trabecular bone-cement samples (open cell rigid polyurethane foam, similar to osteoporotic human trabecular bone and bovine trabecular samples from the iliac crest, Zhao et al. and Tozzi et al. respectively) within a μ CT scanner to assess the evolution of stress in a stepwise fashion. These studies found good agreement with computational and experimental data, with Zhao et al. showing that a composite of bone and cement was considerably lower than what would be expected by cement alone or predicted by a rule of mixtures, such as that used in other studies described above [75]. Tozzi et al. found that greater cement penetration or contact area did not increase the compressive strength contradicting the results of Janssen et al. [78] in a FE study of cemented total hip arthroplasty. This study examined friction coefficients and morphology under both compression and tension, which along with the results of Tozzi et al. potentially suggest that cement penetration is more

important for failure under tension rather than compressive loading. This advocates the idea of buckling trabeculae at the boundary of the cement region and a lack of load transfer into the cement-bone interdigitated region.

In a similar study to those by Zhao et al. and Tozzi et al. [76], [77], Kinzl et al. [79] examined the effect of PMMA shrinkage due to polymerisation and different interface properties. The shrinkage of PMMA upon polymerisation affects the bone-cement interface due to gaps developing between the two materials caused by the volume change of the PMMA. Kinzl et al. showed that the shrinkage affected their models in two ways, firstly, the loss of volume and creation of gaps reduced the load transmission between the materials and secondly it created residual stresses in the PMMA, causing bone damage at trabecular connections from compressive and shear loading.

Sikora [80] used FE models generated from μ CT images of ovine lumbar vertebra to investigate the bone-cement interface. An analytical model of the behaviour of trabecular bone struts embedded in cement was used to impose properties on an interface region defined in the model. The analytical model was used to predict the interdigitation between the two materials and forecast the characteristics of failure between them, it was then used to determine the plastic and elastic properties to apply to the defined interface mesh. The use of the interface layer with explicitly defined properties produced a good agreement with the experimentally determined results for the sample under compression. However, the study was carried out using cylindrical specimens of height 25 mm and diameter 13 mm, hence a further investigation would be required to identify whether this method would prove useful for modelling whole vertebrae.

1.5.5 Discussion of FE studies

Finite element studies modelling vertebrae are currently producing accurate models for single intact vertebrae as described in Section 1.5.3. The methods employed by Chevalier et al. [62] to reduce the need for high resolution μ CT scans and instead use more clinically relevant resolution, present a valuable tool for clinicians. Problems arise however when attentions are turned to modelling vertebroplasty, specifically, modelling how the two materials interact while under compression.

Despite the well validated results for the three μ FE interface studies [76], [77], [79], which examined more accurate methods of modelling the boundary between materials, the computational cost of running such simulations on whole vertebrae restricts further research. For example, the simulations reported by Zhao et al [76] required 700 hours on a high performance computer for a trabecular bone sample. Hence, methods employed by

Sikora [80] require further testing and validation for whole vertebrae in order to strike a balance between agreement with experimental data and computational cost.

1.6 Capture of Population Variation

There is an increasing need for patient variation to be taken into account in pre-clinical testing; however, experimentally there can be difficulties in controlling this variation. Obtaining large quantities of varied tissue can be problematic and time consuming, and there are additional challenges in characterising the variation. Hence, FE studies present a number of potential benefits in terms of controlling and evaluating variables related to the tissue or the treatment. Currently however, the main advantage of FE studies is the ability to run multiple scenarios for the same model (for example using different material properties for implants or different quantities of cement in vertebroplasty studies); this removes specimen variability from the study. In order to use FE studies to examine patient variability across the population then there is a need for large quantities of patient specific models, similarly to experimental work.

Variability in terms of shape, size and density of vertebrae is high between individuals [81], with even greater variability possible for those with various pathologies. The large variation in outcomes of vertebroplasty reported by Kallmes et al. [4] may in some way be linked to the variation in the vertebrae themselves. Hence, large sets of models that cover the whole range of variation across the population and then categorised based on vertebroplasty effectiveness, could be used to examine if there are relationships between vertebral variance and (mechanical) outcomes of the treatment.

As described in Section 1.5, FE models can be generated from image data. This gives the opportunity to build large sets of models describing the population. However, these large datasets are difficult to gain access to, for ethical reasons among others [82] and model construction from such imaging is time-consuming. Instead most studies use vertebra-specific models from a limited selection of experimental specimens [42, 61, 64, 75].

Parameterised finite element models are often used to understand the effect of variables on a system, however, understanding which variables are important and how to set the limits of these variables becomes difficult with the increasing complexity of the system. Statistical approaches that are able to determine the main modes of variation and therefore reduce the number of variation types to those that have the greatest impact can aid such parameterised approaches.

Statistical shape modelling (SSM) and statistical shape and appearance modelling (SSAM)

approaches describe the geometric variation and the material property variation across an input set of models that represent the population [83]. Such models usually describe a mean shape and the main modes of variation found within the population. These approaches and their relevant applications are examined in the following subsections.

1.6.1 Statistical Shape and Appearance Models

The purpose of Statistical Shape and Appearance Models (SSAM) is to capture and describe the variation within a given dataset. It allows an understanding of how different variables interact and what variables control the greatest variation. Here, there is a focus on using Principal Component Analysis (PCA) to perform dimensionality reduction for shape and material property analysis, but the algorithm can be used for many other types of data.

1.6.1.1 Principal Component Analysis

Principal component analysis is a method of identifying patterns in data and expressing the data in such a way as to highlight the similarities and differences. Since patterns in data are difficult to visualise in more than 2 or 3 dimensions, PCA represents a powerful tool for data analysis.

The first step of PCA is subtracting the mean from each of the data dimensions. The mean subtracted is the mean across each of the dimensions, so for a two dimensional dataset, the mean of the x values is subtracted from each of the x values and the same with the y values. This gives a dataset with a mean of zero.

Following this, the covariance matrix is calculated. The covariance is a measure of how much the dimensions differ from the mean with respect to the other dimensions. Hence, the covariance is always a measure between two dimensions. For a data set with more than two dimensions, for example (x, y, z) , then the covariance would be measured between x and y , x and z , and y and z . The covariance is given by:

$$\text{cov}(X, Y) = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{(n - 1)}$$

While the values of the covariance are of less importance, the sign tells us the nature of the relationship (a positive or negative correlation). The covariance matrix presents all of the possible covariance values for each dimension (n) in a matrix with n columns and n rows. For example, a three dimensional dataset (x, y, z) would have a covariance matrix of the form:

$$C = \begin{pmatrix} \text{cov}(x, x) & \text{cov}(x, y) & \text{cov}(x, z) \\ \text{cov}(y, x) & \text{cov}(y, y) & \text{cov}(y, z) \\ \text{cov}(z, x) & \text{cov}(z, y) & \text{cov}(z, z) \end{pmatrix}$$

Next, the eigenvectors and eigenvalues of the covariance matrix are found. The calculation of these for a given matrix is usually carried out using an iterative method for matrices of 3×3 and greater ensuring that the vectors produced are unit eigenvectors. For an n dimensional matrix there are n eigenvalue and eigenvector pairs. These eigenvectors of a matrix are orthogonal to each other regardless of the number of dimensions. So, by taking the eigenvectors of the covariance matrix, lines (or vectors) that describe relationships within the data can be found. Identification of size of the vector, the eigenvalue, tells us the scale of each of those relationships. The eigenvector with the largest eigenvalue becomes the principal component of the dataset.

The next step is often to order the eigenvectors by their eigenvalue, giving the components in order of significance. Dimensional reduction usually occurs next, removing those components (eigenvectors) that have the least significance. This data can be described by using the eigenvectors as the axis, giving an understanding of how each data point contributes to the mode of variation. Alternatively or additionally, the dimensionally reduced data can then be reconstructed onto the original axis, providing a more clear understanding of the modes of variation within the dataset without the noise created by the less significant modes of variation.

1.6.1.2 Application of Principal Component Analysis to Biomechanical Problems

With regard to the biomedical field, PCA has been applied to sets of data containing similar shapes, such as sets of human femurs and knee joints. In these examples it allows an understanding of how the bones vary across the population contained in the given data-set.

Previous studies took a series of measurements, for example, dimensions of bones using specific anatomical features along their length [84]. These measures became the input for PCA, from which relationships and the main modes of variation were found. More recent studies carried out PCA using surface/solid meshes as the inputs, these later studies are described below.

Often the first step for PCA is to normalise scaling, rotational and translation differences

between the specimen meshes within the dataset. This is carried out via generalised Procrustes analysis (GPA) and is often referred to as a preprocessing step for PCA, it is described in detail by Grassi et al. [82] and Väänänen et al. [85]. Once normalised, matrices are created from columns containing specimens and rows containing the nodal coordinates for each mesh. Using this matrix and the mean coordinates from it, the PCA algorithm can be applied, from which the principal components can be acquired. These principal components describe each mode of variation, ordered in terms of the largest variance. To determine how much variation is contained within each principal component, the eigenvalue for each component can be divided by the sum of all eigenvalues. The first three modes of shape variation for a femur (used here as a convenient example) can be seen in Figure 1.8, with the minimum and maximum extremes from each mode of variation seen in the wireframe and shaded model respectively [82].

Similar methods can be used for the material properties of bone and how it varies between specimens when using solid mesh or finite element model based inputs. Models can be generated with brightness values assigned to each element of the mesh (i.e. a finite element model), a covariance matrix can be assembled based on the greyscale value of each element in the mesh. From here a similar method to above can be applied to acquire the principal components and the amount of variation contained within each of the components.

The understanding of the variation of a data-set and its description in terms of a potentially reduced number of components, allows the generation of ‘virtual specimens’ that fit within the ranges of the input data set. For example principal components could be altered in terms of their standard deviation about the mean specimen (based on the input data). This would allow the generation of models where the geometry and material property distribution fit within the bounds of variation of the current data, but do not currently exist due to possible gaps in the data set. It also means that the variation can be quantified, allowing effects caused by variation to be studied in an iterative fashion. For example, in the case of vertebrae, if changes to the diameter of the spinal canal correlate with process length (a correlation that would be difficult to identify without PCA), then the effect this has on specimen strength could be examined in depth by altering the component that controls this relationship. Models that are described by this relationship could be generated and solved through FEA, furthering our understanding of the relationship. However, generating FE models from the statistical models can produce some problems regarding reliability. For example, the generated models heavily rely on the input database, hence the database is required to accurately represent the entire population. There can also be issues with the generating algorithm’s ability to produce shapes and materials not present

in the starting database [82], [85].

The Grassi et al. study [82] examined whether PCA based modelling could be used for representing the femur and concluded that such generated models were able to describe the population of femurs using generated FE models. They found that 50 modes or the first 50 principal components were required to accurately reconstruct the femur while maintaining errors below that originating from pixel size and another 40 to generate accurate density in the models.




Figure 1.8: Femoral shape variations for the first three modes from the PCA. From 115 bones the maximum eigenvalue is shown with the minimum eigenvalue shown in wireframe. Taken from [82].

In a study by Fitzpatrick et al. [86], statistical shape models were developed in order to identify the relationship between the shape and function of the patellofemoral joint. Here, 26 magnetic resonance scans of the knee were taken, from which FE models of the joint in question were created. PCA was applied to the set of models, from which the first 15 principal components describing the shape were taken, these 15 components accounted for 97.2% of the total geometric variation, with the first component (PC1) counting for 47.7%. PC1, as is often the case in biological models, described the variation in the size of the joint components, with the next two components describing the position of the patella and details regarding the conformity and depth of the joint. The following modes of variation described more subtle variations in the geometry of the model.

In order to assess the robustness or ability to accurately predict outcomes of the reconstructed bones or joints, the studies [82, 86] used a leave-one-out approach. This approach neglects one specimen from the development of the shape-function models. The final model, based on the included specimens can be used to predict the mechanical behaviour or properties of the left out specimen by using the generated FE models.

There are also a range of studies that identify the variation found in the spine. These include a characterisation of the shape variation in the cervical spine [87] and studies that aim to identify the variation in the lumbar spine for both spinal curvature variation [88–90] and to identify the variation in single vertebrae [81]. The study by Hollenbeck et al.[81] investigated the effect of single vertebra variation on the curvature of both the whole lumbar section and on specific functional units. The study used a relatively large cohort of 52 lumbar spines, finding that the main mode of variation for single vertebrae was scale variation. Subsequent modes of variation were more intricate geometric variations such as changes to the pedicle angles. Limitations of the study surround the isolation of the statistical shape models to each spinal level and hence failed to describe the gradual level by level variations. Additionally, the model did not include any description of the underlying material properties, which may also share relationships with the geometric modes of variation and hence limit the model’s clinical use.

1.6.2 Discussion

The studies described above present useful workflows of generating virtual subjects from a statistical model acquired from a databases of scans. Such methods allow a variety of population based studies, including those regarding the variation of vertebrae across the population.

The ability to generate models not present in the input dataset and hence create a large set of models representing the population allows us to identify relationships not openly visible. For example in the study by Fitzpatrick et al. [86] a 5 mm change in the patella position caused a 25 % increase in the contact pressure mid-flexion. This relatively easy quantification of the relationship between geometry and function could be adopted to quantify the relationship between vertebral geometry and the mechanical response to vertebroplasty. Equally, such relationships could enable the prediction of vertebral fracture, adjacent level fracture or advise clinicians as to the quantity and location of cement for different patients.

1.7 Conclusion

Vertebroplasty potentially provides a valuable method to treat patients with osteoporotic compression fractures, however, the uncertainty regarding its effects, especially regarding certain subsets of patients, mean that further research is required. While the mechanisms of pain relief are not fully understood, certain assumptions regarding the mechanical stabilisation and restoration of stiffness can be investigated. Such investigations may help to understand how patient groups respond to the treatment in different ways.

Current experimental studies have provided a range of useful techniques for mechanically testing vertebrae and for carrying out the vertebroplasty procedure itself. However, obtaining enough vertebrae to represent variation across a population in order understand its effects on different patients is a difficult task and has not been attempted. A possible solution is FE modelling in conjunction with PCA. FE models have been shown to represent the intact vertebrae accurately using both μ FE and continuum level models.

Problems arise when modelling the augmented vertebrae, with few studies modelling specimen specific augmented vertebrae and instead modelling vertebrae with arbitrary volumes of cement. Studies that have modelled and validated against experimental results have shown poor agreement, owing to the incorrect modelling of the cement-bone interface or incorrect selection of material properties for the cement-bone interdigitated region.

Finally, principal component analysis has shown its value in the literature presented above allowing the variation in a set of specimens to be accurately described and in certain cases models have been spawned at standard variations away from the mean.

1.8 Aims & Objectives

The main aim for the project was to examine whether the variation in outcomes for patients undergoing vertebroplasty is due to underlying biomechanical differences in their vertebrae. To achieve this overall aim, the following objectives were defined:

1. To develop methodologies for testing individual vertebrae in the laboratory pre – and post- vertebroplasty using animal tissue. This study is reported in Chapter 2.
2. To develop methods of generating specimen-specific FE models of the experimental specimens and validate their outcomes pre- and post-vertebroplasty. This study is also reported in Chapter 2.
3. To apply the methods developed in 1. and 2. to test and model human vertebrae. This study is reported in Chapters 3 and 4.

4. To build a statistical shape and appearance model using the models developed in 3. as the inputs. This study is reported in Chapter 5.
5. To understand and characterise the variation found within the statistical shape and appearance model. This study is also reported in Chapter 5.
6. To understand how the variation found within the statistical shape and appearance model affects the mechanical outcomes of augmentation. This study is also reported in Chapter 5.

Chapter 2

Bovine Tail Vertebrae Study

2.1 Introduction

This chapter outlines the methods developed and results obtained for testing and modelling a set of bovine tail vertebrae. It was split into two main sections, with the first discussing the development of experimental methods testing bovine tail vertebrae and the second (section) describing the methods of computationally modelling these vertebrae using FEA. These main sections are split further into method development, sensitivity tests and results. In this chapter the experimental and computational work was limited to bovine tail vertebrae due to their plentiful nature and relatively similar geometry to human vertebrae, while not having many of the problems with increased yield strength and density of porcine or other animal tissue [91]. This development work allowed translation of the same or similar methods to human lumbar vertebrae, that are reported in the following chapter.

2.2 Experimental Methods

2.2.1 Introduction

The experimental methods that were developed and the early results that were acquired in this section allowed easier transition to using human tissue and provided valuable processes and approaches for the development of specimen specific finite element models.

The steps in the developed methods involved dissection of the soft tissue from the vertebrae, potting in PMMA end-caps, scanning using a μ CT scanner, compression testing and augmentation, the order of which can be seen in Figure 2.1. Specimen preparation, fracture generation and initial μ CT scanning were undertaken jointly with Ruth Coe (PhD student, University of Leeds). Vertebroplasty training was carried out with Dr Peter Loughenbury

and Dr Vishal Borse from the Leeds General Infirmary and an initial set of test vertebrae were augmented with Sebastien Sikora and Fernando Cornelio Zapata (Research Fellows, University of Leeds). Testing, μ CT scanning and augmentation of the vertebrae presented in the results section were carried out solely by the author.




Figure 2.1: Flow-chart detailing the experimental process from initial dissection to final load test.

2.2.2 Specimen Preparation

Bovine tails were acquired from a local abattoir and frozen to -20°C prior to use. They were defrosted in a 4°C fridge for approximately 24 hours before the initial dissection. The three most caudal vertebral (CC1 to CC3) were kept, discarding the remainder of the tail due to the elongation of the vertebral bodies in more distal vertebrae. In addition to the elongation of the vertebral bodies, the spinal canal narrows in the distal region, limiting its ability to house a steel rod used for mounting the specimens. Soft tissue was removed

from the vertebrae as thoroughly as possible, including the intervertebral disc material and material occupying the spinal canal. This was carried out in order to remove potential error when comparing experimental results of stiffness to the vertebra models developed from μ CT scans (in which the soft tissues are not modelled) and to allow insertion of a metal rod through the spinal canal to aid alignment.




Figure 2.2: Photograph and diagram depicting the method of creating end-caps for the specimens. Red arrows indicate the location of the removed pedicles.

Once dissected, and in subsequent breaks between procedure steps, the vertebrae were wrapped in phosphate buffered solution (PBS) soaked tissue, to limit the drying of the vertebral bone. The specimens were potted in PMMA end-caps to allow repeated loading of the vertebrae with the same orientation and positioning, while constraining the vertebrae as little as possible and allowing flexion of the upper endplate. Such flexion occurs naturally in the human spine, hence allowing it to occur experimentally was important. The setup for potting the vertebrae can be seen in Figure 2.2. Vertebrae were held using a rod placed through the spinal canal. Depending on the level of the vertebra, the spinal canal was packed with foam around the rod forming a snug fit while the vertebra was held approximately 5 mm above a petroleum jelly lubricated metal surface. Also, depending on the level of the vertebra, any pedicles that protruded past the limits of the metal cylinders were removed with a hacksaw at their base to prevent issues with the loading and scanning tests which followed. Most often this was limited to the most caudal vertebra, shown by the red arrows in Figure 2.2. Lubricated hollow metal cylinders of ~ 10 cm diameter were used to form the endcaps for the vertebrae. A 2:1 powder to liquid component, by weight, PMMA mixture (Cold Cure, WHW Plastics Ltd., Hull, UK) was added until the endplate

of the vertebral body was covered up to the point where the body starts to become concave. After approximately 20 minutes the PMMA had sufficiently set to turn over the vertebra and create the end-cap at the other end using the same process with the addition of a level to ensure the creation of parallel end-caps.

Once the PMMA was set, the vertebrae were wrapped in more PBS soaked tissue before being frozen or stored in a fridge until they were loaded to fracture. The specimens were frozen only if more than 24 hours would pass before the next stage of testing to reduce the number of freeze thaw cycles.

2.2.3 Axial Compression

2.2.3.1 Fracture Creation

All specimens underwent axial compression using a material testing machine in order to generate fractures within the vertebral body. Mounted vertebrae were placed between two steel end-plates, each of which contained four screws to inhibit lateral motion of the specimen when under load. The upper plate also contained a chamfered hole to allow the alignment of the specimen using the marker located above the centre of the vertebral body. This meant that the loading point was directly below the head of the testing machine. A steel ball was used to allow flexion of the upper endplate; this was fitted between the chamfered hole and the fixture on the head of the materials testing machine. The steel ball acted as the centre of rotation about which the upper end-cap could rotate. This permitted rotation mimicked natural loading of the vertebra and increases the likelihood of physiological anterior wedge fractures. Details of the setup can be seen in Figure 2.3.

Loading of the vertebrae started with a preconditioning from 50 N to 300 N for 10 cycles at a rate of 1 mm/minute to remove any viscoelastic effects in both the bone and any remaining soft tissue. Following the preload, displacement was increased by 1 mm/minute until either the load reached 9500 N (a safety limit due to the 10 kN load cell limit) or a visible failure occurred on the real-time load-displacement plot during compression. This failure was observed as a peak in load with the compression being stopped once a clear decrease in force was observed. Examples of both scenarios can be seen in Figure 2.4.

An additional non-fractured control specimen was loaded to only 5000 N during the initial step to gain insight into the effects of fracture on the subsequent process.




Figure 2.3: The experimental setup for axial loading the vertebral specimens. Shown is the screws, securing the PMMA endcaps to the steel plate, the chamfered hole where the steel ball is seated. Space at the other end of the chamfered hole is reserved for the delrin marker, allowing accurate and repeated central loading of vertebrae.




Figure 2.4: The difference between failure (A) and non-failure (B) for bovine tail vertebra compressed to a maximum load of 9500 N or until a peak was observed.

2.2.3.2 Post Fracture and Post Augmentation Stiffness Calculation

In order to find the stiffness of the previously fractured and augmented specimens a similar loading procedure was used. However, following the preload, compression was stopped when the load reached 5000 N as a means to limit additional damage and fractures to the vertebrae. This ensured that the vertebral stiffness across the three stages (intact, post-fracture and post-augmentation) was calculated from the same range of loads (0 - 5000 N).

The stiffness of the specimens throughout their tests was calculated using a Python script, written by the author, on the raw data from the materials testing machine. The script allowed the limits of the range of interest to be set. Using a defined segment size it then calculated the gradient of the segment using a least squares regression. The gradient

calculation was repeated at increments of 0.1 mm over a specified data range and the greatest gradient was defined as the specimen stiffness. The segment size, from which the gradient was calculated was set to 0.3 mm, overlapping increments of 0.1 mm were used and the range of interest was set to 0 - 5000 N. An illustration of this can be seen in Figure 2.5.

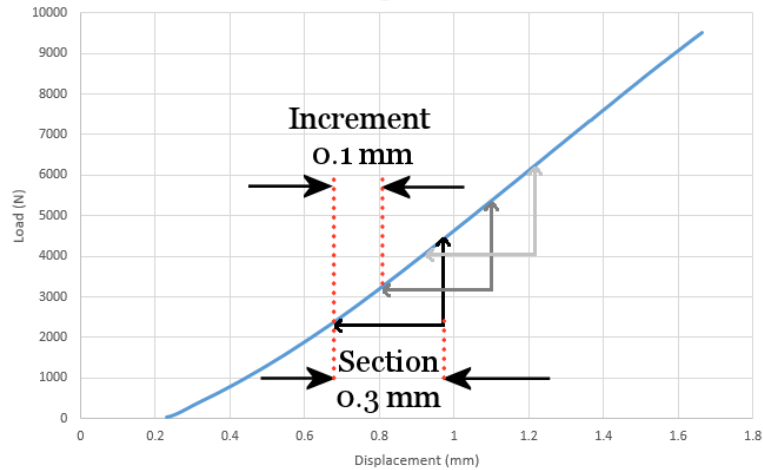


Figure 2.5: A typical load displacement curve showing how the gradient was taken from 0.3 mm long sections incremented at 0.1 mm across the length of the curve.

To examine the effect of the choice of the range of interest, the gradients for six specimens were examined over three different ranges. The results are presented in Figure 2.6. If the load-displacement curves were perfectly linear within the “linear region” the stiffness in the three ranges of interest in Figure 2.6 would give an equal value for the stiffness or maximum gradient. However, given that these values were found to differ for the three ranges, it suggested non-linear behaviour. As shown in Figure 2.6 the recorded maximum stiffness varies greatly depending on what portion of the load displacement graph was being examined. The maximum gradient and therefore stiffness was found to only be measurable with the full 0-9500 N range. However, subsequent damage on the post-fracture tests needed to be avoided. This ruled out testing the fractured vertebrae in the full range, hence 0 to 5000 N was chosen as a compromise. The stiffness measurement was carried out in the 0 to 5000 N range for all load tests, ensuring a uniform stiffness measure between specimens and test stages.

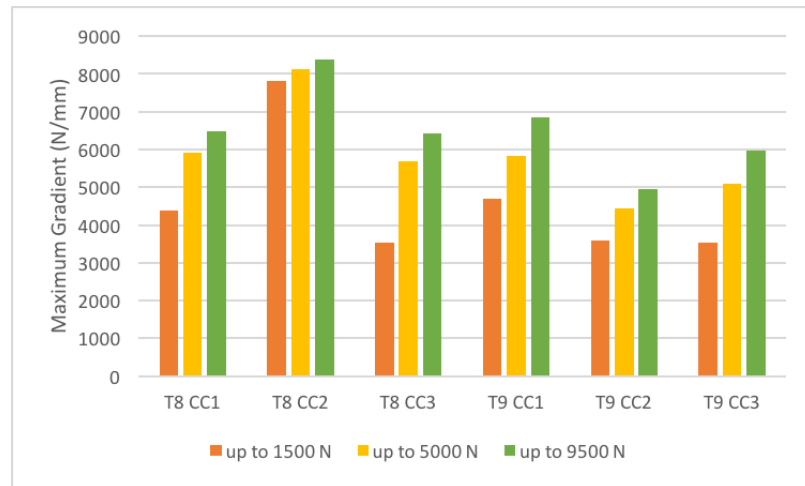


Figure 2.6: The difference seen when measuring the greatest gradient (stiffness) using different portions of the load displacement curve. From 0 to 1500 N, 0 to 5000N and 0 to 9500N.

2.2.4 Vertebroplasty

Due to the differences between human and bovine vertebrae it was not possible to perform bi-pedicular vertebroplasty on the bovine specimens using the methodologies established for human vertebrae. The main difficulty was the greatly increased density of the bovine vertebral bone, meaning that rather than pushing the vertebroplasty needle into the vertebra by hand, a mallet and vice were required. In addition to this, the force required to inject cement into the vertebral body was greatly increased. The vertebroplasty method for bovine tail vertebrae was therefore developed over several iterations due to these difficulties. This sub-section details the initial procedure, the problems encountered and solutions developed to allow a clinically relevant volume of cement to be injected and captured in μ CT scans.

2.2.4.1 Initial Procedure

The procedure was initiated by using bone nibblers to remove the rounded end of both posterior pedicles, providing a surface to start the needle entry. The needle used was a 9 gauge vertebroplasty needle (Figure 2.7), consisting of two parts. The pointed insert, used to create the needle track, could be removed once the needle was at the correct depth creating a hollow tube. A syringe could be attached to the threaded top of the needle, allowing injection of the cement. While holding the vertebra in a table-mounted vice the needle's 1 cm markings were used to estimate the depth and angle needed to reach the anterior quarter of the vertebral body. Care was taken to ensure the pedicle was not damaged as the needle was inserted. A mallet was used to insert the needle until it was at the depth required; the procedure was repeated for the other pedicle, reusing the same

needle.




Figure 2.7: The 9 gauge vertebroplasty needles used for augmentation, showing the hollow needle and the pointed insert.

The PMMA cement was mixed 1:1 monomer to powder, by weight, to ensure that it could be drawn up via the syringe and to allow enough time to inject the cement before it thickened and set. This additional setting time and reduced viscosity is also used by clinicians, who use ratios up to 1:0.74 monomer to powder with no adverse outcomes associated despite the reduced modulus and strength often reported [92, 93]. While the vertebra was held in the clamp of a retort stand, the syringe was attached to the needle, which in turn was inserted into one of the pre-made tracks through the pedicle into the vertebral body. Cement was pushed into the vertebrae using the syringe, until 3-4 mL was inserted into both sides of the vertebrae, with cement being inserted as the needle was removed to back fill the channel created by the needle. The vertebrae were then left for approximately an hour until the cement had set before scanning.

2.2.4.2 Complications and Changes to the Procedure

Various problems were encountered while carrying out the procedure that required the methods to be adapted. These challenges and their solutions are described below.

Vertebral Temperature: With the initial specimens, there was a difficulty in injecting the cement and it was mainly limited to the needle tracks rather than the vertebral body. To counter this, the vertebrae were warmed to 37°C for an hour or until the internal region of the vertebrae had reached this temperature (using a temperature probe in the vertebroplasty needle hole). This meant that the bone marrow inside the vertebra was no longer solid and therefore could be displaced by the cement, making the injection much easier.

Radio-opacity of Cement: A second problem was the opacity of the cement on μ CT scans, which made it difficult to segment the cement region in the vertebral body as can be seen in Figure 2.8:A. Here, the cement was indistinguishable from the bone marrow and can only be seen in the needle channel. The solution to this was to mix barium sulphate (BaSO_4) with the PMMA to achieve the radio opacity seen in Figure 2.8:B, where the bright area in the centre of the vertebral body is the injected cement and BaSO_4 combined. Radio opacifiers are often used clinically to observe the cement distribution through fluoroscopy during the augmentation procedure. Due to the hydrophilic nature of the BaSO_4 powder it was important to use a completely dry beaker when thoroughly mixing it with the PMMA powder to limit aggregation of the BaSO_4 , which can be seen as the bright spots in Figure 2.8:B. The two components were used in a 1:4 BaSO_4 to PMMA powder ratio, mixed 1:1 with the liquid PMMA component. The barium sulphate and PMMA powder were vigorously mixed to minimise the agglomeration found by Sikora [80] both prior to and after the addition of the PMMA liquid component.




Figure 2.8: A: μ CT scan of an augmented vertebrae, with some visible PMMA residing in the needle channel. B: μ CT scan of an augmented vertebrae using PMMA mixed with barium sulphate.

Cement Leaking from Vascular Channels: Preventing the cement from exiting the vertebrae from vascular channels while injecting the cement proved to be another obstacle to achieving a physiologic fill volume for the vertebrae. These channels lead both out of the anterior face and from the vertebral body into the spinal canal, this can be seen in Figure 2.9 and 2.10. *In vivo*, these channels would be filled with vasculature preventing the cement leaking through them. Two main methods were used to stop cement leaking while carrying out the procedure on the bovine tail vertebrae. The first was to use the same rod used for mounting the vertebrae in their end-caps to limit the passage of cement into the spinal canal. The second was to use blu-tac to cover the external vascular channels, wrapped with cling-film to hold it in place. This allowed any bone marrow free passage out of the vertebrae, but provided sufficient resistance to limit the flow of cement, acting similarly to vasculature in the *in vivo* case.




Figure 2.9: A: μ CT scan of an augmented vertebrae showing the cement leaking from vascular channels on the anterior side. B: Photograph of an augmented vertebrae cut into four quarters showing a vascular channel leading into the spinal canal.




Figure 2.10: μ CT scans of two vertebra, showing the cement leaking into the spinal canal and out of the vascular channels and the vertebral surface.

2.2.5 MicroCT Scanning

μ CT scans were taken at three occasions during the experimental process: before and after the initial load to failure, then following the augmentation of the specimens. The process required the vertebrae to be defrosted and at room temperature, given that the radio-opacity of water differs between solid and liquid states, hence vertebrae were usually defrosted overnight in a 4°C fridge. Imaging of vertebrae was carried out using the SCANCO (SCANCO Medical AG, Switzerland) XtremeCT II, a high resolution peripheral quantitative computed tomography (HR-pQCT) scanner. Vertebrae were loaded two at a time in a carbon fibre loading cradle and placed within the μ CT scanner. The settings used for the scans were: an isotropic voxel size of 82 μ m, energy settings 900 μ A, 60 kV and 300 ms exposure time. These settings were based upon previous studies using the same scanner and similar vertebrae carried out in the group [42, 60, 80, 91].

2.2.6 Results

The stiffness values for twelve vertebrae from six bovine tails are shown in Figure 2.11. Of the twelve vertebrae only two, the first and second vertebra of the second tail (T2 CC1 &

T2 CC2), were fractured. The remaining nine (excluding T1 CC3, the non-fractured control vertebra) reached 9500 N without a clear failure, however all showed a slight decrease in the gradient before the 9500 N limit.




Figure 2.11: The maximum stiffness of 12 bovine tail vertebrae between 0 and 5000 N taken from the load - displacement data. Stiffness values for the intact vertebrae, following fracture and following subsequent vertebroplasty are shown. * Indicates those specimens that achieved a clear failure below 9500 N.

The results for the fill volume of cement in the augmented specimens are presented in Table 2.1; this measurement was acquired from the segmented models generated from the μ CT scans. Fill volumes of between 3% and 17% were achieved. There was a lack of a correlation between fill volume and an increase in augmented specimen stiffness over fractured stiffness and only five vertebrae showed an increase in stiffness compared to their fractured state. The images in Figure 2.12 shows the extent of the cement fill for the two vertebrae with the largest fill volume.

The attempt to reduce cement leaking through vasculature during the vertebroplasty procedure can be seen in Figure 2.13. The methods employed were found to greatly reduce the quantity of cement observed in both the spinal canal and around vascular channels at the vertebral body surface when compared to scans in Figure 2.10.

The relationships between the intact stiffness and the amount this changed following fracture and following vertebroplasty are shown in Figure 2.14. In both cases, there was a lack of correlation between the difference in stiffness and the intact stiffness. This result

Table 2.1: The volume of cement and the vertebra volume for the 12 specimens used, along with the percentage cement fill and an indication as to whether the stiffness values of the augmented vertebrae were greater than the fractured stiffness. This information was measured from the down-sampled models generated from μ CT scans of the vertebrae.

Vertebrae	Cement Volume (mm ³)	Vertebra Volume (mm ³)	Cement Percentage of Vertebra Volume (%)	Increase in Augmented Stiffness over Fractured Stiffness
T1 CC1	2260	32440	6.97	*
T1 CC3	465	27039	1.72	
T2 CC1	663	23285	2.85	
T2 CC2	3405	20373	16.71	*
T4 CC3	1363	25446	5.36	
T6 CC1	830	29332	2.83	
T8 CC1	1257	37357	3.36	
T8 CC2	4489	29248	15.35	
T8 CC3	1041	28403	3.67	*
T9 CC1	2922	45681	6.40	
T9 CC2	2210	38894	5.68	*
T9 CC3	2437	35840	6.80	*

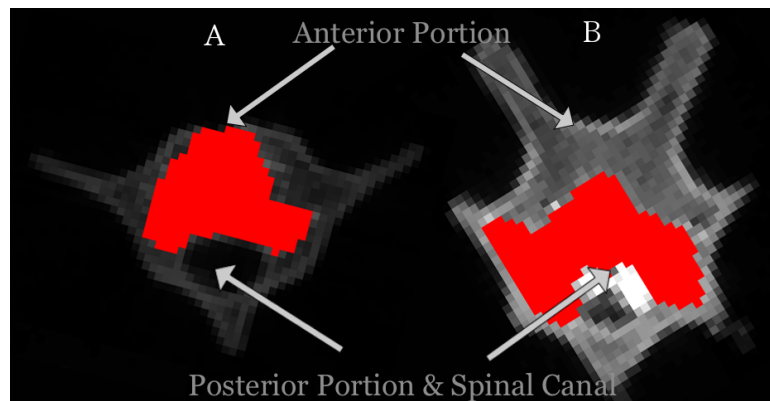


Figure 2.12: Axial μ CT slices of T2-CC2 (A) and T8-CC2 (B), with cement masked in red, showing the extend of cement fill at the point where the cement was most anterior.

indicates that the magnitude of any increase or decrease in the vertebral stiffness following augmentation was not caused, or a feature of the initial, intact vertebral stiffness.




Figure 2.13: μ CT scans of four augmented vertebra using a steel rod to fill the spinal canal and blu-tac to cover the external vascular channels. Shows greatly reduced cement content within the spinal canal with less cement at the surface of vascular channels. (A) T2 CC1, (B) T6 CC1, (C) T9 CC3, (D) T9 CC2.




Figure 2.14: A: The difference between the post augmentation and fractured stiffness against the intact stiffness. B: The difference between the post augmentation and intact stiffness against the intact stiffness.

2.3 Finite Element Modelling

2.3.1 Introduction

As discussed in Chapter 1, finite element models have been used to understand various vertebral treatments and here offer the potential to investigate how vertebral features influence the outcome of augmentation. Such investigations include identification of geometric and material property features, which through the use of FE models can be

changed programmatically to pinpoint the effects of treatments like vertebroplasty.

Here, the main aim was to develop methods that enable the creation of specimen specific models of bovine tail vertebrae, allowing the subsequent creation and generation of the much more clinically relevant human lumbar vertebrae using similar methodologies. Initially the focus was on the generation of models that accurately describe the mechanical behaviour of intact bovine specimens, once this was achieved to a reasonable degree, an attempt to model augmented specimens was made. In addition to these larger goals, certain sensitivity tests were carried out, including those to understand the effects that additional meshes, mesh sizes and mesh interactions have on model stiffness. Finally some preliminary investigations were made into the effect of changing augmented region positions, however the majority of this investigation was reserved for the generated statistical shape models reported in Chapter 5.

2.3.2 Software and Imaging Resources

The main pieces of software and the versions used in the project are described here. The commercial image processing and mesh generation software package Simpleware ScanIP 2016.09 (Synopsys, Mountain View, CA, United States) was used for image segmentation, application of morphological filters and mesh generation. The image processing software and distribution of ImageJ, Fiji version 1.51a [94] was used for image processing, application of certain thresholds and calculation of histogram data. A plugin for ImageJ was used named Bonej, version 1.4.1 [95] for calculations of bone and trabecular properties.

For the finite element modelling, the commercial software Abaqus CAE 6.14-1 (Dassault Systemes, Velizy-Villacoublay, France) was used.

The numerical computing environment MATLAB R2015a (The MathWorks Inc, Natick, MA, USA) was used for the conversion of the proprietary ScanCo ISQ format to DICOM format among other scripts written to analyse data. In most areas of model development and analysis, Python 3.6.3 (Python Software Foundation, <https://www.Python.org>) was used. Python 2.7.3 (Python Software Foundation, <https://www.Python.org>) was used for interfacing with the Abaqus scripting interface and for running scripts within the ImageJ macro environment. Solving of finite element models was undertaken in part on ARC2, part of the High Performance Computing facilities at the University of Leeds, UK.

2.3.3 Model Creation

2.3.3.1 Model Orientation

Model orientation was set such that the z-axis described the vertebrae through the axial plane. This was the axis along which loads were applied during testing. Other axes were left undefined, with the x/y axis being different for each vertebral model due to discrepancies in the rotation of the vertebrae during μ CT scanning.

2.3.3.2 Modelling Methodology for Non-Augmented Vertebrae

The scans acquired from the μ CT scanner were converted from the ISQ file format, generated by the scanner software, into the more portable TIFF image format files using an existing in-house MATLAB script that additionally converts the greyscale of the scan into 256 bins. This conversion from 16 bit TIFF files with 65,536 bins to 8 bit TIFF files was carried out to provide a more manageable number of material properties. Once the scan had been pre-processed, it was imported into ScanIP ensuring that the spacing of voxels was correctly set, in this case to 82 μ m. Once imported, the location of the loading point was identified to simulate the correct experimental load within ABAQUS; the marker (see Figure 2.3) appears bright on the scan and its centre was taken as the load point, calculated by converting the position into mm. This was achieved by multiplying by the native resolution of 82 μ m.

The following parts of model creation were carried out using a Python script from within the ScanIP software. The script carries out the process described below and was generated by the author by carrying out the process manually and in order to understand the steps required and then writing a script based on the log file to perform those actions. The development of the script removed much of the user variation in the segmentation of each vertebral model. The effect of user variation during the segmentation process is examined in Section 2.3.5.5.

It was easier to down-sample the image stack prior to segmentation, due to the time required for the software to generate high resolution masks and increased memory usage at higher resolution. The effect of down-sampling can be seen in Figure 2.15. However, in certain cases, for example when modelling vertebral augmentation, in order to attempt to capture the intricacies of the structure and the boundaries between cement and trabecular bone, it was favourable to generate the mask prior to down sampling. The image stack was down-sampled to voxels of 1 mm \times 1 mm \times 1 mm. Previous studies that had examined sensitivity to mesh size of vertebral FE models have shown a good trade off between




Figure 2.15: Side and top view of a vertebral μ CT scan showing the effect of the downsample from $82 \mu\text{m}$ to 1mm cubed.

computational cost and model accuracy at this resolution [69]. Additionally, previous studies using similar methodology and resolution have build specimen specific models with good agreement to the experiment [42, 60, 96]. It was therefore deemed unnecessary to carry out further mesh sensitivity analysis in this study.




Figure 2.16: Side view of a vertebral model showing segmented vertebra, including the internal void that is filled.

Once down-sampled, the image stack was segmented into the constituent parts - the vertebrae and cement end-caps. The different regions that were required to be segmented have different greyscale values, hence the general shape of the masks was created through a thresholding tool that selected volumes of the image stack between two greyscale bounds.

For the end-caps these bounds were usually between greyscale values of 12 to 22 and, if the specimen was not augmented, the vertebrae between 23 and 255. For augmented specimens these limits change to 23 to 65 for the vertebrae and 66 to 255 for internal cement containing barium sulphate. These values were selected by visually limiting the amount of unwanted material selected within the threshold and maximising the wanted material, for example by selecting as much of the end-caps as possible while limiting the selected background and vertebral material to a minimum. The results of this thresholding operation can be seen in Figure 2.16. It was preferable to avoid internal voids within each mask, due to the potential for errors that may arise from the extra surfaces and contacts these surfaces would create. These voids were therefore removed with the use of the morphological close and cavity fill tools within ScanIP (Figure 2.16).

The following parts of the method were carried out manually, following the completion of the automatic segmentation Python script. The two end-caps were separated into two separate masks by first duplicating the mask and then flood filling each end-cap to form separate masks.

An FE model was created using the previously generated masks and properties for the volume meshing, materials and contacts were set. The grid size for the model was set to $1 \times 1 \times 1$ mm using the FE grid algorithm which uses a mix of tetrahedral and hexahedral elements, C3D4 and C3D8 respectively. Material properties were set to homogenous with a Young's modulus = 2.45 GPa and Poisson's ratio = 0.3 for the endcaps [97]. Material properties for the internal cement regions for the augmented vertebrae are described later. The material properties for the vertebral volume were set to a greyscale based material type using the greyscale background information. The coefficients were set so that both the density and Young's modulus were equal to the greyscale value for that element, allowing the Young's modulus to be set correctly in the following steps and as described in Section 2.3.4.

Contacts were set as placeholders to be edited in Abaqus in the steps following. Contact pairs were defined between each connecting component and another was added between the superior end-cap and the upper boundary on the z-axis. The second contact type was a node set between the inferior end-cap and the lower boundary on the z-axis used to create an enastre boundary condition for the model base.

Following this, the FE model was meshed, then exported in an Abaqus input (INP) file format. This file was imported into Abaqus where the following configuration was completed. A second Python script was used to perform these operations because they




Figure 2.17: Side & top down view of a vertebral model showing the alignment of the analytical rigid plane.

were identical between models. First, an analytical rigid plate was created to represent the upper loading platen of the materials testing machine and was centred at the loading point previously found from the marker, this can be seen in Figure 2.17. Once aligned any previous placeholder interactions were removed and a tied interaction was created between the rigid plate and the superior end-cap, along with tied interactions between the vertebra and both end-caps and, if appropriate, to any internal cement volumes. An encastre boundary condition was created at the bottom surface of the inferior end-cap removing all rotational and translational movement and therefore mimicking the experimental setup. A displacement boundary condition was applied to a reference node at the centre of the rigid plate and therefore aligned with the loading position, the properties were set such that 1 mm of displacement was applied in the negative z direction; lateral motion in the x and y planes was restricted, while rotation about the loading point was allowed, mimicking the experimental steel ball from the experimental setup.

The Python script was written to define the material properties of the greyscale dependent vertebral elements by setting the Young's modulus to the greyscale value multiplied by a conversion factor (which is discussed in section Section 2.3.4). The full model was then submitted for solving through a command in the script. Finally, the predicted stiffness was determined by dividing the axial reaction force (i.e. in the axis of load application) at the reference point by the 1 mm of applied displacement.

2.3.3.3 Augmented Model Generation

In order to capture the detail of interdigitation between the vertebra and the injected cement, the masking process was carried out prior to downsampling, as shown in Figure 2.18. If this process was attempted following downsampling, it was found to be difficult to define the cement boundaries and the masked volume was inaccurate when compared to the full resolution scan. Masking the internal cement region used the same thresholding approach

as described above. Once downsampled the small areas of highlighted voxels, that actually represented bone, were removed.




Figure 2.18: A lateral slice through an augmented bovine tail vertebra, showing the cement mask in red at the full $82 \mu\text{m}$ resolution.

2.3.4 Material Properties

2.3.4.1 Bone Material Properties

Material properties for the bone tissue were modelled elastically using a bone element-specific elastic modulus (E_{ele}) that is dependent on the average greyscale value for the element in question (GS_{ele}) with the conversion factor between the two being α .

$$E_{\text{ele}}(\text{GPa}) = \alpha \text{GS}_{\text{ele}}$$

This was required because each element containing differing quantities of bone and bone marrow due to the continuum level modelling carried out. Hence, a homogenous value for the trabecular bone would not have represented the varying material properties seen between different elements. The conversion factor, α , was used to convert between the greyscale value for each element and the Young's modulus.

A separate set of 24 bovine tail vertebrae underwent the same experimental and computational methods described previously in order to tune the value of α . This additional set was split into two groups of 12 and the work was undertaken in collaboration with Dr Sebastien Sikora, Dr Fernando Zapata Cornelio and Ruth Coe (Research Fellow, Research Fellow and PhD student respectively). The specimen groups consisted of a calibration

group that was used to determine the value of α , and a validation group that was used to determine the level of agreement in predicting the stiffness of an independent set of specimens.

The calibration for the conversion factor, α , was carried out using a golden section search scalar optimisation process. Specifically using the Brent method within the opti4Abq toolbox [98]. The opti4Abq toolbox is an optimisation method written in Python 2.X that runs on a set of abaqus models to minimise the difference (in a least square sense) between the FEA output and a corresponding set of data. The minimisation approach is illustrated in Figure 2.19, where the objective function and outputs (f0, f1) represent the RMS error between computational and experimental results that needed to be minimised to find the optimum conversion factor. The input parameters (p0, p1) are the greyscale conversion factors, of which two were used, an initial and the initial plus a small constant (p0 and p0 + δ), to find the gradient. The value of p1 (the second α value guess) was calculated by the value of the gradient (from p0 and p0 + δ) and another constant, ζ . This constant, ζ , was chosen as a compromise between computational cost (smaller values requiring the models to be solved more times) and accuracy (potential to find other local minima). This was iterated over, first finding the RMS error using p1 and then finding the gradient from p1 and p1 + δ . This process of finding the gradients and new px values was repeated until the cost function achieved a value of 10^{-3} . The cost function is described:

$$|f(px) - f(px + 1)| < cost$$

Where $px + 1$ is the current guess of the parameter and px is the previous guess. The value of the cost function was again chosen for its compromise between accuracy of the result and computational cost. Initial estimates for the value of conversion factor, p0, δ , ζ and threshold value of the cost function were based on previous studies on different animal vertebrae [91].

A limitation of this approach, relating greyscale values to Young's modulus values, is the inability to compare material properties between studies using different μ CT scanners. A machine independent conversion factor (through a baseline of Hounsfield units to Young's modulus) is not available due for two reasons. The first is that, given the calibration between model and experimental data, any change to the properties of the model (boundary conditions, contact interactions, mesh density and element types) would change the greyscale-Young's modulus relationship. The second is that the Hounsfield units calculated are machine dependent [91, 99] and so a calibrated material property would require a further conversion based on the scanner used, utilising the method described by

Zapata et al. [91].

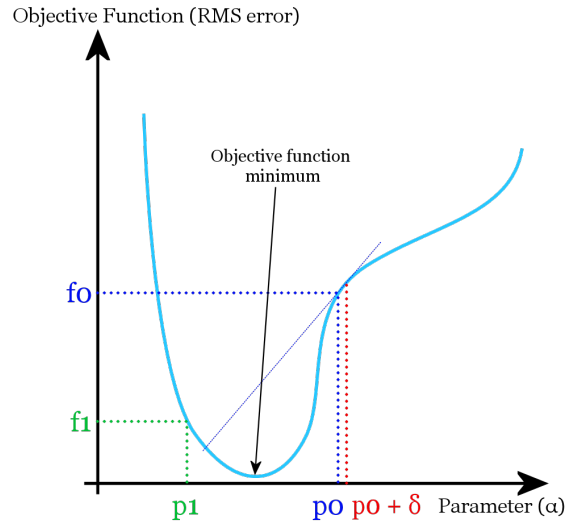


Figure 2.19: An illustration of the gradient descent method used to optimise the greyscale conversion factor for the models. The objective function represents the RMS error between computational and experimental results, with the parameters being the greyscale conversion factors being tested.

2.3.4.2 Augmented Specimen Material Properties

For convenience the values for the Young’s modulus for the interior cement volume were initially set to that of the inferior and superior end-caps. However, previous studies into modelling bone cement interactions [42, 64, 75, 77–80, 100] suggest simple material properties of PMMA do not accurately describe the environment. Due to the rule of mixtures and the results found in the literature [79, 101], the effect of reducing the Young’s modulus was investigated. This was carried out by reducing the Young’s modulus in 10 percent increments from a value of 2.45 GPa to 1.225 GPa.

Additionally, a preliminary investigation was carried out to examine the effect of including a yielding material interface between the bone and the cement, following the work by Sikora [80]. Here, a small interface layer (1 mm in thickness due to the model resolution) represents buckling that occurs in the trabeculae partially captured in the cement volume. This investigation identified the effect of different yield stress values in combination with different Young’s moduli for both the interface and the main cement volume.

The yielding region was created through the duplication and then erosion of the cement mask within Simpleware ScanIP, shown in Figure 2.20. This gave two masks, where one acts as the internal cement volume and the other as an interface layer. Generic material properties were set to the region within ScanIP, allowing the values of the yield stress and Young’s moduli values to be changed within the Python setup script. The properties tested

included changing the yield stress and Young's modulus of the region, whilst maintaining the properties of the internal cement volume at 1.225 GPa. Values for the yield stress in the interface region were varied between 1 MPa and 0.001 MPa, with the material becoming perfectly plastic beyond this. Young's modulus values for the interface region were varied between 1.225 GPa and 0.1 GPa, the tested range of these material properties was based on the results found by Sikora [80]. Sikora found that properties in these ranges gave the most accurate results based on a set of augmented trabecular bone samples. These values were manually tuned to find the optimum values with respect to the level of agreement between the experimental and computational results.




Figure 2.20: Creation of the interface layer: A, showing the initial description of the cement region in pink and B, showing the interface in pink and cement region in red, following the duplication and erosion.

2.3.5 Sensitivity Tests

2.3.5.1 Mesh Size Sensitivity

Element sizes of $1 \times 1 \times 1$ mm were used throughout, following previous convergence studies on specimen specific vertebral models described previously. The results of previous convergence study on porcine vertebrae showed that reducing the element size below $2 \times 2 \times 2$ mm led to changes in the model that were smaller than predicted errors originating from other factors, such as experimental errors and the simplification of boundary conditions [59]. However, reducing the element size to $1 \times 1 \times 1$ mm allows greater resolution when modelling the intricacies of the cement mesh for augmented specimens, the difference between the two resolutions can be seen in Figure 2.21. At resolutions smaller than $1 \times 1 \times 1$ mm, the trabecular bone and trabecular space starts to be captured, therefore creating models which are no longer at the continuum level.




Figure 2.21: Mid-slice through an augmented vertebra, cyan: vertebral body, red: cement. A, element size of 1 x 1 x 1 mm. B, element size of 2 x 2 x 2 mm.

2.3.5.2 Sensitivity to an Additional Mask

The addition of cement into the vertebral body created an extra mesh boundary within the mesh containing the vertebral elements. In order to test what effect this may have on the stiffness of models, an un-augmented specimen was tested with an extra mesh boundary representing edge of a likely cement region, but with the material properties of the elements within it based on their greyscale as with the other bone elements. The mask was created by duplicating the vertebral mask and eroding it until the volume was approximately 20 % of its original volume. This allowed testing to be carried out on the effect of the extra mesh boundary alone. The use of an augmented specimen would have allowed a more accurate cement shape, but it would have hindered setting material properties to that of the bone greyscale and would therefore have created an additional level of uncertainty. Mesh interactions between the two meshes (internal vertebral surface and the cement mask surface) were set using the contact pair interaction and treated similarly to the interaction between the end-caps and the vertebrae using a tied interaction. Following model setup in Abaqus as outlined in Section 2.3.3, the model was loaded in compression to 1 mm and its stiffness was recorded.

There was no difference in the stiffness output between the two models, with and without the internal cement mesh, meaning that any changes to the augmented model stiffness would be due to the material properties of the cement.

2.3.5.3 Mesh Interactions

The effect of mesh interactions between the vertebral body and the internal cement mesh were tested by comparing a) tied interactions between the two surfaces and b) removing any interaction and merely changing the material properties of the internal cement region (neglecting the contact pair steps described earlier). This was carried out for four augmented specimens following the same setup within Abaqus as described earlier.

The results are presented in Table 2.2 and show a negligible difference between tied and non-tied (single mesh) models. This difference falls well below the difference between experimental and computational models, especially for the augmented specimens, hence the effect of this interaction can be neglected from further tests.

Table 2.2: The difference between interaction properties, tied and not tied for four augmented vertebrae specimens.

Vertebrae (Tail Number, Vertebral Level)	Tied Interaction (N/mm)	No Tied Interaction (N/mm)	Difference (%)
T2 CC1	5496	5496	0
T2 CC2	8086	8086	0.001
T6 CC1	3686	3686	0.001
T4 CC3	6059	6059	0.0005

2.3.5.4 Augmentation Location Sensitivity

Another study was carried out to identify the sensitivity of the models to the position of the cement volume. Tests were carried out to identify the effect of moving a 12 % fill cement volume axially and sagittally in 2 mm increments. This utilised the yielding interface described in Section 2.3.4.2 along with the +CAD tools built into Simpleware ScanIP to move a surface based description of the cement volume in the two anatomical planes. A sphere surface of volume equivalent to 12 % fill volume for the T12 CC2 vertebra was created within the +CAD software and imported into the project file for the vertebra, where it was positioned in its first position. This first position was 1 mm from the top of the vertebral body, ensuring 1 mm of bone surrounded the sphere. It was then converted into a mask, duplicated and eroded using the same method described in Section 2.3.4.2 to create the yielding material interface and set with optimum material properties. Following this the mesh was generated and the Abaqus input file was exported. The sphere surface was then moved 2 mm in the inferior direction, the previous masks deactivated and the process repeated through the use of a Python script to carry out the operations until the bottom of the vertebra had been reached. This was also carried out in the sagittal plane.

The results of moving the cement volumes through the T12 CC2 vertebra are shown in Figures 2.22 and 2.23. For the sagittal plane a reduced stiffness was seen when the volume of cement was positioned most anteriorly, where the yielding interface of the cement volume encroaches on the denser bone of the cortical shell. Centrally the stiffness was greatest, where the cement volume replaces the least dense bone, which can be seen in Figure 2.24, where the least dense (blue) parts are “hidden” by the cement volume in certain positions. Little change to the stiffness at the posterior side of the vertebra was seen due to the




Figure 2.22: The effect of moving the cement volumes from the most anterior position to the most posterior position on the recorded stiffness of the vertebra, when using uniform 12 % fill volume of cement.




Figure 2.23: The effect of moving the cement volumes from the most superior position to the most inferior position on the recorded stiffness of the vertebra, when using uniform 12 % fill volume of cement.

additional support from the posterior elements.

Axial movements of the cement volumes showed that the most inferior and most superior positions resulted in the greatest stiffness and central positions resulted in the least stiff models. This was potentially due to the greater quantity of bone surrounding the cement volumes at top and bottom where the vertebral body was wider. The peak in the central loading positions was again likely due to the “hiding” of the weaker bone in the centre of the vertebra.




Figure 2.24: A density map of the T12 CC2 vertebra. Yellow/orange elements are the most dense, while blue elements are the least.

Finally, despite the changes to the stiffness when moving the cement volume, these changes are smaller than five percent between the most and least stiff result, suggesting cement volume movement at these small volumes was not significant. It may also suggest that cement shape may play a more important role, where endcap to endcap distributions show much larger effects as reported in previous literature [75, 102, 103].

2.3.5.5 User Variability Sensitivity

A user variability study was conducted to identify the variation in the modelling approaches, specifically masking the bone and endcap regions. Four users each masked the bone and endcaps for eight vertebrae using thresholding and other morphological filters. This process was later automated using the best approach as described previously. The FE models were then generated and the material properties were optimised using the greyscale optimisation process separately for each users' models.

The results for the variability can be seen in Table 2.3 showing the mean, maximum, minimum and range as well as the mean for each measure. Across the different models, the mean spread in the stiffness of values obtained was 158 N/mm. The maximum difference between users was for specimen six where a range of 246 N/mm was found. This equates to a maximum possible error of 14 % for different users carrying out the same models creation process.

While 14 % error between users is relatively high, the models used in this chapter and subsequent chapters were all build by the author. Users tended to produce consistent differences in the stiffness of generated models, for example User 1 had consistently lower stiffness values and User 3 had consistently higher stiffness values (means of 1876 N/mm

Table 2.3: The variability the modelling approaches of four users with 8 models, each users models undergoing separate greyscale optimisation and FE model solving

Specimen	Stiffness (N/mm)			
	Mean	Maximum	Minimum	Range
1	2457	2567	2379	188
2	2298	2368	2218	150
3	2374	2463	2261	202
4	1600	1659	1545	114
5	2281	2324	2218	106
6	1899	1990	1744	246
7	1330	1398	1236	162
8	1183	1219	1119	100
Mean	1928	1998	1840	158

and 1952 N/mm respectively). Given these results, the Python script previously described was used within Simpleware ScanIP to eliminate the user choice and therefore single user variability.

2.3.6 Statistical Analysis

To quantify the agreement between the computational results and the experimental results for the measured stiffness, the concordance correlation coefficient (CCC) was used. The CCC measures the agreement between two variables and was described by Lin [104] as a method to evaluate reproducibility. The CCC was calculated as:

$$CCC = \frac{2\rho\sigma_x\sigma_y}{\sigma_x^2 + \sigma_y^2 + (\mu_x - \mu_y)^2}$$

Where μ_x and μ_y are the means for the two variables and σ_x^2 and σ_y^2 are the corresponding population variances. ρ is the Pearson correlation coefficient between the two variables. For n independent pairs of samples:

$$CCC = \frac{2S_{12}}{S_1^2 + S_2^2 + (\bar{Y}_1 - \bar{Y}_2)^2}$$

Where,

$$\bar{Y}_j = \frac{1}{n} \sum_{i=1}^n Y_{ij}, \quad S_j^2 = \frac{1}{n} \sum_{i=1}^n (Y_{ij} - \bar{Y}_j)^2, \quad j = 1, 2;$$

and

$$S_{12} = \frac{1}{n} \sum_{i=1}^n (Y_{i1} - \bar{Y}_1)(Y_{i2} - \bar{Y}_2)$$

Specifically, this quantifies the degree to which pairs sit on the $x = y$ line. Departures from this perfect agreement line result in a $CCC < 1$ even in cases where the Pearson correlation coefficient would equal 1. $CCC = 0$ corresponds to no agreement and a value of -1 would indicate perfect negative agreement.

2.3.7 Results

The optimisation process gave a value for the conversion factor of $0.012529 \text{ GPa}^{-1}$, allowing conversion between greyscale values for elements (measured on a scale of 0 to 255) and their elastic modulus values (measured in GPa). This value was used for modelling the bone constituents of all the intact and augmented vertebrae. The agreement between the *in vitro* and *in silico* results for the specimen specific models are shown in Figure 2.25. The agreement of intact vertebrae achieved a concordance correlation coefficient (CCC) of 0.49 compared with 0.14 for the augmented vertebrae with a simple tied interaction (Table 2.4). The non-augmented CCC value increased to 0.60 if the uncharacteristically stiff T8-CC2 was removed.




Figure 2.25: The *in silico* compared with *in vitro* stiffness for intact specimens (triangles) and augmented specimens (circles). The dotted line shows a one-to-one correspondence.

The effect of changing the modulus of the cement volume in the augmented specimens is

Table 2.4: The mean, standard deviation and concordance correlation coefficient (CCC) of the intact and augmented vertebrae (simple tied interaction) for *in vitro* and *in silico* results.

Intact Specimens	Mean Stiffness	Standard Deviation	CCC
<i>in vitro</i>	5684	1196	0.4895
<i>in silico</i>	5610	958	
Augmented Specimens			
<i>in vitro</i>	4246	1371	0.14
<i>in silico</i>	6507	1298	

presented in Figure 2.26. There was a linear decreases in the stiffness of vertebrae with the reduction of the elastic modulus for the internal cement volume. The two vertebrae that show more prominent decreases in stiffness were those vertebrae that contained larger volumes of cement following their augmentation.

The effect that this had on the data in comparison to the *in vitro* stiffness results can be seen in Figure 2.27, where the reduction in *in silico* stiffness moved the data points closer to the $y = x$ line of perfect agreement between the experimental and computational results. This gave a CCC of 0.18.

The effect of the addition of a yielding material interface between the cement and bone further increased the agreement in comparison to the *in vitro* results. The increasing agreement with increasing sophistication of modelling the augmentation can be seen in Figure 2.28, comparing the simple tied interaction, the reduced modulus approach and the combination of this and the yielding material interface. An improvement was seen when combining the reduced modulus for the cement with the yielding interface (Young's modulus 1.225 GPa, yield stress 0.005 MPa), with the CCC improving to 0.23. The effect of changing these interactions and properties was not uniform, also shown in Figure 2.26, the vertebral models respond differently to different interactions, with increases and decreases seen within the same set of vertebrae when changing the same interaction.




Figure 2.26: The percentage decrease in the vertebral stiffness after reducing the elastic modulus of the cement volume within 12 augmented vertebrae.




Figure 2.27: The effect of reducing the elastic modulus of the cement volume within 12 augmented vertebrae. Shows the in silico stiffness for the six elastic moduli tested against their in vitro stiffness.




Figure 2.28: A comparison of the different approaches to modelling augmentation in the bovine tail vertebrae. Comparing a simple tie between the bone and cement, a reduced modulus for the cement region and a reduced modulus in combination with a yielding interface layer between the bone and cement.

2.4 Discussion

2.4.1 Experimental Discussion

The aim of the experimental arm of the study was to develop the various aspects of the methodology on a set of bovine tail vertebrae. Previously developed methods of μ CT scanning were applied and verified for the bovine tail vertebrae. Methods of performing vertebroplasty and dealing with the unique challenges arising from the bovine vertebrae were developed. This part of the chapter provided a good basis for both the continued testing of vertebroplasty and characterisation of vertebrae, factors which will become increasingly important when using human tissue in the following chapter.

Regarding the results of stiffness at the intact, fractured and augmented stages, the expected trends were not always clear. Most commonly the intact vertebrae have the greatest stiffness with the fractured stiffness showing a reduced value following the damage created with the initial load to failure. The variation of the decrease (and increase) in stiffness for the fractured vertebrae may have a variety of reasons, although the most likely cause is the variation in the level of damage caused by the initial loading step. The load-displacement outputs of these tests varied between those that included a failure (Figure 2.4:A) and those that showed no sign of failure up to the limit of the load cell (Figure 2.4:B). There was not a clear correlation between the level of damage seen in the initial loading step and the reduction in the stiffness seen following the reloading. It is not to say that the vertebra that reached 9500 N experienced no damage, with the gradient of the load displacement curve often reducing and plateauing as the 9500 N limit approached. The interesting increase in the fractured stiffness for T1-CC1 compared to the intact stiffness may be explained if it is assumed that the compacted trabeculae following the first load to 9500 N resulted in a stiffer material for the following tests. Other possibilities that may explain the increase in stiffness include a change in the seating of the vertebra within the PMMA endcap. If the initial load changed how the vertebra was positioned in the endcaps it would change the response to loading, even when loaded at the same point. This may be the reason that the control vertebra (T1 CC3), also showed an increase in stiffness following its initial load to 5000 N. The fact that a second vertebrae also presented an increase in stiffness following the initial load and that a small reduction in the stiffness was seen in the post augmentation load of the control specimen means that little information can be drawn from its results. A larger set of control specimens may help to elucidate the effect of the fracture when understanding the post augmentation stiffness, however, given the lack of any relationship between post-fracture and post- augmentation properties, it remains unclear as to whether this would help.

The cement fill volume information shows that a relatively small percentage volume of cement was injected into the vertebrae on average, with only two vertebrae approaching the clinically relevant 20% fill. Unexpectedly, only one of these two vertebrae receiving more than 15% cement fill showed an increase in augmented stiffness over the fractured stiffness. A possible explanation is that it is not only the fill volume that is important in restoring the vertebral stiffness but the placement or shape of the cement fill volumes. This is shown when comparing the segmented scans of the two vertebrae with the greatest fill volume with the T2-CC2 specimen showing cement extending to the anterior wall of the vertebral body, while the cement is limited to the posterior and centre of the vertebral body for T8-CC2. This may help to explain why the stiffness of T8-CC2 did not increase following augmentation. The reduction in stiffness following augmentation for seven of the twelve vertebrae may be due to damage caused by the insertion of vertebroplasty needles. Clinically this damage left behind from the needle channels would heal (depending on the level of osteoporosis), most likely restoring the portion of the vertebrae that received the needles back to its intact properties. Additionally, needle insertion into the dense bone of the bovine tail vertebrae required a mallet to reach the anterior portion of the vertebral body, clinically this would not be required due to the much less bone of the human lumbar vertebrae [41]. Such a violent approach would have likely damaged the bone surrounding the needle entry, causing micro-fractures not visible on the μ CT scans.

Another possible area of inconsistency is the temperature at which the vertebra were mechanically tested. While it is ensured that the specimens were fully defrosted they were tested at both fridge temperature (4°C) and room temperature (20°C). The effect of this variation in temperature needs to be identified or, otherwise, a consistent temperature should be used in future studies. This latter recommendation was adopted for the subsequent tests on human vertebrae.

Despite encouraging results regarding the vertebroplasty methodologies, it was difficult to achieve the desired quantity of the cement in the vertebral body. This was mainly due to the difficulty injecting the cement in a smooth manner, which may have been caused by either the tip of the needle becoming blocked following its reinsertion into the needle track. One option to test in future work would be side opening needles, which would help guide the cement more accurately to the regions required while circumventing issue with the needle becoming blocked. Additionally, the shape of the vertebrae added to the difficulty of performing vertebroplasty, with their narrow, cylindrical shape meaning the cement had to travel large distances axially through the trabecular bone to achieve clinically relevant cement fill volumes. In contrast the much wider human lumbar vertebrae used in latter

chapters provide much more space to inject larger volumes of cement. This, coupled with a greatly reduced average trabecular density for the osteoporotic human vertebrae (see Figure 1.5 and [38, 41]) should allow for greater volumes of cement to be injected.

When attempting to identify trends between the pre and post augmented and pre and post fractured vertebral stiffness, no relationships were found. It was expected that initially weaker vertebrae would show a larger change in stiffness between the fractured state and the augmented state and similarly between the intact and augmented states. However, no such trends were found, suggesting that properties other than the stiffness or bulk material properties were the cause of the variation. These properties may include the vertebral geometry and trabecular structure, which are investigated through the use of statistical shape models in Chapter 5.

The experimental methods reported in this chapter provided a basis for the experimental work using human tissue albeit some adaptation was required due to the differences between the tissue types.

2.4.2 Computational Results Discussion

The aims of the computational part of this chapter were to develop methods to build models of vertebrae using automated approaches and to carry out sensitivity tests. These sensitivity tests were to identify the best methods to use when carrying out similar model creation and simulation of the set of human lumbar vertebrae in Chapter 3. This was achieved, with automation of the segmentation process within ScanIP and automation of the model setup within Abaqus carried out using a series of Python scripts. These scripts require no conversion to be used with the modelling of human lumbar vertebrae, with the exception of the thresholding carried out in ScanIP. The preliminary sensitivity tests provided an initial indication of which parameters the vertebrae were sensitive to, allowing further investigation in subsequent chapters.

The current methods of masking and meshing the internal volumes of cement in the augmented models provided a robust method of creating augmented models. The sensitivity tests carried out on the additional mask inside the vertebrae showed that any effects seen are due to the volume of cement and not a result of the mesh changes. Similarly the inclusion of a tied interaction between the two meshes failed to affect the result.

The sensitivity tests on the location of the cement volume suggested little influence of its position on the reported stiffness. It did highlight a potential source of error, where anteriorly placed volumes of cement would be expected to increase the stiffness of the

vertebrae, given the reduced support and density on the anterior side. However, due to the yielding interface replacing the dense cortical shell, the result was a reduced stiffness. Therefore, care needs to be taken when investigating cement volume movement as to not create unnatural simulations. When modelling the experimental specimens, where the cement region is taken from the μ CT scans, this is not a problem because the cement does not replace the bone when injected, therefore preserving the cortical shell in the vertebral model. Additionally, moving the cement volumes highlights the importance of filling or replacing the weakest, least dense parts of the vertebrae with cement. Although in this case the change of replacing the central “void” in the vertebra was small compared to placing the cement volume elsewhere, as cement volume size increases (and modelling accuracy improves) the effect may become more pronounced.

The intact models agreed well with experimental results, with very similar results for the mean and, excluding anomalous results, a good CCC value, showing that the previously validated conversion factor worked well with this set of data and that segmentation and model setup worked correctly. The poor agreement between the augmented specimen models and their experimental counterparts was an expected result that agrees with similar studies in the literature [42]. In an attempt to produce a better agreement, the elastic modulus of the cement volume was reduced in accordance with the experimental results of Race et al. [101] and similar methods employed by Wijayathunga et al. [42], where the reduction in modulus was expected due to gaps between the bone and cement caused by cement shrinkage and also by pores within the cement itself. The reduction in stiffness formed a linear pattern as the elastic modulus was reduced, with those vertebrae that showed the greatest reduction being those containing the largest volume of cement. While these results did show a reduction in the stiffness, closer to that of the experimental values, it did not explain the disagreement fully. This suggests that a combination of improvements to the augmented models are required. The final increase in the agreement between experiment and model was seen with the addition of the yielding interface. This reduced the stiffness of the set on the whole and also increased the CCC value. Interestingly, this was not a uniform or predictable effect with no correlation between yielding interface surface area or volume and the change in stiffness. It suggests the effect was influenced by the density of the surrounding bone on various regions and therefore was not a bulk effect.

The results from this section suggest that an accurate representation of the bone-cement interface is vital for accurate models of augmentation. The modelling methods provided a repeatable process for building specimen specific models that were utilised for modelling

human vertebrae, as reported in the following chapter.

Chapter 3

Non-augmented Human Lumbar Vertebra Study

3.1 Introduction

Previous *in vitro* studies of vertebroplasty, as well as the work in Chapter 2, have used a range of animal vertebrae, including bovine, ovine and porcine tissue among others. While all of these vertebrae have important uses for developing methodology and understanding certain characteristics, none share all of the characteristics of human vertebrae including the same trabecular structure, thickness of cortical shell, axial strength or stiffness. Hence, to understand the mechanical effects of vertebroplasty in human lumbar vertebrae, human tissue is required.

The primary goal of the work presented in this chapter was to develop methodologies to accurately model non-augmented human lumbar vertebrae. The details of the human tissue used in this study are described along with the experimental methods for testing and imaging the specimens. The finite element methods employed, and the sensitivity tests undertaken are then reported. Comparisons are then made between the experimental and FE results and the implications are discussed.

3.2 Experimental Methods

This section describes the experimental methods used and developed to test and record details of the set of human lumbar vertebrae. This includes an overview of the major changes compared to the methods presented in Chapter 2, the specifics of some of the new additions to these methods and the reasons for, and results of these improvements.

3.2.1 Specimens

Fourteen lumbar vertebrae from four cadaveric spines were used, with the details shown in Table 3.1. These spines were sourced from the Leeds GIFT 2 Research Tissue Project, with consent for the use of tissue being taken by the NHS Blood and Tissue Services.

Table 3.1: Details of the lumbar sections used from four cadaveric spines.

Spine Name	Vertebrae	Sex	Age
Spine 1	L1, L2, L3, L4, L5	F	90
Spine 2	L1	F	94
Spine 3	L1, L2, L3	M	86
Spine 4	L1, L2, L3, L4, L5	M	83

3.2.2 Dissection & Potting

The first step in the experimental setup, as with the bovine tail vertebrae, was the dissection and potting of the vertebrae. As before, potting was required to ensure parallel endcaps to allow repeated axial loading on materials testing machines, as well as to remove excessive soft tissue that could affect the results, given that only the bone would be subsequently modelled. Dissection of the human vertebrae presented added challenges compared to dissecting the bovine tail vertebrae due to the degenerated nature of the vertebrae, with fused facet joints common and difficulties circumventing the bone growths on some of the most degenerated vertebrae.

The geometry of human lumbar vertebrae varies considerably to that of the bovine tail vertebrae for which the methodology was initially developed. In particular, the human vertebrae possess much larger posterior elements with the facets extending much lower, below the inferior surface of the vertebral body. If these protruding processes were potted, it would have led to much more of the posterior elements being constrained, therefore restricting the rotation of the vertebral body endplates under axial load. The solution adopted was to remove the posterior elements, following methods that had also been used previously in the literature [42, 60], where only the vertebral body is modelled. This allows the stiffness of the vertebral body alone to be captured and modelled. The posterior elements were removed by cutting through the pedicles at the narrowest part, limiting damage to the region. This is shown in Figure 3.1.

Specimens were potted similarly to the bovine tail vertebrae in the previous chapter, using a clamp and retort stand (shown in Figure 3.2) instead of the steel rod, given the lack of a spinal canal. Vertebrae were aligned vertically, in the direction of the applied load and held inside the steel potting container. PMMA (cold cure, WHW Plastics Ltd.), was used



Figure 3.1: The removal of the posterior elements, through cutting at the narrowest part of the pedicle.

to pot the specimens ensuring 5 mm of PMMA cement covered the top and bottom of the vertebrae with little around the sides. A cement depth was used such that the vertebrae were not over constrained while making sure the entire inferior and superior endcaps were supported.



Figure 3.2: The holding of the vertebrae during potting in PMMA endcaps.

3.2.3 Imaging

All specimens were imaged under μ CT before the initial loading test and again following augmentation (described in the following chapter). The images were taken with the same scanner and settings as described in Chapter 2, Section 2.2.5.

3.2.4 Loading

The aim for the mechanical testing of the vertebrae was to determine the stiffness of the vertebrae and not to generate fractured in the vertebrae (as was carried out in the previous

chapter). The vertebrae were loaded with an initial conservative maximum load of 800 N, as has been used previously for similarly osteoporotic vertebrae [42, 44], rather than the high loads used by Pneumaticos et al. [49]. However, after loading two of the initial set of vertebrae, it was found that the stiffness continued to increase up to the maximum 800 N. Following loads up to 2000 N, it was found that the stiffness reached a maximum between 1300 and 2000 N, with three of the initial four specimens showing some degree of failure in the final 400 N of loading. Due to this failure, a maximum load of 1600 N was used as a compromise for the remaining vertebrae used, although as presented below, some failure was still observed.

3.2.5 Mechanical Data Processing

The maximum stiffness of the vertebrae was found in the same fashion as with the bovine tail vertebrae - measuring the stiffness of segments at increments over the length of the curve.

An investigation into the effect of segment size (the length of each section from which the stiffness is found, illustrated in Figure 2.5) was carried out. The aim was to optimise the segment size to find the most accurate measurement of maximum gradient while avoiding measuring the noise in the measurement. The script used in the previous chapter was rewritten in Python, allowing iteration and reporting of the maximum stiffness when using different segment sizes. An increment size (space between the start of each segment) of 1 data point was used, this corresponded to 0.0017 mm.

The segment size had large effect on the measured maximum stiffness with a range of over 1000 N/mm in the case of the two Spine 2 L1 tests. The segment size needed to be small enough to capture the maximum stiffness while avoiding the noise (seen to the right in Figure 3.3) when using a segment size below 18 data points. Hence, a value of 20 data points was chosen, a value that avoids the noise while being on the plateau of the lines.




Figure 3.3: The effect of reducing the segment size on the maximum stiffness reported from four human vertebrae loaded to 2000 N. Using an increment size of 1 data point (0.0017 mm) and segment sizes of 100 to 1 data point (0.17 mm to 0.0017 mm).

3.2.6 Image Processing

The trabecular architecture of the vertebrae was characterised in order to understand the relationship between trabecular structure and the vertebral stiffness. To examine these properties, an accurate method of determining the threshold that described the trabecular bone from μ CT scans was required. This threshold describes the limit of the greyscale representing the trabecular bone and the start of either marrow or empty space. To enable a fair comparison between vertebrae, a region of interest (ROI) was selected from the vertebral body. Given the large variation in cortical shell thickness, along with certain vertebrae containing large osteophytes and other extra bone growth, the ROI was selected to be the largest cylinder that could fit within the vertebral body while not capturing any of the cortical shell.

3.2.6.1 Histograms

In order to understand the spread of brightness values from the set of lumbar vertebrae, the histograms of the ROIs were plotted (Figure 3.4).

The lower greyscale values represent the empty space within the ROI which translates to regions where the bone marrow has drained from the trabecular bone, most likely during the unavoidable freeze thaw cycles. The peak in the histogram at an approximate greyscale value of 16 is due to the bone marrow. The remaining portion of the histogram represents bone, where variation is due to the differences in the mineral content of the bone, with




Figure 3.4: The normalised (with respect to the volume of the ROI) histogram data for the 14 lumbar vertebrae.

more mineralised bone appearing brighter on μ CT scans. The specimens ROIs that contain the brighter values are those that exhibit more degeneration for example Spine 2 L1, where the histogram is shifted to the right, suggesting many more bright pixels and therefore more bone spurs and other degeneration. This is because the extra bone growths on the degenerated vertebrae are formed of denser cortical-like material.

3.2.6.2 Threshold Optimisation

The selection of a threshold that defines the limits of the trabecular bone in terms of greyscale values was necessary for the determination of the bone volume fraction (BV/TV) and some of the computational modelling methods employed. The full resolution μ CT scan ($82 \mu\text{m}$) of each specimen was imported into imageJ (Fiji distribution, version 1.51a). The BoneJ plugin for imageJ was used for all of the trabecular structure metrics, with the optimise threshold tool being the focus here. The threshold optimisation feature works by maximising the connectivity (Conn.D) against threshold value [95]. This tool was run on the region of interest using the default settings for the connectivity options shown in Table 3.2. The resulting threshold values are shown in Table 3.3, showing some variation across the range of vertebrae. However, the differences in these suggested thresholds can be grouped with the spine, for example: Spine 1 has a range between 15 and 17, while Spine 4 has a range between 18 and 20. This promotes the idea that these differences are due to the mineralisation of the bone. This is especially salient given that Spine 1 is from a female donor and Spine 4 from a male donor, where, the BMD of post-menopausal women is greatly reduced compared to their similarly aged male counterparts [105].

The chosen value for the threshold directly impacts the reported metrics and therefore the consistent value of 18 was chosen. Using lower or higher values for the threshold usually results in the trabeculae being described as thicker or thinner respectively. With respect to modelling, the difference in the mineralisation between specimens will still be captured when using a fixed threshold, given that the binary image will be down sampled before material properties are acquired. Here, thicker binarised trabeculae correspond to a brighter continuum level voxel. The results in Table 3.5 show the BV/TV data and degree of anisotropy values for the 14 specimens using the fixed threshold of 18. Similar trends to that which was seen from the ‘optimise threshold’ results was again observed in the BV/TV data, with grouping between the spines evident. This suggests that despite the use of uniform thresholds used across the set of vertebrae, the main differences in morphology are still captured.

Table 3.2: Settings used for the ImageJ plugin, BoneJ: Optimise Threshold.

Options	Values
Tests	11
Range	0.2
Subvolume Size	256
Erosian Cycles	0
Dilation Cycles	0

Table 3.3: The suggested threshold from the optimise threshold BoneJ tool for the 14 lumbar vertebrae in the set.

Specimen	Suggested Threshold
Spine 1 L1	17
Spine 1 L2	16
Spine 1 L3	16
Spine 1 L4	15
Spine 1 L5	16
Spine 2 L1	23
Spine 3 L1	18
Spine 3 L2	19
Spine 3 L3	19
Spine 4 L1	20
Spine 4 L2	19
Spine 4 L3	18
Spine 4 L4	19
Spine 4 L5	18

3.3 Computational Methods

The methods reported in this section describe the process used to model non-augmented vertebrae, with associated sensitivity tests to understand the most appropriate approach to adopt for various aspects of the modelling process.

The methods used in this section broadly follow those used in the previous chapter, with a general process including the segmentation and FE model generation in Simpleware ScanIP, and solving the models in Abaqus using setup scripts written in Python. One difference in this chapter is that the models were run on ARC2, part of the High Performance Computing facilities at the University of Leeds. Here, the submission of jobs (simple shell scripts describing the process to be carried out) onto the cluster was carried out using the Sun Grid Engine, a ‘first come first served’ scheduler for the cluster. By running the FE simulation on the HPC cluster (using resources of 10 cores 1024 MB of memory per core), the time to solve each model was reduced to 5 - 10 minutes compared to 20 - 40 minutes on a desktop PC, depending on the model and material properties used. In addition to the time saving benefits, this allowed off loading of large (5 - 10 GB per model) Abaqus output files and through the submission system it allowed the iterative testing of material properties, where models could be re-solved using different material properties read from a simple text file.

The methodology of using ARC2, one of the HPC clusters, remained much the same as when using a desktop PC. Here, Abaqus input files were copied over to a user directory on one of the cluster nodes, along with files detailing the loading positions for the vertebrae, three Python scripts and a shell script. The first Python script set up the model, importing each Abaqus input file and applying boundary conditions, loading positions and material properties. The second Python script was used to run the models with values of the number of cores and memory usage matching those of the requested HPC resources. The models were run with an initial increment size for the step of 0.1, a minimum increment of 0.0001 and a maximum increment of 1. The final Python script was used for post processing; the Abaqus output file was read and the axial reaction force (RF3) recorded. The resulting stiffness (reaction force divided by the displacement to give a stiffness) from each model was written to a results file which could be copied to the desktop PC and analysed. Finally, the shell script described the resources to be requested from the HPC cluster, including a maximum amount of time, and then gives the command to run each of the Python scripts through Abaqus in turn: `abaqus cae nogui=setup.py && abaqus cae nogui=run.py && abaqus cae nogui=post.py`.

3.3.1 Vertebral Modelling Using the BV/TV Method

In this chapter, two methods for modelling the non-augmented vertebrae were compared. The first (the direct-greyscale method) followed a near identical procedure to that used to model non-augmented bovine tail vertebrae, described in Section 2.3. The second was a new

method named the 'bone volume fraction' method used the BV/TV to derive the elastic modulus. In this section the difference between these modelling methods are outlined with the bone volume fraction method being described in detail.

The aim of the new method was to gain greater trabecular definition by using a full resolution scan with the bone regions segmented and binarised using a threshold. This full resolution segmented scan was then down-sampled to voxels with edge length of 1 mm, meaning that each voxel has a greyscale value proportional to the BV/TV value for the region captured by that voxel. Areas that contain more bone would therefore have a higher greyscale value.

The BV/TV method followed similarly to the method used by Robson Brown et al. [60] with a few minor changes. The scans were converted from the SCANCO proprietary ISQ format into a stack of TIFF files using an in house Matlab script, which additionally reduced the 16 bit data to 8 bit images. A threshold was chosen to be a greyscale value of 18, the mean of the values described in Section 3.2.6.2, and was applied to the full resolution scan stacks in ImageJ. Once the binary stack was created, a Gaussian filter was applied, also within imageJ, with $\sigma_{x,y,z} = 1$. This was applied in order to remove speckling found surrounding the end-caps and some of the other noise visible in the scans. The two stacks of images, binary and non-binary, were then imported into ScanIP; carried out by opening one stack of images (the original stack) and then importing a second background (the binary image stack). Both backgrounds were downsampled to voxel sizes of 1 mm³, with the downsampled, thresholded stack being used to create the mask for the vertebrae and the original, non-thresholded stack being used for the segmentation of the end-caps. Greyscale material properties for the vertebral mask were taken from the binary stack, while the remainder of the methods followed the same process as the direct-greyscale method. A visual comparison of the two approaches can be seen in Figure 3.5, where the improved definition of trabecular structure and cortical shell using the bone volume fraction method can be seen in E and F, following the downsample to 1 mm³ resolution.




Figure 3.5: A comparison of the direct-greyscale method (A, B and C) and the bone volume fraction method (D, E and F). A shows the full ($82 \mu\text{m}$) resolution scan, B shows the same image downsampled to 1 mm resolution and C shows the segmented scan. D shows the segmented bone at $82 \mu\text{m}$, E shows this image downsampled to 1 mm and F shows this image after segmentation.

3.3.2 Sensitivity Studies

Sensitivity tests were carried out to identify the sensitivity to different properties, modelling methods and boundary conditions. These tests informed as to which properties gave the best results with regard to the experimental-computational agreement, whether the error associated to a certain property is comparable to other sources of error and gave a measure of how robust the models were.

As with the models in the previous chapter, the commonly carried out mesh convergence studies are not carried out here due, in part, to the greyscale optimisation process used to find the material properties for the bone. There are two changes taking place with a changing mesh size: geometric changes due to the finer mesh capturing more of the geometric features and changes due to the element size altering the calibration coefficient. Therefore, the results of a mesh size optimisation become difficult to unpick. As described previously (Sections 1.5.1 and 2.3.5.1, [69]), the optimum mesh size of $1 \times 1 \times 1 \text{ mm}^3$ elements has been used.

The models used in the sensitivity of the endcap depth were two vertebrae, that varied considerably in their geometry. All 14 vertebral models were used for the load position sensitivity, where the test cases were quick to automate and run.

3.3.2.1 End-cap Depth & Contact sensitivity

In the experimental tests, the cement endcaps are merely moulded to the shape of the top and bottom of the vertebral body allowing slight motion of the vertebrae within them. However, previous FE models used tied contacts between the bone and endcaps, removing any motion between the two and locking neighbouring nodes together. From an initial model output, it was observed that there were high surface stresses at the edges of the

cement endcaps, at both the top and bottom of these regions, shown in Figure 3.6. The depth of the endcaps experimentally also varied between 5 mm and 20 mm depending on the shape of the vertebral endplates, with the more concave vertebral endplates requiring more cement to fill the space, forming a flat cement surface. Therefore, tests were carried out to identify the effect of cement endcap depth and to optimise the contact properties between the bone and PMMA endcaps to find the closer fit to the experiment.




Figure 3.6: The effect of having tied contacts between the endcaps and the bone, with the increases in Von Mises stress indicated.

Different depths of ± 1 , ± 2 and ± 3 mm for the top and bottom endcaps independently and together were tested. Additionally, the effect of changing the cement endcap depth with the currently used tied contact to a frictionless contact was also investigated, along with variations to the properties of the frictionless contacts.

The endcap depth was altered by adding and removing layers within the model file in ScanIP. The endcap contact conditions were changed within the Abaqus Python scripts, by creating a small sliding, node-to-surface contact with frictionless tangential and “hard” normal behaviour. It was found the allowing separation of the components after contact had been established prevented the model solving in some cases, but this option was found to have minimal effect so was not selected. These tests were carried out on two vertebrae, Spine 1 L2 and Spine 4 L5, which were used due to the large shape difference between the L2 and L5 vertebrae.

3.3.2.2 Load Position Sensitivity

The effect of loading the vertebral models at different positions across the superior endcap was tested. This developed the understanding of error in choosing the load position for the model (aiming to match the experimental loading point as closely as possible) and to understand how the vertebrae respond to different loading positions. Sixteen different positions were tested, shown in Figure 3.7. These positions included 1 mm and 2 mm away from the centre in the posterior, anterior, left and right directions along with larger deviations of 10 mm and 20 mm from the centre in the same directions. A 1 mm

axial displacement was applied at the different positions as was used previously, with the remaining boundary conditions kept constant (using the frictionless endcap contacts described above). The two magnitudes of distances from the centre tested, 1 and 2 mm, and 10 and 20 mm, were used to test the error due to the 1 mm resolution and to understand the response to non-central loading respectively.




Figure 3.7: The loading positions used for the load variation tests. Position X is the original central loading position for ap and lr loading, other loading positions are 1 mm, 2 mm, 10 mm and 20 mm away from the centre.

3.3.3 Material Properties

The general method for obtaining material properties for the bone elements in the FE models uses the same approach as used in Section 2.3.4 where the Young's modulus of each element is equal to the greyscale value multiplied by a conversion factor. The conversion factor was optimised by comparing the computational stiffness against the experimentally acquired stiffness. For the optimisation process, the vertebrae were split into two groups at random, giving a calibration and validation set, allowing a test of the validity of the greyscale conversion factor. The split of the vertebrae for these two sets can be seen in Table 3.4, where the greyscale conversion factor was derived from the calibration set and applied to the validation set where the agreement was measured.

Results from the initial optimisation process using the calibration and validation processes gave good results: the agreement of the validation set had CCC values of above 0.5.

Table 3.4: The vertebrae used for the calibration set and the validation set in the optimisation process.

Calibration Set	Validation Set
Spine 1 L4	Spine 1 L1
Spine 1 L5	Spine 1 L2
Spine 2 L1	Spine 1 L3
Spine 3 L2	Spine 3 L1
Spine 4 L1	Spine 3 L3
Spine 4 L2	Spine 4 L3
Spine 4 L4	Spine 4 L5

The validation set for the direct-greyscale method had a CCC value of 0.591 and the CCC value for the validate set from the “bone volume fraction” method gave a value of 0.831. Confirming again that model creation and the method of material derivation was valid.

Following the results of the sensitivity tests, the final value for the greyscale conversion factor was $0.0009625 \text{ GPa}^{-1}$ for the BV/TV method. This was used for the bone properties for all final non-augmented models. The optimised greyscale conversion factor for the models generated using the direct-greyscale method used a value of 0.00416 GPa^{-1} .

The remaining material properties used for all of the models are a Poisson’s ration of 0.3 for the bone and endcaps and a Young’s modulus of 2.45 GPa for the endcaps. Boundary conditions for the models included the frictionless contact between the vertebrae and their endcaps and a tied contact between the analytical platen and the superior endcap. A displacement of 1 mm was applied to all of the models.

3.3.4 Vertebral Density Visualisation

To visualise the density distribution of the trabecular bone within the FE models, Simpleware ScanIP was used. Here, FE models were generated and the clip box was used to remove an axial half of the model from the render. Mass density material properties were then assigned to the model within the visibility options, given that material properties for the region had been set to greyscale based. This provided colour maps of the density of the vertebral mid-slice.

3.4 Results

The results described in this section are split into two main sections: those derived and obtained through experimental means and those measured from the computational, generated models.

3.4.1 Experimental Results

Material properties for the vertebrae can be seen in Table 3.5, with grouping evident within spines for the bone volume fraction measurements. For example, the mean BV/TV value for Spine 1 was 0.159 ± 0.026 , while the mean for Spine 4 was 0.248 ± 0.022 . A relationship was found between the bone volume fraction and the degree of anisotropy, with a correlation coefficient of $r = 0.74$, this is presented in Figure 3.8. Despite this, less grouping within spines could be seen in the degree of anisotropy for the vertebrae.

The vertebral stiffness following experimental loading is presented in the following chapter in Table 4.2, with similar correlations between stiffness and density.

Table 3.5: BV/TV and degree of anisotropy (DA) values found using the ImageJ plugin BoneJ tools Volume Fraction and Anisotropy respectively using a threshold of 18.

Specimen	BV/TV	DA
Spine 1 L1	0.174	0.387
Spine 1 L2	0.17	0.364
Spine 1 L3	0.137	0.359
Spine 1 L4	0.127	0.418
Spine 1 L5	0.187	0.341
Spine 2 L1	0.391	0.199
Spine 3 L1	0.255	0.327
Spine 3 L2	0.267	0.231
Spine 3 L3	0.281	0.301
Spine 4 L1	0.257	0.239
Spine 4 L2	0.241	0.338
Spine 4 L3	0.244	0.348
Spine 4 L4	0.247	0.371
Spine 4 L5	0.249	0.186

Table 3.6: The experimental stiffness results.

Specimen	Non-augmented Stiffness (N/mm)
Spine 1 L1	2991
Spine 1 L2	3456
Spine 1 L3	3244
Spine 1 L4	3223
Spine 1 L5	2891
Spine 2 L1	6149
Spine 3 L1	5153
Spine 3 L2	5357
Spine 3 L3	5338
Spine 4 L1	3277
Spine 4 L2	5064
Spine 4 L3	6098
Spine 4 L4	4957
Spine 4 L5	7185




Figure 3.8: The relationship between the bone volume fraction and the degree of anisotropy with a correlation coefficient of 0.74.

3.4.2 Computational Results

The results of modelling the non-augmented vertebrae are presented in Figure 3.9, comparing the initial direct-greyscale method (the method used with the bovine tail vertebra in the previous chapter) with the best case using the bone volume fraction method and the results of the sensitivity test described previously. The results show a large improvement using the results of the sensitivity tests and the bone volume fraction method, with a concordance correlation coefficient improving from 0.55 to 0.86.

The effect of using the bone volume fraction method on the greyscale background and colour map of the bone density can be seen in Figure 3.10, where the much more clearly defined cortical shell and trabecular structure is evident.

The varying greyscale distribution and differences between the different spines can be seen in Figure 3.11. The general trends match the overall bone volume fraction values described earlier in Table 3.5, where the vertebrae from Spine 1 show much lower bone volume fractions which matches the much more blue heat maps when compared to the denser Spines 2, 3 and 4.




Figure 3.9: The results of using the direct-greyscale method and the bone volume fraction method to model non-augmented human lumbar vertebrae. Red shows the agreement using the direct-greyscale method and blue shows the agreement using the bone volume fraction method, with the orange dotted line showing perfect agreement.




Figure 3.10: A and B: The resultant 1 mm³ greyscale background from the direct-greyscale and bone volume fraction method respectively, where the brightness in A has been increased for visibility. C and D: a heat map of the greyscale values through the mid slice of models from the direct-greyscale and bone volume fraction method respectively.




Figure 3.11: The variation in the greyscale distribution across the mid-slice of the vertebrae from Spine 1, 2, 3 and 4. The general density changes as well as shifting density distributions are visible.

3.4.3 Sensitivity Study Results

3.4.3.1 Loading Position Sensitivity

The results of the load position sensitivity tests are shown in Figures 3.12 and 3.13. Generally, loads a small distance to the posterior (above denser bone) resulted in higher stiffness values while all anterior loads resulted in a reduced stiffness (Figure 3.12). Loads lateral to the central loading point (Figure 3.13) resulted in little mean change to the stiffness of the models for the 1 and 2 mm positions. Larger movements away from the central loading position (10 mm and 20 mm) resulted in almost all vertebrae presenting a reduction in stiffness. Loading positions at 20 mm left and right presented the largest reductions, where in the natural body the expected loads are less compared to anterior/posterior loading. Additionally, more anterior loads resulted in larger reductions in the stiffness, given the less dense nature of anterior portions of the vertebral body.

These results also show that the effect of movements of 1 mm away from the experimental loading position only changed the stiffness by approximately 5 %. Given the resolution of the models is limited to 1 mm, the accuracy of the selection of the experimental loading positions is also limited to 1 mm. The maximum 5 % change to the stiffness may describe some of the error seen when comparing the experimental and computational results. However, as discussed in more detail in Section 3.5.3.2, these errors may be considerably smaller than the errors in the acquisition of the experimental stiffness from load-displacement data.

3.4.3.2 End-cap Depth & Contact Sensitivity

The results of the end-cap depth and contact sensitivity tests are shown in Figures 3.14 and 3.15. It was found that regardless of the contact, increasing the depth of the endcaps increased the stiffness of the vertebrae, while reducing the depth reduces the measured stiffness. The effect this extra or reduced constraint caused by the changing endcap depth broadly followed the same trends for the two vertebrae tested. However, the response did differ most notably at the extremes of endcap depths (± 3 mm), where Spine 4 L5 experienced the greater reduction in stiffness when reducing the endcap depth, while the Spine 1 L2 vertebra saw the greater increase following depth increase. This difference is a consequence of the differences in the bone density, Spine 4 being denser, and the shape, Spine 4 L5 being much wider with greater protruding inferior bone spurs compared to the narrower L2 vertebrae with superior protruding bone spurs. The variation in the change to the stiffness within ± 1 mm change in depth was up to 5 %, for both cases. Given that the maximum error in masking the cement endcaps is 0.5 mm due to the resolution that




Figure 3.12: The effect of the loading position for the human lumbar vertebrae, shown as a percentage change compared to a central loading position, for load positions from the posterior to anterior according to Figure 3.7

is used for the segmentation, it suggests that this error has a minimal effect on the end results. As mentioned previously, this is due to the other errors associated with measuring the experimental stiffness.

The effect of changing the contact from tied to frictionless can be seen by comparing Figures 3.14 and 3.15. Excluding the very large change in stiffness seen in the L5 vertebra when reducing the depth of the inferior endcap, the adoption of the frictionless contact reduces the effect of endcap depth on measured stiffness. This is due to a reduction in the constraining nature of the frictionless contact, allowing a more natural measure of the vertebral stiffness rather than a combination of the vertebral and endcap combined stiffness. The anomalous L5 vertebrae here is a function of having less support and the slipping that this allowed, resulting in the reduced measured stiffness. Optimising the greyscale conversion factor for the set of models with a tied constraint and a frictionless contact gave an improved agreement with a concordance correlation coefficients of 0.848 and 0.865 respectively. Therefore the frictionless contact provided a small improvement in the agreement between the computational and experimental results.




Figure 3.13: The effect of the loading position for the human lumbar vertebrae, shown as a percentage change compared to a central loading position, for load positions from the left to the right according to Figure 3.7.




Figure 3.14: The change in stiffness following a change to the thickness of (a) the inferior, (b) the superior and (c) both cement endcaps. In each case, the effect of $\pm 1, 2$ and 3 mm change in thickness is shown. The endcaps had a tied contact to the vertebrae.




Figure 3.15: The change in stiffness following a change to the thickness of (a) the inferior, (b) the superior and (c) both cement endcaps. In each case, the effect of $\pm 1, 2$ and 3 mm change in thickness is shown. The endcaps have a frictionless contact to the vertebrae.

3.5 Discussion

The outcomes of this human vertebrae study are discussed in the following subsection. It is split into three, discussing the experimental results, limitations and finally the results of the computational aspects of the study.

3.5.1 Experimental Results

Experimentally there was large variation seen in the results, in terms of material properties (BV/TV values), geometry and response to loading. The strong relationship between the bone density and the degree of anisotropy is explained by the thinning and dispersal of the horizontal struts between trabeculae with age [29, 30, 106]. This means that with the age related, general density loss of the bone, the loss of horizontal trabecular struts results in increased anisotropy values. The experimental response to loading showed a large range in the stiffness values of the vertebrae, even for vertebrae with similar volumes and for vertebrae originating from the same spine. The variation seen in the intact vertebrae aids the justification for the study: if such a wide range of vertebral properties exist within such a small data set, then it explains why such variation is found clinically when identifying the response to vertebroplasty. For example, in some of the randomised clinical trials and in studies of the epidemiology of vertebral fractures, no statistical significance in the output measures could be found due to the large standard deviations found in measurements [3, 4, 16].

The repeatability of the individual specimens is somewhat of an unknown, given the requirement of a single loading step to avoid the accumulation of damage prior to augmentation. An attempt to understand the repeatability is described in the following chapter using repeated loads of augmented vertebrae.

3.5.2 Limitations Regarding Specimen Numbers & Vertebral Features

One of the main limitations to the study is the numbers of vertebrae included in the set. Having four human spines with a limited number of lumbar vertebrae available from each spine, meant that only 14 out of a possible 20 vertebrae were able to be tested. This meant that there was a skew towards the top of the lumbar section with limited numbers of L4 and L5 vertebrae. Such skews to the data set have effects to any statistics applied to the data set and will be discussed more thoroughly in the chapter on Statistical Shape and Appearance modelling, Section 5.7. More relevant to modelling the vertebrae is the effect this has on the determination and optimisation of the greyscale conversion factor. The large change in geometry from L4 to L5, rather than the gradual change between L1 to L4,

introduces questions about the response to loading (and augmentation) for this vertebra in particular and how it differs from other levels. A resolution to this limitation would be to increase the specimen numbers and look at vertebral levels in isolation; identifying level specific differences and properties. Variation at different vertebral levels may require level specific greyscale conversion factors, such as those required when modelling different animal species [91]. However, within the current set of vertebrae greater variation was found between spine rather than level (for BV/TV, stiffness and vertebral body volume), even when including L5 vertebrae. This suggests that segregation of vertebral levels is unnecessary for the current specimen set.

Another limitation of the study is the age range and limited information regarding vertebral health. Because of the unknown vertebral quality or health (other than measurable trabecular bone properties), the initial load to failure carried out in the bovine tail study was deemed not appropriate. While no evidence of fractures could be seen in the μ CT scans, other than slight wedge shapes in some specimens, there was also no evidence of the bovine tail vertebrae having succumbed to fracture in μ CT either, despite clear fractures on load displacement curves. Since the state of the vertebrae was relatively unknown in their pre-augmented state, the study was therefore effectively an investigation into the effects of prophylactic vertebroplasty similar to a number of previous studies [44, 96, 107].

3.5.3 Computational Results

Many developments in the computational methods have been made through the course of this chapter, with the resulting improvements showing their importance. The significance of these developments and the results are discussed in the following subsection, split into a discussion of the improvements to modelling non-augmented vertebrae and the results of various sensitivity studies.

3.5.3.1 Non-augmented Models

The improvement in the agreement between computational stiffness and the experimental stiffness results using the bone volume fraction method and the results of the sensitivity tests is significant and provided a good foundation for modelling the more challenging augmented models. The use of the bone volume fraction method in isolation (without the results of the other sensitivity tests) showed a large improvement over the direct-greyscale method. The direct-greyscale method performed similarly on the bovine tail vertebrae as with the human lumbar vertebrae, with a $CCC = 0.49$ for bovine tail vertebrae and $CCC = 0.55$ for the human vertebrae.

The small improved agreement of modelling the human vertebrae over the bovine vertebrae when using the same direct-greyscale method suggests that either the human vertebrae are more receptive to the modelling method used, or that the experimental procedure was more repeatable and therefore provided more accurate results. One possibility is that the wider human vertebral body provided an easier subject to load axially. Slight off axis loads in the bovine tail vertebrae may have caused the vertebrae to not be seated in the endcaps correctly. This sensitivity to endcap contact properties was shown earlier and given the tendency of the bovine vertebrae to be loaded non-axially (due to their narrow columnar shape) they may be more sensitive to such endcap contact properties. So, while experimentally the bovine vertebrae were able to slip in the endcaps, computationally the two materials were tied, restricting this ability and reducing the agreement with the experiment. Conversely, the wider human vertebrae had less of a tendency to slip and were less sensitive to the endcap contact. This change to the endcap contacts seems to be a common transition from similar studies where tied contacts were used [42] to more recent studies [66, 79], however the property is rarely presented in detail in the literature.

The improvement using the bone volume fraction method is most likely due to the added definition of the trabecular bone and cortical shell, especially given how important the correct representation of load sharing is for accurate models [108]. This agreement is much stronger than the agreement found in similar studies that used a comparable methodology to the direct-greyscale method [42, 60, 91] and comparable to methods that used more complex although individual, specimen specific material properties for each model [64, 66]. An advantage the current study has over these latter studies is a uniform material property relationship, that once calibrated over a large set of vertebrae can be used on any additional human lumbar vertebrae that have not been used in the calibration. It also allows the use of scans from any μ CT scanner, negating the need for scanner calibration, provided a satisfactory method for separating the bone from the background exists. Having greater definition of the cortical shell and trabecular bone alignment allows much more accurate descriptions of the load transfer through the vertebrae, rather than the more simple description with the more homogenised material properties of the direct-greyscale method.

This method also removed any effect of the bone marrow, which, due to the freeze/thaw cycles, had leaked from the vertebrae in some regions. Due to the initial segmentation and binarisation at $82 \mu\text{m}$ using the BV/TV method, the bone marrow was not present in the segmented full resolution or down-sampled scans and therefore had no impact on the material properties of the models. Conversely, using the direct-greyscale method, areas

lacking bone marrow would have a lower Young's modulus and areas with bone marrow would have higher Young's modulus, not corresponding to the actual properties of these regions and inaccurately changing the load distribution.

One of the limitations of the bone volume fraction method is that it relies on having trabecular resolution scans of the vertebrae to apply the trabecular bone thresholds to. This is a similar limitation to studies that also achieved a good agreement between computational and experimental results of augmented stiffness [66], where high resolution μ CT scans were required for the acquisition of BV/TV and fabric information for their methodology. This limits the origin of the scans to *in vitro* scans rather than *in vivo* scans where the standard clinical resolution is limited to approximately 1 mm and larger. Given the intention to generate larger populations of vertebra models for use in statistical shape modelling, this limitation on the scan origin is an unwanted consequence. A possible solution to this problem would be to use 16 bit images rather than the currently used 8 bit images. This change gives an increase in the number of greyscale values available from the 256 of 8-bit images to over 65,000 for 16-bit images. The main benefit of the bone volume fraction method is the increased definition of bone structure and a greater contrast between the brighter, denser cortical bone, the trabecular bone and the empty, non-bone areas. Using 16-bit images along with image manipulation to change the contrast of the scans may allow the added detail and contrast to be captured with relatively low (≥ 1 mm) resolution scans, therefore greatly increasing the number of vertebral specimens available for use. An example of how this could be achieved is shown in Figure 3.16, where the background from the direct-greyscale method has been edited through adjusting the contrast to give a similar background to that achieved using the bone volume fraction method. The detail in the centre of the vertebral body is less clear in this example, however, it was achieved through simple contrast manipulation of the 8 bit TIFF stack, meaning that more trabecular detail could be preserved using a curves filter (remapping of the image tonality) and/or as already proposed, 16 bit images.

3.5.3.2 Sensitivity Test Significance and Influence on Errors

Errors in the measurement of the maximum stiffness exist due to the cut off at 1600 N in the experimental loading. This meant that for some of the vertebrae the maximum stiffness was not reached and therefore not captured in the load displacement curve. A prediction of the failure stress and an estimation of the failure load through simulation of a vertebra prior to experimental loading could help to remove this problem. However, this would have required yet another set of human vertebrae.




Figure 3.16: An example showing the ability to create a greyscale background similar to that achieved using the BV/TV method, but instead adjusting the contrast of the background from the direct-greyscale method.

Changing the endcap depth presents an interesting study into the effects of over constraining the models. For the models with the tied constraints changing the endcap depth had a different effect for the two vertebrae (although the same overall trend, where thicker endcaps led to more constraint and therefore higher stiffness values). The two spines tested showed different responses to exposing and covering more of the bone with smaller and larger endcap depths. The large reduction in the stiffness of Spine 4 L5, when removing 3 mm of endcap material from the inferior endcap, of approximately 20 % is caused by exposing a less dense area of bone, which, when no longer supported by the tied endcap resulted in the reduced stiffness. Other than this specific effect the general trend was an increased sensitivity to increasing the cement depth, with a much smaller sensitivity to exposing more bone. This suggests that in any further testing, care must be taken to not over-constrain the vertebrae experimentally by using thick endcaps. Over-constraining the vertebrae experimentally was already reduced through the removal of the posterior elements.

Changing the endcap depths when the frictionless contact was used resulted in (for the most part) much less sensitivity to changing the endcap depth, when compared to the tied endcap models. The exception is the Spine 4 L5 vertebra, where, when the endcap depth is reduced a very large decrease in the stiffness is seen. This is caused by slipping that occurs between the vertebrae and the endcap, given the lack of both physical constraint and constraint in the form of a boundary condition. An illustration of this can be seen in Figure 3.17, showing that with such little constraint the anterior inferior of the vertebral body moves further to the right (anterior). This results in the reduced stiffness recorded. However, such small endcaps were not created experimentally and therefore create no additional errors to the models. Importantly, the use of frictionless contact with the “no separation after contact” option set results in smaller sensitivity to endcap depth modelling errors and resulted in increased agreement between experiment and model for the whole

set of 14 vertebrae.




Figure 3.17: An illustration of the slipping that occurs in Spine 4 L5 when the endcap depth is reduced by 3 mm on the superior and inferior endcaps

The effect of changing the loading position aided the understanding of the mechanics of the models and their response to loading. Experimenting with different loading positions allowed development of our understanding of the response to different natural loading conditions that would be experienced during daily life. This could help to answer the important question of how vertebroplasty affects adjacent level loading and how it may affect the chances of adjacent level fracture. These investigations were carried out with statistical shape modelling and reported in Chapter 5, where the response to loading following augmentation was investigated and, given the more systematic nature of the generated models, allowed for an easier interpretation of the results. Here, the response to loading outside of the ± 1 mm error range follows a mostly expected pattern with loads anterior, posterior, left and right at ≥ 10 mm away from the central loading position resulting in progressively reducing stiffness values. The reduced material beneath the axial loading point explains most of these results, with the greater reduction when loaded anteriorly being due to the generally reduced density of the anterior portion. The smaller reduction in the stiffness when loaded posteriorly is due to the increased density of the posterior region of most of the vertebrae and was seen in the greyscale density maps of the set. The anomalous result is the Spine 3 L2 vertebra, where posterior loading of 10 mm behind the central loading position resulted in an uncharacteristic increase in the measured stiffness. The Spine 3 L2 vertebrae had particularly degenerated superior and inferior endplates, protruding anteriorly at both ends, but with little protrusion posteriorly. This meant that when measuring the central loading position experimentally, a loading point was selected further anterior than what would have been picked without the bone growths. This explains why moving the loading position further to the posterior resulted in an increased stiffness with respect to the central loading point, given that the experimentally placed central loading point was incorrectly placed anteriorly. This presents another reason for carrying out load position tests on the generated models in Chapter 5, where central loading can be ensured and the requirement to match the experimental loading position is

removed.

3.5.3.3 Summary of Computational Results

The use of the bone volume fraction method has clearly provided a more accurate method for the continued modelling of human lumbar vertebrae, providing a much stronger agreement with the experimental results. This methodology will be used in to the following chapter. The results of the sensitivity studies suggest that a frictionless contact between bone and endcaps should be used in subsequent models, again, due to the improved agreement with the experimental results.

Chapter 4

Augmented Human Lumbar Vertebrae Study

4.1 Introduction

This chapter describes the experimental augmentation process and the creation of FE models of these augmented vertebrae, building on the methodology and results of the previous chapter. The fourteen lumbar vertebrae used in the previous chapter were used again in this chapter, undergoing augmentation and testing prior to modelling. This process continued the development of methodology from Chapter 2, adapting to the new challenges created by the human tissue and improving it through performing sensitivity tests.

Presented in this chapter is the augmentation process and the methodology developed to model the augmentation. Experimental results are described along with computational results from the models of augmentation and the sensitivity studies. Finally, the results of the chapter are discussed.

4.2 Experimental Methods

This section describes the process of experimentally augmenting fourteen human lumbar vertebrae. This methodology is adjoined with an investigation into the sensitivity to repeated loading and the errors associated.

4.2.1 Vertebroplasty

The methods for augmenting human vertebrae differed from the methods developed for bovine tail vertebrae due to the different geometry, size and density. Changes include the

method and route of needle insertion, and the type of needle used; these are described below.

The first change to the procedure was that the human vertebrae, being much less dense, did not require the vertebroplasty needle to be inserted with the aid of a mallet. Instead the needle could be pushed by hand through the cortical shell and into the vertebral body, with the vertebra being held in a vice to ensure the correct angle and depth was achieved.

An oblique approach was used, as shown in Figures 4.1 and 4.2. This avoided damage to the pedicle-canal region, especially in L1 and L2 where the pedicles are at their narrowest. Bi-lateral vertebroplasty was used to ensure the maximum fill volume possible.

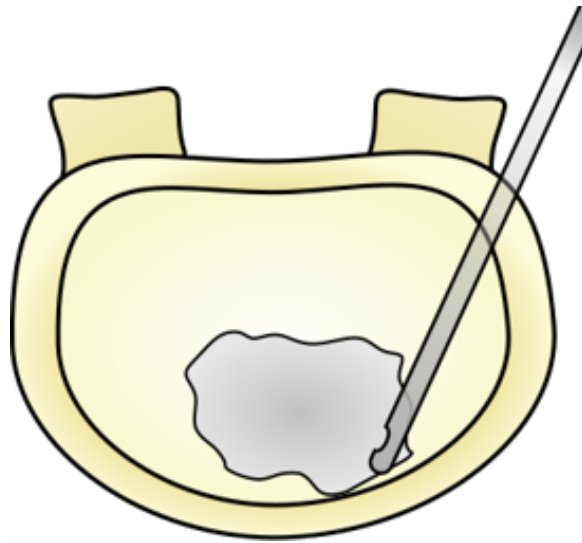


Figure 4.1: An illustration showing the approach to needle insertion and cement position.




Figure 4.2: A μ CT scan showing the injected volume of cement at the anterior of the vertebral body with the cement track from the exiting needle to the right. For this specimen cement leaked through the anterior wall, limiting the quantity of cement injected.

Due to the difficulties encountered in injecting large volumes with an end tipped needle

(Chapter 2), a side tipped needle was used, with the cement directed into the anterior-centre region.

It was attempted to inject the largest possible volumes of cement into the vertebrae, and the injection was stopped upon cement leakage, mimicking the clinical process. Leakage generally occurred through one of the vascular channels at either the posterior or anterior portions of the vertebrae. Cement loss was observed to continue occurring after the injection was stopped due to pressure build up. In addition to this, cement was also lost in the needle itself, where, as the cement set, the amount remaining in the needle became difficult to measure. This meant that the cement fill could not be accurately measured though the amount of PMMA leaving the syringe. Hence the cement fill volume was measured using segmented μ CT scans, using the volumes of the masks created to define the injected cement regions. This methodology is described in the computational modelling section (3.3).

4.2.2 Loading and Repeated Loading Tests

The same 1600 N load was applied following augmentation as was applied to the vertebrae in their non-augmented state.

Single loads were applied to the vertebrae pre-augmentation, in an attempt to limit damage to the vertebrae, however this meant there was no direct measure of whether the changes seen post-augmentation were due to the vertebroplasty procedure or simply due to the effect of repeated loading. To attempt to understand this potential error, vertebrae having undergone augmentation were tested three more times in an iterative fashion, removing each from the load testing machine, testing the next specimen in the set and repeating. Removing the vertebrae from their steel housing between tests allowed error in loading position and setup to be tested along with repeated loading of the vertebrae.

The results of repeated loading can be seen in Figure 4.3. The majority of specimens show a reduced stiffness for the repeated loads following the initial load, for which there are a few possible reasons. The reduction in stiffness could be due to damage being caused during the initial load to 1600 N. In Figure 4.4 it is possible to see slight failure in the three Spine 3 vertebrae, although failure cannot be seen in the Spine 2 L1 vertebrae, which also exhibits a reduction in stiffness in the repeats. An option that may explain why repeats 2 and 3 often share a similar stiffness compared to repeat 1, is that the vertebrae were seated differently in their endcaps for the first repeat, settling into position for repeats 2 and 3. This trend of a different stiffness for repeat 1 compared to 2 and 3 is evident in Spine 1 L3 and L4, Spine 2 L1 and Spine 4 L1, L2 and L3. Another potential cause is the

changing temperatures (frozen to room temperature to above room temperature during the μ CT scan) causing the vertebrae to expand out of the endcap in some cases. However, if freeze thaw cycle damage or the change in temperature was the cause of the variability in repeated loading then more consistent changes would be expected.

The iterative reduction in stiffness for the Spine 3 L2 vertebrae can be explained as damage being caused after each iteration. This can be seen in Figure 4.5, with the three repeats each showing a yielding before the 1600 N limit and a smaller maximum load and stiffness after each repeat.

The figures in Figures 4.4 and 4.5 show the data for the loading, from which the maximum stiffness values are found. This excludes the initial cyclic loading, and starts the loading at 50 N and the displacement at 0 mm. It is most likely that the changes to the stiffness during the repeated loads of the vertebrae are due to a combination of the effects described above. However, these results form a baseline of understanding when comparing the pre and post augmentation vertebral stiffness measurements.




Figure 4.3: The stiffness of the augmented vertebral specimens over the course of an initial load and three repeated loads.




Figure 4.4: The load - displacement results for the initial load for all used vertebra




Figure 4.5: The load - displacement results for the Spine 3 L2 vertebra. Showing results of the intact load and post augmentation load as well as the three repeats.

4.3 Computational Methods

This section describes the methods used to model augmentation in the vertebral models through new and novel methods. Additionally, methods of measuring the vertebral body volume and visualising augmentation are presented.

4.3.1 Modelling Augmentation

To model the augmented region in the vertebral models, a range of different material properties and approaches were investigated. In addition to varying the material properties of the injected volumes of cement, the effect of using a combination of scans from before and after augmentation was also investigated. To achieve this, scans of pre/post augmentation required registration. The registration method was developed for use with the BV/TV method described earlier, where the non-augmented full resolution binary scan and full resolution augmented scan are required to be aligned.

4.3.1.1 Registration of μ CT Scans

Artificial brightening of regions surrounding the internal cement (specifically the concentrated volumes of barium sulphate and the artefacts associated) affect the material properties that are applied to the material, seen in Figure 4.6:B. Given that the material properties are based on the greyscale background, this has the effect of artificially stiffening or changing the properties of elements in these regions. To negate this effect, the intact and post-augmentation models were registered, translating them into the same spacial location. This allowed the cement to be defined, masked and modelled based on the post-augmentation scan and the remaining vertebra was defined from the non-augmented scan, using same background used for the non-augmented models.

Registration of the images was carried out in 3D Slicer (version 4.10, [109]), using scans of the segmented, full-resolution, non-augmented vertebrae and the full-resolution augmented vertebral scans. The method of registration was a landmark-based approach using three landmarks for each vertebra, translating the scan of the augmented vertebrae to that of the non-augmented vertebra. These points were selected at the most superior point of one pedicle, the most inferior point of the other pedicle and the most inferior and anterior point of the vertebral body. The selection of these points proved to provide a repeatable registration of the vertebrae, without selecting superfluous landmarks. Once registered, the resulting translated augmented vertebrae scan was cropped with respect to the non-augmented scan to provide an aligned and registered pair of scans.

4.3.1.2 Registered Scan Model Generation

The augmented models (using the registered scans) were created in ScanIP, following similar methods to the bone volume fraction method described in Section 3.3.1. The initial step was to import the registered augmented scans and apply an initial downsample to a resolution of 0.5 mm³. This resolution provided a compromise of definition and




Figure 4.6: An illustration of the registration process, registering the non-augmented scan (A) with the augmented scan (B) and showing the combined image in C. The images are at 1 mm resolution, showing the artifacts created by the barium sulphate when viewed at this resolution, while to increase accuracy the registration carried out in Slicer 3D was at the full $82 \mu\text{m}$ resolution.

computational cost to segment the scans, with no benefit gained from segmenting at full resolution prior to the downsample to 1mm.

To segment the three regions (background, bone and endcaps, and injected cement) it was found that the most reproducible and accurate method was to use the Otsu thresholding process [110], a filter that is built in to ScanIP. It was used to find four masks in the scan and works on a threshold clustering method. In this case searching for four clusters or masks separated by different thresholds and returning a mask that described the background, bone, endcaps and the cement region. From this, only the cement mask was retained. The cement mask and augmented scan background were then downsampled to 1 mm resolution and cropped to the same dimensions as the non-augmented model. From the non-augmented model, the masks describing the vertebra and endcaps were imported, along with the binarised background of the vertebra. The origin of the masks coming from the two backgrounds can be seen in Figure 4.7.

The track left behind by the vertebroplasty needle was clearly visible on the post-augmentation μCT scan, however it was not explicitly modelled in the previous models of augmentation. While experimentally it was attempted to fill this space with cement when removing the needle, in some cases it was not possible or the needle track was only partially filled. An example of the track left behind can be seen in Figure 4.8:B, with comparisons to the BV/TV background in A, and the correction, through removal of the




Figure 4.7: An illustration showing the origin of the masks from the non-augmented scan (A), the augmented scan (B) and the combination in C.

mask in C. Removal of the vertebral mask in the regions where the needle track was clear in the post-augmentation scan was carried out using the (un)paint line tool in Simpleware ScanIP. This change to the models required further changes to the material properties of the cement and interface regions as described in Section 4.3.1.3.



Figure 4.8: Adding the needle track to registered models. A, shows the masks overlaid onto the BV/TV greyscale background, B shows the masks overlaid onto the augmented scan background and C shows the removal of the mask where the cement track is visible. The pink mask shows the cement region, yellow shows the interface and red is the vertebral mask.

4.3.1.3 Cement Volume Material Property Tests

Material properties for the augmented region and the yielding interface required optimisation from those used in Chapter 2, given the changes following the results of the sensitivity tests. These material properties were tested in an iterative fashion changing the properties in the material properties file read by the Python setup script. Following the results of the above tests (explicit modelling of the needle track and registration methods), the final

material properties for the cement and interface regions are reported in Table 4.1 for the three main methods tested.

Table 4.1: The material properties for the injected cement volumes and its surrounding interface region for the three main modelling methods used. P.P., indicates perfectly plastic

Property	Modelling Method		
	Non-registered	Registered	Registered With Needle Tracks
Cement Young's Modulus	1.7 GPa	1.2 GPa	1.5 GPa
Cement Poisson's Ratio	0.4	0.4	0.4
Interface Young's Modulus	0.005 GPa	0.008 GPa	0.01 GPa
Interface Poisson's Ratio	0.4	0.4	0.4
Interface Yield Stress	50 Pa \rightarrow P.P.	5 Pa \rightarrow P.P.	5 Pa \rightarrow P.P.

Boundary conditions for the cement region include: tie contacts between the vertebrae and interface, and between the interface and the cement volume. The remaining material properties and boundary conditions remain the same as with the non-augmented models, described in Section 3.3.3.

4.3.2 Measuring the Vertebral Body & Augmentation Volume

While the volume of the vertebral mask can be measured directly within ScanIP or Abaqus, this volume includes the remainder of the pedicles and other bone growths and spurs. The size of the pedicles that remain after their removal during dissection varies between level and spine and is especially different with the L5 vertebrae, where the shape and size vary dramatically. To account for this variation and to get a more accurate cement fill percentage for the augmented models, elliptic cylinders were fitted to the vertebral body as shown in Figure 4.9.




Figure 4.9: A diagram showing the fitting of the cylinders to the vertebral body, where the point cloud describing the vertebra comes from an STL file, the red ring shows the points included in the mid-slice and the central red dot indicates the centre of the mid-slice.

This was carried out using a set of scripts written in Python and utilising surface (stl) files describing the vertebral geometry. The axes are defined in Figure 4.9:B. First, the middle

slice of a set of nodes that describe the vertebra was found. From the middle slice, the shape of the ellipse was formed. Given that the alignment of the vertebrae is uniform, the mid point on the y axis was found by searching for the two nodes at the extremes of the x axis, forming a line through the centre of the model, parallel to the x axis. This gives the “ a ” measure of the ellipse description. The “ b ” measure is found by using a similar process - searching for the node furthest away from the centre on the “ x ” axis, forming a line through the centre of the model, parallel to the y axis. The “ a ” and “ b ” measures gave the shape of the ellipse and the height of the elliptic cylinder was found by raising and lowering the ellipse until a node in the middle ± 5 mm of the endplate intersected the ellipse. This defined the height of the ellipse.

The volume of the elliptic cylinder was found ($V = \pi abh$, where a and b have been described and h is the height) and compared to that of the total vertebral volume. This enabled an identification of whether stronger relationships existed between the cement fill volume and vertebral stiffness when the fitted elliptic cylinder was used as the vertebral volume.

4.3.3 Cement Distribution Visualisation

Understanding features of the cement volume, for example determining whether cement volumes are distributed or concentrated was difficult. Two dimensional images of the augmented vertebrae only give an understanding of the layer in question and the 1 mm resolution of FE models makes identifying the nature of the structure from the generated model visualisation difficult. The solution to this problem was to use the 3D project feature in imageJ. This created a 3D rendering from the stack of images and rotated the resulting object through a range of angles, in this case 360 degrees. Variables of transparency bounds, opacity bounds and the projection method were set, in addition to the axis of rotation. These were left at their default using the brightest point projection method and the Y-axis for the rotation axis. This generated a model that was rotated around the Y-axis, providing a view of the cement distribution through the vertebrae and not one limited to a slice by slice view.

4.4 Results

4.4.1 Experimental Results

The vertebral stiffness following experimental loading is presented in Table 4.2 for the non-augmented vertebrae and following vertebral augmentation. The change in stiffness is shown in Figure 4.10 where an increase in stiffness was seen for Spine 1 L2 to L5 and

Spine 2 L1, while the remaining vertebrae presented a reduction in stiffness following augmentation.

Table 4.2: The experimental stiffness results for pre and post augmentation.

Specimen	Non-augmented Stiffness (N/mm)	Augmented Stiffness (N/mm)	Change in stiffness (N/mm)
Spine 1 L1	2991	2339	-652
Spine 1 L2	3456	3980	523
Spine 1 L3	3244	4152	907
Spine 1 L4	3223	4744	1521
Spine 1 L5	2891	4273	1382
Spine 2 L1	6149	7119	969
Spine 3 L1	5153	4444	-709
Spine 3 L2	5357	4848	-510
Spine 3 L3	5338	4779	-559
Spine 4 L1	3277	3231	-46
Spine 4 L2	5064	5042	-22
Spine 4 L3	6098	5749	-349
Spine 4 L4	4957	4488	-468
Spine 4 L5	7185	4403	-2782

A strong relationship was found between the percentage fill volume achieved and the bone volume fraction of the vertebra, where a reduced density of bone resulted in larger quantities of cement being injected into the vertebra before cement leakage occurred. This relationship can be seen in Figure 4.11.

A comparison of the total vertebral volume to the fitted cylinder volume is shown in Table 4.3. The percentage reduction in vertebral body volume, from total vertebral body volume to the fitted cylinder volume, varies significantly and is due to changes to the vertebral shape outside of the main vertebral body portion. For example, Spine 2 L1 was a degenerate vertebrae with extra bone growth at the anterior limits of the vertebral body. Other large changes were due to the size of the pedicles, where larger vertebrae (Spine 4) had proportionally larger pedicles.

The effect of this change on the relationship between cement fill and change in augmentation stiffness was identified and is shown in Figure 4.12. While there was a change in the relationship, a clear trend between the cement fill percentage and change in stiffness for the majority of the vertebrae is still not evident using either approach to measuring the vertebral body volume. The total vertebral body volume was used for all future fill volume comparisons.

The cement distributions following augmentation can be seen in Figures 4.13 to 4.16. From the 3D projection images and slices from μ CT stack, it was determined which vertebrae contained dispersed and which vertebrae contained concentrated volumes of cement, this




Figure 4.10: The experimental stiffness of the vertebrae pre and post augmentation. Showing that 5 of the 14 vertebrae showed an increase in stiffness following augmentation.




Figure 4.11: The relationship between the percentage cement fill of the total vertebral volume achieved through the augmentation procedure and the bone volume fraction value for each vertebra.

Table 4.3: Vertebral body volume compared to the volume of the fitted cylinder.

Specimen	Total Vertebral Body Volume (mm ³)	Fitted Cylinder Volume (mm ³)	Change in Volume (%)
Spine 1_L1	31229	17490	-44
Spine 1_L2	31789	20551	-35
Spine 1_L3	33519	22014	-34
Spine 1_L4	35123	22796	-35
Spine 1_L5	31189	22254	-28
Spine 2_L1	33835	17420	-48
Spine 3_L1	45866	31073	-32
Spine 3_L2	53695	34470	-36
Spine 3_L3	56003	36750	-34
Spine 4_L1	35137	22231	-37
Spine 4_L2	36791	23837	-35
Spine 4_L3	39909	26036	-35
Spine 4_L4	43357	25612	-41
Spine 4_L5	45606	26936	-41




Figure 4.12: The effect of using the fitted cylinder as a measure of vertebral body volume, on the relationship between augmentation fill volume and change in stiffness following augmentation. The red points show the relationship using the total vertebral volume, with the relationship using the fitted cylinder shown in blue.

characterisation is shown in Table 4.4.

Table 4.4: The characterisation of the augmented vertebrae into vertebrae with dispersed and concentrated volumes of cement.

Vertebra	Cement Distribution
Spine 1 L1	Dispersed
Spine 1 L2	Concentrated
Spine 1 L3	Concentrated
Spine 1 L4	Concentrated
Spine 1 L5	Concentrated
Spine 2 L1	Dispersed
Spine 3 L1	Dispersed
Spine 3 L2	Concentrated
Spine 3 L3	Concentrated
Spine 4 L1	Dispersed
Spine 4 L2	Dispersed
Spine 4 L3	Concentrated
Spine 4 L4	Concentrated
Spine 4 L5	Dispersed

Despite the lack of a relationship between fill volume and stiffness for the set of vertebrae on the whole, trends could be seen when splitting vertebrae into groups with: *a)* dispersed volumes of injected cement and *b)* concentrated volumes of cement. This can be seen in Figure 4.17 where a trend ($r = 0.83$) can be seen between the percentage fill of the total vertebral volume and the percentage change in stiffness for those vertebrae with concentrated volumes of cement.




Figure 4.13: Transverse and sagittal 3D projection views of the Spine 1 spine following augmentation, with the cement region being the brighter material in the vertebral body. The L1 vertebra is defined as having a dispersed volume of cement while the remaining L2, L3, L4, L5 vertebrae have concentrated volumes of cement.




Figure 4.14: Transverse and sagittal 3D projection views of the Spine 2 spine following augmentation, with the cement region being the brighter material in the vertebral body. The vertebra is defined as having a dispersed volume of cement.




Figure 4.15: Transverse and sagittal 3D projection views of the Spine 3 spine following augmentation, with the cement region being the brighter material in the vertebral body. The L1 vertebra is defined as having a dispersed volume of cement while the remaining L2 and L3 vertebrae have concentrated volumes of cement.




Figure 4.16: Transverse and sagittal 3D projection views of the Spine 4 spine following augmentation, with the cement region being the brighter material in the vertebral body. The L1, L2 and L5 vertebrae are defined as having a dispersed volume of cement while the remaining L3 and L4 vertebrae have concentrated volumes of cement.




Figure 4.17: The relationship between the percentage fill of the total vertebral volume and the percentage change in the vertebral stiffness following augmentation. The line and r value are for the red points where the cement volume was characterised as concentrated. The remaining blue points indicated the vertebra where the cement volume was characterised as dispersed.

4.4.2 Computational Modelling

The optimum results using the image registration and outcomes of the sensitivity tests, including the explicitly defined needle tracks can be seen in Figure 4.18 with the green triangles. This registration method with the defined needle tracks achieved a $CCC = 0.62$, the registration method not using the needle tracks achieved a $CCC = 0.46$, while the initial method used (similar to the method used with the bovine tail vertebrae) achieved a CCC of 0.18. Comparing the approach using the registration method to the method using the registration with defined needle tracks: computationally stiffer vertebrae were weakened and computationally weaker vertebrae were stiffened due to the higher Young's modulus for the interface and cement regions.




Figure 4.18: The results of using the initial method (red circles), the registration method (blue triangles) and using the defined needle tracks (green triangles), showing the agreement to the perfect $x=y$ line in orange.

4.5 Discussion

The results of both the experimental and computational work have shown a large improvement compared to the outcomes of the similar study on bovine tail vertebrae. These results are discussed below.

4.5.1 Experimental Results

The experimental results of vertebrae having undergone augmentation provide an insight into how the type of augmentation and the type of vertebrae affect the mechanical response to augmentation. Following augmentation, five of the 14 vertebrae showed an increase in the stiffness. Similar results were reported by Wijayathunga et al. [42] and no statistical difference was found in the three studies by Kinzl et al. [64, 66, 79] unless the cement was distributed to both of the endplates, forming a column of cement from the superior to the inferior. Four of the vertebrae that showed an increase in stiffness were from the same spine (Spine 1), which was the spine that had a considerably smaller bone volume fraction compared to the other spines. Given that this spine had the smallest bone volume fraction, it also received the largest quantity of cement, given the relationship between

the percentage fill achieved and the bone volume fraction. This relationship is most likely due to the ease of cement injection before leakage, with more dense vertebrae failing to receive large volumes of cement before the injection was stopped due to vascular channel leakage. Less dense vertebrae, conversely, had less bone tissue and therefore more space to receive cement. The lower BV/TV of Spine 1 therefore was more receptive to an increase in stiffness due to its lower initial stiffness and the larger volume of cement that could be injected. This may suggest that these type of vertebrae are the best candidates for the procedure given the risks of damage through the high intra-vertebral pressures that can be generated [102, 111]. In cases where the vertebral density is high, but high fill volumes are still pursued, damage may be caused by the high pressures “pushing” trabeculae to the side.

Interestingly and converse to findings presented here, the study by Aquarius et al. [102] suggested that cement volumes (like those described here as concentrated) were caused by high pressures “pushing” and damaging trabeculae. In that study, the more dispersed distributions gave more interdigitation due to the lower pressures and therefore resulted in reduced damage and stiffer augmented vertebrae. It may however be the case that these dispersed and concentrated volumes described in this thesis are a different phenomenon to the ball shaped and irregular shaped cement distributions described in that study [102]. An unwanted consequence of the high pressures during the vertebroplasty injection is the filter pressing that can occur [43]. Here, separation of the radio-opacifier from the cement is possible, meaning that experimentally and clinically (through fluoroscopy) the extent of the injected cement volume is potentially unknown. Additionally, a potential separation of the liquid and powder components of the PMMA may result in changing material properties of the cement region and a separation of the radio-opacifier may mean that the computational descriptions of the augmented regions are incorrect.

Cement volume positions were generally to the anterior of the vertebrae and to the middle of the inferior-superior plane, matching the desired position according to the National Institute for Health and Clinical Excellence Guidance [1]. These positions were achieved through the insertion of the needle and angling of the side opening needle tip. Given the lack of obvious variation in cement positions and shapes it was difficult to characterise the cement distributions, hence, a simplified description of the cement was used. This described whether the cement was in a concentrated volume or dispersed in a larger region or whole vertebra. Following this characterisation, strong relationships were found between the change in stiffness and percentage fill, if only those vertebrae with a concentrated volume of cement were included. This result suggests that distributed volumes of cement

are difficult to predict and should be avoided when carrying out the procedure using fluoroscopy, especially with more dense vertebrae where the likelihood or disperse volumes of cement increased. As highlighted above, previous studies only found significant increases in the stiffness of augmented vertebrae when the cement distribution was endplate to endplate stretching from the superior to the inferior of the vertebral body [75, 102, 103]. The sagittal views of the augmented vertebral projections suggest that all vertebrae except Spine 2 L1 and Spine 3 L1 achieved such end-to-end distributions with no noticeable difference in the change in stiffness of these two vertebrae, especially given that the Spine 2 L1 vertebrae received a large increase in stiffness following vertebroplasty.

To summarise the change in stiffness following augmentation based on the material properties of the vertebrae, the following features have the greatest effect:

1. A smaller bone volume fraction results in larger quantities of cement being injected before leakage occurs.
2. If the volume of cement is contained in a concentrated volume then the greater the volume of cement then the larger the increase in stiffness following augmentation.
3. If the volume of cement is dispersed through the vertebral body then there is no relationship between the volume of cement injected and its effect of the vertebral stiffness.
4. Concentrated volumes of cement are more likely if the bone volume fraction of the vertebra is lower (and therefore has a reduced density).

The effects of vertebral geometry and material properties on the response to augmentation will be investigated further using statistical shape modelling in Chapter 5, allowing a more detailed evaluation of the suggested approaches to vertebroplasty based on the tests carried out in this study.

4.5.2 Computational Results

The results of modelling augmented vertebrae showed a reduction in the agreement when compared to the non-augmented vertebral models, mirroring that found by Wijayathunga et al. [42]. Reasons for this are centred around modelling the internal cement region, given the good agreement seen between the non-augmented vertebrae and their models. More specifically the problems in modelling the cement region may be due to both experimental and modelling techniques.

Issues originating from the experimental work come from the difficulty in capturing the

extent of the injected cement volume, often due to clumping of the barium sulphate and separation of the barium sulphate from the other components of the PMMA cement. The clumping of barium sulphate was also found by Sikora [80]; in that study agglomeration was reduced through vigorous mixing of the PMMA monomer and barium sulphate prior to mixing with the PMMA powder. However, others suggested that the separation of the barium sulphate from the cement (calcium phosphate in that study [43]) during the high pressures of the injection were the cause of the subsequent agglomeration. While this problem is minimised through segmentation of the cement region at higher resolutions before the down-sampling, a problem still lies in capturing the intricate interdigitation between the cement and the trabecular bone at a resolution of 1 mm. One issue overcome through the use of the registration method is the removal of the bright halo of the cement region (from clumps of barium sulphate), which solved the problems found in the previous chapter using bovine tail vertebrae and problems encountered by others [42, 80]. This allowed more accurate modelling of the material properties of the vertebral bone and additionally improved accuracy of modelling the interface, given the removal of the false, dense bone surrounding the yielding interface in the bovine tail models. The inclusion of the explicitly modelled needle tracks increased the agreement significantly, with the compromise of interrupting an otherwise automated modelling process. In addition to the difficulties that agglomeration causes for modelling cement augmentation, it may also change the material properties of the PMMA cement. Such changes to the material properties are difficult to measure, given the unique environment within the vertebral body.

An additional experimental challenge is the unknown level of damage caused to the vertebrae when the needle is inserted. Given that after clear fractures (based on load-displacement curves), damage could not be seen on 82 μm scans of bovine tail vertebrae, it is unlikely that micro-fractures caused through needle insertion will be captured in the μCT scans of the human vertebrae. Further tests could investigate whether a change to the Young's modulus as a function of the density in the areas surrounding the needle tracks would improve the agreement. This could represent either damage to the bone in the form of micro-fractures or an increased modulus for the compacted bone created by the needle.

Despite the reduced agreement of the augmented models compared to the pre-augmented specimens, there was a large improvement compared to the initial method which was used for modelling augmentation in the bovine tail vertebrae. As with the non-augmented models, the CCC using the initial method was very similar to the CCC using the comparable method on the bovine tail vertebrae, CCC = 0.18 and 0.15 for the human and bovine models

respectively. This makes the improvement using the registration significant and a beneficial step in making the models of augmented vertebrae and means that the improvement in agreement is not due to the change of tissue or other experimental improvements.

Given the difficulty capturing the detail between the cement and trabecular bone, using a higher resolution for finite element models could be the next step towards improving agreement between augmented models and their experimental counterparts. However, it comes with a number of difficulties which have been shown in the study by Zhao et al. [76], where the increase in resolution presents many other unknowns that need to be modelled. For example modelling the specific interaction between the bone and the cement including the coefficient of friction between them. This would require difficult experimental investigations, given the unique environment inside the vertebrae originally filled with bone marrow. Along with modelling properties that are not applicable in continuum level models, there is the problem with computational cost, where the increased time to solve models would become problematic given the investigatory nature of the study and the fact that the behaviour in simplified bone plug models is different to what is seen in whole vertebrae [80]. Some groups have explored a middle ground between the two approaches. Kinzl et al. [66] used spheres that were randomly inserted into the augmentation region to represent the pores that are present, although not visible, in the augmented regions of the experimentally augmented specimens. This methodology is similar to one of the approaches attempted in the current study which unfortunately performed less favourably than the currently presented approach. It involved an inverse greyscale material property for the augmented region based on the underlying non-augmented background. Empty (marrow containing) voxels were given material properties of, or close to PMMA, voxels that contained trabecular bone were given the material properties of bone and intermediate greyscale values attempted to represent the gaps that form as the cement shrinks during setting. The limited agreement found using this methodology suggests that the middle ground between approaches is just as problematic as the approaches themselves. A large array of unknown material properties exist and the unknown relationships between them make calibration and tuning difficult and somewhat arbitrary.

4.5.3 Quantification of Error

Sources of error from modelling come from discrepancies between the experimental setup and specimen geometry and material properties. Sensitivity tests carried out identified the effect of errors in accurately describing the cement cap depth and loading position accuracy, given the 1 mm accuracy limit. The maximum possible error for load position

in a ± 1 mm anterior/posterior/left/right direction was less than 5%, with the worst case error for the encap depth tests (1 mm thicker for both endcaps) being approximately 5%. Other possible errors would include accurately capturing the model geometry, however simply eroding or dilating the vertebral mesh would not accurately describe the type of meshing errors that may occur (it is unlikely that masks would be created at ± 1 voxel around the whole model). To attempt to answer this, a user variability test was carried out in Section 2.3.5.5. There, the maximum user variability gave errors of 14 %, however as discussed in that chapter, users achieved consistently higher or lower stiffness values compared to their peers. Hence, a 14 % error within a single users segmentation process is both unlikely and further reduced through the use of the Python script carrying out the segmentation. Therefore simply adding the two main sources of error gives a worst case error of approximately 10 %, an error which is relatively small when compared to the possible errors in obtaining the maximum stiffness and the errors between experimental loading repeats.

The specific error in the experimental load is difficult to measure mainly due to the high value and limited supply of human tissue; to measure the error in repeated loads more accurately, these repeats would need to be taken in the non-augmented state given the large array of unknowns the augmentation process introduces. However, this would have led to possible damage or unwanted fracture of the vertebrae, making the results following augmentation more difficult to interpret. As it stands the differences between repeats of the augmented vertebrae show changes of approximately -3000 N/mm (59 % reduction) to +1200 N/mm (39 % increase), with the large reductions originating from repeated testing in plastic regions of the load displacement curve. These large errors between repeated measures are the worst case however and it is unlikely that such errors would occur between non-augmented specimens in a loading range that avoids the potential of causing damage. It is therefore likely that the results presented here, being the consequence of two loads total (one non-augmented, one augmented), would have much less error than that seen in the repeated loads.

4.6 Conclusion

Overall the models produced here can estimate the stiffness of non-augmented human lumbar vertebrae to a high degree. The predictive abilities of models of vertebroplasty is reduced, although a large improvement over studies that used similar approaches and had similar constraints in mind (clinical resolutions for scans). Importantly, the models can represent the large range of outcomes from the augmentation process, suggesting

that the representation of augmentation is robust enough to model and experiment with different augmentation scenarios. Therefore the work presented in this chapter provides a good starting point to use statistical shape and appearance modelling to understand the intricacies of vertebroplasty.

Chapter 5

Statistical Shape & Appearance Modelling

5.1 Introduction

Statistical shape and appearance models (SSAM) give the opportunity to investigate the consequences of variation in a dataset that would otherwise be difficult to isolate. In the current study, it was also not possible to obtain and test sufficient experimental specimens to find trends in the data. The effect of the anterior height of the vertebral body on the response to loading is a simple example of this. To understand the effect of anterior vertebral height on vertebroplasty either experimentally or computationally using specimen specific models, a large range of vertebrae with varying anterior heights would be required to represent (without gaps) the range in the population. Even if this could be achieved, confounding factors and other variation would make the presentation of specific relationships difficult. Using statistical shape and appearance models with principal component analysis (PCA), the modes of variation can be characterised and isolated, allowing an understanding of how specific modes of variation affect the chosen measure. A caveat here (similar to specimen specific models, but potentially less obvious) is that the modes of variation are specific to the input set, i.e. the models used to make the statistical shape models. Hence, any estimations of the types of variation and the effects these modes of variations have is specific to that set of models and not necessarily applicable to the population on the whole. It is dependent on how representative the input set is of the population as a whole.

The general aims of most statistical shape and appearance models are to describe the geometric and material property variation of bones across a sample population. This is

usually carried out by comparing the mean of the sample group to the main modes of variation found within it [83]. As reported in Chapter 1, such approaches have been used to characterise shape variation in the cervical spine [87], pelvis [112] and femur [113, 114]. Like this study, others have used statistical shape models to identify the variation in the lumbar spine for both spinal curvature variation [88–90] and to identify the variation in single vertebrae [81]. However, no-one has yet used these tools to examine the shape and density distribution together for vertebrae, which is the aim of the current study. Regardless of the bone in question in these studies, the aim is most often to determine the mean, which, dependent on the population sample can be described as the population norm. The studies then aim to understand how variations affect the outcomes of different treatments.

5.1.1 General Methodology

Generally, statistical shape and appearance models start with the segmentation and creation of models that form the input set of the statistical model. Following this is a rigid registration step, where rotations and translations are applied to all meshes or models that form the input data set so that they share a common coordinate system. Methods of performing the rigid registration often use an automated approach, where certain degrees of freedom are restricted to ensure registration is achieved. For example, with very cylindrical vertebrae, if rotation is allowed around the coronal axis with vertebrae whose volumes vary significantly, registration may never be achieved as one vertebrae would fit inside the other and rotate without end. Manual techniques can be used similar to those used in the previous chapter (Section 4.3.1.1), with automated rigid registration techniques also using landmark registration.

An input set is usually formed of a set of meshes (surface or volume), that are created from sets of scan data that describe the desired population. These meshes are created using a similar processes to those used in the previous chapters. A mean model is usually generated by first selecting a surface mesh from the input set and registering it to the other models (surface meshes) of the input set using deformable registration, a mean model is calculated from these transformations. Each of the other models are then selected and registered to the other models from the input set until the change between generated mean model is no longer significant (described in detail by Clogenson et al. [83]). Once the mean model has been created the transformations required to deform each input model into that mean are recorded as the variation from the mean shape.

Given the large array of variation captured in those transformations, a method of di-

mensional reduction is usually employed, commonly either Principal Component Analysis (PCA) [83] or Single Value Decomposition [115]. This allows the reduction of the number of variables being identified, allowing a focus on those that contain or describe the largest variations.

5.1.2 Chapter Overview

The work in this chapter focusses on using the 14 human lumbar vertebra FE models described in the previous chapters as the input set to create a statistical shape and appearance model based upon them. Specifically, this chapter describes the methods used to create the statistical models, and find the main modes of variation within the input set; the process is outlined in Figure 5.1.

Methods of creating the statistical shape model and generating models from it are described. These models represent specific modes of variation and standard deviations away from the mean model. Methods of performing validation and characterisation of the generated models are described and the results are presented. Following this, the results are discussed.

The methods of creating artificially augmented models are then described along with methods of testing the consequences of augmentation in them. Results of the effect of augmented region volume, position and the effect of different loading positions are then presented. Finally, there is a discussion of the results regarding augmented vertebral models.




Figure 5.1: A flow chart describing the study design and usage of different generated models.

5.2 Methods

A plugin for Simpleware ScanIP was developed through a collaboration between the University of Leeds and Simpleware (Synopsys). The tool utilises the Insight Tool Kit (ITK) to generate a SSAM model based on an input set of finite element models created in Simpleware ScanIP, and was developed specifically for the analysis of vertebrae. In this case the models used were of non-augmented vertebrae described in the previous chapters, using the bone volume fraction methodology to generate the downsampled greyscale distribution. The process of using the plugin with the 14 human lumbar vertebrae as the input set is described in this section.

5.2.1 Statistical Shape and Appearance Model Creation

The creation of the SSAM through the use of the plugin is carried out in five main steps. The first of these steps is a pre-processing step, converting the masks and background images into necessary formats (.mha & .mhd) for use with the ITK tool-kit.

In the next step, a rigid registration of the models is performed, aligning the masks and backgrounds to a shared coordinate system. This step is carried out using the ITK library with the registration being measured according to a mean squared image-to-image metric using a gradient descent optimiser. A limit to the number of attempts or steps allowed can be set; the default value of 100 steps for each input model was used for this study.

The geometry of each input specimen is described in a deformable registration step, measuring the transform required to morph the mean of the input set onto each of the other input models. This step utilises the ITK FEM registration filter.

Material properties of the input models are described by the difference between the greyscale value of each voxel from the input set to the greyscale value of the corresponding voxel on the mean model. Since each input model will have a different number of elements, voxel to voxel correspondence cannot be ensured. To bypass this problem the greyscale background of each input model is first morphed through the same deformable registration transform that each input model underwent to capture the geometric variation. The change of each voxel's greyscale value from each model can then be measured and added to the geometric variation to describe the total variation of each model.

The transforms between the mean and each of the input vertebrae which describe the vertebral geometries along with the differences in voxel greyscale values are used as inputs for PCA. Here the outputs of the step include the principal components, their eigenvalues, the percentage of variation captured in that component and the cumulative percentage variation captured.

In the final step, new models formed of mask and background combinations are generated using the ITK Image Warp filter. These mask and background combinations are from within the envelope of geometric and material property variation created by the input set. The models generated from the SSAM can be created according to the principal components, choosing the variation desired from standard deviations within each principal component. The plugin allows for the first five principal components (PC1, PC2, PC3, PC4 & PC5) to be used as variables for model generation, altering the value of the principal components by up to three standard deviations (S.D.) away from the mean in both positive and negative directions. The notation of positive and negative standard deviations used

in this chapter refers to a direction of variation and not a specific type of variation. For example the variation in a length measurement may be described by positive and negative standard deviations, yet the positive S.D. does not necessarily describe an increase in length compared to the mean length.

Given the large array of data that could be generated from each of the five PCs, only the first three principal components were used in this preliminary study, and for each of these, 1 S.D. increments away from the mean were considered up to ± 3 S.D.s. These first three PCs provided an investigation into the largest section of the variation, which can be assumed to have the largest effect on the response to loading and therefore augmentation.

5.2.2 The Generation of Vertebral Models to Understand Variation

To attempt to understand what each of the principal components represented in terms of the variation captured by it, models were created at ± 1 , ± 2 and ± 3 standard deviations away from the mean for the first three principal components independently. In addition to these, the mean model was generated, allowing comparisons of the other generated models against it. The models used for the characterisation of variation are listed in Table 5.1. A separate set of generated models are used for validation and were described in Section 5.2.2.1.

Table 5.1: Descriptions of the models used for characterisation of generated models and of the variation described by different principal components. These models were also used for subsequent tests on the effect of augmentation.

Name	PC1	PC2	PC3
Mean	0	0	0
PC1 +3	+3	0	0
PC1 +2	+2	0	0
PC1 +1	+1	0	0
PC1 -1	-1	0	0
PC1 -2	-2	0	0
PC1 -3	-3	0	0
PC2 +3	0	+3	0
PC2 +2	0	+2	0
PC2 +1	0	+1	0
PC2 -1	0	-1	0
PC2 -2	0	-2	0
PC2 -3	0	-3	0
PC3 +3	0	0	+3
PC3 +2	0	0	+2
PC3 +1	0	0	+1
PC3 -1	0	0	-1
PC3 -2	0	0	-2
PC3 -3	0	0	-3

5.2.2.1 Range of Models Generated for Property Validation

There was a need to investigate whether the range of geometries, material properties and resultant stiffness values found in the input set was represented in the SSAM within the first three principal components. Therefore a subset of models that encompassed the majority of the variation in PCs 1, 2 and 3 were generated to validate the process. If PC1, 2 and 3 were represented on three orthogonal axes, with the mean vertebrae at the centre, then a cube of side lengths from +1 to -1 SD should capture approximately 68% of the total variation. This is due to the quantity of the population expected to fit within 1 S.D. of the mean for a normally distributed input set. The models generated for these tests are shown in Table 5.2. These models were used to examine whether the majority the input model variables fitted within the variation limits created by the “cube”, for geometric measurements, greyscale background and resultant stiffness, which would provide confidence that the variation of the input set had been captured and represented by the SSAM.

Table 5.2: Standard deviations of the models generated to validate the geometric, material property and stiffness variation. The vertebrae form a “cube” of variation that should describe all possible variation within 1 S.D of the mean for the first three principal components. 0 describes the mean position for the principle component and the 0,0,0 vertebrae represents the overall mean model.

PC1	PC2	PC3
0	0	0
-1	0	0
1	0	0
0	-1	0
0	1	0
0	0	-1
0	0	1
1	1	1
-1	1	1
1	1	-1
-1	1	-1
1	-1	-1
-1	-1	-1
1	-1	1
-1	-1	1

5.2.2.2 FE Model set-up

Once the vertebrae were generated using the plugin, the rest of the FE model was set up. This included forming endcaps to mimic the experimental set-up and defining how the input set was loaded. Endcaps were created programmatically within the ScanIP Python scripting environment with a diameter of 90 mm, equal to the diameter of the experimental

endcaps. A depth of 6 mm was used, ensuring the superior and inferior tips of the vertebrae were captured with a minimum of 1 mm between the vertebra and end of the endcap. The remaining properties for the FE model were identical to those used for the models of the human vertebra which became the inputs to the PCA plugin. The material properties for the cement and bone were defined in the same way as described in Section 3.3, with the same greyscale-dependent bone properties based on the bone volume fraction method. The boundary conditions were also identical, favouring the frictionless contact between bone and endcap, also described in Section 3.3.

As in the previous chapters, a 1 mm displacement was applied to a loading point and the reaction force was measured to obtain the stiffness. The loading point was determined by finding the centre of the vertebral body in the middle slice axially and translating the point axially to the maximum height of the model. The load was then applied through an analytical platen which was tied to the superior endcap, matching the specimen specific model set-up in the previous chapters.

5.2.3 Measurement of Vertebral Geometry

Accurate measuring of the vertebral geometry was another method of validating shape and geometric features of the outputs from the SSAM generated vertebral models. Previous studies have either taken the vertebral geometry measurements from x-ray scans [116, 117] or μ CT scans [100, 118], where the measurements have been made through moving a cursor to the locations of specific points and reporting the distance between them. While this method may produce accurate results, there is inherent human error associated with it, the number of measurements is limited to the number of planes available and applying this to large sets of data is time consuming which may lead to further error.

The approach used here employed the 1 mm voxel resolution FE models. The mask describing the vertebral body of these vertebrae was exported as a stereo-lithography file (STL), allowing the surface of the vertebra to be measured. Measurement of the vertebra was carried out in Matlab where the STL file was imported forming a point cloud describing the surface from the surface nodes. Once imported, the point cloud was split into thirds in each of the three anatomical planes, with the point clouds describing the vertebra within these thirds being stored in separate variables. Cuboids were then fitted to the nine point clouds, maintaining the correct alignment and therefore prohibiting rotation of the cuboid. The cuboid fits can be seen in Figures 5.2 to 5.4.

Measuring the two longer sides of the nine fitted cuboids provided a total of 18 measurements describing most aspects of the vertebral geometry, these can be seen with their abbreviations




Figure 5.2: The three cuboids fitted to the vertebra point cloud in the axial plane.




Figure 5.3: The three cuboids fitted to the vertebra point cloud in the coronal plane.

in Table 5.3. This includes identifying wedge shaped vertebrae, recorded as reduced anterior height compared to the posterior height, left to right wedge shapes, recorded as sagittal left height being smaller or larger when compared to the sagittal right height. The measurements were used for validation by comparing the measurements of input model geometry to the measurements of the generated models. Here, the input models were compared to the “cube” set of models, describing all possible variation within 1 S.D. of the mean from the first three principal components.




Figure 5.4: The three cuboids fitted to the vertebra point cloud in the sagittal plane.

Table 5.3: The measurements and abbreviations taken from the vertebrae, used to compare and validate the generated models to the input models.

Measurement	Abbreviation
Sagittal Left Height	SLH
Sagittal Left Depth	SLD
Sagittal Mid Height	SMH
Sagittal Mid Depth	SMD
Sagittal Right Height	SRH
Sagittal Right Depth	SRD
Coronal Anterior Height	CAH
Coronal Anterior Width	CAW
Coronal Mid Height	CMH
Coronal Mid Width	CMW
Coronal Posterior Height	CPH
Coronal Posterior Width	CPW
Axial Inferior Depth	AID
Axial Inferior Width	AIW
Axial Mid Depth	AMD
Axial Mid Width	AMW
Axial Superior Depth	ASD
Axial Superior Width	ASW

5.2.4 Validation of Vertebral Variation

Validation of the vertebral properties was carried out for the geometric, material property and stiffness measurements for the generated models in the “cube” set, comparing them to the input set models.

Geometric validation was carried out through the comparison of the 18 vertebral measurements described above, ensuring the range of variation in the input set was represented in the generated models. In addition to the 18 measurements, the mean and range of vertebral volumes was compared from the generated “cube” models to the input set.

Material property variation was validated by comparing the mean and range of the mean greyscale value for the models and by comparing the greyscale histograms of the generated models to the input set. The histograms were normalised to account for the differing size of model by dividing each histogram value by the number of voxels present in that model.

The model stiffness was validated by comparing the range and mean stiffness of the generated models to the input model stiffness values.

5.2.5 Load Position Sensitivity Tests

The sensitivity of the generated models to being loaded anteriorly and posteriorly was tested. This allowed an investigation into how the modes of variation affected anterior and posterior loads for the vertebrae. The methodology used here followed the same set-up used when identifying the effect of loading position in the human lumbar vertebrae in Section 3.3.2.2. Here, the load positions were limited to 10 mm and 20 mm posterior and anterior of the central loading position, generally corresponding with the anterior and posterior walls of the vertebral body at 20 mm.

5.3 Results from the Generated Models

5.3.1 General Results of the Statistical Shape and Appearance Model

Results from the principal component analysis of the input set show that 72 % of the variation within the dataset was captured in the first five principal components. The individual breakdown of the weighting of each is shown in Table 5.4. Of these five available principal components only the first three were used in this study and these described 57 % of the variation within the input dataset.

Table 5.4: Variation captured in the first five principal components.

Principal Component	Percentage of Variance	Cumulative percentage of Variance
1	29.95	29.95
2	14.33	44.28
3	12.29	56.56
4	9.22	65.78
5	5.91	71.69

5.3.2 Validation of Generated Vertebrae

5.3.2.1 Geometry Validation

The variation in geometry for combinations of the first three principal components, with all possible combinations within ± 1 S.D. of the mean is presented here. The 18 geometric measures found within the ± 1 S.D. models are presented in Figure 5.5. The results show a strong agreement between the input and the generated models, with the means of the two sets agreeing closely. The range of the measurements was also found to agree well. Differences in the range of generated vertebrae compared to the range of the measurements seen in the input set are due to the generated models describing only one standard deviation of variation.

Volume differences between the input set and generated vertebrae were also small, with the difference between the mean input set volume and the mean generated model volume being 10 %, the means being 39500 mm^3 and 38245 mm^3 respectively. The range of vertebral volumes was also comparable: the input set range was 31189 to 56003 mm^3 while the generated vertebrae for ± 1 S.D. with all possible combinations of PC 1, 2 and 3, had a range of 28152 to 50462 mm^3 .




Figure 5.5: The variation in the 18 geometric measures of the input set models compared to the variation in the generated models for the first three principal components including all combinations of the +1, mean and -1 standard deviations. The range in the measurement is shown by the line and the mean value is shown by the square.

5.3.2.2 Material Property (Greyscale) Validation

The variation seen in the histograms representing the greyscale differences between generated models and the input can be seen in Figures 5.6 to 5.8. The range of histograms from the input set shown as the grey range and the coloured lines showing the histogram for each of the generated models. This normalised data shows a general agreement with the shape of the curves and relative quantities being very similar. Deviations from the similarity are, for the most part, at the high and low greyscale values, where in general there are fewer voxels in this range for the generated models. Similar trends exist for all of the principal components tested. An example background in Figure 5.9, comparing an input background mid-slice with the mean generated model mid-slice, shows that the small voxel counts in the high and low greyscale regions correlate with having less contrast. Additionally, as shown in Figure 5.9, the important denser cortical shell region, is preserved and accurately described in the background of the mean generated model. While the presence of the cortical shell is clear, its definition and relative contrast with the trabecular bone or empty spaces is not as well preserved. Other features such as the posterior vascular channel was also preserved, albeit not as clearly as on the input set.

The mean greyscale values of the input set and the generated models from the “cube” set were 108 and 97 respectively. The ranges were 78 to 164 for the input set and 72 to 122 for the generated models for all combinations of ± 1 S.D..

5.3.2.3 Finite Element Stiffness Validation

The range of stiffness values from the generated models from “cube” set of models matched closely with the range of stiffness values from the input set. These were 2791 - 6039 N/mm and 2887 - 6172 N/mm for the input and generated models respectively. Additionally the means of the sets also agreed well with the mean computational stiffness of the input set being 4508 N/mm and 4134 N/mm for the mean generated model.




Figure 5.6: Histogram of the models generated from principal component 1, with ± 3 , 2 & 1 standard deviations away from the mean. The values are normalised with respect to the total volume of the vertebrae, allowing comparisons with different sized model. The grey background represents the range of histograms seen in the input set.




Figure 5.7: Histogram of the models generated from principal component 2, with ± 3 , 2 & 1 standard deviations away from the mean. The values are normalised with respect to the total volume of the vertebrae, allowing comparisons with different sized model. The grey background represents the range of histograms seen in the input set.




Figure 5.8: Histogram of the models generated from principal component 3, with ± 3 , 2 & 1 standard deviations away from the mean. The values are normalised with respect to the total volume of the vertebrae, allowing comparisons with different sized model. The grey background represents the range of histograms seen in the input set.




Figure 5.9: Greyscale backgrounds for the axial mid-slice from, A, the Spine 1 L1 input vertebra and B, the mean generated model.

5.3.3 Measuring & Interrogating Vertebral Variation

5.3.3.1 Geometry Variation

The geometric variation found in the first three principal components was identified using the images of the surface point clouds of the mean, ± 3 , ± 2 and ± 1 standard deviations away from the mean, along with the measurements of the 18 variables described previously.

General changes to the volume of the generated vertebral models can be seen in Figure 5.10. This shows a small and similar change to the volume of the vertebrae in both PC1 and PC2, corresponding to general size changes and shape changes respectively. The majority of the volume changes were isolated to PC3, which captured size variation, with little change to the shape.




Figure 5.10: The vertebral body volume variation for the first three principal components, including the variation at ± 1 , 2 & 3 standard deviations away from the mean and including the mean.

Images of the point clouds are shown in Figure 5.11. Simplified figures can be seen in Figures 5.12 to 5.14, showing the mid-slice through each of the anatomical planes. While information is lost in these slice images, it allows a clear visualisation of the modes of variation found in the main parts of the vertebral body. Here, only the +3, mean and -3 standard deviations of the principal components are shown, simplifying and allowing clear visualisation of the main mode of variation present in each principal component.

Principal component 1 contained the least geometric variation of the first three principal

components. The mid slice axially showed an almost identical shape and size between the +3, mean and -3 standard deviations, while the coronal and sagittal planes showed minor changes in the overall size and volume of the vertebrae. The changes in total vertebral volume was limited to 17 % larger and 15 % smaller than the mean model, in the +3 S.D. and -3 S.D. generated models respectively.

Principal component 2 showed the most varied geometric variation with many of the different shapes and some of the degeneration of the input set clearly visible. The large changes to the axial mid slice (Figure 5.12) of the second principal component appear to represent the different levels of the lumbar spine that made up the input set. Positive standard deviations away from the mean, especially the +3 S.D. slice, had the much wider posterior portion aligning with that of the L5 lumbar vertebrae. Conversely the much narrower -3 S.D. model appeared similar to the L1 lumbar vertebrae. This change in width between the +3 and -3 S.D. was especially evident in the coronal views in Figure 5.14. The change from L5-like vertebrae at +3 S.D. to L1-like vertebrae at -3 S.D. was also reflected in the sagittal views (Figure 5.13) where anterior and posterior wedge shapes were seen in the +3 and -3 S.D. models respectively.

Finally, principal component 3 also contained a considerable quantity of the geometric variation. Here, the mode of geometric variation was the size of the vertebrae with the -3 S.D. model being considerably larger than the mean vertebral model (54 % larger) and vice versa with the +3 SD model, where the volume was reduced by 40 % compared to the mean.

Quantification of all of the measurements taken on the generated models can be seen in Table 5.5. The colouration in the table allows easy identification of the changes to the vertebral geometry compared to the mean generated model. For example, while the changes to the measurements were near uniform in PC 1 and 3, the non-uniform nature of PC 2 shows the different shapes that are present and described above. Additionally the numerical values presented in the table allows quantification of the shapes seen in the models, for example the anterior and posterior wedge shapes in PC 2, described with the Coronal Anterior Height and the Coronal Posterior Height having an antagonistic relationship.




Figure 5.11: Three dimensional views of the surface point clouds of the vertebral models from the first three principal components, showing the mean, +3 and -3 standard deviations away from the mean. Showing how the geometry is captured in the first three principal components.




Figure 5.12: Axial views of the mid slice of the vertebral models from the first three principal components, showing the mean, +3 and -3 standard deviations away from the mean. Showing how the geometry is captured in the first three principal components.




Figure 5.13: Sagittal views of the mid slice of the vertebral models from the first three principal components, showing the mean, +3 and -3 standard deviations away from the mean. Showing how the geometry is captured in the first three principal components.

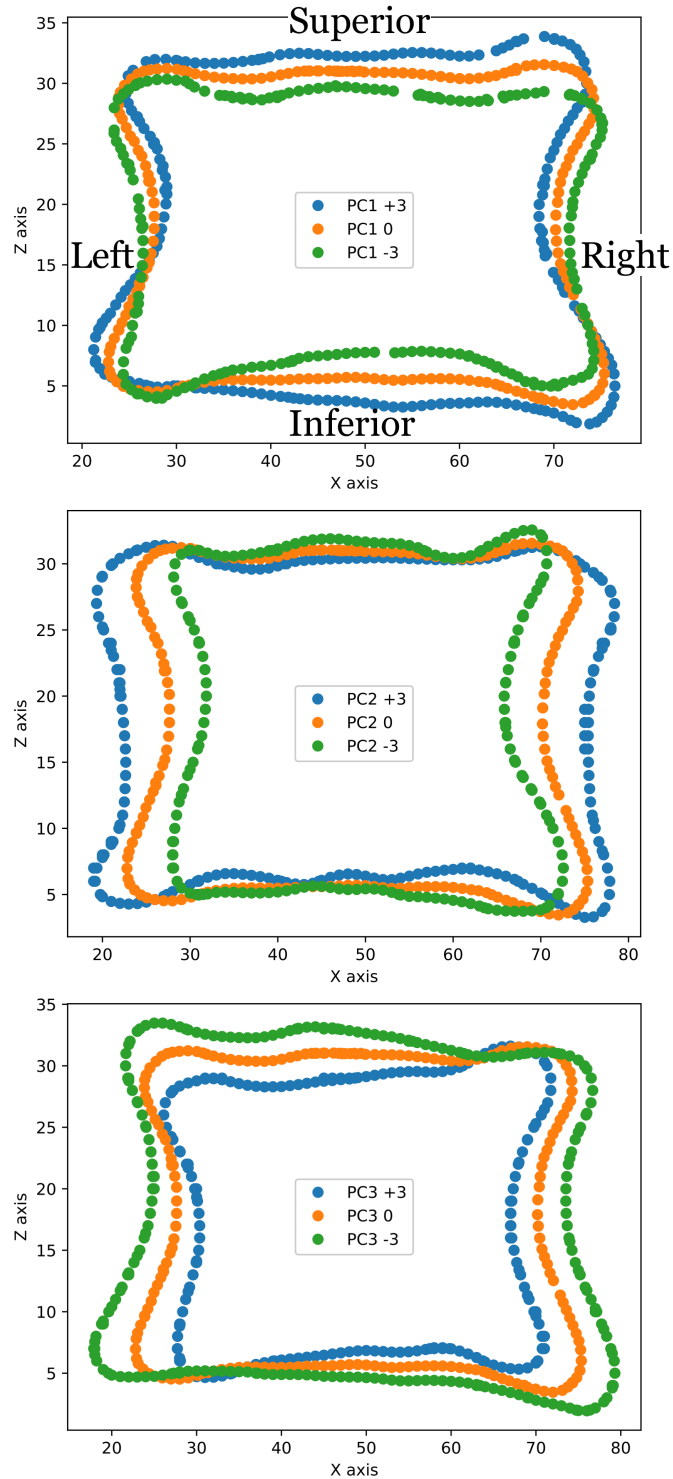


Figure 5.14: Coronal views of the mid slice of the vertebral models from the first three principal components, showing the mean, +3 and -3 standard deviations away from the mean. Showing how the geometry is captured in the first three principal components.

Table 5.5: The measurements of the 18 variables (abbreviations in Table 5.3), shown as a fraction of the mean model's measurements. Colouration shows the reduced measurements in red and larger measurements in green showing the geometric variation quantitatively. The colouration in the table allows easy identification of the changes to the vertebral geometry compared to the mean generated model. For example, while the changes to the measurements are near uniform in PC 1 & 3, the non-uniform nature of PC 2 shows the different shapes that are present and described above. The acronyms for the measurements are given in Table 5.3

	SLD	SLH	SMD	SMH	SRD	SRH	CAH	CAW	CMH	CMW	CPH	CPW	AID	AIW	AMD	AMW	ASD	ASW
PC1 +3	1.119	1.121	1.130	1.157	1.138	1.119	1.143	1.038	1.122	1.057	1.068	1.038	1.147	1.057	1.110	1.004	1.129	0.976
PC1 +2	1.070	1.083	1.078	1.107	1.088	1.074	1.090	1.020	1.082	1.036	1.047	1.024	1.101	1.036	1.068	1.004	1.083	0.985
PC1 +1	1.030	1.040	1.042	1.052	1.051	1.038	1.053	1.000	1.038	1.022	1.022	1.011	1.048	1.022	1.031	0.992	1.040	0.996
Mean	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
PC1 -1	1.002	0.956	0.987	0.952	0.966	0.960	0.949	1.007	0.967	0.984	0.988	0.984	0.956	0.984	0.973	1.013	0.962	1.011
PC1 -2	0.997	0.927	0.977	0.914	0.944	0.922	0.899	1.016	0.956	0.986	0.979	0.988	0.906	0.966	0.944	1.039	0.938	1.021
PC1 -3	1.001	0.906	0.976	0.901	0.918	0.885	0.863	1.031	0.956	0.996	0.983	0.979	0.858	0.948	0.923	1.054	0.918	1.031
PC2 +3	1.089	1.174	1.078	1.208	1.028	1.139	0.953	1.233	1.046	1.173	1.075	1.193	1.101	1.147	1.137	1.352	1.102	1.204
PC2 +2	1.049	1.114	1.042	1.133	1.021	1.095	0.970	1.136	1.017	1.107	1.044	1.137	1.058	1.100	1.091	1.239	1.067	1.124
PC2 +1	1.022	1.063	1.012	1.066	0.992	1.051	0.985	1.066	0.995	1.054	1.011	1.071	1.030	1.054	1.043	1.122	1.034	1.067
Mean	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
PC2 -1	1.014	0.960	1.007	0.967	1.006	0.976	1.024	0.988	0.992	0.967	0.985	0.940	1.018	0.967	0.971	0.939	0.985	0.957
PC2 -2	1.052	0.930	1.049	0.967	1.034	0.956	1.059	0.966	1.005	0.932	0.968	0.872	1.026	0.932	0.960	0.901	0.972	0.902
PC2 -3	1.082	0.914	1.079	0.984	1.065	0.957	1.085	0.935	1.034	0.896	0.958	0.824	1.038	0.901	0.959	0.879	0.969	0.849
PC3 +3	0.912	0.830	0.944	0.828	0.952	0.857	0.853	0.868	0.984	0.868	0.965	0.876	0.794	0.835	0.811	0.905	0.842	0.898
PC3 +2	0.935	0.889	0.949	0.886	0.941	0.906	0.894	0.914	0.963	0.903	0.951	0.914	0.861	0.887	0.879	0.930	0.895	0.932
PC3 +1	0.967	0.944	0.967	0.944	0.965	0.952	0.940	0.960	0.985	0.944	0.968	0.949	0.933	0.944	0.932	0.947	0.947	0.966
Mean	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
PC3 -1	1.065	1.055	1.058	1.063	1.053	1.044	1.070	1.055	1.029	1.060	1.035	1.033	1.075	1.060	1.064	1.051	1.056	1.037
PC3 -2	1.123	1.108	1.118	1.126	1.097	1.094	1.126	1.133	1.067	1.119	1.073	1.083	1.147	1.119	1.135	1.104	1.113	1.074
PC3 -3	1.151	1.169	1.146	1.190	1.140	1.143	1.163	1.210	1.129	1.180	1.114	1.118	1.219	1.180	1.200	1.161	1.170	1.111

5.3.3.2 Greyscale & Material Property Variation

Variation in the mean greyscale was almost completely isolated to PC1, with only a small amount of that variation in PC2 and PC3 (Figure 5.15). The range of the mean greyscale values for PC1 are from the least dense at -3 S.D. away from the mean with a mean greyscale value of 42, to a maximum of 162 at +3 S.D. away from the mean.

However, despite the similarity in the mean of the greyscale variation for PC2 and PC3, the distribution varies significantly, seen in Figure 5.16. In PC2 and PC3 (although more clearly noticeable in PC3) the density distribution shifts from the posterior to the anterior of the vertebral body. In PC2 this shift of densest material is from anterior to posterior in the models from -3 S.D. to +3 S.D., as the vertebrae change in shape from L1-like to L5-like. In PC3 the shift of the densest part is from posterior to anterior with reducing vertebral body volume (from -3 to +3 S.D.). This means that the larger vertebrae have a much denser posterior portion and less dense anterior, while the smaller vertebrae have a much denser anterior portion with a less dense posterior/pedicle region. The density distribution in PC1's models is more uniform, with more general changes to the density as described earlier. The -3 S.D. vertebrae had the least dense properties and +3 S.D. had the most dense. In the +3 S.D. model of PC1 the thickness of the cortical shell increases significantly compared to that of the least dense (-3 S.D.). However, the contrast between the shell and the cancellous regions in the extremes remain similar, while the mean model's cortical shell is less clearly defined.

5.3.3.3 Resulting Stiffness Variation

The variation in the resulting stiffness of the generated models are presented in Figure 5.17 and Table 5.6. The trends within the different principal components match closely with the variation seen in the mean greyscale (Figure 5.15). This suggests a stronger relationship between the greyscale background and the resulting stiffness of the models than any of the other variations in the generated models when loaded centrally.




Figure 5.15: The mean greyscale variation for the first three principal components, including the variation at ± 1 , 2 & 3 standard deviations away from the mean and including the mean.

Table 5.6: The stiffness of the generated models from PC1, 2 and 3, showing the fractional change to the mean model in each case.

Model	Stiffness (N/mm)	Fractional Change from the Mean
PC1 +3	6543.38	1.582
PC1 +2	5887.75	1.424
PC1 +1	5059.89	1.224
Mean	4134.88	1.000
PC1 -1	2962.82	0.717
PC1 -2	1696.47	0.410
PC1 -3	1131.86	0.274
PC2 +3	4159.29	0.980
PC2 +2	4291.53	0.831
PC2 +1	4355.98	0.770
Mean	4134.88	1.000
PC2 -1	4051.57	1.006
PC2 -2	3434.91	1.038
PC2 -3	3182.81	1.053
PC3 +3	2930.69	0.709
PC3 +2	3420.55	0.827
PC3 +1	3846.57	0.930
Mean	4134.88	1.000
PC3 -1	4409.05	1.066
PC3 -2	4240.58	1.026
PC3 -3	4141.31	1.002




Figure 5.16: The variation in the greyscale distribution across the mid-slice of the vertebrae generated from PC1, PC2 and PC3, for each of the $\pm 1, 2$ and 3 standard deviations from the mean. Showing how the changing distribution of the greyscale, even for PC2 and PC3 where the mean greyscale variation is minimal. Red colours indicate denser bone and blue colours indicate the least dense bone.




Figure 5.17: The stiffness variation for the first three principal components, including the variation at ± 1 , 2 & 3 standard deviations away from the mean and including the mean.

5.3.3.4 Loading Position Effect on Model Stiffness

The effect of loading position on the generated vertebral model stiffness was similar to the effect seen with human vertebral models in Chapter 3 and is shown in Figure 5.18, for posterior to anterior loads. Broadly, the stiffness was greatest when the vertebral models were loaded centrally, reducing as loads were applied at greater distances from the centre.

Movement of the load from posterior to anterior for models with PC1 varied showed large reductions in stiffness furthest from central loading points, with the effect greatest for the densest vertebrae (+3 S.D.) and least for the least dense vertebrae (-3 S.D.). The reduction in stiffness at 20 mm away from the centre, both posteriorly and anteriorly, was proportional to the stiffness when loaded centrally, with reductions of approximately 70 % for all standard deviations from the mean.

Variations in PC2 showed a very similar response to loading position for all of the generated models.

Finally, the stiffness was mostly uniform when PC3 was varied, except at the posterior loading points, where the larger vertebral body size models (negative standard deviations), showed an increased stiffness when loaded posteriorly compared to the other models.




Figure 5.18: Effect of loading position on the vertebral stiffness for PC1, PC2 and PC3 from posterior to anterior loading points. Loading point 1 and 2 are 20 mm and 10 mm posterior of the central loading point respectively, loading point 3 is the central position and load points 4 and 5 are 10 mm and 20 mm anterior, respectively. Annotations describe the main modes of variation found in each PC.

5.3.3.5 Summary of Variation Characterisation

The main modes of variation were found to be split between the three principal components:

- PC1 represented mainly density variation. Positive S.D.s were more dense, negative S.D.s were less dense. This had the greatest influence on stiffness, with the more dense vertebrae having a higher stiffness.
- PC2 represented shape variation. Positive S.D.s were L5-like in shape, negative S.D.s were L1-like in shape.
- PC3 represented volume variation. Positive S.D.s were largest in volume, negative S.D.s were smallest in volume.

5.4 Discussion of Non-Augmented Generated Vertebrae

The methods developed in this part of the chapter have provided methods of creating models that represent incremental variations of vertebrae across the input set. The methods also provide ways of measuring and understanding the variation that exists within a set of vertebrae. This discussion identifies the quality of the validation and describes the types of variation found within the input set of vertebrae.

5.4.1 Validation

The models generated using the statistical shape modelling approach represented the input set well with their measurements matching closely to the measurements of the input set. The 18 measures of the geometry described the vertebral shape well, with the most important measurements being ones describing wedge shapes and the general shape changes between L1-like and L5-like vertebrae. Wedge shapes are often an indication of a wedge fracture and therefore this was an important mode of variation to capture. General shape changes, between L1-like and L5-like vertebrae were also an important mode of variation to capture, given that the input set includes vertebrae from these levels. The range of variation within 1 standard deviation (all possible combinations of the PC1, 2 and 3) was found to capture most of the variation seen in the input set for all measurements taken. An expectation that the remaining variation not captured in 1 S.D. would be captured in the other two standard deviations seems valid given that one standard deviation should only describe 68 % of the variation. In a few cases, models produced variation that was not seen in the experimental input set; a possibility here is that the combinations of principal components used to make that model were unnatural. Additionally, with a larger and

therefore more varied input set, such variation may be seen in the input models. Models used in the remaining investigations were from principal components in isolation, which all fell within the range of the input set. In addition to the range being well represented, the measurements of the mean generated model aligned closely with the mean of the input set measurements, suggesting that the models are not skewed to one end of the variation spectrum.

Validation of the material properties of the models through a visual comparison showed that the mean generated greyscale background matched closely with the greyscale backgrounds of the input set. The generated greyscale backgrounds lacked only in contrast and definition of the trabecular structure which is expected given the averaging nature of the approach. This averaging effect on the background is caused by changes to the vertebral bone, that at a trabecular level and resolution used here, do not follow any trends. For example, given two geometrically similar, L1 vertebrae with a similar internal density distribution, the trabecular structure will still vary significantly when the models are registered together. Hence, any model generated using this approach will not have a greyscale background with the contrast seen in the input set. While the clarity in the generated backgrounds is reduced compared to the input set using the BV/TV method (Section 3.3.1), more definition is likely preserved when compared to the normal method, given the background slice images compared in Chapter 3. Also, as seen on the density colour maps, a clear definition of the cortical shell was visible on all of the generated models, both when the mean density was high and low. This clear definition of the cortical shell is important feature given the importance of the cortical shell in load transfer [67] and the results shown in the previous chapter regarding the use of the bone volume fraction method. Overall density was well represented with comparable mean density of the input set and generated models. Ranges of the means were also similar, with less variation seen in the generated models, however, only 1 S.D. was measured explaining the reduced range compared to the input set. The comparisons of the greyscale histograms to the range of histograms determined from the input set showed that the range of variation and quantity of different greyscale values were broadly captured. The histograms reiterate the loss of contrast in the generated models with the histograms having higher, more centralised peaks, lacking the brighter and darker elements.

Given that the response to loading in the form of stiffness measurements is the primary measure in this study, ensuring a correct representation of the stiffness in the generated models was the most important variable to validate. Both the range of stiffness values from the generated models and the stiffness of the mean model were very comparable

to the range of the input set and the mean of the input set. This suggests that despite the reduced contrast and accuracy in the greyscale background of the generated models, enough information is captured to accurately represent the linear loading behaviour.

Despite the first three principal components only describing 57 % of the variation within the input dataset, the majority of the geometric variation and density variation was captured in the models. Additionally, in the case of the geometric variation, most variation was described within a one standard deviation limit. A possible explanation of what the remaining 43 % of the variation describes is found when comparing the sagittal plots of samples from the input set (Figure 5.19) with those of the generated models in the results above. It is clear from this comparison that much of the noise, bumps and osteophytes present on the input set do not exist in the generated models. Additionally, osteophytes that do exist on some of the generated models would not be described by the broader 18 measures of vertebral geometry used for validation. This means that while the 18 measurements suggest all variation is being captured, some is missed.

5.4.2 Variation

The variation found in the statistical shape and appearance model and the models generated from it capture the variation found in the input set well, as described above. While this variation may not describe the variation found in the patient population that would expect to receive vertebroplasty, it does represent a broad range of shape and material property changes given the relatively small size of the input set. This range would not be seen if shape models were created for each vertebral level in isolation, such as in the study by Hollenbeck et al. [81].

PC1 described the largest part of the variation and can be characterised as a description of the density variation. This density variation changed both the mean density and the distribution within the vertebral body, as seen in the colour maps. Of all the variation, changes to the density through the greyscale background changes had the largest effect on the stiffness of the vertebrae, with a variation of ~ 5000 N/mm between stiffest and least stiff. While the change in mean density varied linearly across the principal components, the resulting stiffness did not. The resulting stiffness appeared to reach a plateau at both positive and negative 3 standard deviations from the mean, suggesting that the density distribution (not just the mean density) or minor shape changes have a role in the overall stiffness (given that the volume change is also linear). While variation in the volume and shape of PC1 are small, the variation is present and follows a pattern where the smaller vertebrae have the lowest mean density and the smallest resultant stiffness. Larger

vertebrae would correspond to a larger person and therefore higher loads transmitted through the spine, which would explain the scenario described. However, this trend is not represented in PC2, where larger vertebrae have lower mean greyscale values. In PC3, where the greatest volume variation exists, there is little density variation, unlike what would be expected. This may suggest a limitation of the small sample size of the study, where not enough vertebrae exhibit the trend where larger vertebrae have higher density. Alternatively, the two variables may only be correlated over small volume changes, not larger volume changes as in PC3. As with all of the correlations described here, the sample size needs to be expanded to confirm whether they are population wide relationships or limited to the input set used. The density distribution across PC1 does not show shifting between the posterior and anterior of the vertebral body although it does show differing contrasts between the cortical shell and cancellous bone. For example, the negative S.D., least dense vertebrae showed a much stronger cortical shell to cancellous bone density ratio. Additionally, the negative standard deviations showed a much thinner cortical shell when compared to the mean model and positive standard deviation. This may represent the osteoporotic nature of the ageing vertebrae where the trabecular bone quantity is greatly reduced along with the thinning of the cortical shell [108, 119].

Variation in PC2 contained a similarly small amount of volume variation, however, the shape varied significantly. Axial mid-slice views of the vertebrae generated from PC2 show the variation most clearly with negative standard deviations being more L1-like and the positive standard deviations being more L5-like. More specifically, the negative standard deviations were much narrower and had a much smaller axial and coronal cross sectional area. Sagittal views of the negative standard deviation models show an anterior wedge shape, mimicking the shape of the most superior lumbar vertebra. The positive standard deviations had a posterior wedge shape, again mimicking the shape of the most inferior lumbar vertebra, following the inflection of the spinal curvature at L3. This agrees with the shapes seen in the input set (Figure 5.19) and with results discussed in the literature [9] where opposing wedge shapes exist either side of the L3 vertebra. The positive standard deviations also have a much larger axial cross section with extending pedicles that match those of the L5 vertebrae. The shape changes were accompanied by relatively small mean density changes and similarly small changes to the stiffness. As with PC1, the negative standard deviations had the smallest stiffness, however, the stiffness peaked at +1 S.D. and not +3 S.D.. This may be due to either the shape differences or the density distribution. These are discussed in more detail in Section 5.4.3, where different loading positions help to illuminate the mechanics of all of the principal components.

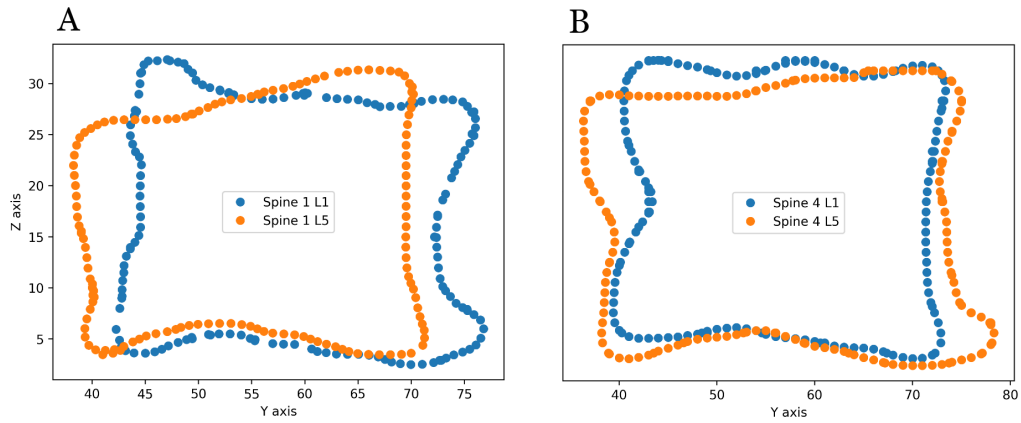


Figure 5.19: Sagittal views of the L1 and L5 vertebrae from spine 1 and spine 4 of the input set in A and B respectively, showing the opposing wedge shapes of the two.

The variation captured in PC3 consisted of mainly volume changes, with little changes to the shape and density. The primary mode of variation here, volume, had a flipped relationship with the positive and negative standard deviations. Conversely to the other two PCs the positive standard deviations have the smallest volume and the negative standard deviations have the largest volume. However, the direction of variation or sign of the standard deviation is arbitrarily assigned and hence has no influence or holds any meaning on the results. The shape variation that did exist showed a similar trend to PC2 in the sagittal plane, where the smaller +3 S.D. model had a posterior wedge shape similar to the smaller L1-like vertebrae. While the mean density remained nearly constant in this PC, a clear change to the density distribution was present. The larger negative standard deviation models had a very high posterior density that shifted to the anterior in the positive S.D. (and smaller) models. Finally, variation in the stiffness of PC3 matched closely to PC2, although reversed, where the peak in the stiffness is in the -1 S.D. model. Again, this relationship is discussed in more detail in Section 5.4.3, where different load positions helps to explain the trends.

The only comparable study in the literature that investigated the variation in individual lumbar vertebrae, rather than the spinal section, was a study by Hollenbeck et al. [81]. This study aimed to characterise individual lumbar vertebrae though building statistical models and to characterise the shape and alignment of the spinal section as a whole, including specific functional units. The two main differences between that study and the one presented here is the inclusion of L1 to L5 vertebrae and the use of material properties of the bone both in this study. In the Hollenbeck study, material properties were not included in the model and the variation in the vertebral levels was investigated in a level by level approach. A very similar quantity of variation was captured in the principal components between the two studies, with approximately 61 % and 66 % variation captured

in the first four principal components in the Hollenbeck study and this study respectively. This is an unexpected result given the 55 input vertebrae in the Hollenbeck study and only 14 in this study. The main mode of variation found by Hollenbeck was a general scaling variation. This difference from the current study is due to both the inclusion of material properties and the inclusion of vertebrae from all lumbar levels, otherwise PC3 (volume change/scaling) would most likely be the primary mode of variation matching the result by Hollenbeck. Hollenbeck et al. also included the posterior elements in the models and so many of the of latter modes of variation included changes to the angle of these. The study also reports the same changing vertebral dimensions from L1 to L5 that are found in PC2 in the current study, although the measurements in that study were from separate SSMs.

5.4.3 Load Position Effect

The effect of load position was investigated for the human tissue, where the possible error in choosing the loading point of the model was ± 1 mm given the resolution. The effect of moving the loading position within the 2 mm range gave changes to the stiffness of approximately 5 % (a small contribution to the total possible error). Here, the changes to the loading points are ± 10 mm and ± 20 mm, giving loading points at approximately the posterior boundary, the midpoint between the posterior boundary and the central loading position and similar points on the anterior side. This gave a reasonable understanding of the effect of anterior and posterior loading for the generated models, both with and without augmentation and more importantly a more detailed understanding of the effect of augmentation on different types of vertebra in the following sections of this chapter.

The response to loading for the different non-augmented models in the PC1 variation set followed expectations, given the incremental change in the mean density. In varying PC2, the near identical response from most of the generated models is explained by their similar volumes and similar mean greyscale values. The -1, -2 and -3 S.D. models that do not fit the same trends showed a reduced posterior density and hence showed a reduced stiffness when loaded posteriorly. A similar effect was seen when varying PC3, where the increased posterior stiffness of -2 and -3 S.D. models compared to the other models is explained through a large increase in the posterior density. The increased anterior density for the +3 and +2 models in PC3 seemed to have little effect on anterior loading.

5.4.4 Summary of the Non-Augmented Generated Vertebra

The models that have been created from the SSAM have provided both a well validated (to the input set) and an interesting range of vertebral variation to test the effect of augmentation in the remainder of this chapter.

5.5 Methods of Characterising Augmentation

The non-augmented vertebral statistical shape and appearance models described above have been used to evaluate the effects of cement augmentation, which is described in this section. The aim of this study was to identify the effect of vertebral variation on the outcomes of augmentation (in terms of changes to the vertebral stiffness). The models used to test the effect of augmentation were those used for the characterisation of variation and are listed in Table 5.1.

5.5.1 Artificial Augmentation of Generated Models

It was desired to test the effect of augmentation on the different modes of vertebral variation found in the generated models. Augmentation in the generated models was carried out through the artificial augmentation process using the approach described below. Different size augmented regions and different positions of the augmented regions were examined. This was to help to understand how variation in the augmentation procedure (position, fill volume) is affected by the vertebral variation.

Two masks were created, one for the cement region itself and one for the interface region. A mask based on an eroded vertebral body mask was used for the shape of the augmented region. For this, a copy of the vertebral mask for each generated models was made, a recursive Gaussian filter 5 pixels deep was applied and then it was converted to a surface. The surface could then be scaled to represent different percentage fill volumes. The percentage fill was determined by dividing the mean cement volume by the volume of the vertebra in question. The desired augmentation volume was then selected. This was either 20, 35 or 50 percent fill, given the range of fill volumes achieved experimentally. The cement volume was then scaled by $\sqrt[3]{\frac{\text{desired \% fill}}{\text{current \% fill}}}$ on each of its axis. This gave three fill volumes for each generated model which was positioned in the centre/anterior of the vertebral body, then converted to a mask. The mask was then duplicated and eroded by one pixel to give the yielding interface layer upon which the correct material properties were applied.

5.5.2 Augmented Region Volume Sensitivity

The sensitivity of the generated models to the size of the augmented region was carried out using the methodology described above to explore the effect of 20 %, 35 % and 50 % fill volumes positioned centrally. This was combined with the methodology described in Section 5.2.5 to understand how the augmented region volume affects the response to alternate loading positions.

5.5.3 Augmented Region Position Sensitivity

Similarly to tests carried out with the bovine tail vertebrae, the position of the augmented region was tested. The mean generated model was used for this test along with the mean generated volume of cement. This mean volume of cement was created by utilising the same PCA plugin for ScanIP as for the vertebral models, however the input models used here were those of the augmented region in the post-augmentation models. The mean volume of cement was moved through the vertebral body, from the anterior to posterior on the Y axis, from left to right on the X axis and from the most inferior point to most superior point on the Z axis. This was carried out by converting the mean generated volume of cement into a surface mesh. This surface mesh could then be positioned using the surface tools in ScanIP. The same limits used in the bovine tail study were implemented where a 1 mm space between the cement volume and the edge of the vertebral body was required, to preserve the cortical shell. The effect of changing the augmented region position in 2 mm increments (x, y and z axis) was measured by way of the axial stiffness. An illustration of the movement on the Y axis (posterior to anterior) is shown in Figure 5.20.

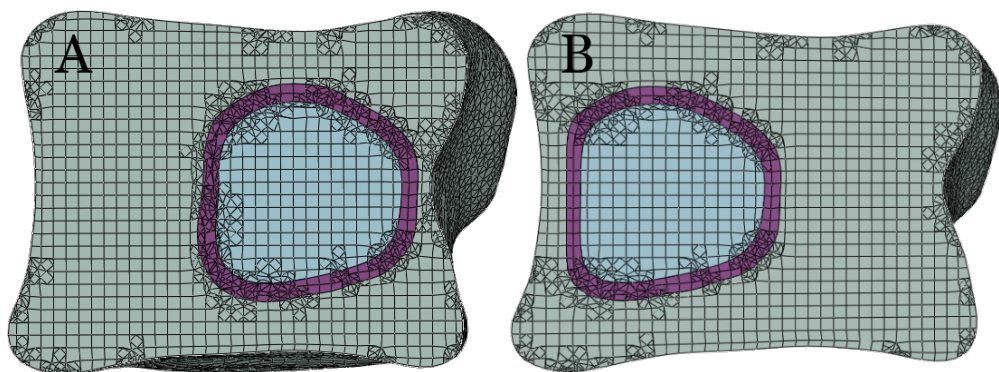


Figure 5.20: An illustration of the loading positions in the mean model. A showing the 0 position on the Y axis, where the volume of cement is positioned posteriorly and B showing the 12 position on the Y axis where the cement volume is positioned anteriorly.

5.6 Results of Augmentation in Generated Models

5.6.1 Influence of Augmentation Volume on Vertebral Stiffness

The general effect of augmentation in the generated vertebrae can be seen in Figures 5.21 to 5.23, with the effect on stiffness isolated to each principal component. The vertebrae described by variations in PC1 showed an increase in stiffness with increasing cement fill volume, except for the more dense negative standard deviations away from the mean, where a reduction in stiffness was seen for all cement fill volumes. A bimodal relationship was also observed in the vertebrae generated from variations in PC2 and PC3 when augmentation was simulated. When PC2 was varied, positive standard deviations from the mean showed a more pronounced increase in the stiffness at the larger fill volumes when compared to the negative standard deviations. When PC3 was varied, increasingly negative standard deviations from the mean had a more noticeable increase in the stiffness at higher cement fill volumes, correlating with the vertebrae with larger vertebral body volumes.




Figure 5.21: The effect of augmentation on vertebral stiffness in the vertebrae generated from principal component 1, for 20, 35 and 50 percent fill volume. The results show the change from the non-augmented stiffness increasing with fill volume, with a more pronounced effect seen in the less dense vertebrae (-1, -2 & -3 S.D. from the mean).




Figure 5.22: The effect of augmentation on vertebral stiffness in the vertebrae generated from principal component 2, for 20, 35 and 50 percent fill volume. The results show the change from the non-augmented stiffness increasing with fill volume, with a more pronounced effect seen in the more L5-shaped, positive standard deviation models.




Figure 5.23: The effect of augmentation on vertebral stiffness in the vertebrae from principal component 3 (describing mainly volume variation), for 20, 35 and 50 percent fill volume. The results show the change from the non-augmented stiffness increasing with fill volume, with a more pronounced effect seen in the larger vertebral models (negative standard deviations away from the mean).

5.6.2 Influence of Augmented Region Position on Vertebral Stiffness

The resulting stiffness values of the mean generated vertebra when the augmented region was moved within it are presented in Figures 5.24 to 5.26. The movement in the x axis, from left to right in the frontal/coronal plane showed little change to the stiffness for the middle five augmented region positions. There was a symmetrical reduction in the stiffness at the furthest left and furthest right positions where the augmented region encroaches on the denser cortical bone. When the augmented region was moved along the Y axis, from the posterior to the anterior, an increase in stiffness was observed when the cement is positioned in the anterior portion of the vertebral body. Finally, when the augmented region was moved axially the greatest stiffness was recorded when the volume of cement was in the most superior 4 mm of the vertebrae.

The percentage fill of the mean generated cement volume in the mean generated vertebral model was 12 %. Regardless of the augmented region's position, a reduction in the stiffness was recorded, from a non-augmented stiffness of 5045 N/mm. The range of stiffness values recorded when moving the augmented region was between 4298 and 4430 N/mm.




Figure 5.24: The stiffness response to augmented region movements on the X axis, from left to right when viewing from the coronal plane.




Figure 5.25: The stiffness response to augmented region movements on the Y axis, from the posterior to the anterior of the vertebral body.

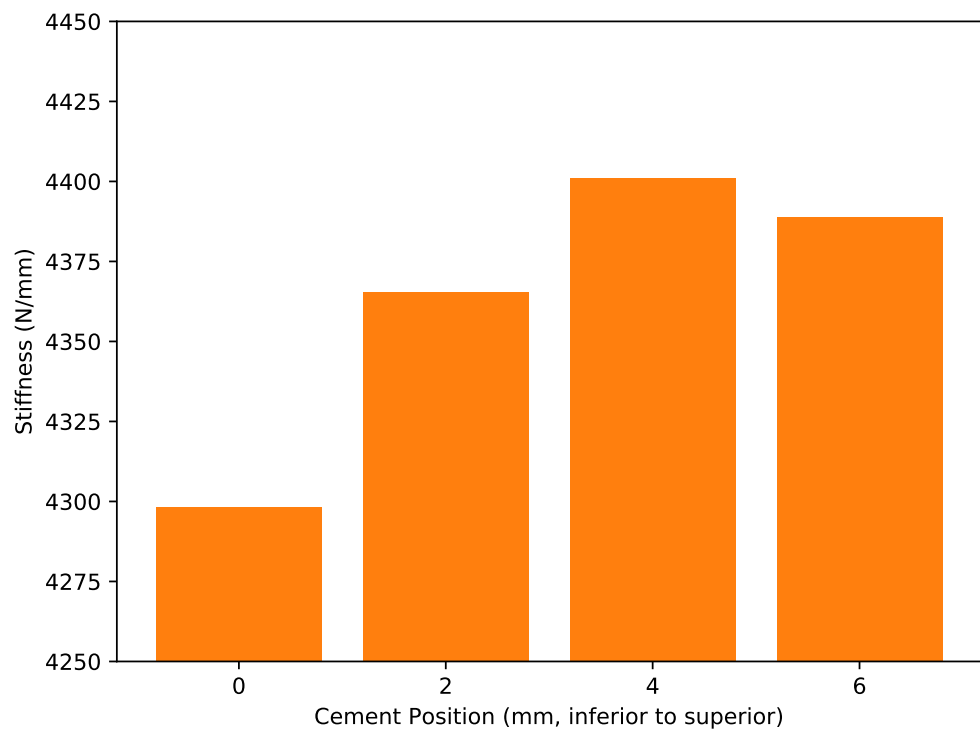


Figure 5.26: The stiffness response to augmented region movements on the Z axis from inferior to superior positions.

5.6.3 Loading Position Effect on Model Stiffness

The effect of loading position on the augmented vertebrae is presented in Figure 5.27, following on from the load position effects on non-augmented vertebral models in Section 5.3.3.4 and discussed in Section 5.4.3. At 20 % fill volume, the loading position had little effect on the response to augmentation, with the percentage change less than 5 % for all but variations in PC1. Here, only the -3 S.D. model showed a change compared to the central loading position, with an increase in stiffness when loading posteriorly.

At 35 % cement fill, trends became apparent which continue and became more noticeable at 50 % fill. Variations in PC1 had a simple relationship, where the augmentation had a smaller effect when loaded at the most posterior or anterior loading positions compared to the centre for the negative standard deviations. However, in the positive standard deviations (the more dense vertebrae), the change in stiffness was greater when loaded anteriorly compared to posteriorly.

Variations in PC2 had a near uniform reduction in stiffness when loading posteriorly compared to the central loading position for 35 and 50 % fill volumes. When loaded anteriorly, augmentation had little effect on the negative standard deviation models, where the response to augmentation was similar for each loading point. For the positive standard deviation models however, there was a large increase in the anterior stiffness following augmentation, with a doubling of the change in stiffness seen in the +3 S.D. model when loaded anteriorly compared to centrally.

Variations in PC3 had a similar response to augmentation at different loading points to changes in PC2, where posterior loads had little effect, or a reduced response compared to central loading, except for the +3 S.D. model, where augmentation has the greatest effect on stiffness when loaded posteriorly. Augmentation increased the stiffness most when loaded anteriorly for the negative standard deviations, representing the largest vertebral body volumes.




Figure 5.27: Effect of loading position on the augmented vertebral stiffness for PC1, PC2 and PC3, and 20, 35 and 50 % fill from posterior to anterior loading points. Loading point 1 and 2 are 20 mm and 10 mm posterior of the central loading point respectively, loading point 3 is the central position and load points 4 and 5 are 10 mm and 20 mm anterior, respectively. Annotations describe the main modes of variation found in each PC.

5.7 Discussion of Augmented Generated Vertebrae

The methods developed in the second part of this chapter have allowed an investigation into the mechanics of vertebral augmentation that is not possible with specimen specific modelling approaches. Few studies have identified the variation in vertebrae from the lumbar spine and the novelty here is the investigation into both the variation across the lumbar section and the response to augmentation. The results acquired may provide suggestions for clinical practice, although a larger quantity of input specimens would be required to apply any results more broadly.

This discussion identifies the consequences of variation on augmentation.

5.7.1 Response to Augmentation

The response to augmentation can be measured and understood much more clearly in this chapter compared to the Bovine Tail Vertebrae chapter and the Human Lumbar Vertebrae chapter. This is mainly due to the isolation of vertebral variation and removal of experimental variation, hence the results described here are in much more detail than the previous chapters allowed.

Broadly, the methods used to augment the generated models were very similar to those used for the experimental specimens with matching material properties and interface layers. A single difference in the methodology (other than being able to choose the position of augmentation) was the description of the needle track. In the previous chapter the description of the needle track was found to provide a large increase in agreement of the models compared to the experiment. Their absence here is due to the effectively parametric nature of including them. For every vertebrae, the needle approach varies and hence the shape and volume of the augmented region vary. Given the lack of defined vertebral levels in these generated models, modelling the needle track correctly would be difficult. However, along with variations in the augmented region shape, it would be a valuable piece of further work to accompany the study. It would allow an investigation of the effect different needle approaches and its effect on both the resulting stiffness and the type of augmented volume expected from the needle position.

The response to small quantities of cement injected into the vertebrae was found to be small. This is shown when comparing 20 % fill volumes to 35 and 50 % fill, where an increase in the stiffness is only seen in the least dense vertebrae. The largest change in the response to augmentation is seen in PC1, where it is clear that the density variation is the driving cause of the response. Here, the change to the non-augmented vertebrae is also the

greatest of the principal components and those that had the lowest non-augmented stiffness showed the greatest increase following augmentation. In these models the very weak, low density and low modulus regions were supported and stiffened by the inserted augmented regions, especially when the fill volume was larger. With the more dense vertebrae the stiffness reduced following augmentation. While the cement region has a higher Young's modulus compared to the majority of the trabecular bone, the yielding interface is most likely to be the cause of the reductions in stiffness recorded, by replacing the dense bone with a material of low Young's modulus and plasticity. The interface captures and describes the damage caused not only by the needle tracks but also damage to the trabeculae from the cement injection as described by Aquarius et al. [102], suggesting that the reduction in stiffness is a valid result.

The response to augmentation in the models where PC2 was varied was reduced when compared to PC1 where the maximum increase in stiffness was 15 % compared to the 100 % in PC1. This corresponds to the change in the non-augmented stiffness, where much less variation was found. With PC2, conversely to PC1 there is no correlation between the non-augmented stiffness and the augmented stiffness, where the least stiff non-augmented vertebra showed the smallest change to the stiffness. This may suggest that the relationship between the pre and post augmentation stiffness only exists when vertebral density is the primary cause of the stiffness variation. The cause of the different response to augmentation in this principal component (PC2) is the changing shapes, where narrower L1-shaped vertebrae show a smaller response to augmentation and vice versa. This is due to the type of bone that surrounds the augmented regions or the type of bone that is replaced by the augmented region. In the L5-like, +3 S.D. models much of the internal structure is very low density and hence when replaced by a higher modulus region the increase in stiffness is greater. The converse occurs for the L1-like vertebrae where more dense bone is replaced by the augmented region.

The change in stiffness seen in the PC3 variation set extended higher than the largest change seen when PC2 was varied, while the smallest change in stiffness was greater than any of the of the other principal components. This means that, regardless of the volume change (the main variation in PC3), if a cement fill volume greater than 20 % is used an increase in stiffness is assured. The reasons for the variation following augmentation in this principal component are explained in and discussed with the addition of load position data below in Section 5.7.2.

Despite some studies suggesting that small quantities of cement ($\leq \sim 20$ % [44, 54, 55, 107]) are enough to increase or restore the stiffness, or other properties of fractured vertebrae to

that of the natural vertebrae, the results presented here suggest otherwise. The results presented here do match those presented in the previous chapter however, with increases in stiffness only seen in the larger fill volumes and in vertebrae with lower bone volume fractions.

5.7.1.1 Response to Augmented Region Position

The reduction in stiffness seen when the cement was moved to lateral extremes was because the yielding interface of the cement volume encroached on the denser cortical shell defined by the greyscale background of the models. The effect of the intermediate cement positions had a near symmetrical appearance which was a consequence of the near left-right symmetry of the mean generated vertebra.

Inferior to superior movements of the cement showed a small change to the stiffness with a 2 % increase to the superior of the vertebra, though a reduction when at the topmost point of the vertebra. This can be explained by understanding the greyscale distribution of the model background or the material properties of the elements. The elements found in the centre of the vertebra are less dense and have a lower Young's modulus; when the cement was positioned in the centre, less of the low Young's modulus elements were exposed, resulting in an increased stiffness.

Moving the cement volume in a posterior to anterior direction showed a reduction in stiffness when at the extremes of the vertebra similarly to the lateral movements. Away from the extremes, the stiffness rose to the anterior of the vertebra, where the bone is less dense in both natural specimens and the generated models. This may suggest that vertebroplasty injections should always aim to the anterior of the vertebra (for purely stiffness increasing reasons, as well as safety). However, the percentage change in stiffness from the least stiff - most posterior, to the stiffest position is a mere 3 %. Given that the most stiff position still had a reduction of 13 % compared to the non-augmented vertebrae, it suggests other factors may be more important and that the cement position, at least for small quantities of cement, has little effect.

As suggested previously the small change in the stiffness when moving the augmented volume within the vertebrae may be due to its small size. Larger changes in stiffness were only recorded when using between 35 and 50 % fill, suggesting that larger changes to the stiffness when moving the cement volume could also be seen if larger augmentation volumes were used. However, problems would be met regarding fitting the augmented volume within the vertebrae and maintaining the 1 mm boundary. A problem that occurred when using the eroded vertebral mask, where at 50 % fill there were approximately 1 to 3 mm between

the augmented region and the vertebral body walls. This would therefore limit the amount of movement that could be experimented with.

The method of artificial augmentation used here (the eroded vertebral body mask method) is not a perfect solution as it is unreasonable to expect that every vertebrae receives a volume of cement that matches its own shape. However, it does remove problems with the cement region extending outside of the vertebrae (if the mean cement shape at 50 % fill was used). Perhaps an improved method for fitting the cement volumes would be to use the mean cement volume but transform it using the reverse of the transforms that were applied to each generated vertebrae turning it into the mean vertebrae shape. The current version of the plugin does not allow access to these transforms however. Despite this, changing shapes and volumes of the augmented region should be the subject of any further parametric analysis, given that an optimum fill percentage would need to be found.

5.7.2 Load Position Effect

The response to loading position when the vertebrae have certain fill volumes of cement show results that are not obvious without the loading position investigation. Changes and patterns that were present in the load position results were merely amplified through the addition of larger volumes of cement. At 35 % fill the effect of the loading position is clear, and mimics (to a smaller degree) what is seen in the 50 % fill volume models. The increased change in stiffness for the positive standard deviation models in PC2 when loaded anteriorly may be a combination of the shape of the models (the L1 to L5 shape shift from -3 to +3 S.D.) and the shifting density distribution from anterior to posterior with -3 to +3 S.D.. Given that the +1, +2 and +3 S.D. models within PC2 have a lower anterior density, the support given to this region through augmentation may explain the increased change in stiffness. Additionally, the positive standard deviations exhibit greater anterior height. The combination of the additional height and density may explain why the increase in stiffness reached a maximum of +33 %, greater than the +25 % maximum seen in anterior loading of the PC3 vertebrae that exhibit the same shifting density distribution. Posterior loads have little effect to the vertebrae in the PC2 set, with all generated models showing a very similar response, potentially adding weight to the idea that the anterior loading response is shape driven. The negative standard deviation models (L1-like) in PC2 showed a near uniform response to the different loading positions, this may be due to the smaller size along with a smaller anterior height. The reduced height suggests an increase in density (given the near uniform mean greyscale values for PC2), which, as with PC1 causes a reduced response to augmentation.

When loading is varied in the PC3 set there was a reverse relationship between the +2 and +3 vertebrae compared to the other models (+2 and +3 S.Ds showed a greater response to posterior loads, others showed greater response to anterior loads and a reduced response to posterior loads). The main variable in PC3 is the size variation, however given that the cement fill volume was a constant 50 %, the cause of the varied response was likely the change in material properties. Like PC2, this change in material properties is the shift in the density distribution where smaller +2 and +3 S.D. vertebrae have a less dense posterior region. This resulted in an increase in posterior stiffness following augmentation for these models by supporting the less dense materials with cement. The remaining vertebrae have very dense posterior regions, explaining the reduction in posterior stiffness following augmentation - the same effect seen in the densest vertebrae of the PC1 set. Conversely, the reduced anterior density explains the increase in stiffness following augmentation when loaded anteriorly, where support is given to the least dense anterior region.

Further isolation of the variation in the principal components would aid the understanding of the true cause of the different responses to augmentation. For example, the combination of shape and volume changes in PC2 and PC3 with the density distribution changes. Potential risks of this however are that such variation types do not exist in isolation, at least in this dataset, for example a L1 to L5 shape variation may always exhibit the density shift. This may be due to the load sharing variation caused by the bending of the lumbar region. The central lumbar vertebra (L3) may experience more central loading, with the L1 and L5 vertebrae experiencing anterior or posterior loads depending on the degree of lordosis in the lumbar spine. The density shift in PC2 to PC3 where the larger vertebra (likely L5 vertebrae) showed an increased posterior density, suggests that that posterior of the vertebrae experience larger loads in the body. Again, like PC2, the smaller vertebrae in PC3 (likely L1 vertebrae) showed greater anterior density and likely higher anterior loads in the body.

5.8 Conclusion

The methodology shows that this approach is able to accurately capture and describe the variation found in an input set. Tests on the vertebrae generated from this methodology show a behaviour similar to the input set with the mean and range of variation well represented. Using these models to investigate augmentation has provided a much deeper understanding of the mechanics of augmentation, especially from the results of using different loading positions. These results emphasise the relationship between the density of the vertebrae and the response to augmentation that was found experimentally. It also

adds detail to this picture, describing how the bone density distribution varies between generated models and likely between spinal levels.

Chapter 6

Discussion and Conclusion

6.1 Introduction

The aim of the work presented in this thesis was to develop methods of experimentally testing and modelling vertebrae, and to understand the mechanical effects of vertebroplasty over variations seen in a patient population. Firstly, a set of bovine tail vertebrae were used to develop methods of experimentally testing vertebral specimens and building models of both the augmented and non-augmented states. These results allowed a swift transition into using the much more limited and valuable human lumbar vertebrae. Here, methods were refined and improved both on the experimental and computational side to produce computational models that were well validated against the experiment in terms of stiffness. Finally, the specimen specific models formed the input set for the development of a statistical shape and appearance model. This allowed an investigation into the modes of variation that exist within the input set. These modes could then be tested to identify how different modes of variation influence the mechanical response to vertebroplasty. The methodologies and results of the individual studies were discussed in each of the previous three chapters. Here, an overview of the results from across the chapters is presented and their meaning discussed from a clinical perspective. Finally recommendations are made for future studies and continuations of this research.

6.2 Discussion

6.2.1 Experimental Methodologies and Results

In vitro tests of vertebrae under axial compression were performed on both bovine tail (Chapter 2) and human (Chapter 3) specimens. Generally, experimental testing of the human lumbar vertebrae was much more repeatable than the bovine tail vertebrae. Many

of the properties of the bovine tail vertebrae made experimental testing more difficult. These include the smaller, more cylindrical shape which made experimental potting and loading more difficult, and the increased density which hindered the augmentation process. With the human lumbar vertebrae, the wider axial cross section allowed more consistency in the alignment of the vertebrae when potting. This reduces the error, ensuring a true axial load rather than the inclusion of shear forces.

Augmentation of the human lumbar vertebrae was more straightforward due to not needing the mallet for needle insertion or requiring the use of a water bath to allow the marrow to be displaced. While the advantages of having bone marrow at the clinically relevant body temperature are clear, the un-measured changes to the hydration of the vertebrae are not known or measurable and therefore introduce a range of variation to the results. An additional limitation of the bovine tail vertebrae was the void region located in the centre of the vertebral body, which made measuring bone properties difficult as the volume of interest usually contained this void region. The limitations of the bovine tail vertebrae described above and previously published results of using other non-human vertebrae, all of which have their own limitations, suggest that no truly representative alternative to the use of human vertebrae exist. While animal vertebrae can perform an excellent test material for different methodologies, their considerably different material properties, geometries and behaviours mean that cadaveric human vertebrae remain the best tool to understand the effects of different treatments.

The stiffness of the two types of vertebrae (bovine and human), despite the large variation in shape and density, remained comparable. Ranges of the stiffness values were very similar, with the mean stiffness being similar for both, 5101 N/mm for the bovine tail vertebrae and 4599 N/mm for the human vertebrae. A similar range of results was found following augmentation, with many vertebrae showing reductions in stiffness following augmentation for both the bovine and human vertebrae. While the augmentation process became easier because of the reduced density of the cortical shell, the similar number of vertebrae that showed a reduced stiffness following augmentation suggests that similar levels of damage were caused in both specimens. Much larger augmentation volume fill percentages were achieved with the human lumbar vertebrae. This was due to both differences in the vertebra themselves, with the human specimens having a reduced density of trabecular bone and procedural changes, with the use of side opening needles greatly aiding the injection quantity. These changes allowed the injection of much more clinically relevant volumes of cement when compared to the bovine tail vertebrae.

The limitations of the bovine tail vertebrae have been discussed above with alterations

made before using the human tissue; limitations of the human tissue methodology surround specimen numbers and range of variation. As with most studies that utilise statistical analysis of the data, increased numbers of specimens aid the identification of trends. Here the limitation comes from the preliminary and developmental nature of the study. This means that results that were directly influenced by the number of specimens may benefit from the inclusion of further specimens. For example the derivation of the greyscale conversion factor and for finding trends between specimen properties and the outcomes of augmentation. The benefit of a smaller number of specimens was that the methods could be developed faster and allowed the progression of the study to include statistical shape and appearance models. Another limitation was the range of spines included, while the likelihood of osteoporotic compression fractures increase with age, a comparison to a younger, possibly less osteoporotic spine would have been a useful comparison. The high levels of degeneration in some of the spines (spine 2 and 3) also meant that not all of the vertebrae could be used, meaning that the set of vertebrae was weighted to the L1 end of the lumbar section.

6.2.2 Computational Methodologies and Results

Specimen-specific models of both bovine and human vertebrae were developed to simulate the experimental test, enabling one to one comparisons with experimental data. The level of agreement increased from bovine to human vertebra in both non-augmented and augmented cases. While the quality of agreement in the bovine tail vertebrae matched closely to that found in similar studies with the same tissue [91], the agreement for both non-augmented and augmented human vertebrae exceeded this. Conclusions were made regarding the likely sources of error/disagreement between the augmented bovine tail vertebrae models and the experimental results, suggesting the problem lay in the bone properties and not the methodology. However, following the same methodology in the human specimens, where much less damage was caused to the bone due to their reduced density, a similar level of agreement was seen between the computational and experimental results. This suggests that the method was the limiting factor in the agreement, not the damage caused to the denser bone. Given that the agreement was increased significantly when using the bone volume fraction method, it suggests that capturing the description of the cortical shell and internal trabecular structure is vital for a good correlation between experimental and computational results. Despite changes to the greyscale background origin and therefore the conversion factor between greyscale value and Young's modulus, the optimum properties for the augmented region and interface region remained very similar between the greyscale and BV/TV methods. Further improvements of the modelling method were made for the human

vertebrae with the inclusion of the needle tracks in the augmented models. Capturing and describing this damage is clearly necessary for an accurate model of augmentation, and there is a need for further studies to identify and model the damage caused to the bone more accurately. Additionally, given the clear reductions to the stiffness following the experimental repeated loading of post-augmentation specimens, fatigue due to cyclic loading should be included into the model to further understand the propagation of damage within the models, similar to the models created by Coe [120].

The computational models based on the SSAM were found to share the range of geometric, structural and stiffness variation with the human input models from which they were derived. Broad measures of the geometry agreed well with the measurements of the input set and while localised geometric variation was not included in the model, the agreement in terms of stiffness suggests that such details play a negligible role in the mechanical behaviour. Greyscale background contrast between the cortical shell, trabecular bone and empty space was less well defined in the generated models from the SSAM, however, enough detail was captured to produce models with comparative stiffness values to the input models.

The main modes of variation appeared to capture the variation of the input set qualitatively, and match that seen in elsewhere in the literature [81], if such studies included a range of vertebrae along with density data. The main mode of variation also matched the main variable that affected experimental augmentation: the density. The only discernible relationship following augmentation between specimens and augmentation behaviour was the density of the model. Specimens with lower BV/TV received larger volumes of cement and were the only specimens that showed a consistent increase in stiffness.

The changing greyscale distribution worked alongside the shape and volume changes in PC2 and PC3 to give different responses to augmentation. The greyscale distribution in these two principal components was almost isolated from variation in the mean greyscale density and hence allowed a more clear visualisation of its effects on augmentation. Comparing the density maps of the human vertebrae with the generated models, it was difficult to identify the same trends in the input set as those seen clearly on the generated models. For example posterior to anterior density trends were not clear between L1 and L5 vertebrae of the input set or between different sizes of the vertebrae. This highlights the main advantage of any dimensional reduction approach or the PCA technique used here, relationships can be identified that would otherwise be hidden.

One limitations of the current computational study, including the SSAM, is that a more

thorough investigation into different types of augmentation volumes was not included. This would help identify the effect of end-to-end augmentation volumes, which was found to have large effects in other studies [103]. Experimentally, little variation was found in the cement distributions, other than dispersed or concentrated volumes and hence no comparisons could be made. Additionally, needle tracks were included in the specimen specific models, therefore a limitation is their non-inclusion and lack of investigation in the generated models.

6.2.3 Application to Clinical Practice

From the experimental results the vertebrae that appeared most suitable for vertebroplasty were identified as having lower density or a smaller bone volume fraction. These were vertebrae that both allowed larger fill volumes and allowed concentrated volumes of cement, which in turn allowed a larger stiffness response. Vertebrae with lower densities also required smaller pressures and resulted in less cement leaking. Both of these factors are of great importance clinically, with leakage being a major concern and higher pressures resulting in damage to the trabecular structure. The differences in mechanical performance of the dispersed or concentrated volumes of cement are not fully understood, especially given the similar observations, yet opposing conclusions in the study by Aquarius et al. [102], discussed in Chapter 4. However, the relationships between the amount of dispersal and the density are clear, as is the relationship between resultant stiffness and dispersion. This perhaps suggests those vertebrae with reduced bone volume fraction should be the primary target for augmentation.

Computationally, the results also suggest that vertebrae with a low bone volume fraction respond more favourably to augmentation. In the results from the SSAM, specifically models which were varied across PC1, where the density variation was greatest, large changes in the stiffness following augmentation were seen in the least dense models. The most dense vertebrae showed the smallest response and in some cases resulted in a reduced stiffness. This, coupled with the experimental evidence of damage due to needle and pressure, all point towards the poorer outcomes for higher density vertebrae.

One of the main limitations regarding the conclusions about clinical outcomes is the lack of understanding of the desired change in stiffness following augmentation. A restoration of the stiffness to the pre-fractured stiffness would be a desired outcome for those vertebrae having undergone a fracture. However, in prophylactic vertebroplasty, where the aim is to prevent fractures in weak vertebrae, the desired change in stiffness is less clear. Perhaps an upper limit would be when damage to adjacent levels due to the increase in stiffness

begins. While studies have investigated the effect of prophylactic vertebroplasty [44, 102] and the consequences on the adjacent levels [96], few suggestions of a desired stiffness have been made. Due to this, it becomes difficult to suggest an appropriate fill volume. The approximately 100 % increase in stiffness of the least dense vertebrae when a fill volume of 50 % was used seems unreasonable, especially when it changed the anterior stiffness by approximately 75 %. Such increases would likely cause adjacent level failure: if the load distribution changes across the vertebra, i.e. the anterior stiffness changes more than posterior stiffness, then this will certainly change the load distribution. This would likely cause an altered load distribution on the adjacent vertebral levels, e.g. more through anterior, less through posterior, very likely causing failure as the bone will not have had a chance to remodel to accommodate these changing loads. Such changes in loading distribution will not necessarily get evened out by the intervertebral disc, given the often degenerated nature of the discs in elderly patients with osteoporotic vertebrae. Further investigations into the stress adjacent levels experience with different cement volumes and locations would be required to understand the desired augmented stiffness.

From the SSAM, the shape of the vertebrae had a much smaller effect on the resultant stiffness than the density. The stiffness response to augmentation in PC2 and PC3, where the general mode of variation was shape (volume or specific shape changes), appeared to be driven by the shape changes. However, comparing the results when loaded in different positions with the density distribution showed a strong relationship between the density distribution and stiffness. This density distribution meant that augmentation could support weaker anterior portions and reduce the density of the posteriorly dense portions. Therefore, guidance based on this finding would suggest focussing on the area of least density when directing the augmentation procedure. Shape changes may influence other factors not investigated here though, where vertebrae with smaller anterior heights may be easier to give endplate to endplate fill volumes and therefore restore or increase the stiffness with smaller volumes of cement. The results of moving cement volumes within the vertebral body also suggest that targeting the least dense (in that case the anterior) region gives the largest stiffness increase.

6.3 Future Recommendations

6.3.1 Further Computational Modelling

The inclusion of damage in the model (in the form of needle track representation and cement interface here), was shown to give a large improvement to the results in terms of

computational to experimental agreement. A further development on this idea would be the inclusion of a larger description of damage into the whole model. This would be in the form similar to that used by Kinzl et al. [66] where damage is accumulated and elements receive a damage score between 0 and 1, where 1 represents complete failure. Methods such as this may help to describe the damage caused by augmentation postulated by Aquarius et al. [102] and other causes of damage such as the high temperatures generated from the cement curing.

Other improvements to the computational methods would include a deeper investigation into the effect of different contacts between the cement endcaps and the vertebrae. While the tied contact between the two materials was removed and replaced by a frictionless contact material, in the experimental model, the contact is not frictionless and hence optimisation of the contact may be required to further improve the agreement.

Also, an investigation into the effects of adjacent level fracture would help elucidate the optimum location for augmentation volumes and the contribution from modes of variation to this optimum. This would be achieved through modelling of a functional unit that includes the intervertebral discs.

6.3.2 Continuation of Statistical Shape and Appearance Modelling

In order to develop the current study from an investigation into the variation of the input set into one that describes the population, an increased number of specimens would be required. Further specimens would help to identify whether the variation that effects the outcomes of augmentation are limited to the current set of vertebrae or extend more broadly. A large cohort of input specimens would better describe the population and the variation within it. This would allow the factors that influence vertebroplasty in this set to be extended to the population, assuming the factors remain similar. While this study has focussed on method development, the application of the approach to a larger input set would enable the ultimate goal of this approach to be achieved - to provide clinicians with guidelines of which vertebrae respond most positively to augmentation from a mechanical standpoint.

With regard to the varying density distribution and the response to loading position, interesting outputs could be drawn. For example, the varying density distribution may affect risk of adjacent level failure if the region superior or inferior of the augmented region is of a particularly low density. Where adjacent level failure is caused by the stiffness change in different vertebral body regions after augmentation [48]. Studies have commented on the changing mean density through spinal level [9, 121], where density generally reduces through

the lumbar sections, attributed to the reduced loads transmitted through these levels. Other studies identified density variation within vertebrae (anterior/superior/posterior/inferior) [32], however did not investigate the effect at different vertebral levels and instead presented means of a set of vertebrae ranging from T9 to L5. Therefore, an interesting further step would be a more thorough investigation into the changing density distribution at different levels. An investigation into the likely relationship between spinal curvature, the load transfer and the resulting trabecular density may add further weight to arguments over which vertebrae should be the targets of augmentation. Studies that have examined the variation in lordosis and kyphosis through SSAMs [88–90] could be incorporated with vertebral geometry and density to find and provide such relationships between spinal level, spinal curvature and vertebral density distributions. These relationships may elucidate which vertebral levels would be better targets for (prophylactic) augmentation without the need for μ CT scans to identify density distributions.

Addition further steps would include a more detailed investigation into different combinations of the principal components and identification of the modes of variation captured in PC4 and PC5. Furthermore, should the percentage of variation captured in the first five principal components be smaller than needed to understand the total variation, the PCA tool could be developed further to investigate the additional modes.

6.4 Conclusion

The aims and objectives for this study were outlined in Section 1.8 and through the course of Chapters 2 to 5 these objectives have been met. The first two objectives described accurately modelling non-augmented vertebrae and subsequent modelling of augmented vertebrae with close attention to the bone-cement interface. The final objective was to use PCA to identify patient subsets and understand the effect these modes of variation have on the outcomes of vertebroplasty.

The main conclusions on the development of methodology are as follows:

- While bovine tail vertebrae are abundant and provide an excellent material for method development, their similarity to human lumbar vertebrae is limited.
- Variation in recording the stiffness of vertebrae through loading is quite large, despite ensuring repeats use near identical loading points. Additionally, attempts to limit damage while loading is difficult due to the range of vertebral strengths.
- Capturing the vertebral structure in terms of density is vital for accurate modelling of both non-augmented and augmented specimens. Using the bone volume fraction

method gave greatly improved results.

- Simple descriptions of the augmented region are not enough to accurately describe it. Utilisation of interface layer, reduced moduli and a description of the needle tracks are needed.
- The use of statistical shape and appearance models can provide an insight into the mechanical properties of vertebrae that are otherwise hidden.
- The desired post-augmentation stiffness needs to be identified in order to find the correct augmentation volume size, position and shape for each type of vertebrae. This may be identified through an examination of adjacent level failure.

The main conclusions regarding clinical application of the results are as follows:

- Experimental, computational and statistical models suggest that vertebroplasty should be reserved for the least dense vertebrae, in terms of bone volume fraction, if an increase in stiffness is desired.
- Anterior placements of cement generally provide the largest increase in stiffness, however have a smaller effect if the vertebral density is weighted towards the anterior portion.
- Volumes of cement greater than 35 % fill are required to achieve a consistent increase in stiffness across all but the most dense vertebrae.
- Cement should be injected so that the augmentation volume remains concentrated and not dispersed through the vertebrae if possible.

The above conclusions regarding clinical practice all assume an increase in stiffness is the desired outcome. They assume that the increase or restoration in stiffness stabilises the vertebrae and relieves pain for the patient. Little can be achieved to validate this assumption given that pain relief cannot be measured on cadaveric or computation based models. Therefore, a thorough investigation into the desired stiffness change following augmentation must be carried out.

In summary, the work presented in this thesis has developed new methods for representing the trabecular structure in continuum level models of bone, and has provided new approaches for modelling cement augmentation. Both of these developments have greatly increased the modelling accuracy when compared to other, similar studies. The outputs from these methods have been used to study of the effect of patient variance on vertebroplasty and uniquely, have been used to identify the bone density distribution variation within human

lumbar vertebrae. This statistical modelling approach allows the identification of trends and variable relationships that are imperceivable in the input set. The approach used here and the tools developed can be applied to wider patient groups and other treatment scenarios to improve patient stratification and to aid the identification of patients who would benefit most from particular treatments.

References

- [1] *Percutaneous vertebroplasty and percutaneous balloon kyphoplasty for treating osteoporotic vertebral compression fractures*. Standard. National Institute for Health and Clinical Excellence, Mar. 2013.
- [2] G. Baroud and M. Bohnert. Biomechanical impact of vertebroplasty. *Joint Bone Spine* 73.2 (2006), pp. 144–150.
- [3] R. Buchbinder, R. H. Osborne, P. R. Ebeling, J. D. Wark, P. Mitchell, C. Wriedt, S. Graves, M. P. Staples, and B. Murphy. A randomized trial of vertebroplasty for painful osteoporotic vertebral fractures. *New England Journal of Medicine* 361.6 (2009), pp. 557–568.
- [4] D. F. Kallmes, B. A. Comstock, P. J. Heagerty, J. A. Turner, D. J. Wilson, T. H. Diamond, R. Edwards, L. A. Gray, L. Stout, S. Owen, W. Hollingworth, B. Ghdoke, D. J. Annesley-Williams, S. H. Ralston, and J. G. Jarvik. A randomized trial of vertebroplasty for osteoporotic spinal fractures. *New England Journal of Medicine* 361.6 (2009), pp. 569–79.
- [5] H. Gray. *Anatomy of the human body*. 20th ed., 1918.
- [6] M. A. Adams and P. Dolan. Spine biomechanics. *Journal of Biomechanics* 38.10 (Oct. 2005), pp. 1972–1983.
- [7] T. Stewart and R. Hall. (iv) Basic biomechanics of human joints: Hips, knees and the spine. *Curr. Orthop.* 20.1 (Feb. 2006), pp. 23–31.
- [8] N. Yoganandan, F. a. Pintar, B. D. Stemper, J. L. Baisden, R. Aktay, B. S. Shender, G. Paskoff, and P. Laud. Trabecular bone density of male human cervical and lumbar vertebrae. *Bone* 39.2 (Aug. 2006), pp. 336–44.
- [9] S. Edmondston, K. Singer, R. Price, R. Day, and P. Bredahl. The relationship between bone mineral density, vertebral body shape and spinal curvature in the elderly thoracolumbar spine: an in vitro study. *The British journal of radiology* 67.802 (1994), pp. 969–975.

- [10] M. M. Panjabi, K. Takata, V. Goel, D. Federico, T. Oxland, J. Duranceau, and M. Krag. Thoracic Human Vertebrae Quantitative Three-Dimensional Anatomy. *Spine (Phila. Pa. 1976)*. 16.8 (1991), pp. 888–901.
- [11] P. M. Manohar, V. Goel, T. Oxland, K. Takata, J. Duranceau, M. Krag, and M. Price. *Human lumbar vertebrae: quantitative three-dimensional anatomy*. 1992.
- [12] F. Denis. The three column spine and its significance in the classification of acute thoracolumbar spinal injuries. *Spine (Phila Pa 1976)* 8.8 (1983), pp. 817–831.
- [13] F. Magerl, M. Aebi, S. D. Gertzbein, J. Harms, and S. Nazarian. A comprehensive classification of thoracic and lumbar injuries. *European Spine Journal* 3.4 (Aug. 1994), pp. 184–201.
- [14] J. M. Mathis, J. D. Barr, S. M. Belkoff, M. S. Barr, M. E. Jensen, and H. Deramond. Percutaneous vertebroplasty: a developing standard of care for vertebral compression fractures. en. *AJNR. Am. J. Neuroradiol.* 22. February (Feb. 2001), pp. 373–381.
- [15] C. Studies, F. The, and M. Clinic. Vertebral Compression Fractures in Elderly Osteoporosis Patients Receiving Glucocorticoid Intra-articular Injections. *American Family Physician* 6.3 (2006), pp. 206–211.
- [16] L. J. Melton III and D. F. Kallmes. Epidemiology of vertebral fractures: implications for vertebral augmentation. *Academic Radiology* 13.5 (May 2006), pp. 538–45.
- [17] C. Cooper, E. J. Atkinson, W. M. O’Fallon, and L. J. Melton. Incidence of clinically diagnosed vertebral fractures: a population-based study in Rochester, Minnesota, 1985-1989. *Journal of Bone and Mineral Research* 7.2 (Feb. 1992), pp. 221–227.
- [18] F. Burstein, S. Cohen, R. Hudgins, W. Boydston, and C. Simms. The Use of Hydroxyapatite Cement in Secondary Craniofacial Reconstruction. *Craniofacial Reconstr.* 104.5 (Jan. 1987), pp. 1270–1275.
- [19] C. Cooper, T. O’Neill, and A. Silman. The epidemiology of vertebral fractures. *Bone* 14 (Jan. 1993), pp. 89–97.
- [20] T. H. Diamond, C. Bryant, L. Browne, and W. a. Clark. Clinical outcomes after acute osteoporotic vertebral fractures: a 2-year non-randomised trial comparing percutaneous vertebroplasty with conservative therapy. *Medical Journal of Australia* 184.3 (2006), pp. 113–7.
- [21] M. Moro, a. T. Hecker, M. L. Bouxsein, and E. R. Myers. Failure load of thoracic vertebrae correlates with lumbar bone mineral density measured by DXA. *Calcified Tissue International* 56.3 (Mar. 1995), pp. 206–209.

- [22] A. A. Bookman. *Musculoskeletal Imaging*. Vol. 38. 4. Elsevier Health Sciences, 2011, pp. 787–787.
- [23] R. Buchbinder, R. H. Osborne, P. R. Ebeling, J. D. Wark, P. Mitchell, C. Wriedt, S. Graves, M. P. Staples, and B. Murphy. A randomized trial of vertebroplasty for painful osteoporotic vertebral fractures. *New England Journal of Medicine* 361.6 (2009), pp. 557–568.
- [24] S. M. Belkoff, J. M. Mathis, and L. E. Jasper. The Biomechanics of Vertebroplasty. *Spine (Phila. Pa. 1976)*. 26.14 (2001), pp. 1537–1541.
- [25] S. M. Belkoff, J. M. Mathis, and L. E. Jasper. Ex vivo biomechanical comparison of hydroxyapatite and polymethylmethacrylate cements for use with vertebroplasty. *AJNR. Am. J. Neuroradiol.* 23.10 (2002), pp. 1647–51.
- [26] P. A. Hulme, J. Krebs, S. J. Ferguson, and U. Berlemann. Vertebroplasty and kyphoplasty: a systematic review of 69 clinical studies. *Spine (Phila. Pa. 1976)*. 31.17 (Aug. 2006), pp. 1983–2001.
- [27] S. Muijs. Treatment of painful osteoporotic vertebral compression fractures: A Brief Review of the Evidence for Percutaneous Vertebroplasty. *The Journal of bone and joint surgery. British volume* (2011).
- [28] B. Jay and S. Ahn. Vertebroplasty. *Semin. Intervent. Radiol.* (2013).
- [29] P. J. Atkinson. Variation in trabecular structure of vertebrae with age. *Calcified Tissue Research* 1.1 (1967), pp. 24–32.
- [30] a. M. Parfitt. Age-related structural changes in trabecular and cortical bone: cellular mechanisms and biomechanical consequences. *Calcified Tissue International* 36 Suppl 1.1 (1984), S123–S128.
- [31] J. Aaron, P. Shore, R. Shore, M. Beneton, and J. Kanis. Trabecular architecture in woman and men of similar bone mass with and without vertebral fracture: II. Three-dimensional histology. *Bone* 27.2 (2000), pp. 277–282.
- [32] P. a. Hulme, S. K. Boyd, and S. J. Ferguson. Regional variation in vertebral bone morphology and its contribution to vertebral fracture strength. *Bone* 41.6 (2007), pp. 946–957.
- [33] L. Mosekilde, L. Mosekilde, and C. C. Danielsen. Biomechanical competence of vertebral trabecular bone in relation to ash density and age in normal individuals. *Bone* 8.2 (1987), pp. 79–85.
- [34] T. S. Keller. Predicting the Compressive Mechanical Behavior of Bone. *Journal of Biomechanics* 27.9 (1994), pp. 1159–1168.

- [35] T. Hildebrand and P. Rüegsegger. Quantification of Bone Microarchitecture with the Structure Model Index. *Computer Methods in Biomechanics and Biomedical Engineering* 1.1 (1997), pp. 15–23.
- [36] J. S. Thomsen, E. N. Ebbesen, and L. I. Mosekilde. Zone-dependent changes in human vertebral trabecular bone: clinical implications. *Bone* 30.5 (2002), pp. 664–9.
- [37] E. N. Ebbesen, J. S. Thomsen, H. Beck-Nielsen, H. J. Nepper-Rasmussen, and L. Mosekilde. Age-and Gender-Related Differences in Vertebral Bone Mass, Density, and Strength. *Journal of Bone and Mineral Research* 14.8 (1999), pp. 1394–1403.
- [38] H.-J. Wilke, A. Kettler, and L. E. Claes. Are sheep spines a valid biomechanical model for human spines? *Spine (Phila. Pa. 1976)*. 22.20 (1997), pp. 2365–2374.
- [39] F. Kandziora, R. Pflugmacher, M. Scholz, K. Schnake, M. Lucke, R. Schröder, and T. Mittlmeier. Comparison Between Sheep and Human Cervical Spines An Anatomic, Radiographic, Bone Mineral Density, and Biomechanical Study. *Spine (Phila. Pa. 1976)*. 26.9 (2001), pp. 1028–1037.
- [40] H.-J. Wilke, A. Kettler, and K. H. Wenger. Anatomy of the Sheep Spine and Its Comparison to the Human Spine. *Anatomical Record* 247.4 (1997), pp. 542–555.
- [41] Aerssens, Jeroen, Boonen, Steven, Geert, and Dequeker. Interspecies Differences in Bone Composition, Density, and Quality: Potential Implications for in Vivo Bone Research. *Endocrinology* 139.2 (1998), pp. 663–670.
- [42] V. N. Wijayathunga, A. C. Jones, R. J. Oakland, N. R. Furtado, R. M. Hall, and R. K. Wilcox. Development of specimen-specific finite element models of human vertebrae for the analysis of vertebroplasty. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine* 222.2 (Feb. 2008), pp. 221–228.
- [43] S. M. Tarsuslugil, R. M. O’Hara, N. J. Dunne, F. J. Buchanan, J. F. Orr, D. C. Barton, and R. K. Wilcox. Development of calcium phosphate cement for the augmentation of traumatically fractured porcine specimens using vertebroplasty. *Journal of Biomechanics* 46.4 (2013), pp. 711–5.
- [44] N. Furtado, R. J. Oakland, R. K. Wilcox, and R. M. Hall. A biomechanical investigation of vertebroplasty in osteoporotic compression fractures and in prophylactic vertebral reinforcement. *Spine (Phila. Pa. 1976)*. 32.17 (2007), E480–E487.
- [45] R. K. Wilcox, D. J. Allen, R. M. Hall, D. Limb, D. C. Barton, and R. a. Dickson. A dynamic investigation of the burst fracture process using a combined experimental and finite element approach. *European Spine Journal* 13.6 (2004), pp. 481–8.

- [46] G. Gurwitz, J. Dawson, and M. McNamara. Biomechanical analysis of three surgical approaches for lumbar burst fractures using short-segment instrumentation. *Spine (Phila. Pa. 1976)*. (1993).
- [47] D. Ananthakrishnan, S. Berven, V. Deviren, K. Cheng, J. C. Lotz, Z. Xu, and C. M. Puttlitz. The effect on anterior column loading due to different vertebral augmentation techniques. *Clin. Biomech. (Bristol, Avon)* 20.1 (2005), pp. 25–31.
- [48] U. Berlemann, S. J. Ferguson, L. P. Nolte, and P. F. Heini. Adjacent vertebral failure after vertebroplasty. A biomechanical investigation. *Journal of Bone and Joint Surgery. British Volume* 84.5 (2002), pp. 748–752.
- [49] S. G. Pneumaticos, G. K. Triantafyllopoulos, D. S. Evangelopoulos, J. a. Hipp, and M. H. Heggeness. Effect of vertebroplasty on the compressive strength of vertebral bodies. *Spine Journal* 13.12 (2013), pp. 1921–1927.
- [50] J. D. Barr, M. S. Barr, T. J. Lemley, and R. M. McCann. Percutaneous Vertebroplasty for Pain Relief and Spinal Stabilization. *Spine (Phila. Pa. 1976)*. 25.8 (Apr. 2000), pp. 923–928.
- [51] A. G. Tohmeh, J. M. Mathis, D. C. Fenton, A. M. Levine, and S. M. Belkoff. Biomechanical efficacy of uni- vs. bi-pedicular vertebroplasty for the treatment of osteoporotic compression fractures. *Trans. Orthop. Res. Soc.* 24.17 (1999), p. 177.
- [52] K. B. Higgins, D. R. Sindall, A. M. Cuitino, and N. a. Langrana. Biomechanical alterations in intact osteoporotic spine due to synthetic augmentation: finite element investigation. *Journal of Biomechanical Engineering* 129.4 (2007), pp. 575–85.
- [53] J. Graham, M. Ries, and L. Pruitt. Effect of bone porosity on the mechanical integrity of the bone-cement interface. *Journal of Bone and Joint Surgery. American Volume* 85-A.10 (2003), pp. 1901–1908.
- [54] S. Belkoff, H. Deramond, J. Mathis, and L. Jasper. Vertebroplasty: the biomechanical effect of cement volume. *Trans Orthop Res Soc* 46 (2000), p. 356.
- [55] S. M. Belkoff, J. M. Mathis, and E. M. Erbe. Biomechanical Evaluation of a New Bone Cement for Use in Vertebroplasty. *Spine (Phila. Pa. 1976)*. 25.9 (2000), pp. 1061–1064.
- [56] S. J. Lee, B. J. Jun, G. R. Tack, S. Y. Lee, and K. C. Shin. “Prediction and assessment of optimal volume for PMMA injection in percutaneous vertebroplasty using image and biomechanical analyses”. *Trans. Annu. Meet. Res. Soc.* 2002, p. 786.

- [57] A. C. Jones and R. K. Wilcox. Finite element analysis of the spine: Towards a framework of verification, validation and sensitivity analysis. *Medical Engineering & Physics* 30.10 (Dec. 2008), pp. 1287–1304.
- [58] M. Viceconti, S. Olsen, L. P. Nolte, and K. Burton. Extracting clinically relevant data from finite element simulations. *Clinical Biomechanics* 20.5 (June 2005), pp. 451–454.
- [59] M. Mengoni, S. Sikora, V. d’Ottreppe, R. K. Wilcox, and A. C. Jones. In-Silico Models of Trabecular Bone: A Sensitivity Analysis Perspective. en (2014), pp. 393–423.
- [60] K. Robson Brown, S. Tarsuslugil, V. N. Wijayathunga, and R. K. Wilcox. Comparative finite-element analysis: a single computational modelling method can estimate the mechanical properties of porcine and human vertebrae. *Journal of the Royal Society, Interface* 11.95 (June 2014), p. 20140186.
- [61] J. M. Buckley. Sensitivity of Vertebral Compressive Strength to Endplate Loading Distribution. *Journal of Biomechanical Engineering* 128.5 (2006), p. 641.
- [62] Y. Chevalier, D. Pahr, and P. K. Zysset. The Role of Cortical Shell and Trabecular Fabric in Finite Element Analysis of the Human Vertebral Body. *Journal of Biomechanical Engineering* 131.11 (Nov. 2009), p. 111003.
- [63] S. K. Eswaran, A. Gupta, and T. M. Keaveny. Locations of bone tissue at high risk of initial failure during compressive loading of the human vertebral body. *Bone* 41.4 (Oct. 2007), pp. 733–739.
- [64] M. Kinzl, L. M. Benneker, A. Boger, P. K. Zysset, and D. H. Pahr. The effect of standard and low-modulus cement augmentation on the stiffness, strength, and endplate pressure distribution in vertebroplasty. *European Spine Journal* 21.5 (May 2012), pp. 920–929.
- [65] G. M. Treece, R. W. Prager, and A. H. Gee. Regularised marching tetrahedra : improved iso-surface extraction. *Comput. Graph.* 23 (1999), pp. 583–598.
- [66] M. Kinzl, J. Schwiedrzik, P. K. Zysset, and D. H. Pahr. An experimentally validated finite element method for augmented vertebral bodies. *Clinical biomechanics* 28.1 (2013), pp. 15–22.
- [67] S. K. Eswaran, A. Gupta, M. F. Adams, and T. M. Keaveny. Cortical and Trabecular Load Sharing in the Human Vertebral Body. *Journal of Bone and Mineral Research* 21.2 (Nov. 2005), pp. 307–314.

- [68] H. Ritzel, M. Amling, M. Pösl, M. Hahn, and G. Delling. The thickness of human vertebral cortical bone and its changes in aging and osteoporosis: a histomorphometric analysis of the complete spinal column from thirty-seven autopsy specimens. *Journal of Bone and Mineral Research* 12.1 (1997), pp. 89–95.
- [69] A. C. Jones and R. K. Wilcox. Assessment of factors influencing finite element vertebral model predictions. *Journal of Biomechanical Engineering* 129.6 (2007), pp. 898–903.
- [70] Y. Zhao. “Finite element modelling of cement augmentation and fixation for orthopaedic applications”. PhD thesis. University of Leeds, 2010.
- [71] D. H. Pahr and P. K. Zysset. A comparison of enhanced continuum FE with micro FE models of human vertebral bodies. *Journal of Biomechanics* 42.4 (Mar. 2009), pp. 455–462.
- [72] M. a. Liebschner, W. S. Rosenberg, and T. M. Keaveny. Effects of bone cement volume and distribution on vertebral stiffness after vertebroplasty. *Spine (Phila. Pa. 1976)*. 26.14 (2001), pp. 1547–1554.
- [73] G. Baroud, J. Nemes, P. Heini, and T. Steffen. Load shift of the intervertebral disc after a vertebroplasty: A finite-element study. *European Spine Journal* 12.4 (Aug. 2003), pp. 421–426.
- [74] A. Polikeit, L. P. Nolte, and S. J. Ferguson. The effect of cement augmentation on the load transfer in an osteoporotic functional spinal unit: finite-element analysis. *Spine (Phila. Pa. 1976)*. 28.10 (2003), pp. 991–996.
- [75] Y. Chevalier, D. Pahr, M. Charlebois, P. Heini, E. Schneider, and P. Zysset. Cement Distribution, Volume, and Compliance in Vertebroplasty. *Spine (Phila. Pa. 1976)*. 33.16 (2008), pp. 1722–1730.
- [76] Y. Zhao, K. a. Robson Brown, Z. M. Jin, and R. K. Wilcox. Trabecular level analysis of bone cement augmentation: a comparative experimental and finite element study. *Annals of Biomedical Engineering* 40.10 (Oct. 2012), pp. 2168–76.
- [77] G. Tozzi, Q.-H. Zhang, and J. Tong. 3D real-time micromechanical compressive behaviour of bone - cement interface: Experimental and finite element studies. *Journal of Biomechanics* 45.2 (Jan. 2012), pp. 356–363.
- [78] D. Janssen, K. a. Mann, and N. Verdonschot. Micro-mechanical modeling of the cement-bone interface: The effect of friction, morphology and material properties on the micromechanical response. *Journal of Biomechanics* 41.15 (Nov. 2008), pp. 3158–3163.

- [79] M. Kinzl, a. Boger, P. K. Zysset, and D. H. Pahr. The mechanical behavior of PMMA/bone specimens extracted from augmented vertebrae: A numerical study of interface properties, PMMA shrinkage and trabecular bone damage. *Journal of Biomechanics* 45.8 (May 2012), pp. 1478–1484.
- [80] S. Sikora. “Experimental and computational study of the behaviour of trabecular bone-cement interfaces.” PhD thesis. University of Leeds, 2013.
- [81] J. F. Hollenbeck, C. M. Cain, J. A. Fattor, P. J. Rullkoetter, and P. J. Laz. Statistical shape modeling characterizes three-dimensional shape and alignment variability in the lumbar spine. *Journal of biomechanics* 69 (2018), pp. 146–155.
- [82] L. Grassi, E. Schileo, C. Boichon, M. Viceconti, and F. Taddei. Comprehensive evaluation of PCA-based finite element modelling of the human femur. *Medical Engineering & Physics* 36.10 (Oct. 2014), pp. 1246–52.
- [83] M. Clogenson, J. M. Duff, M. Luethi, M. Levivier, R. Meuli, C. Baur, and S. Henein. A statistical shape model of the human second cervical vertebra. *International journal of computer assisted radiology and surgery* 10.7 (2015), pp. 1097–1107.
- [84] K. Chase, D. R. Carrier, F. R. Adler, T. Jarvik, E. A. Ostrander, T. D. Lorentzen, and K. G. Lark. Genetic basis for systems of skeletal quantitative traits: principal component analysis of the canid skeleton. *Proceedings of the National Academy of Sciences* 99.15 (2002), pp. 9930–9935.
- [85] S. P. Väänänen, L. Grassi, G. Flivik, J. S. Jurvelin, and H. Isaksson. Generation of 3D shape, density, cortical thickness and finite element mesh of proximal femur from a DXA image. *Medical Image Analysis* 24.1 (Aug. 2015), pp. 125–134.
- [86] C. Fitzpatrick, M. Baldwin, and P. Laz. Development of a statistical shape model of the patellofemoral joint for investigating relationships between shape and function. *Journal of biomechanics* (2011).
- [87] T. L. Bredbenner, T. D. Eliason, W. L. Francis, J. M. McFarland, A. C. Merkle, and D. P. Nicolella. Development and validation of a statistical shape modeling-based finite element model of the cervical spine under low-level multiple direction loading conditions. *Frontiers in bioengineering and biotechnology* 2 (2014), p. 58.
- [88] A. H. Ali, A.-b. Cowan, J. S. Gregory, R. M. Aspden, and J. R. Meakin. The accuracy of active shape modelling and end-plate measurements for characterising the shape of the lumbar spine in the sagittal plane. *Computer methods in biomechanics and biomedical engineering* 15.2 (2012), pp. 167–172.

- [89] J. R. Meakin, J. S. Gregory, R. M. Aspden, F. W. Smith, and F. J. Gilbert. The intrinsic shape of the human lumbar spine in the supine, standing and sitting postures: characterization using an active shape model. *Journal of anatomy* 215.2 (2009), pp. 206–211.
- [90] J. Boisvert, F. Cheriet, X. Pennec, H. Labelle, and N. Ayache. Geometric variability of the scoliotic spine using statistics on articulated shape models. *IEEE Transactions on Medical Imaging* 27.4 (2008), pp. 557–568.
- [91] F. Y. Zapata-Cornelio, G. A. Day, R. H. Coe, S. N. F. Sikora, V. N. Wijayathunga, S. M. Tarsuslugil, M. Mengoni, and R. K. Wilcox. Methodology to Produce Specimen-Specific Models of Vertebrae: Application to Different Species. *Annals of Biomedical Engineering* 45.10 (2017), pp. 2451–2460.
- [92] S. M. Belkoff, J. C. Sanders, and L. E. Jasper. The effect of the monomer-to-powder ratio on the material properties of acrylic bone cement. *Journal of Biomedical Materials Research* 63.4 (2002), pp. 396–399.
- [93] L. E. Jasper, H. Deramond, J. M. Mathis, and S. M. Belkoff. The effect of monomer-to-powder ratio on the material properties of cranioplastic. *Bone* 25.SUPPL. 1 (1999), pp. 27–29.
- [94] J. Schindelin, I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzsch, S. Preibisch, C. Rueden, S. Saalfeld, B. Schmid, et al. Fiji: an open-source platform for biological-image analysis. *Nature methods* 9.7 (2012), p. 676.
- [95] M. Doube, M. M. Klosowski, I. Arganda-Carreras, F. P. Cordelières, R. P. Dougherty, J. S. Jackson, B. Schmid, J. R. Hutchinson, and S. J. Shefelbine. BoneJ: Free and extensible bone image analysis in ImageJ. *Bone* 47.6 (2010), pp. 1076–1079.
- [96] R. J. Oakland, N. R. Furtado, R. K. Wilcox, J. Timothy, and R. M. Hall. Preliminary biomechanical evaluation of prophylactic vertebral reinforcement adjacent to vertebroplasty under cyclic loading. *The Spine Journal* 9.2 (2009), pp. 174–181.
- [97] B. McCormack, P. Prendergast, and B. O’Dwyer. Fatigue of cemented hip replacements under torsional loads. *Fatigue & Fracture of Engineering Materials & Structures* 22.1 (1999), pp. 33–40.
- [98] M. Mengoni. *opti4Abq (v 2.0), a generic python code to run Abaqus in an optimisation loop*. 2017.
- [99] N. Emadi, Y. Safi, A. A. Bagheban, and S. Asgary. Comparison of CT-number and gray scale value of different dental materials and hard tissues in CT and CBCT. *Iranian endodontic journal* 9.4 (2014), p. 283.

- [100] S. H. Zhou, I. D. McCarthy, A. H. McGregor, R. R. H. Coombs, and S. P. F. Hughes. Geometrical dimensions of the lower lumbar vertebrae - analysis of data from digitised CT images. *European Spine Journal* 9.3 (2000), pp. 242–248.
- [101] A. Race, K. A. Mann, and A. A. Edidin. Mechanics of bone/PMMA composite structures: An in vitro study of human vertebrae. *Journal of Biomechanics* 40.5 (2007), pp. 1002–1010.
- [102] R. Aquarius, J. Homminga, A. J. F. Hosman, N. Verdonshot, and E. Tanck. Prophylactic vertebroplasty can decrease the fracture risk of adjacent vertebrae: an in vitro cadaveric study. *Medical engineering & physics* 36.7 (2014), pp. 944–948.
- [103] J. Steens, N. Verdonshot, A. M. Aalsma, and A. J. Hosman. The influence of endplate-to-endplate cement augmentation on vertebral strength and stiffness in vertebroplasty. *Spine* 32.15 (2007), E419–E422.
- [104] L. I.-K. Lin. A concordance correlation coefficient to evaluate reproducibility. *Biometrics* 45 (1989), pp. 255–268.
- [105] W. H. Organization et al. Assessment of fracture risk and its application to screening for postmenopausal osteoporosis: report of a WHO study group [meeting held in Rome from 22 to 25 June 1992] (1994).
- [106] L. Mosekilde. Age-related changes in vertebral trabecular bone architecture—assessed by a new method. *Bone* 9.4 (1988), pp. 247–250.
- [107] K. B. Higgins, R. D. Harten, N. A. Langrana, and M. F. Reiter. Biomechanical effects of unipedicular vertebroplasty on intact vertebrae. *Spine* 28.14 (2003), pp. 1540–1547.
- [108] S. K. Eswaran, A. Gupta, M. F. Adams, and T. M. Keaveny. Cortical and trabecular load sharing in the human vertebral body. *Journal of Bone and Mineral Research* 21.2 (2006), pp. 307–314.
- [109] A. Fedorov, R. Beichel, J. Kalpathy-Cramer, J. Finet, J.-C. Fillion-Robin, S. Pujol, C. Bauer, D. Jennings, F. Fennessy, M. Sonka, et al. 3D Slicer as an image computing platform for the Quantitative Imaging Network. *Magnetic resonance imaging* 30.9 (2012), pp. 1323–1341.
- [110] M. Sezgin and B. Sankur. Survey over image thresholding techniques and quantitative performance evaluation. *Journal of Electronic imaging* 13.1 (2004), pp. 146–166.
- [111] S. Gravius, N. Kraska, U. Maus, T. Mumme, P. Berdel, and M. Weisskopf. Intravertebral pressure during vertebroplasty—an in-vitro study. *Zeitschrift für Orthopädie und Unfallchirurgie* 147.1 (2009), pp. 43–47.

- [112] S. Meller and W. A. Kalender. “Building a statistical shape model of the pelvis”. *International Congress Series*. Vol. 1268. Elsevier. 2004, pp. 561–566.
- [113] D. P. Nicolella and T. L. Bredbenner. Development of a parametric finite element model of the proximal femur using statistical shape and density modelling. *Computer methods in biomechanics and biomedical engineering* 15.2 (2012), pp. 101–110.
- [114] R. Bryan, P. S. Mohan, A. Hopkins, F. Galloway, M. Taylor, and P. B. Nair. Statistical modelling of the whole human femur incorporating geometric and material properties. *Medical engineering & physics* 32.1 (2010), pp. 57–65.
- [115] E. F. Chan, C. L. Farnsworth, S. M. Klisch, H. S. Hosalkar, and R. L. Sah. 3-dimensional metrics of proximal femoral shape deformities in Legg–Calvé–Perthes disease and slipped capital femoral epiphysis. *Journal of Orthopaedic Research* 36.5 (2018), pp. 1526–1535.
- [116] I. Gilad and M. Nissan. A study of vertebra and disc geometric relations of the human cervical and lumbar spine. 11.2 (1986), pp. 154–157.
- [117] I. Gilad and M. Nissan. Sagittal evaluation of elemental geometrical dimensions of human vertebrae. *Journal of Anatomy* 143 (1985), pp. 115–20.
- [118] K. M. Cheung, D. Ruan, F. L. Chan, and D. Fang. Computed tomographic osteometry of asian lumbar pedicles. *Spine (Phila. Pa. 1976)*. 19.13 (1994), pp. 1495–1498.
- [119] H. Ritzel, M. Amling, M. Pösl, M. Hahn, and G. Delling. The thickness of human vertebral cortical bone and its changes in aging and osteoporosis: A histomorphometric analysis of the complete spinal column from thirty-seven autopsy specimens. *Journal of Bone and Mineral Research* 12.1 (1997), pp. 89–95.
- [120] R. Coe. “Simulation of the Fatigue Performance of Spinal Vertebrae”. PhD thesis. University of Leeds, 2019.
- [121] P. Steiger, J. Block, S. Steiger, A. Heuck, A. Friedlander, B. Ettinger, S. Harris, C. Glüer, and H. Genant. Spinal bone mineral density measured with quantitative CT: effect of region of interest, vertebral level, and technique. *Radiology* 175.2 (1990), pp. 537–543.