

Performance of microstructural Finite Element models in predicting the effective modulus of trabecular bone



Claire Lawrence

Department of Mechanical Engineering
University of Cape Town
South Africa

Supervisors:

Mr T. J. Cloete and Prof. G. N. Nurick

December 2019

PhD thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy
in the Department of Mechanical Engineering at the University of Cape Town

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

Declaration

I, Claire Lawrence, hereby:

1. grant the University of Cape Town free licence to reproduce the above thesis in whole or in part, for the purpose of research;
2. declare that:
 - (a) this thesis is my own unaided work, both in concept and execution, and apart from the normal guidance from my supervisor, I have received no assistance except as stated below.
 - (b) neither the substance nor any part of the above thesis has been submitted in the past, or is being, or is to be submitted for a degree at this University or at any other university, except as stated below.

Signed by candidate

Claire Lawrence
Department of Mechanical Engineering
University of Cape Town
Friday 4th September, 2020

Performance of microstructural Finite Element models in predicting the effective modulus of trabecular bone

Claire Lawrence - 18 December 2019

Trabecular bone is made up of an irregular, interconnecting framework of rod- and plate-like struts [1], therefore the mechanical properties of the bone may only be determined through experimental testing or detailed Finite Element modelling. Experimental testing requires a sample to be removed from the body, which is not possible in living patients. As such, there is a drive to move away from experimental testing and focus instead on creating accurate patient-specific Finite Element models from CT scans of the bone.

The computational “gold-standard” Finite Element model used for trabecular bone, namely the voxel-based method, uses solid tetrahedral elements, which are extremely resource intensive. Vanderroost, et al [2] developed an alternative Finite Element code which discretises the structure into a series of beams and shells. This beam-shell approach vastly reduces the size of the mesh and, consequently, the processing time required for the simulation.

In this work, an analysis cycle was developed to determine the apparent modulus of a structure using the beam-shell Finite Element model [2]. The cycle imports micro-CT scans of a structure, discretises the structure into a beam-shell mesh, performs a Finite Element simulation and outputs the apparent modulus of the structure along with a reconstructed image. The analysis cycle was validated by analysing over 3000 artificially generated images, comprising various configurations of cubic lattices, Kelvin cell lattices and octet truss lattices, and comparing the modulus output by the analysis cycle to baseline results obtained through the simulation of known node and element data. The analysis cycle provided predictions within 10% of the baseline value for most lattices, however there were issues associated with the rasterisation of the input images and postprocessing which caused variation in the results. Overall, it was determined that the analysis cycle is capable of capturing the apparent modulus of a variety of different structures.

Micro-CT scans of 127 bone specimens were run through the analysis cycle. The results from the beam-shell analysis were compared to results from experimental testing [3] and an equivalent voxel-based analysis. There was a clear trend in both the beam-shell and voxel-based data, however the voxel-based method produced stiffer results than the beam-shell method overall. The beam-shell method showed more scatter than the voxel-based method, but contained less significant outliers. The effective modulus, i.e. the modulus of an inner core region, was determined for 17 of the bone specimens and compared to equivalent experimental results. The beam-shell method captured the increase in stiffness between the apparent modulus and the effective modulus as regularly as the voxel-based method, given appropriate boundary conditions were applied. The results produced by both methods can be improved by the removal of machining artifacts and improved segmentation of the micro-CT scans.

This work confirms that the beam-shell method is capable of capturing the apparent modulus of a trabecular bone sample, however the scatter in the data must be reduced for it to be considered a viable alternative to the voxel-based method. It was found that the beam-shell method is equally capable of predicting the relationship between apparent modulus and effective modulus as the voxel-based method. In both the beam-shell results and voxel-based results, the accuracy of a particular data point could only be determined by considering the results in reference to additional simulation and experimental data points. In light of these results, researchers should be cautious in reporting simulation results for trabecular bone without additional verification.

Acknowledgements

The financial assistance of the National Research Foundation (NRF) towards this research is hereby acknowledged. Opinions expressed and conclusions arrived at, are those of the author and are not necessarily to be attributed to the NRF.

I would like to thank my supervisor, Trevor Cloete, for many things, but foremost for his enthusiasm. Trevor, thank you for being passionate about my project even when I was not. Your passion is what sets you apart from others in your field, and has been an invaluable part of this project. Thank you for your constant questioning, encouragement and advice - without you this project would not be possible.

To my co-supervisor, Professor Nurick, thank you for many hours of meetings and proofreading, but even more for sage advice dispensed with kindness and wisdom. Prof., the care you take with your students is incomparable, and your steady, calming influence not only convinced me to undertake this thesis, but also provided an anchor point throughout its duration. Thank you.

I would like to thank the BISRU students who learned and struggled along with me in my time as a postgrad. Particularly, thanks must go to Vinay Shekhar for helping with the presentation of the graphs in this thesis, Ruix Qi, for his constant commiseration, Muhsin Osman, for keeping me company on those early Wednesday mornings, Sherlin Gabriel, for being the sounding board for my complaints and Andrew Curry, for making me laugh with memes and Youtube videos whenever I was having a particularly rough day. You all made the long process of writing a thesis a lot more entertaining than it should have been!

Thanks must go to some other residents of Maintenance Place, specifically Shane Ferguson, Bernadine Minnaar and Priyesh Gosai. Shane and Bernadine, thanks for keeping me company in those early mornings when all the sane people were still asleep; you always managed to put a spring in my step, which no one should have at 6AM. Priyesh, thanks for making me laugh and... laugh and laugh. Your humour made the dull days seem a little brighter.

I would like to thank members of the BISRU staff for their constant support. To Dr Richard Curry, thank you with your help with LS-DYNA. To Prof. Genevieve Langdon and Dr Reuben Govender, thank you for taking the time to answer the millions of questions I had about life after a PhD. To Prof. Steeve Chung Kim Yuen, thank you for always putting a smile on my face and for your constant interest in the lives of the students.

This project would not have been possible without access to the Finite Element beam-shell code. My sincere thanks go to Prof. Harry van Lenthe and Dr Jef Vanderroost for sharing their research and their code with me and answering many, many questions via email from the other side of the world.

Thank you to Keelan Moore, who provided the lattice constructing code used in this project, which formed an integral part of my work. On a more personal note, thanks for three years of friendship which included fascinating conversations, shared complaining and a good deal of back-and-forth “you can do it!”.

To Kelsey Hilton, thanks for providing the bone specimens and scans that formed the basis of this project. Thanks, also, for consistently being willing to help, even when that help had to be dispensed during a weekend.

I would like to thank Anton and the team at the Stellenbosch CT scanning facility for all their help with the scanning of the bone specimens. This research formed an interesting part of my project and the swiftness and professionalism with which it was handled was appreciated.

To my best friend, Caitlyn, who has fought her way through 7 years of Engineering with me: Cait, thank you for countless hours of shared thinking, struggling, moaning, laughing, crying, comfort eating and banging our heads against the wall (both literally and figuratively).

To my mom, Desiree Lawrence, my dad, Frank Lawrence, and my brother, Roscoe Lawrence: thank you for being the best family a girl could ask for. To say I could not have done it without you, would be an understatement.

To my fiancé, Gevashkar Rampersadh: Givs, thank you for many, many hours spent listening to me complain about my thesis and for encouraging me to keep going when it is the last thing I wanted to do. You made the good days better and the very worst days seem not so bad.

Contents

Declaration	i
Abstract	ii
Acknowledgements	iii
List of Figures	xiii
List of Tables	xxxii
1 Introduction	1
1.1 Motivation for research	1
1.2 Problem statement	2
1.3 Aim and objectives	3
1.4 Hypothesis	3
2 Literature review	4
2.1 Trabecular bone	4
2.2 Image acquisition and segmentation	7
2.2.1 Micro-CT scanning	7
2.2.2 Segmentation	8
2.3 Microstructural indices	9
2.3.1 Calculation of indices	10
2.3.2 Correlation to mechanical properties	12
2.4 Cellular solids	14
2.4.1 Types of cellular solids	15

2.4.2	Determining mechanical properties	17
2.5	Finite Element Analysis of trabecular bone	18
2.5.1	Modelling methodologies	19
2.5.2	Approximating Young's modulus	23
2.5.3	Determining effective properties	25
2.6	Past research at BISRU	27
2.6.1	Cellular solids	27
2.6.2	Trabecular bone	29
2.7	Synthesis	31
3	Development of analysis cycle	33
3.1	Image processing	34
3.1.1	Specimen cropping	35
3.1.2	Segmentation	36
3.1.3	Conversion to 3D matrix	36
3.2	Discretisation	37
3.2.1	Mesh generation	37
3.2.2	Mesh refinement	37
3.2.3	Creation of ties	38
3.3	Simulation	39
3.3.1	Input decks	39
3.3.2	Finite Element Analysis	40
3.4	Post-processing	40
3.4.1	Modulus calculation	40
3.4.2	Rasterisation	41
4	Validation of analysis cycle	42
4.1	Test configurations	42
4.1.1	Lattice type	44
4.1.2	Lattice topology	45
4.1.3	Constraint and aspect ratio	47

<i>CONTENTS</i>	vii
4.1.4 Image sizes	48
4.1.5 Additional analysis	49
4.2 Validation process	51
4.3 Analysis	52
5 Simulation of bone	53
5.1 General simulations	53
5.2 Factors affecting simulation results	53
5.2.1 Voxel size	54
5.2.2 Segmentation technique	54
5.2.3 Boundary conditions	55
5.3 Microstructural indices	56
5.4 Analysis of results	56
6 Validation results and discussion	58
6.1 Sources of error	58
6.2 Kelvin cell lattice	59
6.2.1 K8-t3	60
6.2.2 Variation due to lattice size	68
6.2.3 Variation due to strut diameter	74
6.2.4 Variation due to aspect ratio	78
6.2.5 Variation due to constraint	78
6.2.6 Correlation to VBM	81
6.2.7 Summary and significance of results	85
6.3 Notable features in the cubic and octet truss lattices	86
6.3.1 Cubic lattice	87
6.3.2 Octet truss lattice	99
6.3.3 Summary and significance of results	107
7 Bone Results	109
7.1 Standard specimens results	109

7.2	Large specimen results	111
7.3	Resource usage	116
7.3.1	Time	116
7.3.2	Memory	120
7.4	Microstructural indices	120
7.5	Voxel size analysis	127
7.6	Segmentation analysis	128
7.7	Boundary condition analysis	130
7.7.1	Standard specimens	130
7.7.2	Large specimens	131
8	Discussion	136
8.1	BSM vs experimental	136
8.1.1	Standard specimens	136
8.1.2	Large specimens	142
8.2	BSM vs VBM	143
8.2.1	Standard specimens	143
8.2.2	Large specimens	152
8.2.3	Resource allocation	154
8.3	Microstructural indices	157
8.4	Factors affecting simulation results	159
8.4.1	Voxel size	159
8.4.2	Segmentation technique	164
8.4.3	Boundary conditions	166
8.5	Difference in results	170
9	Conclusions	172
9.1	General conclusion	172
9.2	Validation	173
9.3	Performance of BSM	173
9.3.1	Success rate	173

9.3.2	Accuracy of results	174
9.3.3	Resource usage	176
9.4	Microstructural indices	177
9.5	Factors affecting performance	178
9.5.1	Voxel size	178
9.5.2	Segmentation technique	179
9.5.3	Boundary conditions	179
9.6	Closing remarks	180
10	Recommendations	182
	References	184
A	Hardware and software	193
B	Removal of free clusters	194
C	Rasterisation	196
C.1	Beam rasterisation	196
C.2	Shell rasterisation	198
C.3	Optimisation	200
D	Validation results	201
D.1	Cubic lattices	201
D.1.1	C6-t1 - unconstrained	202
D.1.2	C6-t1 - constrained	203
D.1.3	C6-t2 - unconstrained	204
D.1.4	C6-t2 - constrained	205
D.1.5	C6-t3 - unconstrained	206
D.1.6	C6-t3 - constrained	207
D.1.7	C8-t1 - unconstrained	208
D.1.8	C8-t1 - constrained	209
D.1.9	C8-t2 - unconstrained	210

D.1.10	C8-t2 - constrained	211
D.1.11	C8-t3 - unconstrained	212
D.1.12	C8-t3 - constrained	213
D.1.13	C10-t1 - unconstrained	214
D.1.14	C10-t1 - constrained	215
D.1.15	C10-t2 - unconstrained	216
D.1.16	C10-t2 - constrained	217
D.1.17	C10-t3 - unconstrained	218
D.1.18	C10-t3 - constrained	219
D.2	Kelvin cell lattices	220
D.2.1	K6-t1 - unconstrained	221
D.2.2	K6-t1 - constrained	222
D.2.3	K6-t2 - unconstrained	223
D.2.4	K6-t2 - constrained	224
D.2.5	K6-t3 - unconstrained	225
D.2.6	K6-t3 - constrained	226
D.2.7	K8-t1 - unconstrained	227
D.2.8	K8-t1 - constrained	228
D.2.9	K8-t2 - unconstrained	229
D.2.10	K8-t2 - constrained	230
D.2.11	K8-t3 - unconstrained	231
D.2.12	K8-t3 - constrained	232
D.2.13	K10-t1 - unconstrained	233
D.2.14	K10-t1 - constrained	234
D.2.15	K10-t2 - unconstrained	235
D.2.16	K10-t2 - constrained	236
D.2.17	K10-t3 - unconstrained	237
D.2.18	K10-t3 - constrained	238
D.3	Octet truss lattices	239
D.3.1	OT6-t1 - unconstrained	240

D.3.2	OT6-t1 - constrained	241
D.3.3	OT6-t2 - unconstrained	242
D.3.4	OT6-t2 - constrained	243
D.3.5	OT6-t3 - unconstrained	244
D.3.6	OT6-t3 - constrained	245
D.3.7	OT8-t1 - unconstrained	246
D.3.8	OT8-t1 - constrained	247
D.3.9	OT8-t2 - unconstrained	248
D.3.10	OT8-t2 - constrained	249
D.3.11	OT8-t3 - unconstrained	250
D.3.12	OT8-t3 - constrained	251
D.3.13	OT10-t1 - unconstrained	252
D.3.14	OT10-t1 - constrained	253
D.3.15	OT10-t2 - unconstrained	254
D.3.16	OT10-t2 - constrained	255
D.3.17	OT10-t3 - unconstrained	256
D.3.18	OT10-t3 - constrained	257
D.4	Additional analysis: angled cubic	258
D.4.1	C8-t1 - angled unconstrained	259
D.4.2	C8-t1 - angled constrained	260
D.4.3	C8-t2 - angled unconstrained	261
D.4.4	C8-t2 - angled constrained	262
D.5	Additional analysis: VBM	263
D.5.1	K8-t3 - cropped and cropped constrained	264
E	Additional validation results	265
F	Rasterised lattices	268
F.1	Cubic	268
F.1.1	Size 6	269
F.1.2	Size 8	272

F.1.3	Size 10	275
F.2	Kelvin cell	278
F.2.1	Size 6	279
F.2.2	Size 8	282
F.2.3	Size 10	285
F.3	Octet truss	288
F.3.1	Size 6	289
F.3.2	Size 8	292
F.3.3	Size 10	295
F.4	Angled cubic	298
F.4.1	Size 8	299
G	Inherent VBM boundary condition	301
H	Bone specimens	303
I	Segmented Bone Scans	305
J	Choice of line fit	317

List of Figures

2.1	Structure of bone.	5
2.2	Trabecular bone samples showing the characteristic rod-plate microstructure.	5
2.3	Close-up view of bovine femur.	6
2.4	Maximal spheres method used for the calculation of local thickness.	11
2.5	Classification of voxels into end-points, junctions or slab voxels using the <i>Analyse Skeleton</i> function in BoneJ.	11
2.6	Examples of honeycombs.	15
2.7	Open-cell versus closed-cell foams.	15
2.8	Unit cells of the three lattice types used in the current work.	16
2.9	Unit cell for octet truss lattice.	16
2.10	Examples of rod-like and plate-like trabecular bone.	17
2.11	Unit cells of lattices used as idealised models for trabecular bone in literature.	18
2.12	Voxel-based Finite Element meshing of a trabecular bone sample using (a) hexahedral elements, or (b) tetrahedral elements.	19
2.13	Finite Element mesh using the beam-shell method.	20
2.14	Beam-shell discretisation process developed by Vanderoost, et al.	22
2.15	Back-calculation of Young’s modulus.	24
2.16	Inner core or embedding technique suggested by Ün, et al. and used by Harrison and McHugh.	26
2.17	Calculation of the “in situ” tensor value from the results of the PMUBCs and KUBCs simulations.	27
2.18	Boundary conditions used in the work on honeycombs.	28
2.19	Polar diagram showing the effect of confinement on angled cubic honeycombs.	29
2.20	Polar diagrams showing the effect of confinement on angled cubic lattices.	30
2.21	Custom testing platform developed by Hilton.	31

3.1	Flow chart of the developed analysis cycle.	34
3.2	Example of initial images from the bone scans.	35
3.3	Example of a bone substack including full bone cross sections.	35
3.4	Image processing steps applied in Fiji.	36
4.1	Flow chart of the test cases used for validation.	43
4.2	Representations of the three lattice types chosen for use in the validation study: cubic, Kelvin cell and octet truss.	45
4.3	Comparison of 3D renderings of size 2 and 6 octet truss lattices with a trabecular bone sample.	46
4.4	3D renderings of the size 8 cubic lattices with strut diameters of 0.1, 0.2 and 0.3, respectively.	46
4.5	3D renderings of the size 8 Kelvin cell lattices with diameters of 0.1, 0.2 and 0.3, respectively.	47
4.6	C8-t1 lattices in the uncropped and cropped conditions.	48
4.7	Angled cubic lattice.	50
4.8	Flow chart of the validation process.	51
6.1	Apparent modulus error results for an uncropped, unconstrained K8-t3 lattice at varying image sizes.	61
6.2	Inner modulus error results for an uncropped, unconstrained K8-t3 lattice at varying image sizes.	62
6.3	Modulus ratio results for an uncropped, unconstrained K8-t3 lattice at varying image sizes.	63
6.4	Image slices taken from the 110 voxel rasterisation of the K8-t3 lattice.	64
6.5	3D renderings of uncropped K8-t3 rasterised images at 100, 140, 200 and 400 voxels.	64
6.6	FE meshes produced for the uncropped K8-t3 lattices at image sizes of 100 and 400 voxels.	65
6.7	Enlarged image slices from the 290, 300 and 310 voxel TIFF stack for the K8-t3 lattice.	66
6.8	Enlarged view of graph of apparent modulus error and average beam cross- sectional area error for an uncropped, unconstrained K8-t3 lattice at varying image sizes.	67
6.9	Enlarged view of the error plots for the moments of inertia in the primary and secondary directions for an uncropped, unconstrained K8-t3 lattice at varying image sizes.	67

6.10	Illustration of the nodes used to calculate the inner modulus for the K8-t3 lattice.	69
6.11	Apparent modulus error plots for the uncropped, unconstrained K6-t3, K8-t3 and K10-t3 lattices at varying image sizes.	70
6.12	3D renderings of the K6-t3, K8-t3 and K10-t3 lattices at an image size of 250 voxels.	71
6.13	Modulus ratio results for the uncropped, unconstrained K6-t3 and K10-t3 lattices at varying image sizes.	72
6.14	Illustration of the nodes used to calculate the inner modulus for the K6, K8 and K10 lattices at an image size of 350 voxels.	73
6.15	Apparent modulus error results for the uncropped, unconstrained K8-t1 lattice at varying image sizes.	74
6.16	3D renderings of K8-t1 rasterised images at 100, 200, 300 and 400 voxels.	75
6.17	Apparent modulus error results for the uncropped, unconstrained K8-t2 lattice at varying image sizes.	76
6.18	Enlarged view of graph of apparent modulus error and average beam cross-sectional area error for an uncropped, unconstrained K8-t2 lattice at image sizes of 300 to 400 voxels.	76
6.19	Enlarged view of the error plots for the moments of inertia in the primary and secondary directions for an uncropped, unconstrained K8-t2 lattice at image sizes of 300 to 400 voxels.	77
6.20	3D renderings of the K8-t1, K8-t2 and K8-t3 lattices at an image size of 350 voxels.	77
6.21	Comparison of the error in apparent modulus for the uncropped, unconstrained and cropped, unconstrained K8-t3 lattices at varying image sizes.	78
6.22	Modulus ratio results for a cropped, unconstrained K8-t3 lattice at varying image sizes.	79
6.23	Apparent modulus error plot for an uncropped, constrained K8-t3 lattice at varying image sizes.	80
6.24	Modulus ratio results for an uncropped, constrained K8-t3 lattice at varying image sizes.	80
6.25	Apparent modulus error plot for a cropped, constrained K8-t3 lattice at varying image sizes.	82
6.26	Modulus ratio results for a cropped, constrained K8-t3 lattice at varying image sizes.	82
6.27	Enlarged view of the apparent modulus error plot for a cropped, unconstrained K8-t3 lattice at varying image sizes discretised using the VBM.	83

6.28	Enlarged view of the inner modulus error plot for a cropped, unconstrained K8-t3 lattice at varying image sizes discretised using the VBM.	83
6.29	Modulus ratio results for a cropped, unconstrained K8-t3 lattice at varying image sizes discretised using the VBM.	84
6.30	Modulus ratio results for a cropped, constrained K8-t3 lattice at varying image sizes discretised using the VBM.	84
6.31	Example of under- and over-estimations in strut rasterisation.	85
6.32	Apparent modulus error results for an uncropped, unconstrained C8-t1 lattice at varying image sizes.	88
6.33	Enlarged view of the apparent modulus error and element cross-sectional area error results for an uncropped, unconstrained C8-t1 lattice at varying image sizes.	88
6.34	Enlarged view of the error in the primary and secondary moments of inertia for an uncropped, unconstrained C8-t1 lattice at varying image sizes.	89
6.35	Apparent modulus error results for an uncropped, unconstrained C8-t2 lattice at varying image sizes.	90
6.36	Number of shell elements present in the FE mesh for an uncropped, unconstrained C8-t2 lattice at varying image sizes.	90
6.37	3D renderings of the C8-t2 lattice at image sizes of 230, 240, 300 and 310 voxels.	91
6.38	Image slices through a horizontal strut in a cubic lattice.	91
6.39	Representation of the mesh created for a C8-t2 lattice at an image size of 350 voxels.	92
6.40	Apparent modulus results for a cropped, constrained C8-t1 lattice at varying image sizes.	94
6.41	Modulus ratio results for a cropped, constrained C8-t1 lattice at varying image sizes.	94
6.42	Apparent modulus results for an uncropped, unconstrained angled C8-t2 lattice at varying image sizes.	96
6.43	Apparent modulus results for an uncropped, constrained angled C8-t2 lattice at varying image sizes.	96
6.44	Apparent modulus results for a cropped, unconstrained angled C8-t2 lattice at varying image sizes.	97
6.45	Apparent modulus results for a cropped, constrained angled C8-t2 lattice at varying image sizes.	97
6.46	Modulus ratio results for an uncropped, unconstrained angled C8-t2 lattice at varying image sizes.	98
6.47	Modulus ratio results for an uncropped, constrained angled C8-t2 lattice at varying image sizes.	98

6.48	Modulus ratio results for a cropped, unconstrained angled C8-t2 lattice at varying image sizes.	99
6.49	Modulus ratio results for a cropped, constrained angled C8-t2 lattice at varying image sizes.	100
6.50	Apparent modulus results for an uncropped, unconstrained OT8-t2 lattice at varying image sizes.	101
6.51	Inner modulus results for an uncropped, unconstrained OT8-t2 lattice at varying image sizes.	101
6.52	Apparent modulus and beam cross-sectional area results for an uncropped, unconstrained OT8-t2 lattice at varying image sizes.	102
6.53	Enlarged view of primary and secondary moments of inertia results for an uncropped, unconstrained OT8-t2 lattice at varying image sizes.	102
6.54	Error plot of the number of nodes in the FE mesh for an uncropped, unconstrained OT8-t2 lattice at varying image sizes.	103
6.55	Enlarged region of plots of the first layer of cells in a cubic, angled cubic, Kelvin cell and octet truss mesh.	104
6.56	Small elements caused by the presence of the extra nodes in the octet truss lattice.	105
6.57	Modulus ratio results for an uncropped, unconstrained OT8-t2 lattice at varying image sizes.	106
6.58	Modulus ratio results for an uncropped, constrained OT8-t2 lattice at varying image sizes.	107
7.1	Comparison of BSM and experimental results for standard bone specimens. . .	110
7.2	Comparison of VBM and experimental results for standard bone specimens. . .	110
7.3	Comparison of BSM and experimental results for large bone specimens. . . .	112
7.4	Comparison of BSM and experimental results for inner core of large bone specimens.	113
7.5	Comparison of BSM and experimental results for modulus ratio of large bone specimens.	113
7.6	Comparison of VBM and experimental results for large bone specimens. . . .	114
7.7	Comparison of VBM and experimental results for inner core of large bone specimens.	115
7.8	Comparison of VBM and experimental results for modulus ratio of large bone specimens.	115
7.9	Time taken for discretisation of the standard bone specimens using the BSM and VBM. Results shown in ascending order of BSM time.	116

7.10	Time taken for simulation of the standard bone specimens using the BSM and VBM. Results shown in ascending order of BSM time.	117
7.11	Total time taken for discretisation and simulation of the standard bone specimens using the BSM and VBM. Results shown in ascending order of BSM time.	118
7.12	Time taken for discretisation of the large bone specimens using the BSM and VBM.	118
7.13	Time taken for simulation of the large bone specimens using the BSM and VBM.	119
7.14	Total time taken for discretisation and simulation of the large bone specimens using the BSM and VBM.	119
7.15	Comparison of apparent modulus results and bone volume fraction.	121
7.16	Comparison of apparent modulus results and mean trabecular thickness.	122
7.17	Comparison of apparent modulus results and mean trabecular separation.	123
7.18	Comparison of apparent modulus results and connectivity.	124
7.19	Comparison of apparent modulus results and connectivity density.	125
7.20	Comparison of apparent modulus results and shell-to-beam ratio.	126
7.21	Comparison of micro-CT scans of specimen 477 at varying voxel sizes.	127
7.22	Comparison of micro-CT scans of specimen 518 at varying voxel sizes.	127
7.23	Comparison of BSM and experimental results for standard bone specimens segmented using the <i>Optimise Threshold</i> function.	128
7.24	Comparison of BSM and experimental modulus ratio results for large bone specimens segmented using the <i>Optimise Threshold</i> function.	129
7.25	Comparison of BSM and experimental apparent modulus results for standard 10 mm bone specimens segmented using the <i>Optimise Threshold</i> function with a lateral constraint on the top and bottom surfaces.	129
7.26	Comparison of experimental and BSM modulus ratio results for large bone specimens segmented using the <i>Optimise Threshold</i> function with a lateral constraint on the top and bottom surfaces.	130
7.27	BSM results for apparent modulus of standard specimens simulated using a variety of boundary conditions.	132
7.28	BSM results for inner modulus of large specimens simulated using a variety of boundary conditions.	133
7.29	BSM modulus ratio results for large specimens simulated using a variety of boundary conditions.	135
8.1	Comparison of BSM and experimental results for standard bone specimens with line fit.	138

8.2	Comparison of BSM and experimental results for defatted, marrow and confined bone specimens.	139
8.3	Comparison of a defatted and marrow image slice with corresponding segmentation.	140
8.4	Comparison of confined specimen images with corresponding segmentation. . .	141
8.5	Comparison of BSM and experimental results for modulus ratio of large bone specimens with data labels.	142
8.6	Comparison of BSM and experimental results for large bone specimens.	144
8.7	Comparison of VBM and experimental results for standard bone specimens with line fit.	145
8.8	Comparison of VBM and experimental results for defatted, marrow and confined bone specimens.	147
8.9	Comparison of outlying VBM specimen images with corresponding segmentation.	148
8.10	Comparison of BSM and VBM results to experimental results for standard bone specimens.	149
8.11	Comparison of BSM results to VBM results for standard bone specimens.	150
8.12	Comparison of BSM results to VBM results using a log-log scale.	151
8.13	Comparison of VBM and experimental results for large bone specimens.	153
8.14	Comparison of VBM and experimental results for modulus ratio of large bone specimens with data labels.	154
8.15	Comparison of experimental, BSM and VBM results for modulus ratio of large bone specimens.	155
8.16	Comparison of apparent modulus results and bone volume fraction with fitted curve.	158
8.17	Comparison of apparent modulus results and trabecular separation with fitted curve.	160
8.18	Comparison of BSM and experimental results for standard bone specimens with high S2B ratios highlighted.	161
8.19	Comparison of micro-CT scans of boundaries of specimen 518 at varying voxel size.	162
8.20	Noise in 20 μm scan of specimen 518.	164
8.21	Comparison of BSM and experimental apparent modulus results for standard bone specimens segmented using two different segmentation methods.	165
8.22	Percentage increase in apparent modulus results using the <i>Optimise Threshold</i> function versus the <i>Make Binary</i> function.	165

8.23	Comparison of BSM and experimental modulus ratio results for standard bone specimens segmented using two different segmentation methods.	166
8.24	Increase in stiffness caused by the application of lateral and rotational constraints on the top and bottom surfaces of the specimen simultaneously.	167
8.25	Comparison of BSM and experimental results for standard bone specimens with line fit.	168
8.26	Comparison between micro-CT scans used in the Vanderoost study and this thesis.	171
C.1	Perpendicular distance criteria for beam rasterisation.	197
C.2	Parallel distance criteria for beam rasterisation.	197
C.3	Bounding planes used in thick shell rasterization.	198
D.1	Error plots for apparent modulus of unconstrained C6-t1 lattices at varying image sizes.	202
D.2	Error plots for inner modulus of unconstrained C6-t1 lattices at varying image sizes.	202
D.3	Modulus ratio results of unconstrained C6-t1 lattices at varying image sizes.	202
D.4	Error plots for apparent modulus of constrained C6-t1 lattices at varying image sizes.	203
D.5	Error plots for inner modulus of constrained C6-t1 lattices at varying image sizes.	203
D.6	Modulus ratio results of constrained C6-t1 lattices at varying image sizes.	203
D.7	Error plots for apparent modulus of unconstrained C6-t2 lattices at varying image sizes.	204
D.8	Error plots for inner modulus of unconstrained C6-t2 lattices at varying image sizes.	204
D.9	Modulus ratio results of unconstrained C6-t2 lattices at varying image sizes.	204
D.10	Error plots for apparent modulus of constrained C6-t2 lattices at varying image sizes.	205
D.11	Error plots for inner modulus of constrained C6-t2 lattices at varying image sizes.	205
D.12	Modulus ratio results of constrained C6-t2 lattices at varying image sizes.	205
D.13	Error plots for apparent modulus of unconstrained C6-t3 lattices at varying image sizes.	206
D.14	Error plots for inner modulus of unconstrained C6-t3 lattices at varying image sizes.	206
D.15	Modulus ratio results of unconstrained C6-t3 lattices at varying image sizes.	206

D.16	Error plots for apparent modulus of constrained C6-t3 lattices at varying image sizes.	207
D.17	Error plots for inner modulus of constrained C6-t3 lattices at varying image sizes.	207
D.18	Modulus ratio results of constrained C6-t3 lattices at varying image sizes. . . .	207
D.19	Error plots for apparent modulus of unconstrained C8-t1 lattices at varying image sizes.	208
D.20	Error plots for inner modulus of unconstrained C8-t1 lattices at varying image sizes.	208
D.21	Modulus ratio results of unconstrained C8-t1 lattices at varying image sizes. . .	208
D.22	Error plots for apparent modulus of constrained C8-t1 lattices at varying image sizes.	209
D.23	Error plots for inner modulus of constrained C8-t1 lattices at varying image sizes.	209
D.24	Modulus ratio results of constrained C8-t1 lattices at varying image sizes. . . .	209
D.25	Error plots for apparent modulus of unconstrained C8-t2 lattices at varying image sizes.	210
D.26	Error plots for inner modulus of unconstrained C8-t2 lattices at varying image sizes.	210
D.27	Modulus ratio results of unconstrained C8-t2 lattices at varying image sizes. . .	210
D.28	Error plots for apparent modulus of constrained C8-t2 lattices at varying image sizes.	211
D.29	Error plots for inner modulus of constrained C8-t2 lattices at varying image sizes.	211
D.30	Modulus ratio results of constrained C8-t2 lattices at varying image sizes. . . .	211
D.31	Error plots for apparent modulus of unconstrained C8-t3 lattices at varying image sizes.	212
D.32	Error plots for inner modulus of unconstrained C8-t3 lattices at varying image sizes.	212
D.33	Modulus ratio results of unconstrained C8-t3 lattices at varying image sizes. . .	212
D.34	Error plots for apparent modulus of constrained C8-t3 lattices at varying image sizes.	213
D.35	Error plots for inner modulus of constrained C8-t3 lattices at varying image sizes.	213
D.36	Modulus ratio results of constrained C8-t3 lattices at varying image sizes. . . .	213
D.37	Error plots for apparent modulus of unconstrained C10-t1 lattices at varying image sizes.	214
D.38	Error plots for inner modulus of unconstrained C10-t1 lattices at varying image sizes.	214

D.39	Modulus ratio results of unconstrained C10-t1 lattices at varying image sizes.	214
D.40	Error plots for apparent modulus of constrained C10-t1 lattices at varying image sizes.	215
D.41	Error plots for inner modulus of constrained C10-t1 lattices at varying image sizes.	215
D.42	Modulus ratio results of constrained C10-t1 lattices at varying image sizes.	215
D.43	Error plots for apparent modulus of unconstrained C10-t2 lattices at varying image sizes.	216
D.44	Error plots for inner modulus of unconstrained C10-t2 lattices at varying image sizes.	216
D.45	Modulus ratio results of unconstrained C10-t2 lattices at varying image sizes.	216
D.46	Error plots for apparent modulus of constrained C10-t2 lattices at varying image sizes.	217
D.47	Error plots for inner modulus of constrained C10-t2 lattices at varying image sizes.	217
D.48	Modulus ratio results of constrained C10-t2 lattices at varying image sizes.	217
D.49	Error plots for apparent modulus of unconstrained C10-t3 lattices at varying image sizes.	218
D.50	Error plots for inner modulus of unconstrained C10-t3 lattices at varying image sizes.	218
D.51	Modulus ratio results of unconstrained C10-t3 lattices at varying image sizes.	218
D.52	Error plots for apparent modulus of constrained C10-t3 lattices at varying image sizes.	219
D.53	Error plots for inner modulus of constrained C10-t3 lattices at varying image sizes.	219
D.54	Modulus ratio results of constrained C10-t3 lattices at varying image sizes.	219
D.55	Error plots for apparent modulus of unconstrained K6-t1 lattices at varying image sizes.	221
D.56	Error plots for inner modulus of unconstrained K6-t1 lattices at varying image sizes.	221
D.57	Modulus ratio results of unconstrained K6-t1 lattices at varying image sizes.	221
D.58	Error plots for apparent modulus of constrained K6-t1 lattices at varying image sizes.	222
D.59	Error plots for inner modulus of constrained K6-t1 lattices at varying image sizes.	222
D.60	Modulus ratio results of constrained K6-t1 lattices at varying image sizes.	222

D.61	Error plots for apparent modulus of unconstrained K6-t2 lattices at varying image sizes.	223
D.62	Error plots for inner modulus of unconstrained K6-t2 lattices at varying image sizes.	223
D.63	Modulus ratio results of unconstrained K6-t1 lattices at varying image sizes. . .	223
D.64	Error plots for apparent modulus of constrained K6-t2 lattices at varying image sizes.	224
D.65	Error plots for inner modulus of constrained K6-t2 lattices at varying image sizes.	224
D.66	Modulus ratio results of constrained K6-t2 lattices at varying image sizes. . . .	224
D.67	Error plots for apparent modulus of unconstrained K6-t3 lattices at varying image sizes.	225
D.68	Error plots for inner modulus of unconstrained K6-t3 lattices at varying image sizes.	225
D.69	Modulus ratio results of unconstrained K6-t3 lattices at varying image sizes. . .	225
D.70	Error plots for apparent modulus of constrained K6-t3 lattices at varying image sizes.	226
D.71	Error plots for inner modulus of constrained K6-t3 lattices at varying image sizes.	226
D.72	Modulus ratio results of constrained K6-t3 lattices at varying image sizes. . . .	226
D.73	Error plots for apparent modulus of unconstrained K8-t1 lattices at varying image sizes.	227
D.74	Error plots for inner modulus of unconstrained K8-t1 lattices at varying image sizes.	227
D.75	Modulus ratio results of unconstrained K8-t1 lattices at varying image sizes. . .	227
D.76	Error plots for apparent modulus of constrained K8-t1 lattices at varying image sizes.	228
D.77	Error plots for inner modulus of constrained K8-t1 lattices at varying image sizes.	228
D.78	Modulus ratio results of constrained K8-t1 lattices at varying image sizes. . . .	228
D.79	Error plots for apparent modulus of unconstrained K8-t2 lattices at varying image sizes.	229
D.80	Error plots for inner modulus of unconstrained K8-t2 lattices at varying image sizes.	229
D.81	Modulus ratio results of unconstrained K8-t1 lattices at varying image sizes. . .	229
D.82	Error plots for apparent modulus of constrained K8-t2 lattices at varying image sizes.	230
D.83	Error plots for inner modulus of constrained K8-t2 lattices at varying image sizes.	230

D.84	Modulus ratio results of constrained K8-t2 lattices at varying image sizes. . . .	230
D.85	Error plots for apparent modulus of unconstrained K8-t3 lattices at varying image sizes.	231
D.86	Error plots for inner modulus of unconstrained K8-t3 lattices at varying image sizes.	231
D.87	Modulus ratio results of unconstrained K8-t3 lattices at varying image sizes. . .	231
D.88	Error plots for apparent modulus of constrained K8-t3 lattices at varying image sizes.	232
D.89	Error plots for inner modulus of constrained K8-t3 lattices at varying image sizes.	232
D.90	Modulus ratio results of constrained K8-t3 lattices at varying image sizes. . . .	232
D.91	Error plots for apparent modulus of unconstrained K10-t1 lattices at varying image sizes.	233
D.92	Error plots for inner modulus of unconstrained K10-t1 lattices at varying image sizes.	233
D.93	Modulus ratio results of unconstrained K10-t1 lattices at varying image sizes. .	233
D.94	Error plots for apparent modulus of constrained K10-t1 lattices at varying image sizes.	234
D.95	Error plots for inner modulus of constrained K10-t1 lattices at varying image sizes.	234
D.96	Modulus ratio results of constrained K10-t1 lattices at varying image sizes. . .	234
D.97	Error plots for apparent modulus of unconstrained K10-t2 lattices at varying image sizes.	235
D.98	Error plots for inner modulus of unconstrained K10-t2 lattices at varying image sizes.	235
D.99	Modulus ratio results of unconstrained K10-t1 lattices at varying image sizes. .	235
D.100	Error plots for apparent modulus of constrained K10-t2 lattices at varying image sizes.	236
D.101	Error plots for inner modulus of constrained K10-t2 lattices at varying image sizes.	236
D.102	Modulus ratio results of constrained K10-t2 lattices at varying image sizes. . .	236
D.103	Error plots for apparent modulus of unconstrained K10-t3 lattices at varying image sizes.	237
D.104	Error plots for inner modulus of unconstrained K10-t3 lattices at varying image sizes.	237
D.105	Modulus ratio results of unconstrained K10-t3 lattices at varying image sizes. .	237

D.106	Error plots for apparent modulus of constrained K10-t3 lattices at varying image sizes.	238
D.107	Error plots for inner modulus of constrained K10-t3 lattices at varying image sizes.	238
D.108	Modulus ratio results of constrained K10-t3 lattices at varying image sizes. . .	238
D.109	Error plots for apparent modulus of unconstrained OT6-t1 lattices at varying image sizes.	240
D.110	Error plots for inner modulus of unconstrained OT6-t1 lattices at varying image sizes.	240
D.111	Modulus ratio results of unconstrained OT6-t1 lattices at varying image sizes. .	240
D.112	Error plots for apparent modulus of constrained OT6-t1 lattices at varying image sizes.	241
D.113	Error plots for inner modulus of constrained OT6-t1 lattices at varying image sizes.	241
D.114	Modulus ratio results of constrained OT6-t1 lattices at varying image sizes. . .	241
D.115	Error plots for apparent modulus of unconstrained OT6-t2 lattices at varying image sizes.	242
D.116	Error plots for inner modulus of unconstrained OT6-t2 lattices at varying image sizes.	242
D.117	Modulus ratio results of unconstrained OT6-t2 lattices at varying image sizes. .	242
D.118	Error plots for apparent modulus of constrained OT6-t2 lattices at varying image sizes.	243
D.119	Error plots for inner modulus of constrained OT6-t2 lattices at varying image sizes.	243
D.120	Modulus ratio results of constrained OT6-t2 lattices at varying image sizes. . .	243
D.121	Error plots for apparent modulus of unconstrained OT6-t3 lattices at varying image sizes.	244
D.122	Error plots for inner modulus of unconstrained OT6-t3 lattices at varying image sizes.	244
D.123	Modulus ratio results of unconstrained OT6-t3 lattices at varying image sizes. .	244
D.124	Error plots for apparent modulus of constrained OT6-t3 lattices at varying image sizes.	245
D.125	Error plots for inner modulus of constrained OT6-t3 lattices at varying image sizes.	245
D.126	Modulus ratio results of constrained OT6-t3 lattices at varying image sizes. . .	245

D.127	Error plots for apparent modulus of unconstrained OT8-t1 lattices at varying image sizes.	246
D.128	Error plots for inner modulus of unconstrained OT8-t1 lattices at varying image sizes.	246
D.129	Modulus ratio results of unconstrained OT8-t1 lattices at varying image sizes.	246
D.130	Error plots for apparent modulus of constrained OT8-t1 lattices at varying image sizes.	247
D.131	Error plots for inner modulus of constrained OT8-t1 lattices at varying image sizes.	247
D.132	Modulus ratio results of constrained OT8-t1 lattices at varying image sizes.	247
D.133	Error plots for apparent modulus of unconstrained OT8-t2 lattices at varying image sizes.	248
D.134	Error plots for inner modulus of unconstrained OT8-t2 lattices at varying image sizes.	248
D.135	Modulus ratio results of unconstrained OT8-t2 lattices at varying image sizes.	248
D.136	Error plots for apparent modulus of constrained OT8-t2 lattices at varying image sizes.	249
D.137	Error plots for inner modulus of constrained OT8-t2 lattices at varying image sizes.	249
D.138	Modulus ratio results of constrained OT8-t2 lattices at varying image sizes.	249
D.139	Error plots for apparent modulus of unconstrained OT8-t3 lattices at varying image sizes.	250
D.140	Error plots for inner modulus of unconstrained OT8-t3 lattices at varying image sizes.	250
D.141	Modulus ratio results of unconstrained OT8-t3 lattices at varying image sizes.	250
D.142	Error plots for apparent modulus of constrained OT8-t3 lattices at varying image sizes.	251
D.143	Error plots for inner modulus of constrained OT8-t3 lattices at varying image sizes.	251
D.144	Modulus ratio results of constrained OT8-t3 lattices at varying image sizes.	251
D.145	Error plots for apparent modulus of unconstrained OT10-t1 lattices at varying image sizes.	252
D.146	Error plots for inner modulus of unconstrained OT10-t1 lattices at varying image sizes.	252
D.147	Modulus ratio results of unconstrained OT10-t1 lattices at varying image sizes.	252

D.148	Error plots for apparent modulus of constrained OT10-t1 lattices at varying image sizes.	253
D.149	Error plots for inner modulus of constrained OT10-t1 lattices at varying image sizes.	253
D.150	Modulus ratio results of constrained OT10-t1 lattices at varying image sizes. . .	253
D.151	Error plots for apparent modulus of unconstrained OT10-t2 lattices at varying image sizes.	254
D.152	Error plots for inner modulus of unconstrained OT10-t2 lattices at varying image sizes.	254
D.153	Modulus ratio results of unconstrained OT10-t2 lattices at varying image sizes.	254
D.154	Error plots for apparent modulus of constrained OT10-t2 lattices at varying image sizes.	255
D.155	Error plots for inner modulus of constrained OT10-t2 lattices at varying image sizes.	255
D.156	Modulus ratio results of constrained OT10-t2 lattices at varying image sizes. . .	255
D.157	Error plots for apparent modulus of unconstrained OT10-t3 lattices at varying image sizes.	256
D.158	Error plots for inner modulus of unconstrained OT10-t3 lattices at varying image sizes.	256
D.159	Modulus ratio results of unconstrained OT10-t3 lattices at varying image sizes.	256
D.160	Error plots for apparent modulus of constrained OT10-t3 lattices at varying image sizes.	257
D.161	Error plots for inner modulus of constrained OT10-t3 lattices at varying image sizes.	257
D.162	Modulus ratio results of constrained OT10-t3 lattices at varying image sizes. . .	257
D.163	Error plots for apparent modulus of unconstrained angled C8-t1 lattices at varying image sizes.	259
D.164	Error plots for inner modulus of unconstrained angled C8-t1 lattices at varying image sizes.	259
D.165	Modulus ratio results of unconstrained angled C8-t1 lattices at varying image sizes.	259
D.166	Error plots for apparent modulus of constrained angled C8-t1 lattices at varying image sizes.	260
D.167	Error plots for inner modulus of constrained angled C8-t1 lattices at varying image sizes.	260
D.168	Modulus ratio results of constrained angled C8-t1 lattices at varying image sizes.	260

D.169	Error plots for apparent modulus of unconstrained C8-t2 lattices at varying image sizes.	261
D.170	Error plots for inner modulus of unconstrained C8-t2 lattices at varying image sizes.	261
D.171	Modulus ratio results of unconstrained C8-t2 lattices at varying image sizes. . .	261
D.172	Error plots for apparent modulus of constrained C8-t2 lattices at varying image sizes.	262
D.173	Error plots for inner modulus of constrained C8-t2 lattices at varying image sizes.	262
D.174	Modulus ratio results of constrained C8-t2 lattices at varying image sizes. . . .	262
D.175	Error plots for apparent modulus of cropped unconstrained and cropped constrained K8-t3 lattices at varying image sizes obtained using the VBM.	264
D.176	Error plots for inner modulus of cropped unconstrained and cropped constrained K8-t3 lattices at varying image sizes obtained using the VBM.	264
D.177	Modulus ratio results of cropped unconstrained and cropped constrained K8-t3 lattices at varying image sizes obtained using the VBM.	264
E.1	Enlarged view of graph of apparent modulus error and average beam cross-sectional area error for an uncropped, unconstrained K6-t3 lattice at varying image sizes.	266
E.2	Enlarged view of the error plots for the moments of inertia in the primary and secondary directions for an uncropped, unconstrained K6-t3 lattice at varying image sizes.	266
E.3	Enlarged view of graph of apparent modulus error and average beam cross-sectional area error for an uncropped, unconstrained K10-t3 lattice at varying image sizes.	267
E.4	Enlarged view of the error plots for the moments of inertia in the primary and secondary directions for an uncropped, unconstrained K10-t3 lattice at varying image sizes.	267
F.1	Rasterisation of C6-t1 lattice at varying image sizes.	269
F.2	Rasterisation of C6-t2 lattice at varying image sizes.	270
F.3	Rasterisation of C6-t3 lattice at varying image sizes.	271
F.4	Rasterisation of C8-t1 lattice at varying image sizes.	272
F.5	Rasterisation of C8-t2 lattice at varying image sizes.	273
F.6	Rasterisation of C8-t3 lattice at varying image sizes.	274
F.7	Rasterisation of C10-t1 lattice at varying image sizes.	275
F.8	Rasterisation of C10-t2 lattice at varying image sizes.	276

F.9	Rasterisation of C10-t3 lattice at varying image sizes.	277
F.10	Rasterisation of K6-t1 lattice at varying image sizes.	279
F.11	Rasterisation of K6-t2 lattice at varying image sizes.	280
F.12	Rasterisation of K6-t3 lattice at varying image sizes.	281
F.13	Rasterisation of K8-t1 lattice at varying image sizes.	282
F.14	Rasterisation of K8-t2 lattice at varying image sizes.	283
F.15	Rasterisation of K8-t3 lattice at varying image sizes.	284
F.16	Rasterisation of K10-t1 lattice at varying image sizes.	285
F.17	Rasterisation of K10-t2 lattice at varying image sizes.	286
F.18	Rasterisation of K10-t3 lattice at varying image sizes.	287
F.19	Rasterisation of OT6-t1 lattice at varying image sizes.	289
F.20	Rasterisation of OT6-t2 lattice at varying image sizes.	290
F.21	Rasterisation of OT6-t3 lattice at varying image sizes.	291
F.22	Rasterisation of OT8-t1 lattice at varying image sizes.	292
F.23	Rasterisation of OT8-t2 lattice at varying image sizes.	293
F.24	Rasterisation of OT8-t3 lattice at varying image sizes.	294
F.25	Rasterisation of OT10-t1 lattice at varying image sizes.	295
F.26	Rasterisation of OT10-t2 lattice at varying image sizes.	296
F.27	Rasterisation of OT10-t3 lattice at varying image sizes.	297
F.28	Rasterisation of angled C8-t1 lattice at varying image sizes.	299
F.29	Rasterisation of angled C8-t2 lattice at varying image sizes.	300
G.1	Discretisation of a strut using the BSM and VBM.	301
G.2	Rotation of beam-shell mesh with translational z -constraint.	302
G.3	Rotation of voxel-based mesh with translational z -constraint.	302
I.1	Segmentation of bone scans (13–33).	306
I.2	Segmentation of bone scans (38–124).	307
I.3	Segmentation of bone scans (129–165).	308
I.4	Segmentation of bone scans (171–247).	309
I.5	Segmentation of bone scans (248–283).	310

I.6	Segmentation of bone scans (285–366).	311
I.7	Segmentation of bone scans (371–424).	312
I.8	Segmentation of bone scans (427–455).	313
I.9	Segmentation of bone scans (456–494).	314
I.10	Segmentation of bone scans (496–529).	315
I.11	Segmentation of bone scans (530–538).	316
J.1	Least-squares line fit for results from BSM analysis of standard bone specimens.	318
J.2	Histogram of instantaneous gradients.	318
J.3	Histogram of natural logarithms of instantaneous gradients.	319
J.4	Log-normal line fits for results from BSM analysis of standard bone specimens.	321

List of Tables

2.1	Summary of literature relating mechanical properties to microstructural indices.	14
4.1	Properties of the cubic, Kelvin cell and octet truss lattices.	45
5.1	Parameters used for nano-CT scanning at varying voxel size.	54
6.1	Reason for the error surges seen in the uncropped, unconstrained K6-t3, K8-t3 and K10-t3 lattices.	71
6.2	Baseline results for the apparent modulus of the uncropped and cropped K8-t3 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained apparent modulus from unconstrained apparent modulus.	81
6.3	Baseline results for the modulus ratio of the uncropped and cropped K8-t3 lattices in the unconstrained and constrained conditions, including the percentage increase in ratio from the unconstrained to constrained condition.	81
6.4	Average number of voxels representing strut cross-sections in the C8-t1 lattice rasterisations.	89
6.5	Baseline results for the apparent modulus of the uncropped and cropped C8-t1 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained apparent modulus from unconstrained apparent modulus.	93
6.6	Baseline results for the modulus ratio of the uncropped and cropped C8-t1 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained modulus ratio from unconstrained modulus ratio.	93
6.7	Baseline results for the apparent modulus of the uncropped and cropped angled C8-t2 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained apparent modulus from unconstrained apparent modulus.	95
6.8	Baseline results for the modulus ratio of the uncropped and cropped angled C8-t2 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained modulus ratio from unconstrained modulus ratio.	95
6.9	Baseline results for the apparent modulus of the uncropped and cropped OT8-t2 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained apparent modulus from unconstrained apparent modulus.	106

6.10	Baseline results for the modulus ratio of the uncropped and cropped OT8-t2 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained modulus ratio from unconstrained modulus ratio.	106
7.1	Microstructural indices for different voxel sizes.	127
7.2	Apparent modulus results at different voxel sizes.	128
7.3	Average increase in stiffness experienced by the bone specimens due to the addition of lateral and rotational constraints. Increase is determined relative to the “Standard” boundary condition.	131
7.4	Average change in inner stiffness experienced by the large bone specimens due to the addition of lateral and rotational constraints.	131
7.5	Average change in modulus ratio of the large bone specimens due to the addition of lateral and rotational constraints.	134
8.1	Percentage difference in microstructural indices when compared to 20 μm indices.	163
8.2	Increase in apparent modulus results at different voxel sizes.	163
A.1	Hardware used for discretisation.	193
A.2	Hardware used for simulation.	193
A.3	Software used in this work.	193
H.1	List of bone specimens.	304

Chapter 1

Introduction

1.1 Motivation for research

Research into the mechanical properties of trabecular bone has been ongoing for at least 40 years [4, 5]. Despite this long history, there is a continuous drive to understand this complex tissue and reliably predict its mechanical properties.

There are two major complications which arise during the study of trabecular bone. Firstly, there is uncertainty associated with the Young's modulus of the tissue. Young's moduli for trabecular bone have been reported in the literature, however these values span a large range [1]. Moreover, it is not known whether there is a single Young's modulus for the bone or if this value varies in different regions based on how the bone was formed [6].

The second complication is that the structure of trabecular bone is extremely complex [7] and can vary between individuals [1], anatomic sites [1] and over time [8]. Consequently, in order to determine the mechanical properties of the trabecular bone of a particular person, the structure of the bone for that person must be considered. Due to this constraint, there are only two methods which may reliably be used to determine the mechanical properties of the bone, namely experimental testing and detailed patient-specific Finite Element (FE) modelling of the bone microstructure.

Experimental testing has its benefits, but it requires a sample of bone to be removed from the body, which is not possible for patients, particularly those who have already experienced bone degradation and loss due to diseases like osteoporosis. As a result, numerical FE models have become increasingly popular as they can be constructed from micro-CT scans of a bone sample [9]. Currently, micro-CT scans also require a bone sample to be removed from the body, however CT scanning technology is continuously improving, making it likely that general CT scans, which do not require a sample to be removed from the body, could soon be capable of capturing the microstructure in as much detail as current micro-CT scans.

Using FE models to determine mechanical properties has an additional advantage over experimental testing, namely repeatability. A bone sample may only be experimentally tested once as samples are generally tested to the point of failure. In contrast, there is no limit to the number of times a FE model may be simulated, meaning that multiple properties may be determined for a single sample, thereby more fully describing the behaviour of the bone.

The voxel-based method (VBM) [10] is widely-used to construct FE models of trabecular bone from micro-CT scans and is considered to be the computational “gold-standard” [11, 12] for determining trabecular bone properties. In this method, the structure is discretised by creating a solid tetrahedral element for each bone voxel¹ in the micro-CT scan. Defining the elements in this way leads to extremely large meshes, which are computationally expensive to simulate, often requiring specialised high-speed computing resources. The time- and memory-intensive nature of the VBM makes it impractical in a clinical setting.

In an effort to reduce the resources required for the simulation of trabecular bone, Vanderroost, et al [2] developed an alternative Finite Element code which allows a micro-CT scanned structure to be discretised into a series of beams and shells. It was reported that this beam-shell method (BSM) vastly reduces the size of the mesh and, consequently, the processing time required for the simulation, while maintaining excellent agreement to the VBM [2].

Despite the popularity of Finite Element models in trabecular bone research, there are complications when using this methodology that should not be underestimated. Most notably, the FE model is only as good as the micro-CT scan from which it was created. In order to obtain an accurate representation of the structure, small voxel sizes must be used during micro-CT scanning which is both time-consuming and expensive [13]. Furthermore, the image processing applied to the micro-CT scan can have a significant effect on the FE model [14] and, consequently, the apparent properties determined from the simulation. As a result, it is necessary to monitor the choice of image acquisition and processing techniques closely when using FEA to investigate trabecular bone behaviour.

In addition to determining the apparent properties of trabecular bone, there have been attempts, both experimental [3] and numerical [15], to determine the effective properties of the bone, which have greater significance than apparent properties. Effective properties describe the behaviour of the bone as a continuum, and allow for the creation of a material model of the bone. A material model, which ideally captures the behaviour of the bone in all loading scenarios, can be applied to FE meshes and would negate the need for modelling of the complex trabecular microstructure. Although the voxel-based method has been used to predict the effective properties of trabecular bone, the capability of the more efficient beam-shell method in determining effective properties has yet to be explored.

1.2 Problem statement

Most research into trabecular bone is performed using Finite Element Analysis of micro-CT scanned bone specimens using a voxel-based method. Three limitations have been identified in the literature:

- The methodology used to simulate the response of the bone is often applied directly to bone, without first validating the methodology by applying it to simpler structures with known mechanical responses.
- The voxel-based approach has been shown to correlate to experimental results in earlier research². As a result, the majority of current research does not compare the simulation results to experimental results.

¹A voxel is a small volume of a CT scan. It is similar to a pixel in a 2D image, however has three coordinates (x, y, z) associated with it as opposed to only two (x, y).

²See Section 2.5.1.1

- Much of the current research investigates the apparent properties of the bone, rather than the effective properties which have broader applications.

This work seeks to address the aforementioned limitations by using a sufficiently validated microstructural beam-shell FE approach to determine the effective properties of trabecular bone before comparing the simulation results to previously obtained experimental results.

1.3 Aim and objectives

The aims of this work are as follows:

1. Develop methodology to numerically determine the apparent modulus of trabecular bone samples using microstructural Finite Element models.
2. Investigate the performance of microstructural beam-shell FE models in predicting the response of trabecular bone, particularly the apparent and effective moduli.

This aim will be accomplished by completing the following objectives:

1. Create an automated analysis cycle which imports micro-CT scans of a structure, discretises the structure using the beam-shell method and returns the apparent modulus and a recreated image of the structure.
2. Validate the developed analysis cycle using periodic 3D lattices with known mechanical responses.
3. Run trabecular bone samples through the developed analysis cycle and evaluate the degree to which the beam-shell approach captures the response of trabecular bone.

1.4 Hypothesis

The hypothesis of this project is as follows:

The beam-shell model is capable of determining the apparent and effective moduli of trabecular bone, given appropriate material properties and boundary conditions are applied to the FE mesh.

Chapter 2

Literature review

This section provides an overview of the literature relevant to the proposed project, including information pertaining to the structure of trabecular bone and the challenges faced when modelling this complex tissue.

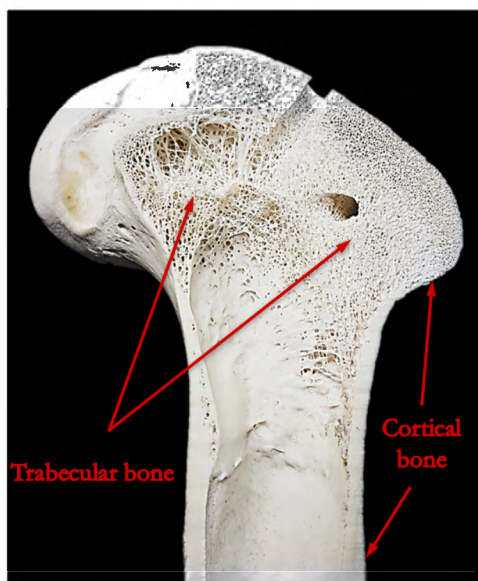
2.1 Trabecular bone

At the macroscopic scale, bone is composed of two major load-bearing tissues, namely cortical and trabecular bone (shown in Figure 2.1a). The outer layer of the bone is composed of dense, hard cortical bone which acts as a protective barrier and adds stiffness and strength to the whole bone. The trabecular (or cancellous) bone fills the interior spaces in the heads of long bones like the femur and is mainly responsible for energy absorption and load distribution in the body [7]. Recent studies have found that trabecular bone carries a significant portion of the load in some regions of the human femoral neck [16] and vertebrae [17]. Although these two tissues are easily distinguishable by their vastly different porosity, it is the arrangement of the lamellae at the microscale which truly sets them apart [1]. As shown by Figure 2.1b, in mature human cortical bone, the lamellae (layers of mineralised collagen fibrils) arrange themselves into concentric cylinders called osteons, with axes aligned along the long axis of the bone; in trabecular bone, the lamellae are arranged into rod-like structures, termed “trabeculae” [7]. These trabeculae combine in a cellular¹ arrangement of rod-rod, rod-plate, or plate-plate connections to form an interconnected framework with marrow-filled cavities [1] as shown in Figure 2.2, which lends stiffness and strength to the structure while minimising mass [18, 19].


A difficulty in the study of trabecular bone is the large variance in microstructure between individuals and anatomic sites. Trabecular bone adapts to the stresses it experiences [18], causing low density rod-like trabecular bone to form in areas of low stress and high density plate-like trabecular bone to form in high stress areas [8]. Because the mechanical properties² of trabecular bone show a direct correlation to its structure [20, 21], the mechanical properties of trabecular bone can vary by a factor of 2–5 based on anatomic site [1]. Additionally, the mechanical properties within the same anatomic site vary significantly based on location, with

¹Here the term “cellular” is used to describe a spatial arrangement, as in cellular solids (see Section 2.2), not a biological cell.

²Due to the multi-scale aspect of trabecular bone, mechanical properties must be considered in relation to the length scales. Three major properties are considered in this work, namely material properties, apparent properties and effective properties, all of which are described in detail in Section 2.5.3.



(a) Cortical and trabecular bone in a bovine femur.



(b) Arrangement of lamellae in trabecular bone and cortical bone. Adapted from [7].

Figure 2.1: Structure of bone.




Figure 2.2: Trabecular bone samples showing the characteristic rod-plate microstructure [8].




Figure 2.3: Close-up view of bovine femur. Differing structures can clearly be seen in regions (a), (b) and (c). (a) is near the load-bearing surface of the bone, (b) is at the free boundary of the trabecular bone and (c) is a transition region.

vastly different structures seen along the load-bearing surface and at the free boundary [1], as shown in Figure 2.3.

Another structural factor that needs to be taken into account is the continuous remodelling of the bone in response to ageing and disease [8]. Trabecular bone is substantially more metabolically active than cortical bone, i.e. it is remodelled more regularly [1], which causes a significant change in microstructure over the course of a lifetime. Furthermore, osteoporosis and other bone diseases cause deterioration in the microstructure of the trabecular bone and lead to a reduction in bone mass [21], which is consistent with an overall reduction in strength and increase in the risk of bone fracture [22].

In addition to the inherent structural complexity of trabecular bone, the analysis of this tissue is further complicated by the uncertainty associated with its material properties, specifically the material modulus or Young's modulus. Although attempts have been made to determine the material modulus of trabecular bone, results vary significantly between studies, with reported moduli generally in the range of 1–20 GPa [1]. Detailed tables summarising the reported values for Young's modulus in the literature may be found in [1], [18], [23], [24] and [25].

More recent research has suggested that the Young's modulus varies within the specimen depending on the density of the bone at a particular point [1]. Many mathematical relationships based on the local density have been proposed and implemented as a way of predicting the Young's modulus [6], however it has been found that even the choice of density-modulus relationship is subject-specific [26]. Using these density-modulus relationships may lead to a significant variation in Young's modulus within the bone, with some specimens showing variations as large as 67 GPa [6].

The Poisson's ratio of trabecular bone is another material property with large uncertainty attached to it. The lack of consensus on the Young's modulus of trabecular bone means that experimental studies investigating the Poisson's ratio of the tissue are rare. In 2007, Hong, et al [27] reported experimentally determined Poisson's ratios of 0.199 (longitudinal), 0.152 (anterior-posterior) and 0.154 (latero-medial). Despite these reported experimental values, the

majority of subsequent numerical studies still implement an isotropic, linear elastic assumption when assigning material properties to the bone and, as a result, a Poisson's ratio of 0.3 is commonly used [6, 11, 15, 16, 26, 28, 29]. Although the linear elastic assumption is inaccurate, Yosibash, et al showed in two previous numerical studies [30, 31] that the influence of Poisson's ratio on the apparent mechanical properties is very small, which encourages continued use of this assumption.

Overall, trabecular bone research presents a unique challenge, with material properties which possibly vary across a specimen [1] and a structure which is irregular [7], inhomogeneous [7], anisotropic [18], time- [8, 18] and site-dependent [1]. As a result, the only way to conclusively determine the mechanical properties of trabecular bone is to analyse the microstructure of the bone at a particular anatomic site for an individual patient.

2.2 Image acquisition and segmentation

Micro-computed tomography (micro-CT) scanners have long been used to obtain images of trabecular bone microstructure. Presently, micro-CT scanning is a leading technique in non-destructive structural evaluation and can be used to visualise and quantify the three-dimensional structure of a bone sample [32].

2.2.1 Micro-CT scanning

Micro-CT scanners function similarly to normal CT scanners, but can produce images with much higher resolutions than a normal CT scanner. X-rays are passed through a structure onto an X-ray detector to produce a 2D projection of the structure. The sample is rotated slightly and the process is repeated to obtain another 2D projection. After the sample has been rotated 180° or 360°, the 2D projections are reconstructed into images of cross-sectional slices of the sample [33]. The stack of cross-sectional images which make up the structure are saved with a .tif extension, and hence are commonly referred to as a TIFF stack.

The TIFF stack produced by the micro-CT scanner consists of a number of grayscale images, where the grayscale value assigned to each voxel represents the relative linear attenuation of the corresponding region in the sample [34]. As such, the grayscale values of the voxels may be used to infer the density of the bone by calibrating the grayscale values of the sample with grayscale values produced by a tissue of known density [34].

The quality of the image obtained from a micro-CT scanner depends on a number of factors, most notably the voxel size of the scan. Scan resolutions are often defined in terms of voxel size in trabecular bone literature, where a resolution of 80 μm indicates that each voxel in the image represents an 80 $\mu m \times 80 \mu m \times 80 \mu m$ region. However, the true spatial resolution of the scan is not equivalent to the voxel size, as it is dependent on a number of other factors [13]. Consequently, it is more appropriate to refer directly to "voxel size" when quoting this parameter, rather than "resolution". Much work on trabecular bone has been performed with the aid of micro-CT scans with small voxel sizes (5 $\mu m - 30 \mu m$) [35–37], however the time taken to scan at these voxel sizes is high and processing these images is computationally expensive [13]. Consequently, some studies digitally coarsen the image after acquisition to increase voxel size and decrease processing time [37, 38].

There have been various studies which have investigated the influence of voxel size on the analysis of trabecular bone micro-CT scans. Due to differing experimental techniques and image processing procedures, there has been no consensus regarding the voxel size required to obtain an accurate representation of trabecular bone. Ulrich, et al [37] reported that a voxel size of $168 \mu m$ was small enough to provide useful information for some bone morphometries, whereas Ladd and Kinney [39] proposed that a voxel size of $20 \mu m$ is required to obtain an accurate result for the elastic modulus. Chen, et al [40] analysed specimens scanned at five voxel sizes in the range $19.5 \mu m - 312 \mu m$ and reported that for a linear simulation the third principal strain, third principal stress and displacement all showed negligible error for voxel sizes up to $78 \mu m$.

2.2.2 Segmentation

Segmentation is a vital step in the analysis of trabecular bone samples which involves dividing the micro-CT scans into bone and non-bone voxels. Generally, segmentation is performed by identifying and applying a threshold; if the grayscale value of a particular voxel falls above the threshold, the voxel is designated as bone, otherwise the voxel is designated as non-bone. The result of segmentation is a binary image, where bone voxels are assigned a value of 1 (white) and non-bone voxels are assigned a value of 0 (black). There are two types of thresholding, which are commonly used in literature [14]:

1. Global thresholding
2. Local thresholding

Global thresholding is the most commonly used segmentation technique [41]. A single grayscale value is chosen as the threshold and all the images in the TIFF stack are segmented based on this value. The threshold may be determined visually by an operator or through a histogram analysis of the image [14]. Automated global thresholding techniques, such as the Otsu method [42], are relatively popular as they do not require the input of an operator, thereby removing some of the uncertainty associated with operator subjectivity [43]. Additionally, these methods are efficient [43], which makes them ideal for processing large micro-CT datasets which can have sizes larger than $1 GB$ [44]. Despite the popularity of global thresholding, problems arising from the choice of a single threshold can significantly reduce the quality of the segmentation [14]. The combination of beam hardening³ and partial volume effects⁴ means that the optimal threshold can vary within a bone specimen. As such, global thresholding may result in the loss of thin trabeculae and oversizing of thicker trabeculae [14].

Local thresholding makes use of multiple locally-determined thresholds throughout the image such that each voxel is segmented based on its neighbourhood [14]. Although local thresholding is more sensitive to variation within a bone specimen [14], it requires the use of extensive computational resources which are not available on conventional desktop computers [41]. Additionally, local segmentation involves setting a number of parameters, which may make it difficult to distinguish changes caused by the segmentation from variations inherent to the sample [13],

³In the polychromatic X-ray beam used in micro-CT scans, the lower energy photons are absorbed more readily than the higher energy photons as they pass through a material. Consequently, the average energy of the X-ray beam increases as it passes through a specimen, resulting in a progressive increase in the attenuation and, consequently, grayscale value. Beam hardening causes artefacts on the scan, and erroneously increases the grayscale values of the edges of the specimen [13].

⁴Partial volume effects refer to voxels which are only partially filled by bone. These voxels have higher grayscale values than a non-bone voxel, but lower grayscale values than a bone voxel [14].

thereby skewing analyses around tissue density.

Numerous studies have investigated the effect of thresholding techniques on the results produced by trabecular bone analyses [14, 41, 43–45]. Waarsing, et al [14] investigated the difference between the performance of local and global segmentation methods at various resolutions and reported a number of interesting phenomena:

- There was no difference between the performance of the local and global thresholding methods when good-quality scans with small voxel sizes were used.
- As the voxel size of the scans increased, both the global and local methods began to overthicken the trabeculae and overestimate the volume of the scan.
- At a voxel size of $53 \mu m$, neither the global nor the local thresholding produced reliable results.
- Global segmentation resulted in the overthickening of the trabeculae on the outer edge of the specimen and underthickening of trabeculae inside the specimen.

Additional interesting points were raised in other studies. Hara, et al [45] found that bone specimens with bone volume fractions less than 0.15 showed a variation of 5% in bone volume fraction and 9% in stiffness when the threshold value was changed by only 0.5%. For specimens with volume fractions above 0.2, these variations decreased to 2% and 3%, respectively. Recently, Rovaris, et al [43] showed that there is no significant difference between the Otsu method and the subjective operator-based method at high resolutions. Finally, Parkinson, et al [41] showed that the Otsu method produced a segmented image which showed excellent agreement between bone volume fraction calculated from the images and the ashed samples.

Overall, it is clear that the choice of segmentation method forms an important aspect of numerical trabecular bone research, with required accuracy and available resources being the deciding factors. Regardless of the thresholding method used, image slices from the segmented scan must be compared to slices from the original scan to ensure accurate segmentation [13].

2.3 Microstructural indices

Microstructural indices, morphometric properties or architectural parameters are measures taken directly from a micro-CT scan which quantify the structure of a trabecular bone sample. It has been proposed that these indices can provide insight into the behaviour of trabecular bone [46] and even provide a basis for the prediction of the mechanical properties [47]. As such, it has been suggested that research involving trabecular bone should report microstructural indices as a way of describing the morphometry of the bone [13].

This section provides an overview of the types of microstructural indices, how they are calculated and the relevant research investigating the link between microstructural indices and the mechanical properties of trabecular bone.

2.3.1 Calculation of indices

The most commonly used indices in literature⁵ are: bone volume fraction (BV/TV), trabecular thickness (Tb.Th), trabecular separation (Tb.Sp), trabecular number (Tb.N), structure model index (SMI), degree of anisotropy (DA) and connectivity. All of these indices may be determined using BoneJ [53], which is free, open-source software created specifically for the image-based analysis of trabecular bone. BoneJ [53] is a plugin for Fiji [54], which is a distribution of the broader image analysis software ImageJ [55].

In the following paragraphs, a brief description of each of the indices is provided, followed by a short description of how the parameter is calculated. This is not an exhaustive list of the microstructural indices, but rather a description of the indices most commonly used in the literature.

Bone volume fraction is a measure of the structural density of the bone, i.e. how much space is occupied by bone instead of air in a particular region of interest (ROI). In order to calculate the bone volume fraction, a three-dimensional ROI is chosen within a segmented TIFF stack, and the number of bone voxels and non-bone voxels in that region are counted. The bone volume fraction is calculated as:

$$BV/TV = \frac{N_B}{N_T}, \quad (2.1)$$

where N_B is the number of bone voxels and N_T is the total number of voxels (bone and non-bone) in the region of interest.

As the name suggests, **trabecular thickness** is a measure of the average thickness of the trabeculae in a bone sample. The local thickness at a point is defined as the diameter of the largest sphere which fits inside the trabecula and contains the given point [56] as shown in Figure 2.4 [57]. The trabecular thickness is reported as the mean of the local thicknesses within the ROI. The **trabecular separation** is calculated in a similar manner, except the local measures are performed on the spaces between the trabeculae as opposed to within the trabeculae.

The **trabecular number** describes the number of trabeculae in the bone sample. There are several ways to calculate trabecular number; BoneJ [53] calculates this index through a function called *Analyse Skeleton*, which is run on a TIFF stack after it has been skeletonised using the *Skeletonise 3D* function. Each voxel in the skeletonised image is labelled as either an end-point, junction or slab voxel depending on if it has less than 2, more than 2 or exactly 2 neighbouring voxels, respectively. Figure 2.5 shows the voxel classification for a simple structure after skeletonisation. The number of trabeculae is counted by considering the slab segments which connect end-point to end-point, junction to end-point or junction to junction [58].

The **structure model index** [21] is intended to measure how rod-like or plate-like the trabecular structure is. To achieve this, the surface area of the structure, S , is calculated by fitting a mesh of triangles to the surface and adding all the areas of the triangles. The mesh is dilated by a small value, r , in the direction of the triangles' vertex normals and the new surface area is determined. The SMI is calculated as:

$$SMI = 6 \frac{\frac{S'}{S^2} V}{S^2}, \quad (2.2)$$

where S' is the change in the surface area and V is the bone volume. Values vary from 0 (plate-like) to 3 (cylindrical rods) and 4 (sphere-like).

⁵Observation based on data collated from [13, 19, 34, 46–52].




Figure 2.4: Maximal spheres method used for the calculation of local thickness. The displayed sphere is the largest sphere which fits within the trabecula and contains point p [57].




Figure 2.5: Classification of voxels into end-points (blue), junctions (purple) or slab voxels (orange) using the *Analyse Skeleton* function in BoneJ [53].

The directional dependence of the bone is captured in the **degree of anisotropy** index and is calculated using the mean intercept length (MIL) method [59]. A random point is chosen within the sample from which a large number of vectors of equal length are drawn. Each time a vector crosses between a bone voxel and a non-bone voxel, an intercept is counted for that vector. The MIL is defined as:

$$MIL = \frac{l}{N_i}, \quad (2.3)$$

where l is the length of the vector and N_i is the number of intercepts. The vector is multiplied by its corresponding MIL and a point cloud is constructed from the multitude of points. An ellipsoid is fitted to the point cloud and eigendecomposition [60] is used to find the lengths of the ellipsoid axes. Finally, the degree of anisotropy is calculated as:

$$DA = 1 - \frac{\lambda_{short}}{\lambda_{long}}, \quad (2.4)$$

where λ_{long} is the eigenvalue of the long axis and λ_{short} is the eigenvalue of the short axis. The structure is considered isotropic if the DA is 0 (i.e. the point cloud is a sphere) and is considered fully anisotropic if the DA is 1.

The **connectivity** is a measure of the number of connected structures in the bone network [61]. The Euler number [62] of each bone voxel is calculated by considering its immediate neighbourhood, i.e. the number of voxels which share a vertex or an edge with the target voxel. The Euler numbers of the individual voxels are added to obtain the Euler number of the bone sample. The contribution of the bone sample to the bone it was cut from is calculated by checking the intersections of voxels and stack edges. Finally, the connectivity is calculated as:

$$\beta_1 = 1 - \Delta X, \quad (2.5)$$

where ΔX is the contribution of the bone sample to the bone it was cut from. The connectivity density may be calculated as:

$$CD = \frac{\beta_1}{TV}, \quad (2.6)$$

where β_1 is the connectivity and TV is the volume of the sample.

2.3.2 Correlation to mechanical properties

Much research has focused on determining whether there is a link between the microstructural indices and the mechanical properties of bone. The purpose of this research is two-fold:

1. Correlations between microstructural indices and mechanical properties may shed light on the behaviour of the bone and provide a better understanding of the dependence of the behaviour on various structural traits.
2. If a correlation exists between the microstructural indices and the mechanical properties, there is a possibility that the indices can be used to predict the mechanical properties of the bone.

The latter of these is more aligned with the focus of this thesis, thus research relevant to this goal is presented in this section and summarised in Table 2.1.

In work on microstructural indices, Ding, et al [63] compared the bone volume fraction, structure model index, degree of anisotropy, connectivity density and trabecular thickness to experimentally determined apparent moduli for 160 human trabecular bone specimens. The connectivity density showed no correlation to the apparent modulus, however relatively weak correlations were seen between the apparent modulus and the other indices, with correlation coefficients ranging from $R^2 = 0.31$ – 0.45 . The strongest correlation stemmed from the structure model index and bone volume fraction, however it was noted that the correlation between the microstructural indices and the apparent modulus appeared to vary depending on the age of the bone specimens. For young individuals, connectivity provided the strongest correlation to the modulus, whereas trabecular thickness provided the strongest correlation in middle-aged individuals and bone volume fraction along with degree of anisotropy provided the strongest correlation for old individuals. In recent work, Moore [64] found the degree of anisotropy to be an inaccurate reflection of the mechanical anisotropy of a structure; this is discussed further in Section 2.6.

Subsequent to the study in [63], Salmon, et al [65] conducted work which showed that the structure model index (SMI) is not an accurate measure of plate- versus rod-like structure in trabecular bone. The calculation of SMI assumes that a structure is made up of entirely convex surfaces, and thus the change in surface area as shown in Equation 2.2 is always positive. Salmon, et al [65] showed that a significant portion of trabecular bone surfaces are concave, which results in a negative change in surface area. As a result, a positive SMI of 3 and negative SMI of -3 in a particular area may result in an overall SMI of 0, indicative of a plate-like structure when, in fact, this is not the case. Salmon, et al [65] conclude that SMI is not a suitable measure for trabecular bone and should not be used in future work.

Uchiyama, et al [47] experimentally determined the apparent modulus of trabecular bone specimens in different anatomical planes and compared the results to various microstructural indices. No correlation was found between the microstructural indices and the apparent modulus in the craniocaudal direction, which implies that any link between microstructural indices and mechanical properties is dependent on the loading direction.

Perilli, et al [66] investigated the correlation between microstructural indices and experimentally determined ultimate stress of trabecular bone specimens. It was reported that the ultimate stress showed an excellent correlation to the bone volume fraction ($R^2 = 0.84$), a good correlation to trabecular separation ($R^2 = 0.77$) and a weaker correlation to both trabecular thickness ($R^2 = 0.60$) and trabecular number ($R^2 = 0.66$).

Cui, et al [67] used Finite Element Analysis to numerically determine the apparent moduli of 15 trabecular bone specimens and compare them to various microstructural indices. The apparent modulus showed a good correlation to the bone volume fraction ($R^2 = 0.76$) and a weaker correlation to the trabecular separation ($R^2 = 0.42$), structure model index ($R^2 = 0.51$) and trabecular number ($R^2 = 0.39$). The study suggests that the best prediction of the modulus stems from a combination of the SMI, trabecular number and trabecular thickness, producing a correlation of $R^2 = 0.80$. However, considering the later findings by Salmon, et al [65], this conclusion may need to be reviewed. Ulrich, et al [49] similarly found that a combination of the bone volume fraction, trabecular separation and degree of anisotropy yielded the best correlation to numerically determined apparent moduli, showing a correlation coefficient of $R^2 = 0.92$.

Table 2.1: Summary of literature relating mechanical properties to microstructural indices.

Reference	Mechanical property	Microstructural indices	Conclusion
Ding, et al [63]	Experimental apparent modulus	<ul style="list-style-type: none"> • BV/TV • SMI • DA • Conn.D • Tb.Th 	Combination of SMI and BV/TV produces best correlation, but this varies with age.
Salmon, et al [65]	-	<ul style="list-style-type: none"> • SMI 	SMI is not an accurate measure of plate- versus rod-like structure in trabecular bone.
Uchiyama, et al [47]	Experimental apparent modulus	<ul style="list-style-type: none"> • BV/TV • Tb.N • Tb.Sp 	Link between microstructural indices and mechanical properties dependent on loading direction.
Perilli, et al [66]	Experimental ultimate stress	<ul style="list-style-type: none"> • BV/TV • Tb.Th • Tb.Sp • Tb.N 	Ultimate stress shows excellent correlation to BV/TV and good correlation to Tb.Sp.
Cui, et al [67]	Numerical apparent modulus	<ul style="list-style-type: none"> • BV/TV • SMI • Tb.Sp • Tb.N • DA • Tb.Th 	Apparent modulus shows good correlation to BV/TV. Best correlation to apparent modulus occurs when a combination of SMI, Tb.Th and Tb.N is used.
Ulrich, et al [49]	Numerical apparent modulus	<ul style="list-style-type: none"> • BV/TV • BS/TV • Tb.Th • Tb.Sp • Tb.N • DA 	Best correlation obtained from combination of BV/TV, Tb.Sp and DA.

Overall, there appears to be some dissent in the literature regarding which index, or combination of indices, provides the best correlation to the mechanical properties. Bone volume fraction consistently shows a correlation to the apparent modulus, however the strength of that correlation is variable. No other index consistently showed a good correlation to the mechanical properties of the bone.

2.4 Cellular solids

Trabecular bone is one of many natural materials which is classified as a cellular solid. As such, it is useful to investigate properties and behaviour of cellular solids in order to better understand the properties and behaviour of trabecular bone. This section briefly discusses the different types of cellular solids, the mechanical properties of these cellular solids and the mechanisms which govern their behaviour.

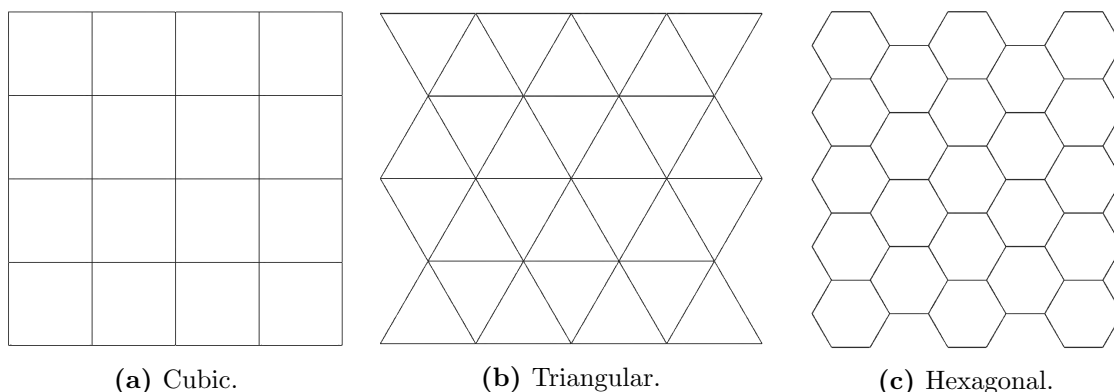


Figure 2.6: Examples of honeycombs: (a) cubic honeycomb, (b) triangular honeycomb and (c) hexagonal honeycomb [68].




Figure 2.7: Examples of an open-cell polyurethane foam and a closed-cell polyethylene foam [18, 69].

2.4.1 Types of cellular solids

Cellular solids are two- or three-dimensional structures composed of an interconnected network of struts and/or plates which form the edges and faces of cells [18]. Two-dimensional cellular solids, referred to as honeycombs, include cubic, triangular and hexagonal cellular structures, as shown in Figure 2.6. Three-dimensional cellular solids, referred to as lattices or foams, are further divided into two groups, namely open-cell and closed-cell foams. The edges of the cells in open-cell foams are composed of a series of struts, which join at vertices, whereas in closed-cell foams, the cell is constructed from solid plates which close off the faces of the cell. Examples of open- and closed-cell foams may be found in Figure 2.7.

A significant amount of research has been performed around creating and manufacturing lattices which are stiff and strong, yet lightweight. Three particular lattices are relevant to the current work, namely the cubic lattice [18], Kelvin cell lattice (or tetrakaidecahedron) [70] and octet truss lattice [71], the unit cells for which are shown in Figure 2.8. The cubic lattice is composed of 6 rectangular faces, whereas the Kelvin cell lattice comprises 14 faces, of which 6 are rectangular and 8 are hexagonal. The octet truss lattice has the most complex unit cell, which is composed of a central octahedral cell which is bounded by 8 edge tetrahedrons as shown in Figure 2.9 [72]. The properties of these lattices are discussed in more detail in Section 4.1.1.




Figure 2.8: Unit cells of the three lattice types used in the current work.




Figure 2.9: Unit cell for octet truss lattice. The central octahedral cell is highlighted in red, with the bounding tetrahedrons shown in grey [72].




Figure 2.10: Examples of rod-like and plate-like trabecular bone. In humans, rod-like trabeculae dominate in the lumbar spine as shown in (a), whereas plate-like trabeculae dominate in the femoral head, as shown in (b) [46].

Although trabecular bone is classified as a three-dimensional cellular solid, it cannot be neatly categorised as either an open- or closed-cell foam due to its inherent plate-rod structure. In areas where the trabeculae are predominantly rod-like as shown in Figure 2.10a, the bone may be viewed as an open-cell foam, with struts making up the edges of the cells. However, in regions where the trabeculae are more plate-like as shown in Figure 2.10b, the structure is more reminiscent of a closed-cell foam. Regardless of the categorization, understanding the behaviour of other three-dimensional cellular solids, either open- or closed-cell foams, can only improve understanding of the behaviour of trabecular bone.

2.4.2 Determining mechanical properties

The behaviour of any cellular solid is governed by two factors, namely the structure of the cellular solid and the material from which the structure is made. Cellular solids have local material properties associated with the material of which they are composed, but also have global, macroscopic properties which are a function of the structure or physical architecture of the cellular solid as well as the material⁶. As a result, the mechanical properties of cellular solids depend heavily on the relative density, which takes both the material and structural properties into account. For cellular solids, the relative density is generally lower than 0.3 and is calculated as:

$$\frac{\rho^*}{\rho_s}, \quad (2.7)$$

where ρ^* is the density of the structure and ρ_s is the density of the material of which the structure is made [18]. The apparent moduli of open-cell foams and closed-cell foams have been shown to correlate to the relative density [18].

Although cell size has some effect on the properties of the cellular solid, the mechanical properties of the cellular solid depend more strongly on the cell shape. If cells are equiaxed, then the properties of the cellular solid are isotropic, whereas if the cells are elongated, the properties

⁶More information on this distinction may be found in Section 2.5.3




Figure 2.11: Unit cells of lattices used as idealised models for trabecular bone in literature.


can show significant directional dependence [18].

As mentioned in Section 2.4.1, many types of cellular solids exist, which span a range of structural complexity. Because the material and structural properties of trabecular bone are complex and variable, attempts have been made to predict the mechanical properties of trabecular bone by making use of simpler cellular solids, generally lattices. These lattices can be manufactured from a known material, which removes the uncertainty inherent in the Young's modulus of trabecular bone, and can be manufactured to a similar relative density as the bone. Additionally, the structure of lattices is much more regular than trabecular bone, which makes it easier to determine the mechanical properties. Goda, et al [73] developed a 3D continuum model of vertebral trabecular bone based on an idealised cellular solid with a repeating unit cell as shown in Figure 2.11a. Diamant, et al [74] built cubic lattices with patient-specific values of trabecular thickness, trabecular separation and bone volume fraction in order to predict the elastic modulus and yield strength of canine vertebral trabecular bone specimens. Yeh and Keaveny [75] used an irregular cubic lattice, the unit cell of which is shown in Figure 2.11b, in their work on the effect of a variation of trabecular thickness on the elastic modulus of the bone.


Although simulating idealised models can never provide a prediction as accurate as the simulation of a sample of trabecular bone, the connection between trabecular bone and cellular solids can provide important information regarding the general behaviour of the bone.

2.5 Finite Element Analysis of trabecular bone

Finite Element Analysis (FEA) has long been used in various streams of trabecular bone research due to its non-destructive, cost-saving nature and repeatability as mentioned in Section 1.1. The following section focuses on the modelling methodologies, material properties and boundary conditions relevant to the current work and briefly discusses the previous work performed in this area.



(a) VBM with hexahedral elements [77].



(b) VBM with tetrahedral elements [78].

Figure 2.12: Voxel-based Finite Element meshing of a trabecular bone sample using (a) hexahedral elements, and (b) tetrahedral elements.

2.5.1 Modelling methodologies

The majority of trabecular bone research is based on micro-CT scans of the microstructure of the bone. The most important function of the chosen modelling methodology is thus to accurately and consistently discretise the 3D image into a FE mesh which is representative of the given microstructure. There are two major discretisation methodologies in use, namely the voxel-based method (VBM) and the beam-shell method (BSM).

2.5.1.1 Voxel-based method

The most widely-used discretisation method for FEA of trabecular bone is the voxel-based method [10]. In the VBM, each voxel of the scan that falls beyond a specified grayscale threshold is converted into an eight-node isoparametric hexahedron element, more commonly known as an eight-node “brick” [76] as shown in Figure 2.12a. The VBM is currently the preferred approach, because it creates the FE mesh directly from the CT scan without the need for geometrical interpretation [76]. Additionally, it allows for the density and, hence, a measure of the Young’s modulus to be determined for each element by considering the grayscale value of the corresponding voxel and applying density-modulus relationships⁷ [6].

Some variations on the classic VBM have been investigated. In 1995, Müller and Rügsegger [78] proposed a method based on the marching cubes algorithm to create a mesh with linear tetrahedrons as opposed to hexahedrons, as shown in Figure 2.12b. Although this method yielded smoother surfaces than the original VBM, the use of tetrahedrons required more computing time and storage capacity while yielding no distinct advantage over the accuracy of the original VBM [76]. Some studies have made use of the tetrahedral elements in subsequent years [37, 79, 80], however the overwhelming focus of current research is on the standard hexahedral VBM [9, 15, 38, 81–84]. Because previous work using the VBM showed excellent correlation to experimental results [35, 85, 86], the VBM is considered to be the computational “gold standard” for FEA of trabecular bone [87].

While the VBM may be the most widely used discretisation method, it creates FE meshes which contain hundreds of thousands of nodes and elements, leading to simulations which are extremely computationally expensive [88]. As a result, only small domains can be simulated using this approach. Additionally, by converting each voxel of the scan into an element, the structure-mechanics relationship of the trabecular bone micro-architecture is lost [2], meaning

⁷See Section 2.5.2.

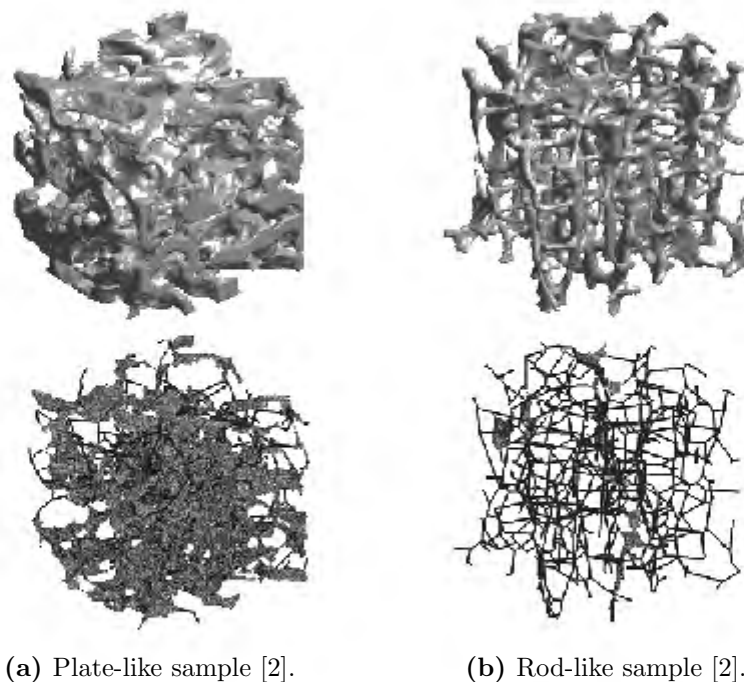


Figure 2.13: Finite Element mesh using the beam-shell method developed by Vanderroost, et al. [2] for a (a) plate-like sample, and (b) rod-like sample. The thickness of the beams and shells is not shown in these images.

that individual contributions to stiffness by the rods and plates in the microstructure cannot be easily assessed.

2.5.1.2 Beam-shell method

More recently, a new method [2, 20], has been proposed which seeks to address the resource-intensive nature of the VBM. This method takes advantage of the inherent plate-rod geometry of trabecular bone to discretise the structure into a series of plates and rods which are allocated as beams and shells for FEA as shown in Figure 2.13. Vanderroost, et al. [2] developed an optimised skeletonisation and meshing algorithm to achieve the beam-shell discretisation, whereas Liu, et al. [20] independently achieved beam-shell discretisation via complete volumetric decomposition using an arc-skeletonisation process. Although the methodology is different, both approaches have shown excellent results when compared to VBM simulations and have been shown to vastly reduce the processing time required for FE simulation [2, 20, 87].

The methodology developed by Vanderroost, et al [2] is used extensively in the current work, therefore a more in-depth description of the process is required. Figure 2.14 gives a visual representation of the various steps employed in the discretisation process, which is described in the list on the next page [2].

1. 2D skeletonisation (Figure 2.14b) is performed on the 6 boundary planes (Figure 2.14a) of the sample and the skeletonised voxels are stored.
2. 3D skeletonisation is performed on the sample (Figure 2.14c) using the shape-preserving algorithm developed by Saha, et al [89,90]. The 2D skeletonised voxels produced by step 1 are protected in step 2 so that the boundaries of the specimen are preserved. The output of this step is a fully skeletonised image.
3. The voxels in the skeletonised image are classified as rod-like or plate-like using an extension of the work by Saha, et al [90].
4. Individual trabeculae are identified by identifying rod-rod, plate-plate and plate-rod junctions (Figure 2.14d). The voxels between junctions or between junctions and end-points are grouped together and classified as an individual trabecula⁸.
5. A thickening algorithm based on the inverse skeletonisation algorithm developed by Pothuaud, et al [91] was used to attribute the voxels in the original 3D volume to individual trabeculae as defined in step 4.
6. Nodes are defined at the intersection of the rod-like trabeculae. If the node is near the boundary of the sample, the node is positioned on the boundary plane.
7. Elements are defined based on the nodes and the classification of rod-like and plate-like trabeculae.
 - a) Rod-like trabeculae are specified as Timoshenko beam elements. The cross-sectional area and moments of inertia of each beam are determined by considering slices through the trabecula perpendicular to the beam axis. The cross-sectional area is determined by counting the voxels in the cross-sectional slice, whereas the moments of inertia are determined relative to two perpendicular vectors, which are output along with the values for the moment of inertia.
 - b) Plate-like trabeculae are specified as 4-node⁹ thick shell elements. The thickness of the trabecula at each node of the shell element is determined using the trabecular thickness algorithm [57] discussed in Section 2.3.1.
8. Tie elements are added between the end-nodes of beam elements and several nodes of the neighbouring shell elements (Figure 2.14f). This is done to provide a more realistic load distribution between the beams and the shells. Shell elements which are within the radius of the intersecting beam are tied to the intersecting beam node.

Vanderoost, et al [2] report that an initial testing phase of the beam-shell method was conducted on artificial structures, but the data from this analysis is not included in the results. Further analysis was conducted using bone specimens, and these results are available in [2]. 257 human trabecular bone samples were simulated using the BSM and the results were compared to the “gold-standard” VBM results to determine the accuracy of the BSM. The samples were $4\text{ mm} \times 4\text{ mm} \times 4\text{ mm}$ in dimension and were scanned at a voxel size of $30\ \mu\text{m}$. In the initial testing phase, it was discovered that the BSM slightly underpredicts the apparent modulus relative to the VBM and thus a correction factor of 1.605 was applied to the BSM results. It is stated that this correction factor is “uniformly applicable for all types of trabecular bone structures” [2]. Consequently, a back-calculation approach (see Section 2.5.2) was used and a

⁸This is similar to the way the trabecular number is determined using the *Analyse Skeleton* function in Fiji. See Section 2.3.1 for more information.

⁹For complex shells, the four-node element was collapsed into a three-node triangular element [2].




Figure 2.14: Beam-shell discretisation process developed by Vanderroost, et al [2]: (a) 3D rendering of micro-CT scanned trabecular bone sample, (b) 2D skeletonisation on the boundaries of the sample, (c) 3D skeletonisation and identification of trabecular rods and plates, (d) identification of trabecular rods and plates in 3D volume, (e) Finite Element beam-shell mesh with colour map indicating element thickness, (f) tie elements employed between beam and shell elements.

Young's modulus of 10 *GPa* was assigned to the VBM simulations, whereas a Young's modulus of 16.05 *GPa* was assigned to the BSM simulations.

A strong correlation was reported between the apparent moduli generated by the BSM and the VBM ($R^2 = 0.97$) [2]. Most BSM values are within 5% of the VBM values, with a maximum error of 11% occurring at low apparent moduli. It was found that the BSM vastly reduced the processing time required for FE simulations, because it produced meshes which had significantly fewer nodes and elements. Vanderoost, et al. [2] found that the beam-shell method caused, on average, a 7-fold decrease in number of elements and 33-fold reduction in the CPU time for a linear FE simulation.

Although the BSM study by Vanderoost, et al [2] shows promising results, the research undertaken was entirely numerical, and no correlation to experimental results is presented. This correlation is essential in providing confidence in the beam-shell methodology.


2.5.2 Approximating Young's modulus

As discussed in Section 2.1, there is some debate about the Young's modulus of trabecular bone, which makes it difficult to assign a value for Young's modulus (and other mechanical properties, like Poisson's ratio) to a Finite Element simulation of trabecular bone. To combat this uncertainty, two approaches are commonly used, which are briefly described in this section.


The first approach is referred to as "back-calculation" [9, 24, 36, 39] and is demonstrated in Figure 2.15. In this method, an arbitrary Young's modulus is assigned to the bone specimens for simulation. The results from the simulations are then compared to experimental results as shown in Figure 2.15a and a scaling factor is determined which scales the simulation results in such a way that they provide a 1:1 correlation with the experimental results as shown in Figure 2.15b. This scaling factor is then applied to the arbitrary Young's modulus initially assigned to the simulations to adjust the results (Figure 2.15b), and this back-calculated modulus is quoted as the Young's modulus of the specimens.

In order for the back-calculation method to be effective, the simulation must be an excellent computational replica of the experimental test set-up, i.e. the mesh must be representative of the structure and the boundary conditions must be well-replicated [92]. Any variation in these factors will influence the moduli produced by the simulations and result in the incorrect value being back-calculated for Young's modulus [93]. This is well-demonstrated by the work of Ladd, et al [39], which reports low back-calculated Young's moduli (compared to experimentally determined values) due to the influence of the end-artifact in the platen-based compression testing setup [24].

The second approach commonly used in the literature, termed the "density-modulus approach" in this work, assigns a unique Young's modulus to every element in the FE mesh based on the grayscale value of the equivalent voxels in the micro-CT scan [6, 26]. The grayscale value of a voxel may be used to determine the density of the bone at that point, and many mathematical relationships have been proposed to relate the density to Young's modulus. This approach allows for variation of the Young's modulus within an individual bone sample, which has been reported to increase the accuracy of the simulation [6, 26].



(a) Generated data.



(b) Adjusted data.

Figure 2.15: Back-calculation of Young's modulus. Generated data in (a) is used to back-calculate a Young's modulus. The generated data is adjusted according to this back-calculated modulus to obtain a 1 : 1 correlation as shown in (b).

The density-modulus approach has two main drawbacks: sensitivity to partial volume effects [14] and choice of density-modulus relationship [26]. There are two possible reasons for a grey voxel (i.e. a voxel with a lower grayscale value than other bone regions) in a micro-CT scan. The density-modulus approach assumes that the bone is less dense in this area, which is valid for some regions of the image. However, a grey voxel could also stem from a voxel which is only partially filled with bone (hence, partial volume effects) [14]. In these cases, the density-modulus relationship is falsely applied, which will negatively influence the results. Secondly, it has been suggested that the best density-modulus relationship to use for trabecular bone varies based on the sample. As such, one relationship is not universally applicable [26] and an incorrect choice may lead to erroneous results.

Both of the methods currently being used to assign Young's modulus to FE simulations have serious limitations and care should be taken when using either approach in trabecular bone simulations.

2.5.3 Determining effective properties

In the analysis of any cellular solid, it is important to understand the responses given by the structure at different scales. For a cellular solid, the mechanical properties can be classified into three main categories: material properties, apparent properties and effective properties. These properties are discussed here in reference to trabecular bone.

Material properties are the mechanical properties (stiffness, strength, etc.) inherent in the material forming each trabecula [18]. The material properties do not depend on the structure and are thus generally considered constant once they have been determined for a given "batch" of material. Although there is some uncertainty regarding the variability of the material properties of trabecular bone [1], Young's modulus is typically assumed to be constant in FEA of trabecular bone to simplify the simulation.

The **apparent properties** of a sample take into account the material properties and the structural properties [94]. Although the material has an inherent Young's modulus, the structure (consisting of a particular arrangement of edges, vertices and faces) also has a modulus [18]. The combined effect of these two moduli gives the apparent modulus of the sample.

The **effective properties** refers to the mechanical properties of an infinite structure, as opposed to a small sub-domain of the structure. As such, the effective properties are free from boundary effects. Usually, these properties are determined via a continuum mechanics approach [84], where a small sub-domain of the structure called a representative volume element (RVE) is chosen for analysis [95]. The RVE is used to link the heterogeneous structural characteristics of the material at the microscale to the effective properties and is useful because it allows smaller domains to be considered, thereby reducing computational time [95]. This continuum mechanics approach has been widely investigated and extensively used in the study of other cellular solids and composites [18,96–100]. There have also been attempts to use the RVE technique in trabecular bone research [18,84,101], however the inherent properties of trabecular bone make it complicated to define an RVE. As determined by Ostoja-Starzewski [102], the continuum mechanics technique may only be used for homogeneous materials, where the RVE is defined as:

1. A sample which is so large that the boundary effects are rendered negligible, or
2. A unit cell which is representative of the structure, where the structure is periodic.

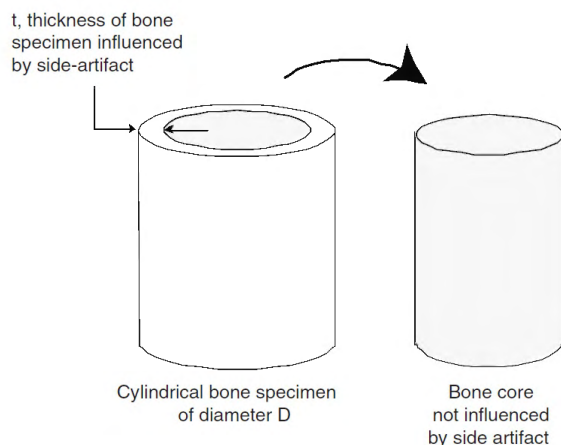


Figure 2.16: Inner core or embedding technique suggested by Ün, et al. [83] and used by Harrison and McHugh [38]. Figure taken from Ün, et al. [83].

Trabecular bone is inhomogeneous [7], and although it shows some periodicity, the variation in structure (even within the same region), makes the unit cell approach extremely difficult to apply. Moreover, because of the location and size of trabecular bone regions in the body, it is often impossible to extract a sample of bone which is large enough to approximate an infinite region [15]. As such, trabecular bone does not meet any of the requirements for an RVE. This means that any sample of trabecular bone which is tested, will be heavily influenced by the applied boundary conditions [84] and, hence, will not yield the effective properties, but rather the apparent properties [94].

Many attempts have been made to calculate the effective properties of trabecular bone through the use of artificially created unit cells [84, 101, 103] or application of various boundary conditions [38, 77, 84], however the recent work by Daszkiewicz, et al. [15], which focuses on minimising the error associated with the so-called “side-artifact” shows the most promising results to date. In 2006, Ün, et al. [83] noted that machining a trabecular bone specimen results in hanging trabeculae around the sides of the specimen, termed “side-artifacts”. Because these trabecular connections are cut off, the trabeculae lose their axial load-bearing capacity, which can cause errors of over 50% in the measured apparent modulus [83]. Harrison and McHugh [38] further investigated the error associated with the side-artifact by using an inner core technique suggested by Ün, et al. [83], where a larger bone sample is used for the simulation, but results are only considered for the inner core of the sample as shown in Figure 2.16. In essence, the trabecular bone sample in question is “embedded” within a larger trabecular region, which allows the loads and boundary conditions to be transferred to the core in a more natural way [15].

Daszkiewicz, et al. [15] used a variation on the embedding approach to determine the apparent elastic properties of a trabecular sample and, later, to calculate the effective properties via a continuum mechanics approach. Trabecular bone may be considered as a linear elastic, orthotropic material [104, 105], meaning that its behaviour may be captured in a second-rank compliance tensor, consisting of 9 independent variables [106]. In [15], an RVE embedded in regions of increasing size is simulated using both kinematic uniform boundary conditions (KUBCs) and periodicity-compatible mixed uniform boundary conditions (PMUBCs). The values for the compliance tensor are calculated from the FEA results and the values for the “in situ” tensor are chosen to be the values that the KUBC and PMUBC values tend to when the thickness of the outer layer is increased as shown in Figure 2.17. It was found that the embedded approach produces effective moduli which are close to computationally determined “in situ” values, although they consistently underestimate the “in situ” values [15].

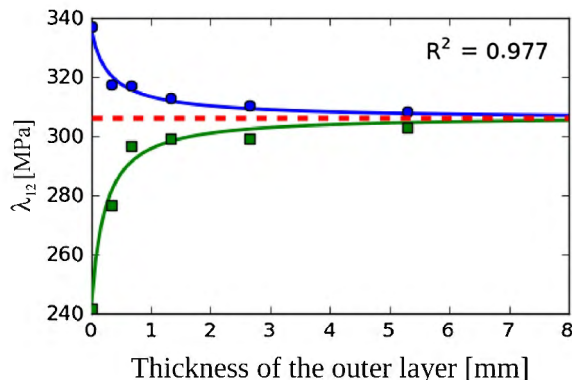


Figure 2.17: An example of the calculation of an “in situ” tensor value from the results of the PMUBCs and KUBCs simulations [15]. In the figure, the blue circles represent the results from the KUBCs simulations, the green squares represent the results from the PMUBCs simulations and the dotted red line represents the chosen “in situ” value.

Although the work by Daszkiewicz, et al. [15] tackles important questions regarding embedded RVEs and boundary conditions, it is entirely computationally based, therefore the correlation to experimental values has yet to be investigated.

2.6 Past research at BISRU

The Blast Impact and Survivability Research Unit (BISRU) has increased its activity in the study of biomechanics in recent years, with several students undertaking research in this field. The aim of the past research was to build a solid understanding of trabecular bone and the techniques used to determine its mechanical properties so as to lay the groundwork for an extensive, numerical study of this tissue. Foundational knowledge was gained through two means:

1. Performing thorough numerical investigations into simpler cellular solids like honeycombs and lattices to investigate the variation in mechanical properties of cellular solids and the numerical techniques that can be used to determine these properties.
2. Obtaining, machining, micro-CT scanning and compression testing a large number of bovine bone specimens to provide a basis for comparison for any future numerical work.

The recent work of three students is of importance to the current work and is briefly discussed in this section. Although the past work has yet to be published, it is important to note that the work by Lawrence [68] was reviewed by an internal examiner and the dissertations by Moore [64] and Hilton [3] have each been reviewed by two international examiners.

2.6.1 Cellular solids

In 2016, Lawrence [68] investigated the apparent properties of two-dimensional cellular solids, termed honeycombs. The aim of this research was to gain a better understanding of basic cellular solids theory with the aim of extending the acquired skills and processes to trabecular bone.

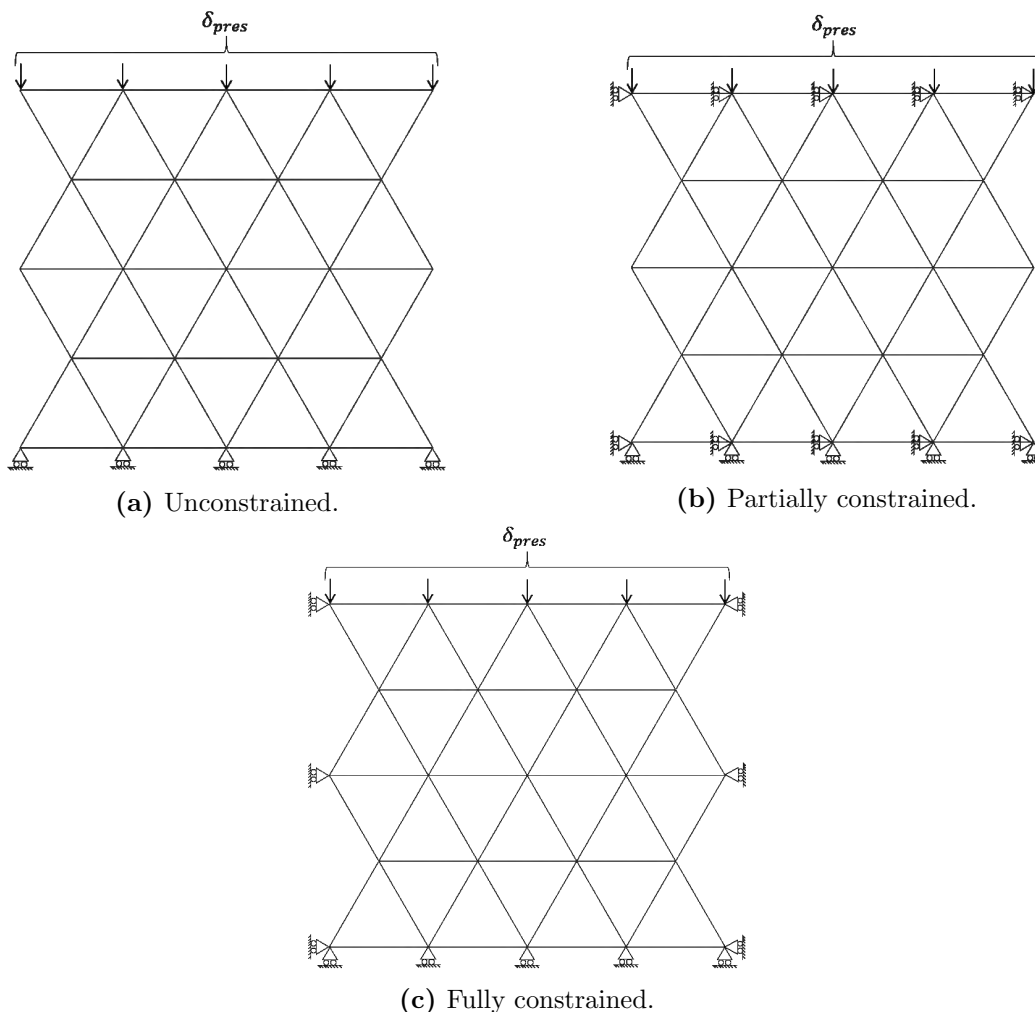


Figure 2.18: Boundary conditions used in the work on honeycombs [68], consisting of (a) unconstrained, (b) partially constrained and (c) fully constrained boundary conditions.

In the work by Lawrence [68], three different types of honeycombs, namely the cubic, triangular and hexagonal honeycombs, were analysed at different loading angles and three different levels of constraint as shown in Figure 2.18.

Through the use of analytical solutions and Finite Element simulations, it was demonstrated that the honeycombs show significantly different responses when subject to constraint at different loading angles. Figure 2.19 shows the results of the cubic simulations which demonstrated the largest variation in apparent stiffness due to confinement.

Following on from the work on honeycombs, Moore [64] investigated the effect of constraint on three-dimensional cellular solids at varying angles. Custom code was written in Python to construct cubic, Kelvin cell and octet truss lattices at the required angle and Finite Element simulations were performed to determine the apparent stiffness of the structures. Similar to the honeycombs, the lattices showed significantly different responses when subject to constraint, as demonstrated in Figure 2.20.

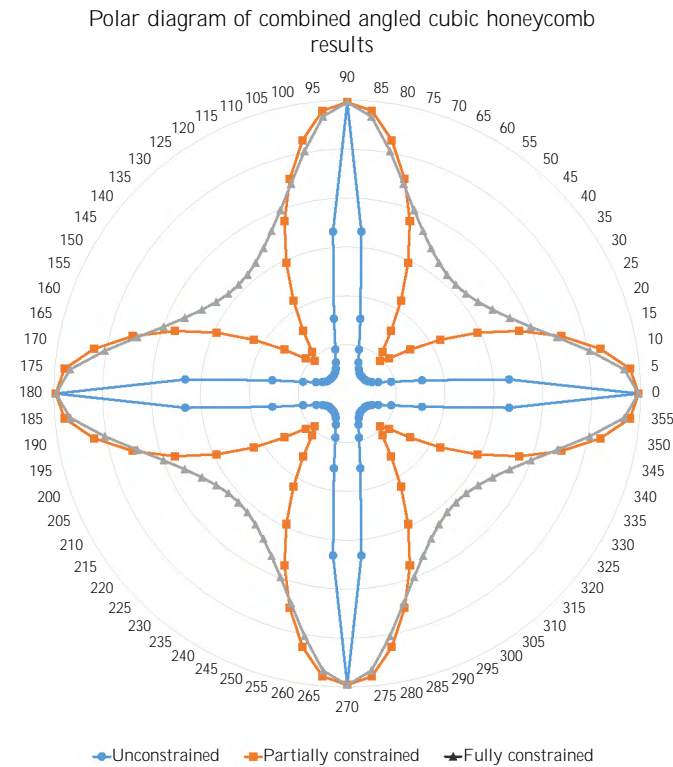


Figure 2.19: Polar diagram showing the effect of confinement on angled cubic honeycombs [68].

In addition, Moore [64] investigated the correlation between geometric anisotropy and mechanical anisotropy of the cubic, Kelvin cell and octet truss lattices. Geometric anisotropy was determined via a microstructural index, the degree of anisotropy (DA), whereas mechanical anisotropy was determined by Finite Element Analysis of the computationally-generated lattices at different angles. It was found that the mechanical anisotropy was not accurately reflected in the DA, with the Kelvin cell lattice showing the highest degree of geometric anisotropy of the three lattices despite its isotropic behaviour. This finding illustrates the need for validation studies in any work involving trabecular bone, and highlights how lattices may be used to obtain valuable information in such a study.

2.6.2 Trabecular bone

In order to gain a greater understanding of the effect of confinement and marrow on the apparent stiffness of trabecular bone, Hilton [3] experimentally tested 127 trabecular bone samples in various conditions. In this work, trabecular bone samples were machined from nine bovine femurs, with sizes as specified:

- $110 \times \phi 10$ mm specimens
- $9 \times \phi 20$ mm specimens
- $8 \times \phi 28$ mm specimens

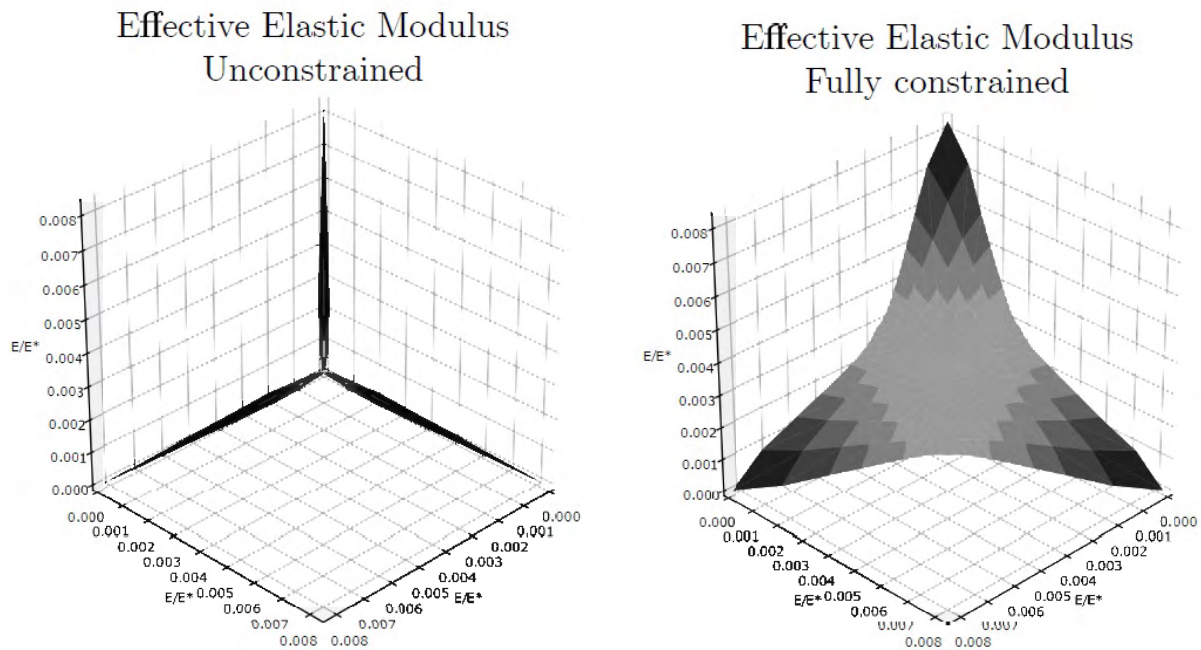


Figure 2.20: Polar diagrams showing the effect of confinement on angled cubic lattices [64].

Each specimen was micro-CT scanned (Phoenix, voxel size $80 \mu m$, $70 kV$, $350 \mu A$, acquisition time of $500 ms$ per image) before being subjected to a quasi-static compression test using a ZWICK 1481 Universal Tester. The standard $\phi 10 mm$ specimens were divided into three groups:

1. **Defatted** specimens, which were defatted (i.e. marrow removed) prior to scanning and testing.
2. **Marrow** specimens, which were not defatted.
3. **Confined** specimens, which were not defatted and were placed in an aluminium confinement collar during testing.

Additionally, the effect of embedded confinement as shown in Figure 2.16 was investigated using the $\phi 20 mm$ and $\phi 28 mm$ bone specimens. For these specimens, a custom testing platform (shown in Figure 2.21) was manufactured which was capable of monitoring the force response for the inner $\phi 10 mm$ bone core, while the ZWICK tester monitored the force response of the specimen as a whole. Using this platform, the apparent moduli of the $\phi 10 mm$ bone core and full bone sample were determined.

The results from the standard specimen tests showed that the defatted specimens were slightly less stiff than the marrow-inclusive specimens (i.e. the marrow and confined specimens), with the marrow-inclusive specimens achieving apparent moduli which were 9%–52% stiffer than the defatted specimens. Additionally, it was noted that the aluminium collar used in the confined specimen tests did not achieve the lateral confinement that was required, as the diameter of the bone specimens was slightly variable, creating a gap between the bone and the collar in some specimens.




Figure 2.21: Custom testing platform developed by Hilton [3] used to determine the inner modulus of bone specimens. The strain gauge on the inner cylinder monitors the response of the inner $\phi 10$ mm bone core, while the overall response is monitored by the ZWICK tester.

For the large specimen tests, Hilton [3] observed that the apparent modulus of the bone core was, on average, 59% higher than the apparent modulus of the sample as a whole, which is consistent with the findings of Ün, et al. [83]. It was concluded that the embedding technique successfully reduces the edge effects of the lateral boundary conditions on trabecular bone specimens, allowing for a more accurate representation of in vivo conditions.

2.7 Synthesis

The beam-shell method (BSM) recently proposed for use in Finite Element Analysis of trabecular bone presents a unique opportunity to perform faster, more efficient FE simulations and easily assess the contribution of the structural features in the bone to its mechanical properties. Despite the potential of the method, some areas for improvement and expansion were identified in the literature and are addressed in the current work.

Validation study

Due to the complex structure of trabecular bone, it is exceptionally important that any FE methodology applied to the bone structure be thoroughly validated using simpler structures with known mechanical responses in order to properly assess the accuracy and efficiency of the method. Although a validation study is mentioned in [2], neither the methodology of this validation study, nor the results are presented in the paper. A validation study of an alternative BSM was performed and reported in [88]. This study considered one rod-rod structure (a cubic lattice) and one plate-rod structure (a cubic lattice with plates on two faces) at varying thicknesses and voxel sizes. A more varied range of structures would have been beneficial in this work as the cubic lattice is a highly idealised structure with mechanical properties which

are solely dependent on axial deformation, unlike trabecular bone.

In the current work, the identified limitations are addressed by performing a validation study which uses three different types of cellular solids, thereby allowing a greater range of behaviour, such as directional dependence or sensitivity to constraint, to be investigated. Additionally, the size, strut thickness and image size of the cellular solids are varied to further investigate and quantify the accuracy and efficiency of the BSM.

Experimental results

In their work using an alternative BSM, Wang, et al [87] compared the BSM results to VBM results and experimental results, however the BSM results produced by Vanderroost, et al [2] were only compared to the VBM results generated for the same trabecular bone specimens - no experimental results were presented. Despite the well-reported correlation of the VBM to experimental results, it remains important for numerical results to be compared to experimental results to provide confidence in the numerical methodology.

In the current work, the BSM and VBM results generated for the scanned bone specimens are compared to experimental results generated for the same specimens. This allows for an investigation into how well the BSM correlates to the experimental work, how well the VBM correlates to the experimental work and how well the BSM and VBM relate to each other, providing a more cohesive analysis of the results.

Effective properties

The majority of FE research around trabecular bone places focus on the apparent properties, rather than the effective properties of the bone. However, determining the effective properties is the first step towards creating an accurate material model for trabecular bone, which could be used to significantly decrease simulation complexity and computational time. Recent work has demonstrated that an embedding approach may be used to determine the effective properties of the bone, but these studies have relied solely on the VBM as a numerical methodology and have lacked experimental validation.

In the current work, a novel investigation into the effective properties of trabecular bone is presented, where experimental effective moduli are compared to numerical effective moduli generated by both the BSM and the VBM. This work seeks to answer the question: “Can a microstructural beam-shell model be used to accurately determine the effective properties of trabecular bone?”.

Chapter 3

Development of analysis cycle

The first objective of this work is to develop an automated analysis cycle which is capable of determining the apparent modulus of a specimen given micro-CT scans of its structure (see Section 1.3). The methodology is referred to as an “analysis cycle” because it imports an image of the structure, and returns an image of the structure (reconstructed from the FE mesh) along with the modulus results, thereby making the process cyclical in nature. The analysis cycle can be broken down into four main sections:

- **Image processing** of the micro-CT scan.
- **Discretisation** of the structure into a beam-shell mesh.
- **Simulation** using Finite Element Analysis.
- **Post-processing** of the results.

These sections may be expanded into ten distinct steps which are detailed below and summarised in Figure 3.1.

Image processing

1. The micro-CT scan (TIFF stack) of the structure is imported into Fiji.
2. The TIFF stack is cropped and segmented in Fiji.
3. The binary TIFF stack is converted into a 3D binary matrix using code written in Python employing the Python Imaging Library (PIL).

Discretisation

4. The 3D matrix is imported into MATLAB, where Finite Element mesh discretisation is performed using the beam-shell code obtained from and developed by Vanderroost, et al. [2].
5. The mesh is refined using custom MATLAB code so as to minimise memory usage and allow for more successful simulations to be performed.
6. Information about the mesh, loads and boundary conditions is written to a series of input decks for FEA.

Simulation

7. A quasi-static compression simulation of the sample is performed in LS-DYNA (via the command prompt) using the input decks described in 6.

Post-processing

8. Results from the simulation are used to calculate the apparent modulus of the sample by employing custom code written in MATLAB.
9. A 3D binary matrix is reconstructed from the node and element data defining the FE mesh using custom code written in MATLAB. This is an important validation step as it allows the user to visually assess whether the discretisation code is building an accurate representation of the structure.
10. The reconstructed 3D matrix is converted into a binary TIFF stack using custom code written in Python.

All the steps of the analysis cycle are performed automatically when the user runs a batch file, which can be written for any number of specimens using MATLAB code developed by the author. A discussion of how each step was achieved is presented in Sections 3.1-3.4 and the hardware and software used in the work is listed in Appendix A.

3.1 Image processing

A number of the steps in the analysis cycle may be classified as image processing, where the TIFF stack is manipulated in some way in order to facilitate further analysis. This includes the image processing performed in Fiji, as well as the conversion from a TIFF stack to a matrix. The processes and reasoning behind the processes are discussed in this section.

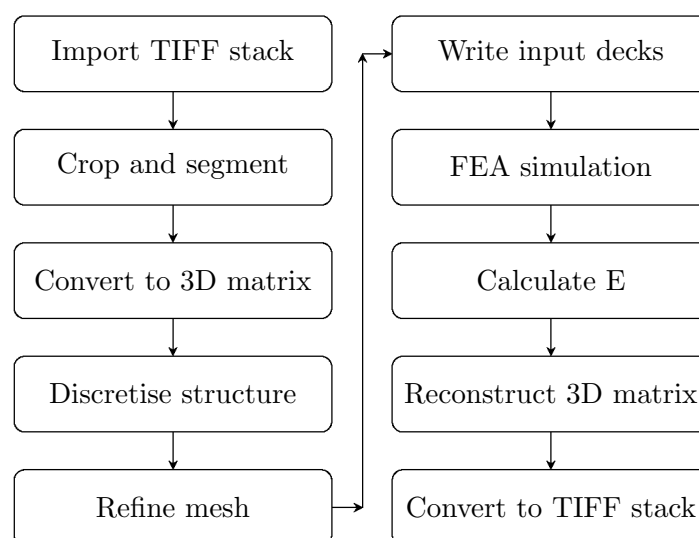


Figure 3.1: Flow chart of the developed analysis cycle.




Figure 3.2: Example of initial images from the bone scans.




Figure 3.3: Example of a bone substack including full bone cross sections, so as to provide “flat” top and bottom surfaces for 2D skeletonisation.

3.1.1 Specimen cropping

The bone scans obtained from Hilton [3] constitute a set of grayscale images representing consecutive slices of the bone. When the faces of the bone specimens are machined, it is impossible to acquire perfectly flat surfaces, therefore the top and bottom faces of the bone are not perfectly parallel to each other or to the micro-CT scanner, leading to initial and final image slices similar to Figure 3.2, where only partial cross-sections of the bone are identified. Because the discretisation code applies 2D skeletonisation on the top and bottom surfaces of the specimen, it is important that the first and last slice in the stack represent full cross-sections of bone (a “flat” face), unlike those shown in Figure 3.2. As such, the specimen needs to be cropped to exclude the partial bone scans.

Cropping was achieved by importing the TIFF stack into Fiji and creating a substack of the largest set of image slices containing full bone cross-sections as illustrated in Figure 3.3. The region to be cropped was determined visually for each specimen¹.

¹The variation in results caused by the subjective cropping process is likely to be very small compared to other factors such as the varying bone structure or image size, therefore this subjective process was considered acceptable for the purposes of the current work. A recommendation for future work using this analysis cycle may include the development of an image registration process to automatically detect the specimen boundaries and crop them appropriately.




Figure 3.4: Image processing steps applied in Fiji. The unedited image is shown in (a) and the two image processing steps are shown in (b) and (c).

3.1.2 Segmentation

The discretisation code [2] used in the analysis cycle requires a binary 3D matrix as input, therefore the TIFF stack needs to be segmented before being converted to a 3D matrix.

In order to achieve an accurate segmentation, the images are sharpened to provide clearer boundaries between the bone and the background (including marrow, air and polystyrene) as shown in Figure 3.4a and 3.4b. After sharpening, the *Make Binary* function in the BoneJ plugin in Fiji is used to segment the images using the Otsu method², so that the bone voxels are given a value of 1 (white) and all other voxels are allocated a value of 0 (black) as shown in Figure 3.4c.

In-house MATLAB code was developed which creates a Fiji macro for a given specimen. This macro allows the importing, cropping, sharpening, segmenting and saving processes to occur without additional user input.

3.1.3 Conversion to 3D matrix

The final step in the initial image processing is to convert the binary TIFF stack into a 3D binary matrix. Code was written in Python using the Python Imaging Library (PIL) which imports the cropped, segmented TIFF stack and converts it into a 3D binary matrix. Similar code was also written which converts the 3D binary matrix back into a TIFF stack; this code is used at the end of the analysis cycle so that the reconstructed images can be compared to the original images.

²Fiji provides many different segmentation options in its interface. The option chosen for this work is *Make Binary*, a global thresholding technique, which calculates a threshold for each image in the stack and individually segments each image based on its corresponding threshold. The *Optimise Threshold* function, another global thresholding technique, calculates the threshold value based on the entire bone volume and uses said threshold to segment all the images in the stack, while trying to reduce connectivity. Although the *Make Binary* option was used as the default segmentation method in the analysis cycle, both the *Make Binary* and *Optimise Threshold* methods were investigated in the current work, a decision which is discussed further in Section 5.2.2

3.2 Discretisation

In order to perform a Finite Element simulation, the imported 3D binary matrix needs to be discretised into a mesh. This work focuses on a microstructural discretisation (i.e. a mesh consisting of beams and shells) rather than a voxel-based discretisation (i.e. solid tetrahedral elements), however the voxel-based method (VBM) is used for comparison purposes in this work (see Sections 4.1.5 and 5.1). Although the beam-shell discretisation is performed using code developed by Vanderroost [2], edits and refinements to the code and the mesh were implemented by the author.

3.2.1 Mesh generation

Upon completion of the image processing, the specimen is imported into MATLAB and discretised into a Finite Element mesh. The discretisation code used in this work was obtained from Professor Harry van Lenthe and Dr Jef Vanderroost at KU Leuven [2]. The code converts a 3D binary matrix into a microstructural Finite Element mesh using beam and shell elements, and outputs information about the nodal coordinates, element connectivity and element geometry. More information about the code may be obtained in Section 2.5.1.2.

The original code obtained from van Lenthe analysed rectangular structures, where all six faces of the sample were labelled as boundaries and thus 2D skeletonisation was performed on these surfaces. An equivalent surface definition for the rounded face of a cylindrical sample is much more difficult to obtain, therefore edits were applied to the code by the developer, Dr Jef Vanderroost, to accommodate cylindrical specimens. Further modifications were applied by the author to allow the code to run for cylindrical specimens. The edited code only defines the top and bottom faces of the sample as boundaries, thereby limiting the 2D skeletonisation to these two surfaces as opposed to all the surfaces of the specimen. As previously mentioned, the function of the 2D skeletonisation is to provide clear nodal definitions on the boundaries so as to facilitate the application of boundary conditions. In the scope of this project, it is only necessary to apply boundary conditions on the top and bottom surfaces, therefore this new boundary definition was deemed to be sufficient for the purposes of this research. With the exception of the boundary definitions and linked skeletonisation, the cylindrical code is identical to the rectangular code.

3.2.2 Mesh refinement

After an initial set of simulations was run, it became clear that the mesh generated by Vanderroost's code could be refined to allow for more efficient and consistent simulations. Two areas of improvement were identified: the allocation of cross-sectional properties to elements and the removal of unnecessary elements from the mesh.

The motivation for and implementation of each of these refinements is discussed in Sections 3.2.2.1-3.2.2.2.

3.2.2.1 Cross-sectional property allocation

In LS-DYNA, mesh information is specified by means of “cards” which contain keywords in a particular format. There are three cards associated with the definition of an element:

- ***SECTION**, which defines the cross-sectional properties and element type.
- ***PART**, which allocates a particular section and material to a part.
- ***ELEMENT**, which identifies the nodes of the element and allocates a particular part to this element.

If every element has different cross-sectional properties, a unique section and part has to be defined for each element. This not only leads to large input decks³ and increased simulation time, but also significantly increases the memory required by MATLAB during any operations on the mesh data. Although trabecular bone is a highly complex and varied structure, it was found that not every element has unique cross-sectional properties assigned to it after discretisation. If all the duplicate sections are removed, the number of section definitions decreases substantially, sometimes by as much as 80%.

As a result of the above observations, duplicate section definitions are removed immediately after discretisation and ***SECTION** and ***PART** cards are only created for each unique section, as opposed to each element. This is achieved using in-house MATLAB code developed by the author.

3.2.2.2 Removal of unnecessary elements

It was noted that in some of the meshes, shell elements with zero thickness were being defined which caused a number of simulations to fail due to a “zero nodal thickness” error. Although some meshes showing this phenomenon were successfully simulated, these shell elements could not contribute towards the apparent stiffness of the structure and were, therefore, unnecessarily increasing the size of the mesh. As a result, in-house code was developed which would remove all shell elements having zero nodal thickness at one or more of the nodes.

Additionally, elements were identified which were connected locally to each other, but were not connected to the bulk of the mesh. These free element clusters do not contribute towards the apparent stiffness of the structure and, moreover, may cause errors in the simulation because they are ill-constrained. In-house code was developed to identify any free elements or free element clusters and remove them from the mesh. A more detailed description of the breadth-first algorithm used for this purpose may be found in Appendix B.

3.2.3 Creation of ties

Part of the novelty of the work by Vanderoost, et al [2] is the use of ties at the intersection of beam and shell elements. As described in Section 2.5.1.2, these ties distribute the load more evenly amongst the small shell elements, which is reported to give a more accurate result.

³It was found that the large input decks associated with the $\phi 20\text{ mm}$ and $\phi 28\text{ mm}$ bone specimens completely prevented the simulation as the combination of ***SECTION**, ***PART** and ***ELEMENT** cards were too large to be read into the pre-processor.

For the purposes of this project, it was decided that each tie would be incorporated as a beam element having the same cross-sectional properties as the beam that it connects to. In-house MATLAB code was developed which adds these tie elements to the element list, along with their corresponding nodes and cross-sectional properties.

3.3 Simulation

After discretisation, the mesh information must be transferred to LS-DYNA (a Finite Element software package) for simulation. This section details how the mesh information was formatted and delivered to LS-DYNA, and describes important aspects of the simulation such as boundary conditions and output requests.

3.3.1 Input decks

The BSM discretisation code [2] originally produced input decks for analysis in MARC, however the FE package available for use in the current work is LS-DYNA. As such, code was written in MATLAB which takes the information output by the discretisation code, reformats it and writes it to a series of input decks (.k files), which can be imported into and read by LS-DYNA.

It is possible to write the mesh information to a single input deck, but because of the size of the trabecular bone meshes, a decision was made to divide the information between a number of input decks as follows:

1. Nodal coordinates and element connectivity (*NODE, *ELEMENT_BEAM, *ELEMENT_SHELL)
2. Section definitions (*SECTION_BEAM_TITLE, *SECTION_SHELL_TITLE)
3. Part definitions (*PART)
4. Nodal sets (*SET_NODE_LIST_TITLE)
5. Material properties (*MAT_ELASTIC_TITLE)
6. Boundary conditions, loading and requests for output (*CONTROL, *DATABASE⁴, BOUNDARY, *DEFINE_CURVE_TITLE)
7. Main card, which imports the information from the other input decks (*INCLUDE).

Dividing the input decks in this way allows for quick access to particular information and negates the need to recreate input decks that do not change from specimen to specimen, particularly the input decks controlling material properties (item 5), boundary conditions and output requests (item 6).

⁴Both the *CONTROL and *DATABASE commands are used extensively in the input deck, however the specific details are not listed here for the sake of brevity.

3.3.2 Finite Element Analysis

Once the input decks have been created, they are imported into LS-DYNA where the FE simulation is performed. The rod-like trabeculae are specified as element type 13, 3D Timoshenko resultant-based beam elements with two nodes for small displacement, linear isotropic elasticity; the plate-like elements are specified as element type 16, fully integrated shell elements with 5 through thickness integration points. Elements are assigned material properties of $E = 1 \text{ GPa}$ and $\nu = 0.3$. The bottom surface of each specimen is constrained from translation in the z -direction, while a prescribed displacement of -1 mm in the z -direction is applied on the top surface⁵ to produce a small strain, quasistatic, linear deformation.

It is necessary to allocate additional memory to LS-DYNA at the beginning of the simulation process due to the size of the mesh. The command prompt is used to allocate the additional memory and run the simulation through LS-DYNA using an implicit integration scheme. Upon completion, a file containing the resultant forces at each of the nodes on the bottom surface is produced.

3.4 Post-processing

Post-processing of the simulation results was required in order to determine the apparent modulus of the structure and construct a binary image of the discretised structure for comparison purposes.

3.4.1 Modulus calculation

In-house MATLAB code was developed which imports the nodal forces and calculates the total resultant force and the apparent modulus of the specimen as shown in Equations 3.1 - 3.3:

$$\sigma_{app} = \frac{4F_{total}}{\pi d^2} \quad (3.1)$$

$$\varepsilon = \frac{\delta}{h} \quad (3.2)$$

$$E_{app} = \frac{\sigma_{app}}{\varepsilon}, \quad (3.3)$$

where F_{total} is the total resultant force on the bottom boundary of the specimen, d is the diameter of the specimen, h is the height of the specimen, δ is the prescribed displacement applied to the specimen, σ_{app} is the apparent stress in the direction of loading, ε is the strain in the direction of loading and E_{app} is the apparent modulus in the direction of loading.

In addition to the apparent modulus, the effective modulus and modulus ratio are also calculated for some of the specimens, where the effective modulus is defined as the modulus of an embedded inner core region as described in Section 2.6.2. The effective modulus calculation takes a similar form to Equations 3.1 - 3.3, however Equation 3.1 is altered as follows:

$$\sigma_{eff} = \frac{4F_{inner}}{\pi d_{inner}^2}, \quad (3.4)$$

⁵Note: These are the default boundary conditions specified by the input decks, however a number of boundary conditions were tested on the bone samples, which are described in Section 5.2.3.

where F_{inner} is the total resultant force experienced by the nodes lying within the inner core region and d_{inner} is the diameter of the inner core of the specimen (10 mm). The modulus ratio is calculated as shown in Equation 3.5.

$$r = \frac{E_{eff}}{E_{app}} \quad (3.5)$$

After calculation, the apparent and effective moduli and moduli ratio are saved in a text file for easy access.

3.4.2 Rasterisation

In order to evaluate the discretisation performed by the beam-shell code, a method was required which would allow the original structure to be compared to the discretised structure. To this end, in-house rasterisation (or image construction) code was developed which creates a TIFF stack using the node, element and cross-section information of the mesh. The recreated TIFF stack, which is constructed to be the same image size as the original TIFF stack, may then be compared to the original TIFF stack to determine if the original structure is being accurately represented by the discretised structure.

The developed code was separated into beam rasterisation and plate rasterisation, respectively. A detailed description of the rasterisation code may be found in Appendix C.

A note on terminology: In this work, the term “image size” is used to describe the number of voxels in the TIFF stack, rather than “image resolution” which is widely used in the literature. It was found that the term “resolution”, although generally understood to relate to the quality of the image, can have a number of different meanings, depending on the context. “Scan resolution” differs from “image resolution”, and even “image resolution” can point to a variety of metrics, including pixel resolution or spatial resolution. As a result, the term “image size” was used as it succinctly and clearly conveys the metric which is being discussed and can be quantified for all TIFF stacks, including those of the virtually generated lattices described in the next chapter.

Chapter 4

Validation of analysis cycle

After the creation and automation of the analysis cycle, a validation study was performed to ensure that the cycle was producing accurate results. Three cellular solids in various configurations were chosen for this purpose, namely the cubic lattice, Kelvin cell lattice and octet truss lattice. These lattices are regular, repeating structures which can be computationally constructed, as demonstrated by Moore [64]. The apparent modulus of the structure may be obtained via Finite Element Analysis. Using the lattices in this way provides a baseline result, against which any results from the analysis cycle may be compared to determine the accuracy of the cycle.

This chapter will describe the validation process, the various test configurations used during the process and the comparison metrics used to quantify the accuracy of the analysis cycle.

4.1 Test configurations

A wide variety of lattices, configurations and constraints were used during the validation process. A summary of the varied parameters is provided below, and a discussion of how these parameters were chosen follows.

- Lattice type: Cubic (C), Kelvin cell (K), octet truss (OT)
- Lattice size: 6 cells, 8 cells, 10 cells
- Aspect ratio¹: Uncropped, cropped
- Constraint: Unconstrained, constrained
- Strut diameter: 0.1 (t1), 0.2 (t2), 0.3 (t3)
- Image size: 100 voxels to 400 voxels² in increments of 10

In total, more than 3000 individual validation tests were performed on lattices of various type, size, aspect ratio, strut diameter and image size using different constraints. An outline of the validation tests performed for a single lattice type³ is provided in Figure 4.1.

¹The aspect ratio in this context refers to the ratio of the height to the width of the specimen.

²100 voxel image = 100 voxels × 100 voxels × 100 voxels

³Some other analyses were performed on a select few lattice configurations, which are not included in Figure 4.1. These additional analyses are described in Section 4.1.5.

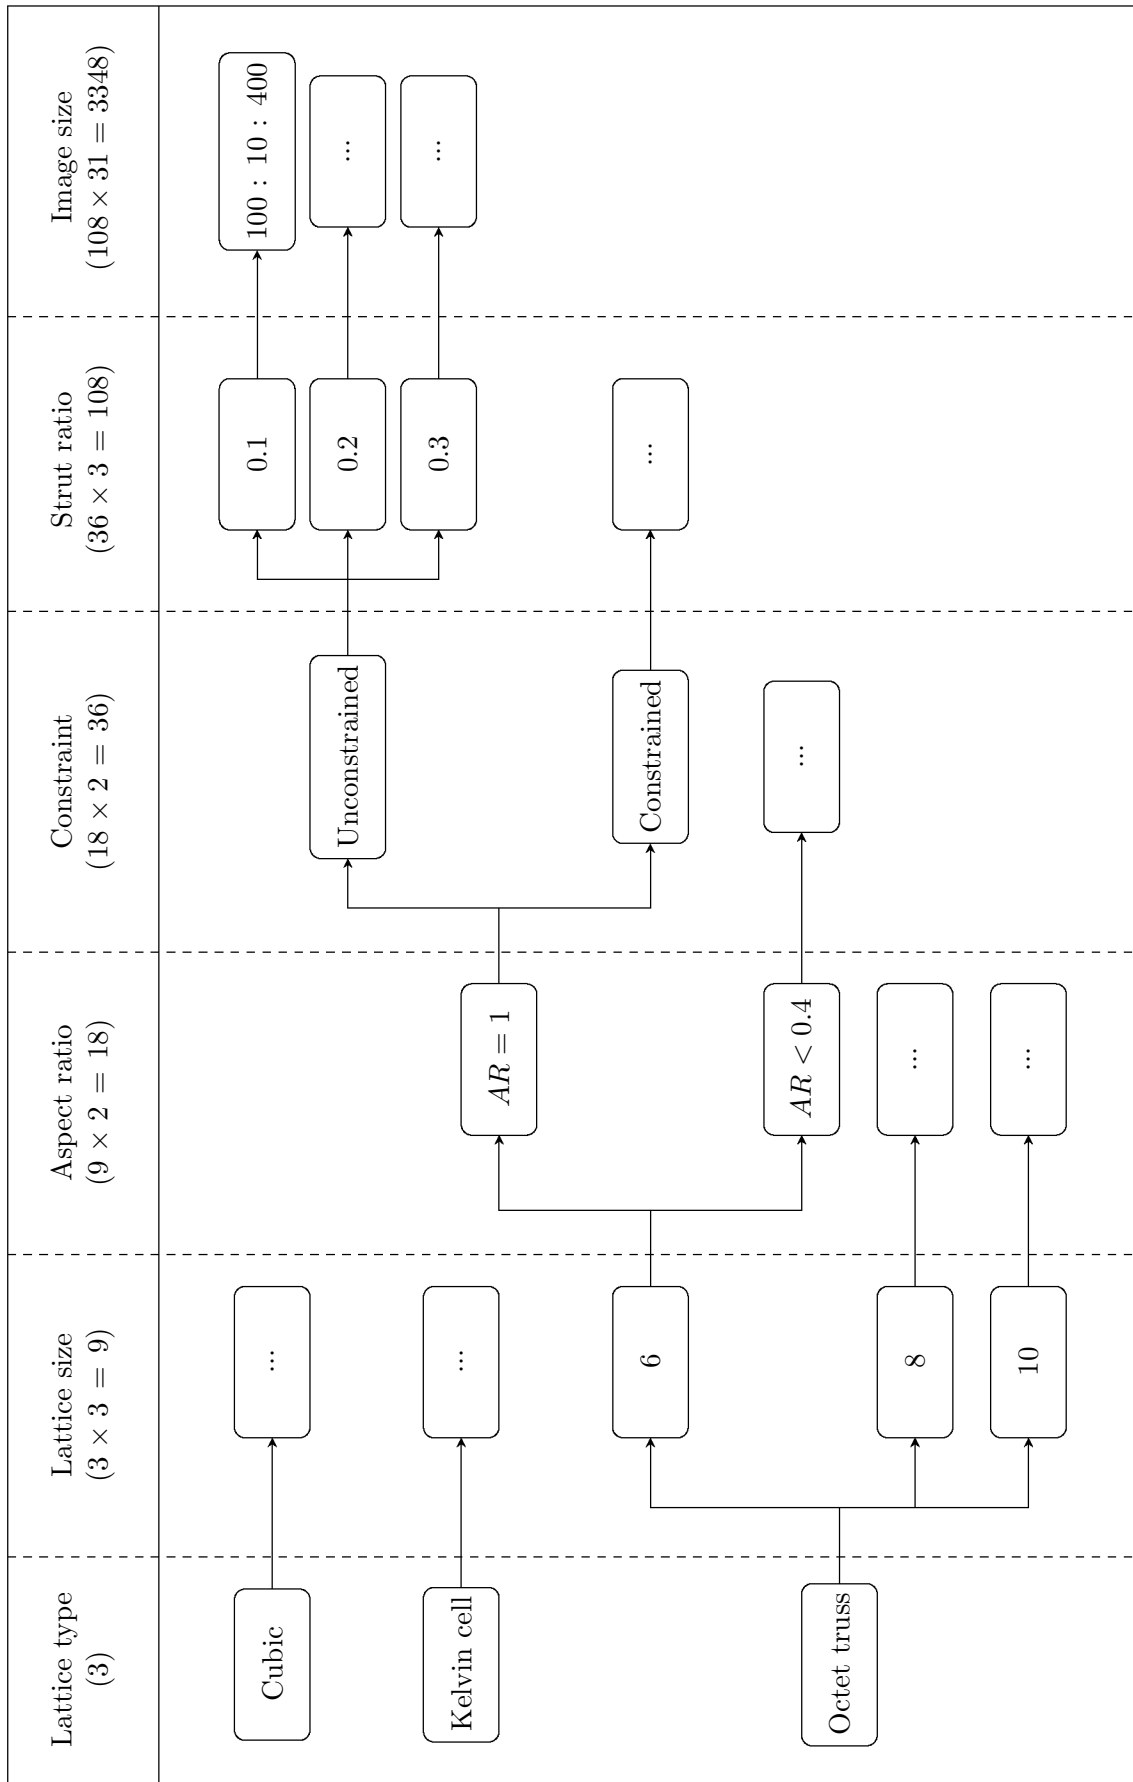


Figure 4.1: Flow chart of the test cases used for validation. The calculation for the total number of tests performed is stepped through underneath each test case title.

Due to the large number of varying parameters, a lattice naming convention was used to succinctly specify lattice type and topology. This consisted of the lattice name, followed by the lattice size and strut diameter. E.g.: A size 6 cubic lattice with a strut diameter of 0.2 is specified as C6-t2. This convention is followed for all the work dealing with lattices.

4.1.1 Lattice type

As discussed in Section 2.4, trabecular bone is classified as a cellular solid. This distinction is important as it implies that the properties of the bone are directly dependent on two main factors, namely: the material from which the bone is made and the topology of the cells which make up the structure [18]. In order to properly validate the analysis cycle, a test regimen was required which included structures showing a similar dependence on material and structure to the bone and a similar structure to the bone itself. The test specimens also required a regular structure, so that they could be artificially created and simulated to determine baseline properties. As such, it was decided that regular, three-dimensional cellular solids, or lattices, should be used for the validation study.

The structure of trabecular bone varies significantly even within the same anatomical region, as is demonstrated in Figure 2.3. In Figure 2.3, the trabecular bone structure changes from being relatively equiaxed at the load-bearing surface, to showing a more elongated structure at the free surface. This change in structure is synonymous with a change in the directional dependence of the properties [18], hence it is important that the analysis cycle be able to capture the response of structures of varying directional dependence. Consequently, the lattices included in the test regimen must show varied directional dependence, from fully isotropic to anisotropic. Additionally, the changing structure can affect how the bone responds to constraint [64, 68], which is important because in vivo trabecular bone is constrained differently at different locations. For example, in Figure 2.3, the trabecular bone is constrained by cortical bone at the load-bearing surface and unconstrained at the free surface. As such, in addition to ranging directional dependence, the chosen lattices should also show a range of constraint sensitivity, from insensitive (i.e. no change in properties when subjected to constraint) to highly sensitive.

Cubic, Kelvin cell and octet truss lattices, shown in Figure 4.2, were chosen for this work because they are analogous to the various trabecular bone structures found within a bovine femur, and show the full range of directional dependence and constraint sensitivity required for the validation study. The Kelvin cell lattice is isotropic, the cubic lattice is highly anisotropic and the octet truss lattice has a lower degree of anisotropy than the cubic lattice, but a higher degree of anisotropy than the Kelvin cell lattice [64]. Additionally, the Kelvin cell lattice shows a high sensitivity to constraint (i.e. the lattice shows a large change in apparent properties when it is constrained versus when it is unconstrained), the cubic lattice is insensitive to constraint and the octet truss lattice shows low sensitivity to constraint [64]. The wide variety of structures and properties represented by the cubic, Kelvin cell and octet truss lattices, summarised in Table 4.1, creates a compelling analogy to trabecular bone and, therefore, the collection of these lattices was determined to be an effective test regimen for the developed analysis cycle.




Figure 4.2: Representations of the three lattice types chosen for use in the validation study: cubic, Kelvin cell and octet truss.

Table 4.1: Properties of the cubic, Kelvin cell and octet truss lattices.

Lattice name	Directional dependence	Constraint sensitivity
Cubic	Highly anisotropic	Insensitive
Kelvin cell	Isotropic	Highly sensitive
Octet truss	Anisotropic	Sensitive

4.1.2 Lattice topology

Although cubic, Kelvin cell and octet truss lattices show structural similarities to trabecular bone, these similarities only become apparent at a particular lattice size⁴. For example, a size 2 octet truss lattice ($2 \times 2 \times 2$) does not provide an effective representation of trabecular bone - the struts are too thick and the spaces between struts are too large - however, a size 6 lattice generated at the same image size has thinner struts and smaller spaces between struts, making it much more representative of bone, as shown in Figure 4.3. It was found that lattices of sizes 6, 8 and 10 were partially representative of trabecular bone, while still being small enough to generate and simulate in a timely manner. As such, these three lattice sizes were used during the validation study.

The lattice size was not the only aspect of the topology which was varied during the validation study. In their pioneering book on cellular solids, Gibson and Ashby [18] report that the “single most important structural characteristic of a cellular solid is its relative density⁵”. As the trabecular bone structure varies within an anatomical region, so does the thickness of the trabeculae and, therefore, the relative density of the structure. In order to investigate how well the analysis cycle captures the response of structures of varying relative density, lattices with varying strut diameters, namely 0.1, 0.2 or 0.3, were generated for use in the validation study as shown in Figure 4.4. Because these lattices were artificially generated, the assignment of measurement units to the diameter of the struts is arbitrary, which is why no units are mentioned. However, it is important to note that these diameters correspond to strut diameter-to-length ratios of 0.1, 0.2 and 0.3, respectively.

⁴Note: Lattice size is defined as the value input to the *numCells* variable in the lattice generation code [64]. Because of the different ways the lattices are generated in the code, there may be some variation in the size of the produced lattices when compared to other lattice types.

⁵The relative density of a lattice is the density of the structure (i.e. the porosity between struts) divided by the density of the material from which the struts are made.




Figure 4.3: Comparison of 3D renderings of size 2 and 6 octet truss lattices with a trabecular bone sample.




Figure 4.4: 3D renderings of the size 8 cubic lattices with strut diameters of 0.1, 0.2 and 0.3, respectively.




Figure 4.5: 3D renderings of the size 8 Kelvin cell lattices with diameters of 0.1, 0.2 and 0.3, respectively.

Changing the diameter served another purpose. When TIFF stacks were generated for the size 8 Kelvin cell lattice, it was found that the 0.1 and 0.2 diameter struts were extremely thin because they are rasterised at a 45° angle, meaning that only one or two voxels were being used to represent the struts in each image slice. These extremely thin struts create a lattice which is not clearly discernible, as shown in Figures 4.5a-4.5b and, consequently, the discretisation code fails to discretise these structures. Increasing the strut diameter results in better quality rasterisation and insures that representative lattices are being generated for the validation study, as illustrated in Figure 4.5c.

4.1.3 Constraint and aspect ratio

One of the main aims of this work is to investigate whether the developed analysis cycle can accurately predict the effective stiffness of a structure, where effective stiffness is defined as the stiffness of an inner core region in the sample (see Figure 2.16 for an illustration). In the experimental work by Hilton [3], it was found that the effective modulus of the bone samples was, on average, approximately 60% stiffer than the apparent modulus, which was due to side artifacts caused by machining. In order to recreate this experimental test set-up with lattices for the purposes of this validation study, two factors need to be considered, namely:

- The experimental boundary conditions of the bone specimens.
- The height-to-diameter ratio of the bone specimens.

The boundary conditions present in the experimental test set-up [3] are difficult to define. In reality, there is some friction between the plates of the testing rig and the bone specimens, which constrains the lateral motion of the top and bottom surfaces, however it is unlikely that lateral motion is hindered completely. As such, neither a laterally unconstrained FE model, nor a laterally constrained FE model is a wholly accurate representation of the experimental set-up. Although work has been done to more accurately replicate the experimental boundary conditions [82, 107], these complex constraints were outside the scope of this project and, therefore, were not employed. Instead, specimens were simulated using laterally unconstrained boundary conditions as well as laterally constrained boundary conditions, so as to understand the effect of constraint on the specimen, while acknowledging that the true response of the bone falls between these two extremes.




Figure 4.6: C8-t1 lattices in the uncropped and cropped conditions.

In addition to the constraint, it was also important to consider the overall dimensions of the bone specimens. The bone specimens used in the experimental tests [3] had height-to-diameter ratios (hereafter referred to as aspect ratios) of 0.750, 0.375 and 0.268, respectively. Because the specimen heights are small compared to their diameters, it means that any boundary condition present in the experimental set-up would have a large effect on the response of the bone. For example, if there was a large amount of friction between the plates of the ZWICK tester and the bone, the bone would be laterally constrained on its top and bottom surfaces. The constraint would not only effect the two surfaces of the specimen, but also permeate into the specimen. For a tall specimen, there would come a point where the effects of the boundary conditions would be negligible and a bulging effect would be seen. However, if the specimen is short in comparison to its diameter, the boundary effects would permeate further into the sample, deterring bulging and effectively creating a more encompassing constraint. Consequently, any specimen which is sensitive to constraint would show a greater difference in constrained versus unconstrained response when cropped to an aspect ratio less than one, than it would if it had an aspect ratio greater than or equal to one.

In order to investigate the effect of constraint in the lattices (which have a standard aspect ratio of 1), two different aspect ratios were considered in the validation study. For each simulated lattice, a second round of validation tests was performed on a cropped lattice having an aspect ratio of approximately 0.375⁶. E.g. all tests performed on an C8-t1 lattice were also performed on a cropped C8-t1 lattice, as shown in Figure 4.6. The terms “uncropped” and “cropped” are used to distinguish between the two aspect ratios for brevity.

4.1.4 Image sizes

One of the first steps in the validation process, described in Section 4.2, is to generate an image of the lattice in the form of a TIFF stack so that it may be imported into the analysis cycle. The quality of the rasterised image depends on the image size, lattice size and lattice topology; if a complex lattice is rasterised at a small image size, the quality of the rasterisation, i.e. how well the generated image represents the lattice, will be poorer than if the same lattice is rasterised at a larger image size. Various lattice sizes and topologies have already been included in the validation study based on other factors, therefore the only parameter left to investigate is the image size.

⁶Lattices were cropped in such a way that whole cells were maintained, i.e. no struts were severed. As a result, the aspect ratio of the cropped lattices varies slightly between lattice types.

The image size causes variability in the rasterisation process because decreasing image size means a decreasing number of voxels in the image. As a result, when the struts are rasterised, very few voxels are used to represent the cross-section of the struts, which leads to incorrect cross-sectional shape, cross-sectional area and moments of inertia. In some cases, the rasterised struts are so thin that a “skeleton” image is produced, where the struts have no significant thickness and, in places, no thickness at all as shown in Figure 4.5a. These images are not representative of the lattice, and therefore will not produce viable results when discretised by the analysis cycle.

The rasterisation of each lattice type only becomes representative at a certain combination of strut diameter and image size, however this combination is unique for the different lattice types and sizes. Therefore, instead of arbitrarily choosing an image size for each lattice, TIFF stacks of each lattice were generated at a variety of different sizes, from 100 to 400 voxels. This range was chosen based on the image sizes of the bone micro-CT scans, which are approximately 140 voxels, 270 voxels and 370 voxels for the $\phi 10\text{ mm}$, $\phi 20\text{ mm}$ and $\phi 28\text{ mm}$ specimens, respectively. Additionally, after initial testing it was found that the rasterisation of almost all the lattices becomes representative within this range, which makes it a satisfactory test criteria.

4.1.5 Additional analysis

Although the configurations described in Figure 4.1 cover a wide range of structures and properties, two additional analyses were conducted as part of the validation process. These additional configurations were not privy to all the test variations described in Figure 4.1, but rather were chosen to provide a specific comparison which, once achieved, did not require further variations of the structure or image to be simulated.

Lawrence [68] and Moore [64] reported that a cubic lattice angled at 45° , which is shown in Figure 4.7, is highly sensitive to constraint, unlike a regular cubic lattice, which is completely insensitive to constraint. Although three levels of constraint sensitivity are provided by the cubic (insensitive), Kelvin cell (highly sensitive) and octet truss (sensitive) lattices used in the validation process, an angled cubic lattice is even more sensitive to constraint than a Kelvin cell lattice, which allows the ability of the analysis cycle to capture constraint sensitivity to be more accurately quantified. As such, the following variations of the angled cubic lattice were analysed as part of the validation process, leading to a total of 248 simulations:

- Lattice type: Angled cubic
- Lattice size: 8
- Strut diameter = 0.1, 0.2
- Aspect ratio: Uncropped, cropped
- Constraint: Unconstrained, constrained
- Image sizes: 100 voxels to 400 voxels in increments of 10

The second additional analysis that was performed as part of the validation cycle did not consist of adding a different structure to the test regime, but rather analysing some of the existing structures with a different method. For comparison purposes, it is useful to determine not only how accurately the BSM predicts the apparent stiffness of the lattices, but also how accurately




Figure 4.7: Angled cubic lattice.

the VBM predicts these values. As such, a collection of lattices was chosen and the analysis cycle was rerun using the VBM instead of the BSM. This was achieved by creating an 8-node hexahedral element (LS-DYNA element type 1) for every “bone” voxel in the binarised image using in-house MATLAB code written by the author. Because of the time- and memory-intensive nature of the VBM simulations, only a select few of the lattices were used for this analysis. The following configurations were simulated using the VBM, for a total of 62 simulations:

- Lattice type: Kelvin cell
- Lattice size: 8
- Strut diameter: 0.3
- Aspect ratio: Cropped
- Constraint: Unconstrained, constrained
- Image sizes: 100 voxels to 400 voxels in increments of 10

Although the additional analyses described in this section are not as extensive as the full validation study presented in Figure 4.1, the additional results will provide key insights into the functioning and accuracy of the analysis cycle.

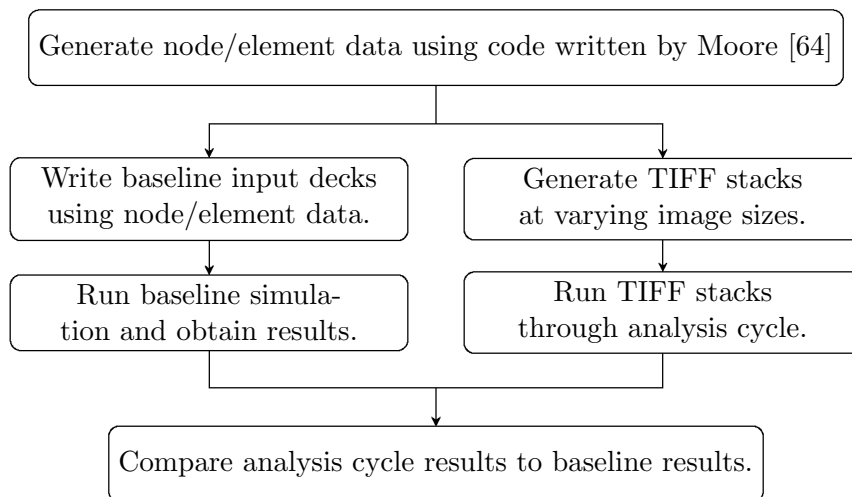


Figure 4.8: Flow chart of the validation process.

4.2 Validation process

The validation process may be divided into two distinct sections: the calculation of the baseline results and the processing of the sample by the analysis cycle. Although the validation study involves, in part, checking if the analysis cycle is capable of dealing with a variety of different structures in a variety of different configurations, the main purpose of the study is to provide information about the accuracy of the results produced by the analysis cycle. Therefore, it is vital to compare the results from the analysis cycle to predetermined values which accurately capture the response of the specimen, the so-called “baseline results”.

In order to determine the baseline results, cubic, Kelvin cell and octet truss lattices were generated using code written by Moore [64], which outputs a text file containing nodal coordinates and element connectivity of the required lattice. Custom MATLAB code was used to generate a set of input decks directly from the node and element information⁷. These input decks were imported into LS-DYNA where a Finite Element simulation was performed, after which the baseline results for apparent modulus, effective modulus and modulus ratio were calculated from the simulation results. In his work, Moore [64] found an excellent correlation between the FEA results generated from his code and analytical results in the literature for the varying lattices. As a result, the baseline results are considered to be a good standard for comparison.

Once the baseline results were obtained, TIFF stacks of the lattices were constructed at varying image sizes using the node and element information output by the lattice generation code [64]. The images were rasterised using custom beam rasterisation code, which is discussed in detail in Appendix C. The TIFF stacks were each run through the analysis cycle to determine the apparent modulus, effective modulus and modulus ratio of the sample.

Finally, the results from the analysis cycle were compared to the baseline results to quantify the accuracy of the analysis cycle. A summary of the validation process is depicted in Figure 4.8.

⁷These input decks differ from the input decks generated as part of the analysis cycle, because the node and element information contained within them is exact, not inferred from an image.

4.3 Analysis

Although the comparison of moduli generated by the analysis cycle to the baseline moduli is the most important metric to consider, there are other values which may be compared to further quantify the accuracy and efficiency of the analysis cycle. For each of the lattices, the following comparisons to baseline values were performed:

- Apparent modulus⁸.
- Effective modulus.
- Modulus ratio
- Average cross-sectional area of beam elements.
- Average moment of inertia of the beam elements in the primary direction.
- Average moment of inertia of the beam elements in the secondary direction.
- Number of nodes in the mesh.
- Number of nodes on the bottom surface of the mesh⁹.
- Number of nodes on the top surface of the mesh.
- Number of elements in the mesh.
- Number of beam elements in the mesh.
- Number of shell elements in the mesh¹⁰.

For each of these values, a percentage difference from the baseline value was calculated as shown in Equation 4.1:

$$P = \frac{x_{cycle} - x_{base}}{x_{base}} \times 100, \quad (4.1)$$

where P is the percentage difference between the numerical and baseline values, x_{cycle} is the value obtained from the analysis cycle and x_{base} is the baseline value.

Additionally, the images produced at the end of the analysis cycle were visually compared¹¹ to the images that were imported into the analysis cycle, i.e. the image which was constructed from the Finite Element mesh data produced by the analysis cycle was compared to the image upon which the discretisation was based. This comparison is important because it provides a quick check of how well the discretisation code works. If the image constructed at the end of the analysis is completely different from the original image, then clearly the discretisation code has not performed well and the results produced are invalid.

⁸See Section 2.5.3 for definitions of apparent and effective modulus.

⁹The number of nodes on the top and bottom surface of the mesh was taken into account because boundary conditions are applied on these surfaces, making these nodes non-trivial.

¹⁰Ideally, there should be no shell elements generated in the lattices, because they are made entirely from rod-like struts.

¹¹An objective way to compare the images would be to subtract the recreated binary matrix (post-analysis cycle) from the original binary matrix (pre-analysis cycle). An attempt was made to use this methodology, however it proved difficult to create images of the lattices that were the same size and had the same positioning as the original images. This is due to the complex manipulation of the binary matrix that occurs as part of the discretisation process created by Vanderoost, et al [2]. For the purposes of the current work, a visual comparison is sufficient to determine whether or not the analysis cycle produces representative meshes, although the more analytical approach to the comparison could be included in future work to further quantify the accuracy of the analysis cycle.

Chapter 5

Simulation of bone

Micro-CT scans of bovine trabecular bone specimens obtained from Hilton [3] were used as a basis for the bone simulations in the current work. This section describes the various simulations performed on the bone specimens.

5.1 General simulations

Micro-CT scans of 110 $\phi 10$ mm bovine bone specimens, nine $\phi 20$ mm specimens and eight $\phi 28$ mm specimens¹ all with a height of 7.5 mm [3] were run through the developed analysis cycle. As the true material properties of the bone specimens were unknown, standard properties of $E = 1$ GPa and $\nu = 0.3$ were assigned to the specimens². Apparent moduli were calculated for all the specimens and effective moduli were calculated for the large specimens by considering the response of an inner 10 mm core as discussed in Section 2.6.2. Equivalent voxel-based simulations were performed on the same 127 specimens for comparison purposes.

5.2 Factors affecting simulation results

Three major contributors to simulation accuracy were identified, namely:

- **Voxel size** of the micro-CT scans
- **Segmentation technique** used during the image processing of the scans
- **Boundary conditions** applied to the mesh

In order to quantify the effect of these variations on the results, multiple iterations of the analysis cycle were run for each specimen. The various iterations are described in Sections 5.2.1-5.2.3.

¹To differentiate between the sizes of bone without constantly quoting diameters, $\phi 10$ mm specimens are referred to as “standard specimens”, whereas $\phi 20$ mm and $\phi 28$ mm specimens are referred to as “large specimens”.

²This arbitrary material property assignment allowed the BSM results to be directly compared to the VBM results. However, once results were obtained, back-calculated moduli were quoted for the BSM and VBM results, respectively.

Table 5.1: Parameters used for nano-CT scanning at varying voxel size.

Voxel size (μm)	X-ray tube potential (kV)	X-ray intensity (μA)	Integration time (ms)
20	70	350	500
40	70	200	250
80	70	70	125

5.2.1 Voxel size

The voxel size used for the micro-CT scans can have a significant effect on the microstructural indices and the apparent moduli determined from FEA [45]. In this work, one standard bone specimen and one large bone specimen were scanned at three different voxel sizes, namely 20 μm , 40 μm and 80 μm in order to investigate the effect of a change in voxel size on the results.

The specimens were scanned using a Phoenix nano-CT scanner, which works similarly to a micro-CT scanner, but is capable of providing much smaller voxel sizes. Because the nano-CT scanner was incapable of producing voxel sizes larger than approximately 30 μm for specimens of the given size, a binning approach³ was used to obtain the 40 μm and 80 μm scans. Additionally, the parameters of the scan had to be changed for each voxel size. The parameters used for the various scans are listed in Table 5.1.

The files output for the scans at voxel sizes of 20 μm and 40 μm were extremely large and could not be processed on the available computer. In order to minimise the required computational resources, a region of interest 5 mm \times 5 mm \times 5 mm was chosen for each specimen. The specimen was cropped to include only the relevant ROI and then imported into the analysis cycle⁴. This cropping technique is prevalent in the literature and was an acceptable method to use for these scans as the simulation results were not being compared to experimental results.

To determine the effect of the voxel size on the scans, three comparisons were made:

1. Visual comparison of segmented scans.
2. Comparison of microstructural indices, including bone volume fraction, trabecular thickness and trabecular separation.
3. Comparison of apparent moduli output by the analysis cycle.

5.2.2 Segmentation technique

The segmentation method used can play a major role in the thickness of trabeculae and, therefore, the apparent modulus of a bone sample. In order to quantify the variation, two built-in BoneJ [53] segmentation methods were considered, which are listed on the next page.

³Binning is a micro-CT scanning mode, where a matrix of voxels (2 \times 2 \times 2, 3 \times 3 \times 3, etc.) is taken in and combined to create a single larger voxel during scanning. The larger voxel has more contrast and less noise than the voxels in the voxel matrix due to the increased signal level measured from the larger effective sensor voxel, however the resolution of the images is lower. Binning occurs during the scanning process, and therefore reduces the amount of data that needs to be transferred to the computer, leading to reduced acquisition times [108].

⁴Equivalent VBM simulations were not performed for these bone specimens as the memory required for the simulations exceeded the available computing resources.

- *Make Binary*, which determines a threshold value for each image in a stack using the Otsu method [42].
- *Optimise Threshold* which determines a single threshold value for the stack based on the average threshold value of all the images in the stack.

Each specimen was segmented using both of these segmentation methods, and the resulting apparent moduli were compared to determine the variation in the results. Additionally, a second set of simulations was performed where a lateral constraint was applied to the top and bottom surfaces of the specimens. The goal of this analysis was to investigate whether a change in segmentation technique causes a change in the sensitivity of the specimens to constraint.

5.2.3 Boundary conditions

In the experimental testing rig used by Hilton [3], the boundary conditions at the top and bottom of the sample are difficult to define. Although no external constraints are applied to the majority of the specimens⁵, there is some friction between the surface and the specimen, which inevitably constrains the specimen, but also allows some motion. There has been some research into accurately simulating the boundary conditions in an experimental test [82, 107], but that was not the aim of this work and so these complex boundary conditions were not employed.

Although this work did not seek to accurately represent experimental boundary conditions, it was considered important to broadly understand the influence of boundary conditions on the simulated apparent moduli of trabecular bone. As such, four different types of boundary conditions were considered:

- **Unconstrained:** the minimum constraint described in Section 3.3.2, i.e. a constraint in the z -direction on the bottom surface of the specimen.
- **Laterally constrained:** in addition to the minimum constraint, lateral translation (x and y) of the specified surface is constrained.
- **Rotationally constrained:** in addition to the minimum constraint, rotation (R_x , R_y and R_z) of the specified surface is constrained.
- **Fully constrained:** in addition to the minimum constraint, lateral translation and rotation of the specified surface are constrained.

With the exception of the “unconstrained” baseline condition, each of these boundary conditions was applied to the top surface, bottom surface and simultaneously to the top and bottom surfaces of the specimen. This resulted in a total of ten boundary condition simulations being performed for each specimen: unconstrained (1 simulation), laterally constrained (3 simulations), rotationally constrained (3 simulations) and fully constrained (3 simulations).

⁵The confined specimens were confined in an aluminium collar, however this confinement was found to be ineffective due to undersizing of the bone specimens, as discussed in Section 2.6.2. As a result, the confined specimens were simulated using the same boundary conditions as the defatted and marrow specimens, despite the variation in experimental set-up.

5.3 Microstructural indices

Several microstructural indices were determined for each specimen, namely:

- Bone volume fraction (BV/TV)
- Trabecular thickness (Tb.Th)
- Trabecular separation (Tb.Sp)
- Connectivity
- Connectivity density

These indices were chosen based on the existence of comparative values in the literature as discussed in Section 2.3.2. In addition to these established indices, the shell-to-beam (S2B) ratio in the beam-shell mesh was calculated for each specimen. Although this is not a traditional index since it is calculated from a FE mesh and not a micro-CT scan, it was important to present an index which provided a measure of the shell-beam split in the bone. The S2B ratio serves as an alternative measure to the structure model index, which has been shown to produce questionable results for trabecular bone [65]. The S2B ratio was calculated by finding the ratio of the volume of the shell elements to the volume of the beam elements in the mesh as shown in Equation 5.1, where A^S is the area of a shell element, \bar{t} is the average thickness of a shell element, A^B is the area of a beam element and l is the length of a beam element.

$$S2B = \frac{\sum_{i=1}^m A_i^S \bar{t}_i}{\sum_{i=1}^n A_i^B l_i} \quad (5.1)$$

5.4 Analysis of results

The performance of the microstructural analysis cycle was judged based on three parameters:

- **Success rate** of simulations, i.e. how often was the discretisation and simulation successful? It is unlikely that all the simulations will run successfully either due to problems with discretisation or simulation of the specimen, however the ultimate goal in the field of trabecular bone FEA is to create a robust FE methodology which consistently produces accurate results, therefore the success rate of the simulations is of interest in this work.
- **Accuracy** of results, which was determined by comparing the BSM results to the experimental [3] and VBM results. It is not expected that the BSM results will show a 1 : 1 correlation with the experimental results due to the arbitrary material properties, specifically Young's modulus, assigned to the FE mesh (see Section 5.1). As a result, any straight-line correlation with the experimental results should be considered a positive outcome as subsequent back-calculation of Young's modulus will provide a 1 : 1 correlation as discussed in Section 2.5.2. When comparing the BSM results to the VBM results, it is expected that the VBM results will be 1.605 times higher than the BSM results, as noted in the work by Vanderroost, et al [2].

- **Resource usage** during discretisation and simulation, specifically with regards to memory requirements and time taken. This comparison will focus on the difference between the resource usage for the BSM and the VBM and how these values compare to the values reported by Vanderroost, et al [2] (see Section 2.5.1.2).

Chapter 6

Validation results and discussion

In total, more than 3000 validation tests were performed on the analysis cycle. The challenge associated with producing so much data is presenting the results in a way that illustrates the trends, without displaying every graph produced during the process. To this end, it was decided that the results of a single lattice type would be displayed and thoroughly discussed to exhibit the performance of the analysis cycle to the reader. The Kelvin cell lattice was chosen to be this exemplary case because of its high sensitivity to confinement; if the analysis cycle is capable of capturing this confinement sensitivity, it is a good indication that the analysis cycle is producing accurate results.

Due to the multitude of simulations performed¹, it is not possible to discuss all the results obtained during the course of the validation study. However, it would be negligent to ignore the results produced by the other lattices. As such, key features from the analysis cycle results for the cubic and octet truss lattices will be pointed out and discussed in Section 6.3.

A full set of results, including apparent modulus, inner modulus and modulus ratio, for all the lattices in the various configurations may be found in Appendix D.

6.1 Sources of error

The aim of the validation study was to quantify the accuracy of the analysis cycle, which depends largely on the discretisation process. However, it is important to note that another major source of error exists outside of the analysis cycle, namely the rasterisation of the initial TIFF stacks. As described in Section 4.2, TIFF stacks of the lattices are rasterised from node and element data and are used as the input to the analysis cycle. If the rasterised images are not representative of the lattices, it is impossible for the analysis cycle to produce representative results, no matter how accurate the method.

¹As a result of the number of graphs presented in this chapter, some figures appear relatively far away from the related text. It is suggested that the reader makes use of the hyperlinks in the text to navigate to the relevant figure. After viewing the figure, the “Alt + Left” combination can be used to return to the previous position in the document (i.e. the related text).

As a consequence of the dual sources of error, the presented results will often be discussed in terms of:

- The quality of the rasterisation, i.e. was the rasterised image representative of the lattice?
- The effectiveness of the analysis cycle, i.e. given a representative starting image, did the analysis cycle produce accurate results?

For each lattice, a representative image size is identified; this is the smallest image size for which a representative lattice is produced during the rasterisation process. Any discussion after the identification of the representative image size will focus on the images which are larger than the representative image size, as these results convey the accuracy of the analysis cycle and not the accuracy of the rasterisation process. After the representative image is identified, the reader should take “all images” to mean “all images larger than the representative image size”, even if this is not explicitly stated.

3D renderings of all the rasterised lattices in the uncropped condition are shown in Appendix F.

6.2 Kelvin cell lattice

Choosing a single lattice type substantially decreases the number of results to be discussed, however there are still multiple test configurations within each lattice which need to be analysed. In order to keep the analysis of these results as non-repetitive and digestible as possible, the uncropped, unconstrained K8-t3 lattice was chosen as a starting point for the discussion. The apparent modulus, inner modulus and modulus ratio results for this lattice are displayed and discussed in detail, after which the variation of the results with lattice size, strut diameter, cropping and constraint is examined in relation to this starting lattice.

The starting lattice was chosen to provide convenient comparisons for the variation in results due to various factors. An uncropped, unconstrained lattice was used so that the variation of the analysis cycle results with cropping and constraint could be discussed. The choice of a size 8 lattice allowed for comparisons to be drawn between the performance of the analysis cycle for smaller (size 6) and larger (size 10) lattice sizes. Finally, a strut diameter of 0.3 was chosen because the rasterisation at this diameter produced representative images, which provided a solid starting point for the analysis cycle (more about this in Section 6.2.3)².

The Kelvin cell lattice results in this chapter are extensively discussed, with corroborating evidence being provided where appropriate. These results and their significance are summarised in Section 6.2.7 for convenience.

²Note: The results for each configuration are shown for image sizes of 100 to 400 voxels, therefore the variation of results with image size is discussed within each section, instead of in a dedicated section.

6.2.1 K8-t3

The error in the apparent modulus results produced by the analysis cycle for the uncropped, unconstrained K8-t3 lattice is shown in Figure 6.1. The error was calculated by finding the percentage difference between the apparent modulus predicted by the analysis cycle and the baseline apparent modulus, as discussed in Section 4.3. In the graph, the analysis cycle results are shown as black circles, the failures³ are shown as red crosses, the 10% error region is shown as a shaded green box with a solid outline and the 20% error region is shown as a shaded orange box with a dashed outline. A similar graphing system is used throughout this chapter to illustrate the results produced by the analysis cycle.

Figure 6.1 shows that there is a large error in the predicted apparent modulus at small image sizes, as well as two discretisation failures, which occur for the 110 and 130 voxel images. As the image size increases, the error in predicted apparent modulus decreases until it stabilises at an image size of approximately 200 voxels. The results produced for image sizes of 200 voxels and above are excellent, with all results falling within the 10% error region, except one. At an image size of 300 voxels, the error jumps to 14.4%, a significant increase from both the 290 and 310 voxel images, which both produce errors below 1%. At 390 voxels, there is a similar surge⁴ in error, though to a lesser degree.

The inner modulus exhibited in Figure 6.2 shows a similar trend to the apparent modulus, with the analysis cycle mostly producing results with less than 10% error for image sizes of 200 voxels and above. Similar to the apparent modulus results, distinct surges in the error may be seen at 300 voxels and 390 voxels. However, in the inner modulus results there are also surges in the error at 220 voxels, and 310 voxels and, overall, there is slightly more variation in the inner modulus results than there was in the apparent modulus results.

Considering the error associated with the prediction of both apparent modulus and inner modulus was low, it makes sense that the analysis cycle provided excellent predictions of the modulus ratio, with errors of less than 10% for all but one of the image sizes above 200 voxels as shown in Figure 6.3. Error surges are seen at 220 voxels and 310 voxels, which correspond to the error surges in the inner modulus graph. Interestingly, no error surge is seen at 300 voxels or 390 voxels, due to the fact that there were error surges in both the apparent modulus and inner modulus predictions. Additionally, it is worth noting that the predicted modulus ratios are consistently greater than 1, indicating that the inner modulus is stiffer than the apparent modulus of the specimen, which is the expected result.


The poor results at small image sizes, the error surges and the variation of the inner modulus are discussed in detail in the following sections.

6.2.1.1 Error at small image sizes


The error associated with the results at small image sizes may be understood by considering the rasterised images which are used as input to the analysis cycle. At small image sizes, the cross-section of a strut is represented by only one or two voxels in each image slice, as is demonstrated in Figure 6.4a, or is entirely missing from the image slice as is demonstrated in Figure 6.4b. This rasterisation results in extremely thin struts, to the point where the struts are indiscernible in a 3D rendering of the images, as shown in Figure 6.5a and 6.5b. Additionally,

³Any specimen which could not be discretised or simulated is indicated as a “Failure” on the graphs.

⁴In this report, “error surge” is used to refer to a sharp increase in the absolute error.




(a) Full set of results.




(b) Enlarged view.

Figure 6.1: Apparent modulus error results for an uncropped, unconstrained K8-t3 lattice at varying image sizes.




(a) Full set of results.




(b) Enlarged view.

Figure 6.2: Inner modulus error results for an uncropped, unconstrained K8-t3 lattice at varying image sizes.



(a) Full set of results.



(b) Enlarged view.

Figure 6.3: Modulus ratio results for an uncropped, unconstrained K8-t3 lattice at varying image sizes.




Figure 6.4: Image slices taken from the 110 voxel rasterisation of the K8-t3 lattice. Slice 17 in (a) shows examples of single voxel cross-section rasterisation, while slice 18 in (b) shows an example of missing cross-section rasterisation.




Figure 6.5: 3D renderings of uncropped K8-t3 rasterised images at 100, 140, 200 and 400 voxels, respectively. At 100 voxels and 140 voxels, the rasterised image is not representative of the lattice, whereas at 200 voxels and 400 voxels, the structure of the lattice is well-represented by the rasterised image.

the missing cross-sections cause discontinuities in the rasterised structure, which are not present in the lattice. When the rasterised images are imported into the analysis cycle, it is likely that the thinning algorithm used during the discretisation process⁵ will fail, as the struts are already skeletonised and any further thinning will result in the elimination of the struts altogether. This kind of failure is demonstrated by the 110 and 130 voxel rasterisations.

If the thinning algorithm does not fail and the discretisation process is successful, the constructed Finite Element mesh will not be representative of the lattice as the rasterised image used as input was, itself, not representative. Additionally, the thinning algorithm may produce unexpected results because of the already “skeletonised” nature of the struts, which could further distort the mesh as shown in Figure 6.6. Although these samples will successfully run through the analysis cycle, they will not produce accurate results, as demonstrated clearly in the 100-150 voxel region in Figure 6.1. As such, any error produced for small image sizes may be considered to be caused by poor image rasterisation, not a fault in the analysis cycle.

⁵See Section 2.5.1.2.

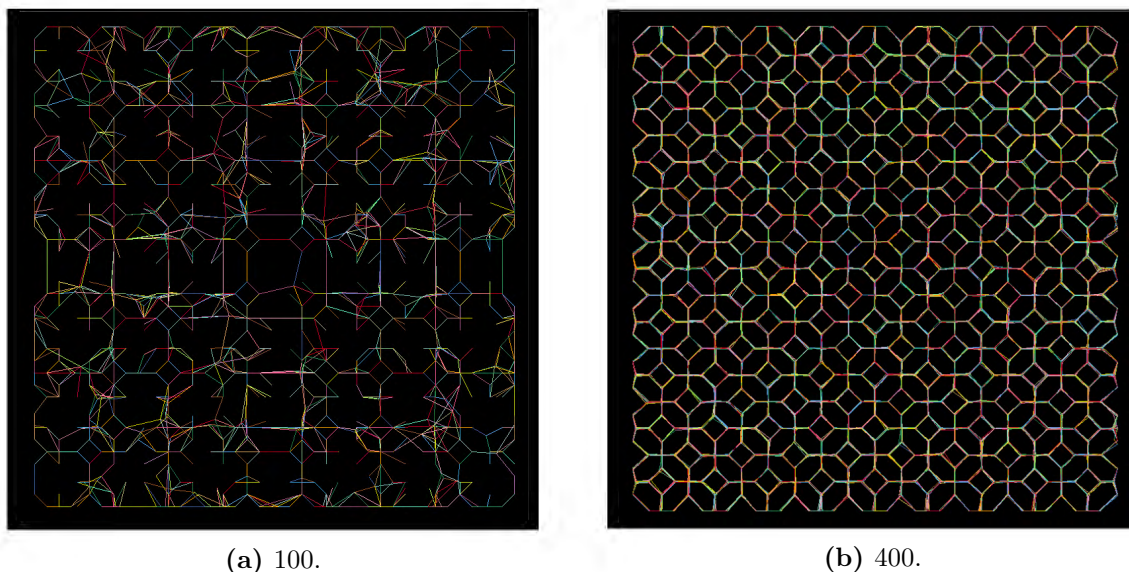


Figure 6.6: FE meshes produced for the K8-t3 lattices at image sizes of 100 and 400 voxels. The 100 voxel mesh is clearly distorted and, overall, not representative of the Kelvin cell lattice.

Considering the quality of images has a direct effect on the accuracy of the results produced by the analysis cycle, it is important to determine at what image size the rasterised structure becomes representative of the lattice. This is done visually by considering the 3D renderings of the images at various image sizes, as well as the image slices associated with each rendering. For the K8-t3 lattice, the images appear to become representative of the lattice at an image size of 200 voxels⁶ as shown in Figure 6.5c. At 200 voxels, the lattice is identifiable from the 3D rendering of the rasterised images and the struts are consistently well-represented, both in the 3D rendering and the image slices. This image size correlates to the start of the excellent results (i.e. $e < 10\%$) seen in Figure 6.1.

6.2.1.2 Error surges

Both the apparent modulus results and inner modulus results show distinct error surges at image sizes of 300 and 390 voxels, with the inner modulus results showing additional surges at 220 and 310 voxels. Considering the consistency of the results from 200 to 400 voxels, it seems unlikely that this error is caused by an issue in the analysis cycle, therefore the rasterised images were investigated as the source of the error.

Figure 6.7 shows enlarged areas in an image slice from the 290, 300 and 310 voxel images. Most of the struts in the Kelvin cell lattice are angled at 45° to the xy -plane (or z -axis), therefore in the rasterised image slices, the strut cross-sections should be represented as ellipses⁷ In the 290 and 310 voxel images, it is clear that the strut cross-sections are elongated along either the x - or y -axis and, therefore, approximate ellipses. However, in the 300 voxel image, the strut cross-sections are equiaxed and, therefore, more representative of circles than ovals, clearly in-

⁶Deciding when an image becomes representative of a lattice is a subjective process. An argument could be made that the 190 voxel image is also representative of the lattice.

⁷The strut cross-sections are circular, and they are sliced at an angle to their axis, meaning that the projected area in the image slice should be an ellipse.




Figure 6.7: Enlarged image slices from the 290, 300 and 310 voxel TIFF stack for the K8-t3 lattice.

dicating that there is some error associated with the rasterisation of the strut cross-sections⁸, which has propagated into the FE mesh. The most likely factor to be affected by this rasterisation error is the cross-sectional area of the beam elements; if the error associated with this area is considered, the spike in the modulus error corresponds to a clear spike in the cross-sectional area, as shown in Figure 6.8. Although this increase in cross-sectional area will certainly increase the apparent modulus, there are two problems that arise from this explanation:

1. The increase in the average cross-sectional area of the beam elements does not produce a corresponding increase in the error associated with the area. In fact, the increase in area causes a decrease in the absolute value of the error, bringing it down to only 0.228%. Why, then, is there an increase in the error of the apparent modulus?
2. Even if the previous point is disregarded, a 5% increase in the cross-sectional area should not cause a 13% increase in the apparent modulus. Rather, there should be an increase in the apparent modulus which is approximately equivalent to the percentage increase in beam cross-sectional area⁹.

Considering these two observations, it is clear that there must be some other factor which is causing the large error in apparent modulus. The rasterised images for the 300 voxel sample showed circular cross-sections in the image slices. Not only can this incorrect rasterisation affect the cross-sectional area of the elements in the FE mesh, but it is also likely to affect the moments of inertia of the beam elements as the shape of the strut is not properly represented. If the error associated with the moments of inertia in the beam elements is considered, the reason for the error in the apparent modulus becomes clear. Figure 6.9 shows the error in the average moments of inertia in the beam elements in the primary and secondary directions, respectively. There is a 14.6% increase in primary moment of inertia at 300 voxels and, similarly, an 8.61% increase in the secondary moment of inertia. These increases in error correspond to the increase in error in the apparent modulus. Additionally, a clear difference in the values of the moments of inertia is seen, implying that the shape of the cross-section of the struts is not circular, as it should be. These errors combine to produce the surge in error seen in the apparent modulus at 300 voxels.

⁸It is important to realise that there is always some error associated with the rasterisation of circular features, as it is impossible to accurately represent these features with a collection of squares, particularly at the image sizes discussed in this report. As such, the error in the rasterisation at 300 voxels is not a flaw in the rasterisation procedure used, but rather an “occupational hazard” associated with rasterisation as a concept.

⁹A baseline FE mesh with erroneous cross-sectional area was simulated to determine how much the apparent modulus should change based on a 5% increase in beam cross-sectional area. This simulation indicated that the apparent modulus should only increase by approximately 5%.




Figure 6.8: Enlarged view of graph of apparent modulus error and average beam cross-sectional area error for an uncropped, unconstrained K8-t3 lattice at varying image sizes.




Figure 6.9: Enlarged view of the error plots for the moments of inertia in the primary and secondary directions for an uncropped, unconstrained K8-t3 lattice at varying image sizes.

A similar phenomenon, though to a lesser degree, may be seen at an image size of 390 voxels. Here there is a slight decrease in the cross-sectional area of the elements ($e = -5.96\%$), which causes a corresponding decrease in the predicted apparent modulus ($e = -9.01\%$), as shown in Figure 6.8. Although there is some error associated with the moments of inertia of the elements, as shown in Figure 6.9, the error in the primary and secondary directions are almost equivalent at -8.78% and -8.69% , respectively. These nearly equivalent errors mean that although the magnitudes of the moments of inertia are slightly smaller than they should be, the shape of the struts is preserved in the FE mesh. As a result, the erroneous moments of inertia have less of an effect on the apparent modulus results than they did at 300 voxels and, therefore, the apparent modulus error is much closer to the cross-sectional area error than was seen at 300 voxels. The same reasoning may account for the error surges at 300 and 390 voxels in the inner modulus results.

6.2.1.3 Inner modulus variation

As mentioned previously, the inner modulus results show more variation than the apparent modulus results. The increased variation in results is expected, as the inner modulus calculation is based on a smaller portion of the mesh than the apparent modulus, meaning that any variation in structure (such as a strut with an incorrect placement, cross-sectional area or shape) will have a more significant effect on the inner modulus, than the apparent modulus. In some cases, such as at 310 voxels, the inner modulus may show an increase in error even if the apparent modulus does not. Here the structure as a whole may be representative of the Kelvin cell lattice, hence the accurate apparent modulus, however locally there may be variations in the structure which will be captured by the inner modulus.

An additional source of error for the inner modulus was identified. Figure 6.10 shows an illustration of the nodes on the bottom surface of the baseline mesh and the 220 voxel mesh created by the analysis cycle. In these images, the nodes used for the calculation of the inner modulus are shaded red. It is clear that there are some structural deviations in the analysis cycle mesh (Figure 6.10a), where only three nodes are grouped together on the bottom surface instead of four. As such, although the 220 voxel mesh captures the essence of the lattice, these deviations, which lie particularly in the inner nodes region, will effect the inner modulus and cause the error surge seen in Figure 6.2.

6.2.2 Variation due to lattice size

Although much of this section focuses on the apparent modulus results, an interesting variation in the modulus ratio was observed in the K6 and K8 lattices which requires discussion. As such, this section will detail the variation in the apparent modulus as well as the modulus ratio with a change in size of the Kelvin cell lattice.

6.2.2.1 Apparent modulus

The variation in apparent modulus error for the K6-t3, K8-t3 and K10-t3 lattices is shown in Figure 6.11. Excellent results are produced for all three lattice sizes, with an error of less than 1% at large image sizes¹⁰. A pattern may be seen in the results produced for each lattice

¹⁰Excluding error surges.

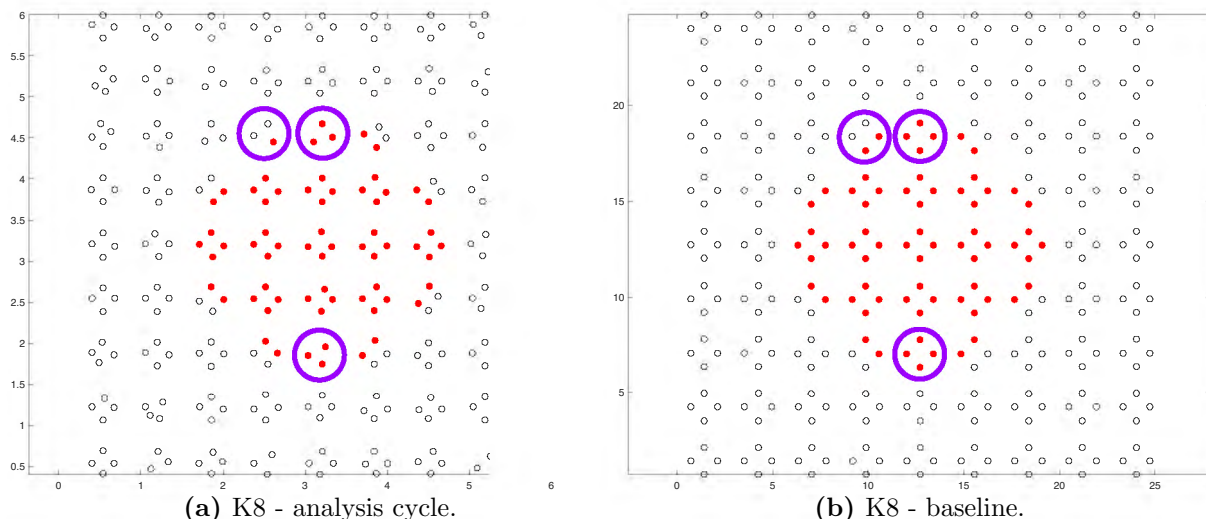


Figure 6.10: Illustration of the nodes used to calculate the inner modulus for the K8-t3 analysis cycle simulation and baseline simulation for an image size of 220 voxels.


size: the error starts off being a high positive value, after which there is a sharp decrease to a maximum negative error, followed by a trend towards 0% error. If the rasterised images are considered, it is seen that representative images are only produced at 160 voxels for K6, 200 voxels for K8 and 240 voxels for K10, which corresponds well to the start of the upwards trend to 0% error in Figure 6.11. This clear pattern shows that the analysis cycle is producing consistent results, and supports the idea that the high error at low image sizes is due to poor rasterisation, not a fault in the analysis cycle.

The discretisation failures for each lattice are not shown in Figure 6.11, but still merit discussion. As the lattice size increases, the number of discretisation failures increases, with the K6 lattice showing no failures, the K8 lattice showing 2 failures and the K10 lattice showing 7 failures. If a particular image size is considered, the K8 and K10 rasterisation is poorer than the K6 rasterisation because the same number of voxels are being used to represent a more complex structure, as shown in Figure 6.12. As such, the image quality decreases, which causes the increase in the number of failures at small image sizes.


Another feature which is consistent in all the lattice sizes, is the presence of error surges, which were discussed in Section 6.2.1.2. In the K6 lattice, these error surges occur at 280 and 350 voxels, whereas the K8 lattice shows error surges at 300 and 390 voxels and the K10 lattice shows a small error surge at 340 voxels. All of these error surges may be attributed to incorrect image rasterisation resulting in erroneous cross-sectional area, cross-sectional shape and/or moments of inertia¹¹. The reasons for the error surges are summarised in Table 6.1¹².

¹¹Note: Erroneous “cross-sectional shape” indicates that the primary moment of inertia is significantly different in magnitude to the secondary moment of inertia, whereas erroneous “moments of inertia” indicates that both moments of inertia show significant error.

¹²These observations were made from the cross-sectional area and moment of inertia graphs for each lattice. The results for the K8-t3 lattice are shown in Figures 6.8 and 6.9; the results for the K6-t3 lattice and K10-t3 lattice are included in Appendix E, in Figures E.1-E.2 and Figures E.3-E.4, respectively.



(a) Full set of results.



(b) Enlarged view.

Figure 6.11: Apparent modulus error plots for the uncropped, unconstrained K6-t3, K8-t3 and K10-t3 lattices at varying image sizes. Failures for the various lattice sizes have been excluded from the graph for the sake of clarity.




Figure 6.12: 3D renderings of the K6-t3, K8-t3 and K10-t3 lattices at an image size of 250 voxels. Clearly, as the lattice size increases, the quality of the image decreases.



Table 6.1: Reason for the error surges seen in the uncropped, unconstrained K6-t3, K8-t3 and K10-t3 lattices. Issues include: erroneous cross-sectional (CS) area, cross-sectional (CS) shape or moments of inertia.

Lattice size	Image size (voxels)	CS area	CS shape	Moments of inertia
K6	280		X	
	350	X		X
K8	300		X	X
	390	X		X
K10	340	X		X

6.2.2.2 Modulus ratio

Enlarged regions of the modulus ratio results for the K6 and K10 lattices are shown in Figure 6.13. The K10 results (Figure 6.13b) show more variation than the K8 results, displayed in Figure 6.3, however the most surprising results stem from the K6 lattice (Figure 6.13a), which shows a relatively consistent offset in the predicted ratio when compared to the baseline ratio. The apparent modulus results for all three sizes of the lattice showed consistently good results, which means that the error in the modulus ratio must be caused by the calculation of the inner modulus.

Figure 6.14 shows illustrations of the nodes on the bottom surface of a K6, K8 and K10 lattice, respectively, where the nodes used for the calculation of the inner modulus (so-called “inner nodes”) are represented by filled red markers. Figures 6.14a, 6.14c and 6.14e show the nodes output by the analysis cycle, whereas Figures 6.14b, 6.14d and 6.14f show the nodes used for the baseline simulations of the corresponding lattices. In the K10 lattice (Figure 6.14e and Figure 6.14f), the inner nodes output by the analysis cycle are almost identical to the inner nodes used for the baseline simulation. In the K8 lattice (Figure 6.14c and Figure 6.14d), there are clusters of 5 nodes instead of the correct 4 nodes, which may cause some variation in the results, but the cylindrical shape of the inner region is maintained. However, for the K6 lattice (Figure 6.14a and Figure 6.14b), the region used to determine the inner modulus in the analysis cycle is rectangular due to missing nodes as indicated, whereas the region used for the baseline simulation is the correct circular shape. A decrease in the number of inner nodes results in a decreased inner modulus and, consequently, the decreased modulus ratio seen in Figure 6.13a.



(b) K10.

Figure 6.13: Modulus ratio results for the uncropped, unconstrained K6-t3 and K10-t3 lattices at varying image sizes.




Figure 6.14: Illustration of the nodes used to calculate the inner modulus for the K6, K8 and K10 lattices at an image size of 350 voxels.




Figure 6.15: Apparent modulus error results for the uncropped, unconstrained K8-t1 lattice at varying image sizes.

The variation in the assignment of the inner nodes stems from a variation in the nodal coordinates. The inner node region is determined by finding the centre and radius of the inner core of the lattice, which are calculated directly from the nodal coordinates. If there is a small error in the nodal coordinates, it will affect the placement of the centre point which, in turn, may cause some of the inner nodes to be excluded from the calculation, particularly because the lattice nodes are so sparse. The radius of the inner core may be similarly affected, which could cause a decrease in the size of the inner region, causing some inner nodes to be missed. Although these issues may cause a deviation in the lattice results, it is unlikely that the same problem will be experienced with more complex structures, like bone, because the meshes will contain more nodes which will be more densely packed on the bottom surface. In this case, if a few nodes are missed during the determination of the inner node region, the inner modulus should not be affected as significantly as in the lattices. Nevertheless, it is important to take note of this possible source of error.

6.2.3 Variation due to strut diameter

Figure 6.15 shows the apparent modulus error plot for the K8-t1 lattice at varying image sizes. There is a significant increase in the number of specimens which failed to discretise during the analysis cycle, with the K8-t1 lattices showing 23 failures where the K8-t3 lattices showed only two. Additionally, the error in the K8-t1 results is over 20% for the majority of the successful specimens, even at large image sizes. Although these results are poor, they are easily explained by the quality of the corresponding rasterised images.

Examples of rasterised K8-t1 images, ranging from 100 to 400 voxels, are displayed in Figure 6.16. It is clear that at this particular lattice size and strut thickness, the rasterised images




Figure 6.16: 3D renderings of K8-t1 rasterised images at 100, 200, 300 and 400 voxels, respectively. Although the quality of the rasterisation increases with increasing image size, at a strut diameter of 0.1, the rasterisation never becomes representative of the K8 lattice.

are not representative of the Kelvin cell lattice at any image size as they are indiscernible in the 3D rendering of the lattice. The extremely thin struts in the rasterised images lead to an increase in the number of failures for the lattice and vastly increases the error associated with any successful specimens, as discussed in Section 6.2.1.1. As such, the K8-t1 lattice does not provide a useful validation test for the analysis cycle, as any error associated with the results is clearly due to the inadequate rasterisation.

Figure 6.17 demonstrates that the K8-t2 results for apparent modulus are better than the K8-t1 results (Figure 6.15), both in frequency of successful discretisation, as well as error in apparent modulus results. Although the error at small image sizes is high, this error decreases as the image size increases, and stabilises at an image size of approximately 300 voxels in the 10% error zone. This trend is similar to the trend seen in the K8-t3 results (Figure 6.1), where an increasing image size resulted in better rasterisation of the lattices and, therefore, more accurate predictions of apparent modulus.

Another familiar feature which may be found in the K8-t2 graph (Figure 6.17), is the two surges in error at 360 and 390 voxels. At these image sizes, the error increases from approximately 1-2% to -12.5% and -10.0% , respectively. Once again, these error surges may be attributed to the rasterisation process, where a slight error in the strut cross-section rasterisation carries through to become an error in the average cross-sectional area and moments of inertia of the beam elements in the FE mesh, as shown in Figures 6.18 and 6.19. Although the cross-sectional shape is unaffected, there is a decrease in the cross-sectional area of the struts and the magnitude of the moments of inertia, which will cause a corresponding decrease in the predicted apparent modulus, as seen in Figure 6.17. Despite these variations in cross-sectional properties, the surges in the absolute error of the apparent modulus are only slightly more than 10%, which is a commendable result.

The K8-t2 results (Figure 6.17) are significantly better than the K8-t1 results (Figure 6.15), however the K8-t3 results (Figure 6.1) are the best of the three sets. When K8-t1, K8-t2 and K8-t3 rasterised images are compared to one another, as shown in Figure 6.20, it is clear that the variation in the results at the different strut diameters is due to the image rasterisation and not the analysis cycle. The K8-t3 lattice images are significantly more representative than either the K8-t1 or K8-t2 images, leading to a greater number of successful simulations and increased accuracy in the results.




Figure 6.17: Apparent modulus error results for the uncropped, unconstrained K8-t2 lattice at varying image sizes.




Figure 6.18: Enlarged view of graph of apparent modulus error and average beam cross-sectional area error for an uncropped, unconstrained K8-t2 lattice at image sizes of 300 to 400 voxels.




Figure 6.19: Enlarged view of the error plots for the moments of inertia in the primary and secondary directions for an uncropped, unconstrained K8-t2 lattice at image sizes of 300 to 400 voxels.




Figure 6.20: 3D renderings of the K8-t1, K8-t2 and K8-t3 lattices at an image size of 350 voxels.




Figure 6.21: Comparison of the error in apparent modulus for the uncropped, unconstrained and cropped, unconstrained K8-t3 lattices at varying image sizes.

6.2.4 Variation due to aspect ratio

Figure 6.21 shows a comparison of the apparent modulus error for the K8-t3 uncropped lattice and the K8-t3 cropped lattice. It is clear from the graph that the analysis cycle provided excellent predictions of the apparent modulus for the cropped lattice, similar to the excellent results produced for the uncropped lattice. Additionally, the cropped lattice results show a similar surge in error to the uncropped lattice results at 300 voxels and, to a lesser extent, at 390 voxels, reinforcing the conclusion that the rasterised images are the reason for the increased error, not any aspect of the analysis cycle itself.

In addition to the excellent prediction of apparent modulus, the analysis cycle also consistently predicted the modulus ratio of the cropped specimens to within 10% of the baseline value as shown in Figure 6.22. Moreover, all except one of the predicted ratios was consistently above 1 indicating that the inner core is stiffer than the lattice as a whole, which is the expected result. Although the modulus ratio at 180 voxels is less than 1, this result is below the image size of 200 voxels, which was determined to be the smallest representative image in the set.

6.2.5 Variation due to constraint

Figure 6.23 shows the apparent modulus error plot for the uncropped, constrained K8-t3 lattice. This graph closely resembles the unconstrained apparent modulus graph exhibited in Figure 6.1, even displaying the same error surges at 300 and 390 voxels. Similarly, Figure 6.24, which shows the results for the modulus ratio of the uncropped, constrained lattice, closely resembles Figure 6.3. Both the apparent modulus and modulus ratio results predicted by the analysis cycle are




Figure 6.22: Modulus ratio results for a cropped, unconstrained K8-t3 lattice at varying image sizes.

within 10% of the baseline results for image sizes of 200 voxels and above¹³, indicating that the analysis cycle is capable of accurately capturing the response of the K8-t3 lattice under different constraints.

Tables 6.2 and 6.3 show the baseline results for apparent modulus and modulus ratio of the uncropped unconstrained, cropped unconstrained, uncropped constrained and cropped constrained K8-t3 lattices. As expected, the cropped lattice is significantly more sensitive to the constraint than the uncropped lattice, showing a 140% increase in apparent modulus and a 34.7% increase in modulus ratio when constrained. In contrast, the uncropped lattice only showed an 18.6% increase in apparent modulus and showed a decrease in modulus ratio of 14.3% in the constrained condition. Considering there was a much larger change in baseline response for the cropped constrained lattices than there was in the uncropped constrained lattices, it is important to investigate how successfully, if at all, the analysis cycle managed to capture the behaviour of the cropped lattices in the constrained condition.

Figures 6.25 and 6.26 show the apparent modulus and modulus ratio results for the cropped, constrained K8-t3 lattice. Despite the fact that the apparent modulus for the constrained lattice was significantly higher than that of the unconstrained lattice, the analysis cycle performed extremely well and predicted the apparent modulus to within 10% for image sizes of 250 voxels and above. Furthermore, the modulus ratio of the cropped, constrained lattice was also accurately predicted, with the analysis cycle producing ratios within 10% of the baseline ratio for all images sizes greater than 200 voxels. Moreover, the error in the modulus ratio decreases to less than 1% for image sizes greater than 360 voxels, an excellent result in any circumstances,

¹³The modulus ratio at 310 voxels lies slightly outside the 10% error region. This is due to improper image rasterisation as described in Section 6.2.1.2, therefore should not be taken as a reflection on the accuracy of the analysis cycle.




Figure 6.23: Apparent modulus error plot for an uncropped, constrained K8-t3 lattice at varying image sizes.




Figure 6.24: Modulus ratio results for an uncropped, constrained K8-t3 lattice at varying image sizes.

Table 6.2: Baseline results for the apparent modulus of the uncropped and cropped K8-t3 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained apparent modulus from unconstrained apparent modulus.

	Unconstrained (MPa)	Constrained (MPa)	Increase (%)
Uncropped	2.680	3.178	18.58
Cropped	2.448	5.886	140.4

Table 6.3: Baseline results for the modulus ratio of the uncropped and cropped K8-t3 lattices in the unconstrained and constrained conditions, including the percentage increase in ratio from the unconstrained to constrained condition.

	Unconstrained	Constrained	Increase (%)
Uncropped	1.090	0.9345	-14.27
Cropped	1.067	1.437	34.68

but especially remarkable if one bears in mind that the analysis cycle generated these results from a series of images only.

6.2.6 Correlation to VBM

Figure 6.27 shows the apparent modulus results for the cropped, unconstrained K8-t3 voxel-based simulation¹⁴. The VBM results show high negative error at small image sizes, similar to that seen in the BSM results. The error steadily increases as the image size increases, passing through the 10% and 20% error zones, and stabilising at an error of approximately 20%. The VBM apparent modulus error is much higher than the 5% or less error readings seen in the BSM simulations of the same cropped, unconstrained lattice¹⁵. A similar trend is seen in the unconstrained inner modulus VBM results shown in Figure 6.28, however there is no clear stabilisation in the results at high image sizes, as the error continues to increase as the image size increases.

Despite the fact that the apparent modulus and inner modulus VBM results show a higher error compared to the corresponding BSM results, the modulus ratios predicted by the VBM are excellent, generally showing an error of less than 10% at representative image sizes as shown in Figure 6.29. A similar phenomenon is seen when the cropped specimens are constrained; the apparent and inner moduli for the cropped constrained lattice fall outside the 20% error zone, however the corresponding modulus ratio results are excellent, consistently showing errors less than 10% for image sizes greater than 150 voxels, as shown in Figure 6.30.

How is it that the modulus ratio can be accurate even when the apparent modulus and inner modulus show errors above 20%? As mentioned previously in Section 6.2.1.2, there are inherent errors in the rasterisation process, particularly where circular features are concerned. This phenomenon is illustrated in Figure 6.31. In a binary TIFF stack, each voxel must be designated

¹⁴Note: Additional rotational constraints, R_x and R_y , were added to the baseline simulations for the VBM simulations to mimic the rotational constraints enforced by the VBM. The reasoning for this change in boundary conditions is discussed in Appendix G.

¹⁵It may be tempting to attribute the high error in the VBM results to incorrect baseline results, however the code by Moore [64] which was used to generate the baseline results was thoroughly tested as discussed in Section 4.2, therefore it is unlikely that the baseline results are incorrect.




Figure 6.25: Apparent modulus error plot for a cropped, constrained K8-t3 lattice at varying image sizes.




Figure 6.26: Modulus ratio results for a cropped, constrained K8-t3 lattice at varying image sizes.



Figure 6.27: Enlarged view of the apparent modulus error plot for a cropped, unconstrained K8-t3 lattice at varying image sizes discretised using the VBM.




Figure 6.28: Enlarged view of the inner modulus error plot for a cropped, unconstrained K8-t3 lattice at varying image sizes discretised using the VBM.




Figure 6.29: Modulus ratio results for a cropped, unconstrained K8-t3 lattice at varying image sizes discretised using the VBM.




Figure 6.30: Modulus ratio results for a cropped, constrained K8-t3 lattice at varying image sizes discretised using the VBM.

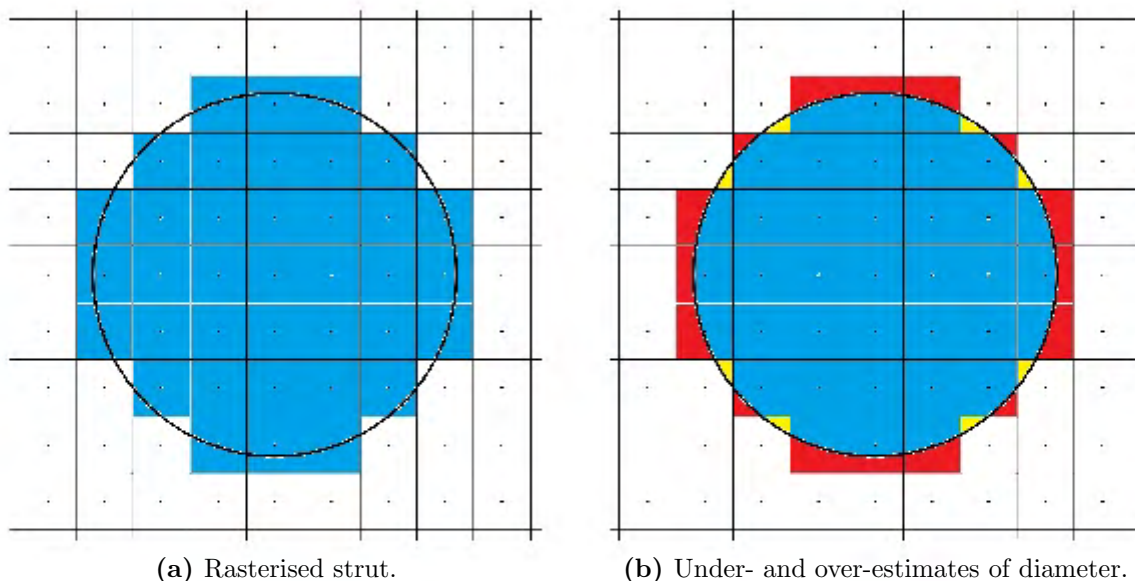


Figure 6.31: Example of under- and over-estimations in strut rasterisation. (a) shows the rasterisation (blue voxels) of a strut with circular cross-section, whereas (b) highlights the regions of under- (yellow) and over-estimation (red) of the strut diameter as a result of the rasterisation.

a value of 0 (non-bone) or 1 (bone), even if the voxel in question is only partially filled by a structure (Figure 6.31a). As such, rasterised images contain both under- and overestimates of cross-sectional area (Figure 6.31b), where partially filled voxels have been designated values of 0 and 1, respectively. Because the VBM involves designating every bone voxel as a solid element, these under- and overestimations may make a significant difference to the apparent modulus and inner modulus results, however the inner modulus relative to the apparent modulus (i.e. the modulus ratio) should not be affected because the rasterisation error is contained within both sets of results.

Interestingly, the results seem to indicate that the BSM diminishes the effect of the rasterisation error seen in the VBM results. Where the VBM uses each voxel in the TIFF stack as a building block for the FE mesh, the BSM infers the structure of the mesh by considering the group of voxels that make up a strut and assigning corresponding cross-sectional properties based on the overall cross-sectional representation of the strut. This more cohesive approach minimises the individual voxel rasterisation errors, which produces a more consistent and, in this case, accurate result.

6.2.7 Summary and significance of results

Overall, the analysis cycle provided excellent results for the K8 lattice in all its various topologies and test configurations. The error associated with the prediction of apparent modulus decreased as the image size, and therefore resolution, increased and generally stabilised within the 10% error region.

In all of the configurations, the error in the analysis cycle predictions was high at small image sizes. However, this can be shown to result from unrepresentative images being used as input to the analysis cycle. If the rasterised images are considered, and an image size chosen which produces a representative image, the results tend to show reasonable error (less than 20%) and,

at higher image sizes, very little error (within 10%). As such, it can be concluded that the error at low image sizes is not due to an error in the analysis cycle, but rather improper rasterisation.

Noticeable surges in error, where the absolute error increases significantly in relation to the adjacent results, appear in most of the K8 graphs. These error surges have been shown to correlate with incorrect rasterisation of the cross-sectional area of the struts, which may lead to incorrect beam cross-sectional properties in the FE mesh. Once again, this error may be attributed to variations in the rasterisation of the images, rather than the analysis cycle.

The results from the various test configurations are summarised as follows:

- **Variation due to lattice size:** All three lattice sizes showed a similar pattern in their apparent modulus error plots. Excellent results were consistently produced for the various sizes once a representative image was generated, with the error in the apparent modulus decreasing to less than 1% at large image sizes.
- **Variation due to strut diameter:** The rasterisation of the lattices with small strut diameters (namely 0.1 and 0.2) was not adequate, resulting in a significant increase in discretisation failures and erroneous results. If the rasterised image was representative of the lattice, the analysis cycle provided good predictions for the apparent modulus, where the error was within 10% at large image sizes.
- **Variation due to aspect ratio:** The analysis cycle produced excellent results for the cropped lattices, which were almost identical to the equivalent uncropped lattice.
- **Variation due to constraint:** The analysis cycle captured both the apparent modulus and modulus ratio of the constrained lattices for the uncropped and cropped conditions. Results were generally within 10% error and, for large image sizes, within 2% error.
- **Correlation to VBM:** The VBM produced less accurate apparent and inner moduli than the BSM, but captured the modulus ratio as accurately as the BSM for both the cropped unconstrained and cropped constrained lattices.

The modulus ratio results for the cropped, constrained K8-t3 lattice are possibly the most significant results produced in the validation study. These results show that the analysis cycle is capable of predicting the change in apparent modulus and inner modulus when the lattice is constrained versus when it is unconstrained, even when the change in results is significant. **The conclusion that can be drawn from these results is that the microstructural Finite Element model used in the analysis cycle is definitively capable of predicting the ratio of the inner modulus to the apparent modulus of the Kelvin cell lattices.**

6.3 Notable features in the cubic and octet truss lattices

Although the Kelvin cell lattice in all its various topologies provided important information regarding the functioning of the analysis cycle, the cubic and octet truss lattice results provided new information about both the rasterisation process and the analysis cycle. A full analysis of the results will not be exhibited for these lattices, however an overview of the various key features in each result set, as well as the reason for these results, is provided.

As with the Kelvin cell lattice, a “starting lattice” was chosen for the cubic and octet truss lattices, and used as a basis for comparison when discussing the variation in results due to lattice size, constraint, etc. The choice of a starting lattice was based on the quality of the image rasterisation, as it is of no value analysing results from the analysis cycle that were produced based on unrepresentative images. As such, the strut thickness which produced the best image rasterisation was determined and the corresponding uncropped, unconstrained size 8 lattice was used as the starting lattice. The chosen starting lattices are: C8-t1 (cubic) and OT8-t2 (octet truss).

6.3.1 Cubic lattice

Overall, the cubic lattice showed excellent agreement to the baseline results, however there were some notable results produced when the strut diameter of the lattices was increased and when a constraint was applied to the lattice. A brief discussion of these phenomena is provided in Sections 6.3.1.1-6.3.1.4.

6.3.1.1 General

Figure 6.32 shows the apparent modulus results produced for the uncropped, unconstrained C8-t1 lattice. Similar to the Kelvin cell results, the error in the apparent modulus is significant at small image sizes, but quickly tends down to the 10% error zone, where it remains for the majority of image sizes greater than 170 voxels. The first good result, seen at 170 voxels, correlates to the first representative rasterisation of the lattice (see Appendix F for rasterised images).

Familiar surges in the error may be seen at 210, 280 and 350 voxels, where the error increases from a value less than 10% to -18.2% , -20.8% and 11.8% , respectively. The error in the apparent modulus is due to the error in cross-sectional area of the beam elements, as shown in Figure 6.33, and a corresponding error in the moments of inertia¹⁶, as shown in Figure 6.34. At 210 and 280 voxels, the cross-sectional area is underestimated, while at 350 voxels, the cross-sectional area is overestimated. The underestimation of the cross-sectional area at 210 and 280 voxels and the overestimation of the area at 350 voxels can be shown to arise from the rasterised images by considering the average number of voxels representing the cross-sections in each image slice¹⁷, as exhibited by Table 6.4. Clearly, just as in the Kelvin cell lattices, the error surges are due to the rasterised images, not the analysis cycle.

6.3.1.2 Variation due to strut diameter

Of the three lattices used for the validation study, the cubic lattice is the simplest, therefore it was expected that it would show an excellent correlation to the baseline results in all cases. However, when the strut diameter of the lattice was increased, there was a significant increase in the error of the apparent modulus at large image sizes, as shown in Figure 6.35. This error in apparent modulus directly correlates to the erroneous inclusion of shell elements in the mesh, particularly on the top and bottom surfaces, as shown in Figure 6.36.

¹⁶Here there is some error in the moments of inertia in both the primary and secondary directions, however these errors are approximately equivalent, meaning that the shape of the cross-section is correct.

¹⁷The average was calculated based on the four struts in the top left corner of the first image slice through the vertical struts.




Figure 6.32: Apparent modulus error results for an uncropped, unconstrained C8-t1 lattice at varying image sizes.




Figure 6.33: Enlarged view of the apparent modulus error and element cross-sectional area error results for an uncropped, unconstrained C8-t1 lattice at varying image sizes.




Figure 6.34: Enlarged view of the error in the primary and secondary moments of inertia for an uncropped, unconstrained C8-t1 lattice at varying image sizes.

Table 6.4: Average number of voxels representing strut cross-sections in the C8-t1 lattice rasterisations for image sizes of 200, 210, 220, 270, 280, 290, 340, 350 and 360. The image sizes corresponding to error surges clearly show an underestimation (in the case of 210 and 280 voxels) or overestimation (350 voxels) in the cross-sectional area of the struts.

	One image size smaller $x - 10$	Image size x	One image size larger $x + 10$
$x = 210$ voxels	5.25	5	6.25
$x = 280$ voxels	10.25	9	11.25
$x = 350$ voxels	17.25	21	17.25

Two questions arise from these results:

1. Why are shell elements included in the mesh, when the lattice is made up, entirely, of beam-like struts?
2. Why does the inclusion of shell elements cause such a large error in the apparent modulus results?

The first question may be answered by considering the rasterised images used as input to the analysis cycle. Figure 6.37 shows the first image slice for the C8-t2 lattice at image sizes of 230, 240, 300 and 310 voxels. The first slice of the TIFF stack for the 240 and 310 voxel images show struts which are substantially thinner than the struts in the 230 and 300 voxel images. This occurs because as the image size increases, the image resolution increases, which means that there are more image slices through the lattice and, consequently, through each strut. For




Figure 6.35: Apparent modulus error results for an uncropped, unconstrained C8-t2 lattice at varying image sizes.




Figure 6.36: Number of shell elements present in the FE mesh for an uncropped, unconstrained C8-t2 lattice at varying image sizes.




Figure 6.37: 3D renderings of the C8-t2 lattice at image sizes of 230, 240, 300 and 310 voxels.

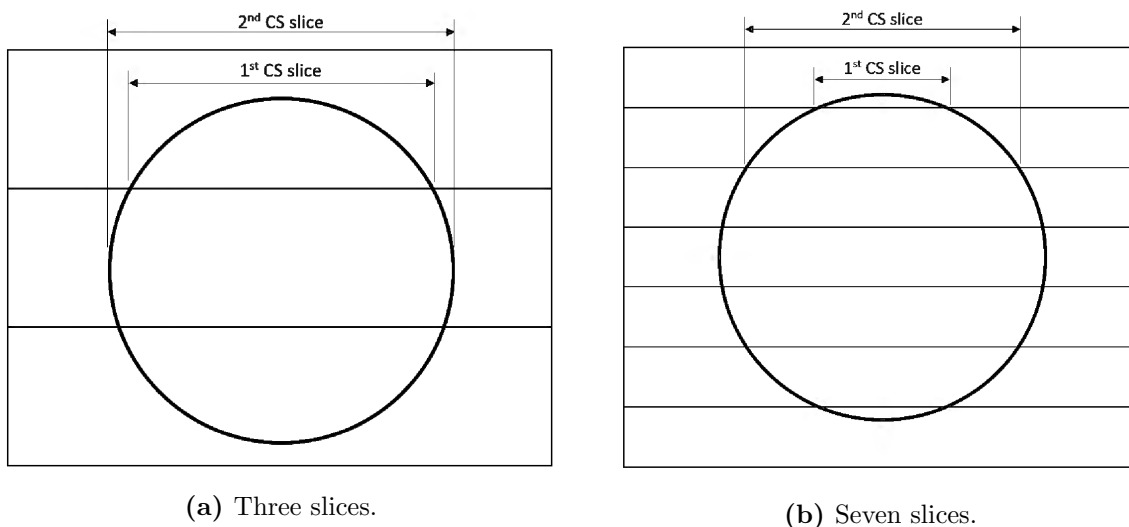


Figure 6.38: Image slices through a horizontal strut in a cubic lattice.

example, consider the rasterisation of the top row of struts in the cubic lattice. In the 230 voxel image, these struts are rasterised using five image slices, but in the 240 voxel image, the rasterisation covers six image slices. Similarly, the 300 voxel rasterisation takes up seven image slices, whereas the 310 voxel rasterisation covers eight. In both of these cases, when the number of image slices through the top strut increases, the volume of the strut is spread over more image slices and, therefore, the thickness of the strut in the first image slice will decrease as illustrated in Figure 6.38. So as the image resolution of the cubic lattice increases, the rasterisation in the first slice gets progressively thinner until, at a given point, the first slice is so thin in comparison to the other slices that the discretisation code identifies this section of the strut as a separate, very thin plate-like strut. As a result, shell elements with negligible thickness are mistakenly allocated to this structure and, in some lattices, this error carries through so that the entire top layer of struts is discretised into shell elements.

Now that the presence of the shell elements can be explained, it is important to discuss why the shell elements create such a large error in the apparent modulus results. This can be attributed to two root causes, namely incorrect structural representation and the removal of the shell elements with zero thickness. When shells are included in the mesh, the lattice structure is fundamentally misrepresented, meaning that a simulation using this mesh could not possibly yield the correct apparent modulus. This incorrect mesh is exacerbated by the removal of the shell elements with zero thickness, as discussed in Section 3.2.2.2, particularly because most, if not all, of the shell elements created on the top surface have zero thickness. When these shells are removed, a final mesh is created which is missing entire sections of the original lattice




Figure 6.39: Representation of the mesh created for a C8-t2 lattice at an image size of 350 voxels. The bottom surface of the mesh is made up entirely of shell elements, and the top surface of the mesh is non-existent due to the removal of the shell elements with zero thickness.

structure, as shown in Figure 6.39. This mesh is not representative of the structure and, hence, cannot accurately predict the apparent modulus of the structure. Moreover, the ties included on the top surface of the mesh will experience large deformation, which will erroneously lower the apparent modulus.

The code for the removal of zero thickness shell elements was created after several of the cubic lattices failed to simulate. By removing these elements, the simulation runs successfully, but produces the erroneous results seen in Figure 6.35. As such, the initial failure in simulation may be taken as a sign that the lattice has not been properly discretised. Hence, the shell elements with zero thickness should not be removed from the mesh, because the removal facilitates the simulation of improperly discretised structures¹⁸.

6.3.1.3 Variation due to constraint

Tables 6.5 and 6.6 show the baseline results for apparent modulus and modulus ratio for the C8-t1 lattice in various configurations. As expected, the addition of a constraint had no effect on either the apparent modulus or modulus ratio. This insensitivity to confinement is captured in the apparent modulus results produced by the analysis cycle for both the uncropped constrained and cropped constrained lattices, the latter of which is shown in Figure 6.40. The modulus ratios predicted by the analysis cycle also capture this insensitivity to confinement, with both the uncropped constrained and cropped constrained lattices providing excellent correlation to the baseline result. The cropped constrained modulus ratio results are consistently within the 10% error zone at large image sizes, as shown in Figure 6.41.

¹⁸The code used to remove the zero thickness shells was not discarded until after the bone specimens had been analysed. However, no zero thickness shells were produced during the discretisation of the bone specimens, so the removal of the code did not affect the bone results.

Table 6.5: Baseline results for the apparent modulus of the uncropped and cropped C8-t1 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained apparent modulus from unconstrained apparent modulus.

	Unconstrained (MPa)	Constrained (MPa)	Increase (%)
Uncropped	10.26	10.26	0
Cropped	10.26	10.26	0

Table 6.6: Baseline results for the modulus ratio of the uncropped and cropped C8-t1 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained modulus ratio from unconstrained modulus ratio.

	Unconstrained	Constrained	Increase (%)
Uncropped	0.9549	0.9549	0
Cropped	0.9549	0.9549	0

Although these results may not seem as significant as, for example, the poor results produced by the inclusion of shell elements in the mesh, the consistent results produced by the analysis cycle should not be overlooked. The stiffness of the cubic lattice is governed solely by its perfectly vertical struts; if these vertical struts are angled only slightly from the vertical, the stiffness of the cubic lattice drops significantly and the lattice becomes sensitive to constraint [64, 68]. The fact that the analysis cycle is capable of consistently capturing both the apparent modulus of the lattice as well as its insensitivity to confinement means that a good quality mesh is being produced during the discretisation process and, consequently, lends a great deal of confidence to the accuracy of the analysis cycle.

6.3.1.4 Variation due to angle

The angled cubic lattice, as with the Kelvin cell lattice, is highly sensitive to constraint, therefore it is important to check whether the analysis cycle is capable of capturing the change in response of this lattice when it is constrained. Tables 6.7 and 6.8 show the baseline results for apparent modulus and modulus ratio for the angled C8-t2 lattice¹⁹ in various configurations. Clearly, the angled cubic lattice is significantly more sensitive to constraint than the Kelvin cell lattice, showing a 119% increase in apparent modulus when the uncropped lattice is constrained, and a massive 675% increase when the constraint is applied to the cropped lattice. The change in the modulus ratio with constraint is not as significant as the apparent modulus, however there is an 80.3% decrease in the modulus ratio of the uncropped lattice when the constraint is applied, which should not be underestimated.

The apparent modulus results of the uncropped, unconstrained angled C8-t2 lattice are excellent with almost all the results above the representative image size of 150 voxels falling within the 10% error zone as shown in Figure 6.42. Error surges may be seen at 230 voxels and 340 voxels, which may be attributed to incorrect cross-sectional properties of the struts stemming from rasterisation. An additional increase in error is experienced at 360 and 370 voxels, where shell elements are erroneously included in the meshes. Despite these departures in accuracy, the

¹⁹C8-t2 was chosen here instead of C8-t1, because the rasterisation for the angled C8-t1 lattice is relatively poor as a result of thin struts caused by the angle.




Figure 6.40: Apparent modulus results for a cropped, constrained C8-t1 lattice at varying image sizes.




Figure 6.41: Modulus ratio results for a cropped, constrained C8-t1 lattice at varying image sizes.

Table 6.7: Baseline results for the apparent modulus of the uncropped and cropped angled C8-t2 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained apparent modulus from unconstrained apparent modulus.

	Unconstrained (MPa)	Constrained (MPa)	Increase (%)
Uncropped	1.856	4.066	119.1
Cropped	1.790	13.87	674.9

Table 6.8: Baseline results for the modulus ratio of the uncropped and cropped angled C8-t2 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained modulus ratio from unconstrained modulus ratio.

	Unconstrained	Constrained	Increase (%)
Uncropped	1.061	0.2086	-80.34
Cropped	0.7797	0.9266	18.84

results, overall, are very good. The results for the uncropped, constrained angled C8-t2 lattice shown in Figure 6.43 are similar to the unconstrained lattice, and fall mostly within the 10% error region. In both the unconstrained and constrained results, there is a very slow decrease in the apparent moduli with image size, which is due to an increasing number of unnecessary nodes in the mesh. This phenomenon is discussed in detail in Section 6.3.2.1. Overall, the analysis cycle has captured the 120% increase in the apparent modulus results of the uncropped lattices.

The apparent modulus results of the cropped, unconstrained angled C8-t2 lattice, displayed in Figure 6.44, show similar accuracy to their uncropped counterparts, which were shown in Figure 6.42. In contrast, the cropped, constrained, angled C8-t2 lattice results shown in Figure 6.45 seem to oscillate around -15% error, which is higher error than was seen in the unconstrained simulations. Although this higher error is larger than errors previously reported in this chapter, the fact that the analysis cycle captured the cropped, unconstrained lattice response to within 10% and the cropped, constrained lattice response to within 15% is still commendable, considering the difference in results is in the order of 600% (see Table 6.7).

The modulus ratio results for the uncropped, unconstrained lattice show more variation than the ratios seen in, for example, the Kelvin cell lattice. However, most of the results still fall within the 10% error region and, therefore, provide good agreement to the baseline result as shown in Figure 6.46. The uncropped, constrained lattice results, for the most part, do not fall within the 10% error zone, as shown in Figure 6.47. However, the modulus ratios are not significantly different from the baseline result of 0.209, with predicted ratios falling between 0.2 and 0.3 for representative images. Although the exact values of the modulus ratio are not predicted accurately by the analysis cycle, the significant decrease from the unconstrained to the constrained results is captured.

The modulus ratio results for the cropped, unconstrained lattice are mostly satisfactory, as shown in Figure 6.48. Although there are some ratios which do not fall within the 10% error zone, these ratios are in the region of 0.9, which is still relatively close to the baseline value of 0.780. The modulus ratio results for the cropped, constrained lattice are shown in Figure 6.49. Because the apparent modulus results for this lattice showed higher error than the other lattices, it was expected that the modulus ratios predicted by the analysis cycle for this lattice




Figure 6.42: Apparent modulus results for an uncropped, unconstrained angled C8-t2 lattice at varying image sizes.




Figure 6.43: Apparent modulus results for an uncropped, constrained angled C8-t2 lattice at varying image sizes.




Figure 6.44: Apparent modulus results for a cropped, unconstrained angled C8-t2 lattice at varying image sizes.




Figure 6.45: Apparent modulus results for a cropped, constrained angled C8-t2 lattice at varying image sizes.




Figure 6.46: Modulus ratio results for an uncropped, unconstrained angled C8-t2 lattice at varying image sizes.




Figure 6.47: Modulus ratio results for an uncropped, constrained angled C8-t2 lattice at varying image sizes.




Figure 6.48: Modulus ratio results for a cropped, unconstrained angled C8-t2 lattice at varying image sizes.

would not be as accurate. Although some of the predicted ratios fall within the 10% error zone, most of the predicted ratios are above 1, meaning that the inner core is stiffer than the whole specimen, which does not correlate with the baseline result. However, there is an increase from unconstrained to constrained results of approximately 20%, which agrees with the baseline values.

Despite the departure of the results from the 10% error zone, it can be said that the analysis cycle performed well. The analysis cycle accurately predicted the apparent moduli of all of the configurations, except the cropped constrained lattice, and captured the change in apparent moduli between unconstrained and constrained configurations. The modulus ratios predicted by the analysis cycle generally showed more error than in the Kelvin cell lattice, however the overall change in ratio between the unconstrained and constrained conditions was captured. Considering the massive variety in behaviour seen in the angled cubic lattice, the degree to which the analysis cycle captured the behaviour is commendable.

6.3.2 Octet truss lattice

The octet truss lattice provided extremely consistent results, albeit at a higher error than was seen in the cubic or Kelvin cell lattices. This phenomenon, as well as the variation of the results due to constraint will be discussed in Sections 6.3.2.1-6.3.2.2.




Figure 6.49: Modulus ratio results for a cropped, constrained angled C8-t2 lattice at varying image sizes.

6.3.2.1 General

The apparent modulus error plot for the OT8-t2 lattice is shown in Figure 6.50. The analysis cycle results are consistently at an error of approximately -20% , which is a significantly higher error than seen in either the Kelvin cell or cubic lattices. A similar offset is seen in the inner modulus results, shown in Figure 6.51. At first, it was suspected that the cross-sectional properties of the elements in the mesh contained high error, which then adversely affected the modulus results. However, Figure 6.52 shows that there was only an error of approximately -5% associated with the cross-sectional area of the elements for representative images. Additionally, the moments of inertia of the elements in both directions showed an error of at most 10% for representative images as shown in Figure 6.53.

The error associated with the cross-sectional properties of the elements was low, however there is a chance that these small errors may result in a high error in the apparent modulus. To ascertain whether this is the case, a baseline simulation was run which included a -5% error in cross-sectional area, and -2% error in the moments of inertia. The apparent modulus result produced by this simulation was only approximately -5% off from the true baseline apparent modulus. This result indicated that the cross-sectional properties are not responsible for the -20% error in apparent modulus and there must be another factor which is causing the higher error.

The only other factor that may be causing the error is the structure of the mesh. 3D renderings of the OT8-t2 lattice indicate that representative images are produced from an image size of approximately 200 voxels onward, therefore it is unlikely that the problem stems from the image rasterisation, hence the discretisation of the mesh was investigated. Figure 6.54,




Figure 6.50: Apparent modulus results for an uncropped, unconstrained OT8-t2 lattice at varying image sizes.




Figure 6.51: Inner modulus results for an uncropped, unconstrained OT8-t2 lattice at varying image sizes.




Figure 6.52: Apparent modulus and beam cross-sectional area results for an uncropped, unconstrained OT8-t2 lattice at varying image sizes.




Figure 6.53: Enlarged view of primary and secondary moments of inertia results for an uncropped, unconstrained OT8-t2 lattice at varying image sizes.




Figure 6.54: Error plot of the number of nodes in the FE mesh for an uncropped, unconstrained OT8-t2 lattice at varying image sizes.

which displays an error plot of the number of nodes in the OT8-t2 mesh, shows that there is a clear increase in the number of nodes in the mesh as the image size increases. Although this phenomenon was detected in both the cubic and Kelvin cell lattices, the error in the node number was relatively small ($< 20\%$) even at high image sizes, which is not the case for the octet truss.

In order to visualise the problem, the first layer of cells²⁰ in a cubic, angled cubic, Kelvin cell and octet truss mesh were plotted as shown in Figure 6.55. In this image, where the nodes are shown as open circle markers, it is clear that there are extra nodes created in the angled cubic, Kelvin cell and octet truss lattices. Interestingly, these extra nodes all appear on the elements lying in the yz -plane angled at 135° to the y -axis, which would explain why there are no extra nodes seen in the cubic lattice. The Kelvin cell lattice contains minimal extra nodes, which correlates with its excellent results. The angled cubic lattice contains more extra nodes than the Kelvin cell, but a significant portion of the struts do not have extra nodes on them. Finally, the octet truss lattice contains a very high number of extra nodes, with almost every strut at the appropriate angle having at least one extra node.

In an ideal mesh, having extra nodes along a strut would not cause an error, as the additional nodes would be placed perfectly along the strut. However, in the BSM, the mesh is inferred from the image, therefore it is highly unlikely that the extra nodes will be perfectly placed. Instead, the extra nodes are slightly offset from the strut, creating small elements which do not maintain the correct angle and, as a result, do not connect properly to the junctions as illustrated in Figure 6.56. These incorrectly angled elements and poor junction connections cause deformation mechanisms which are not seen in the baseline simulation. Although a few of these

²⁰A portion of the lattice which is 1 cell thick, i.e. $x \times x \times 1$ cells.

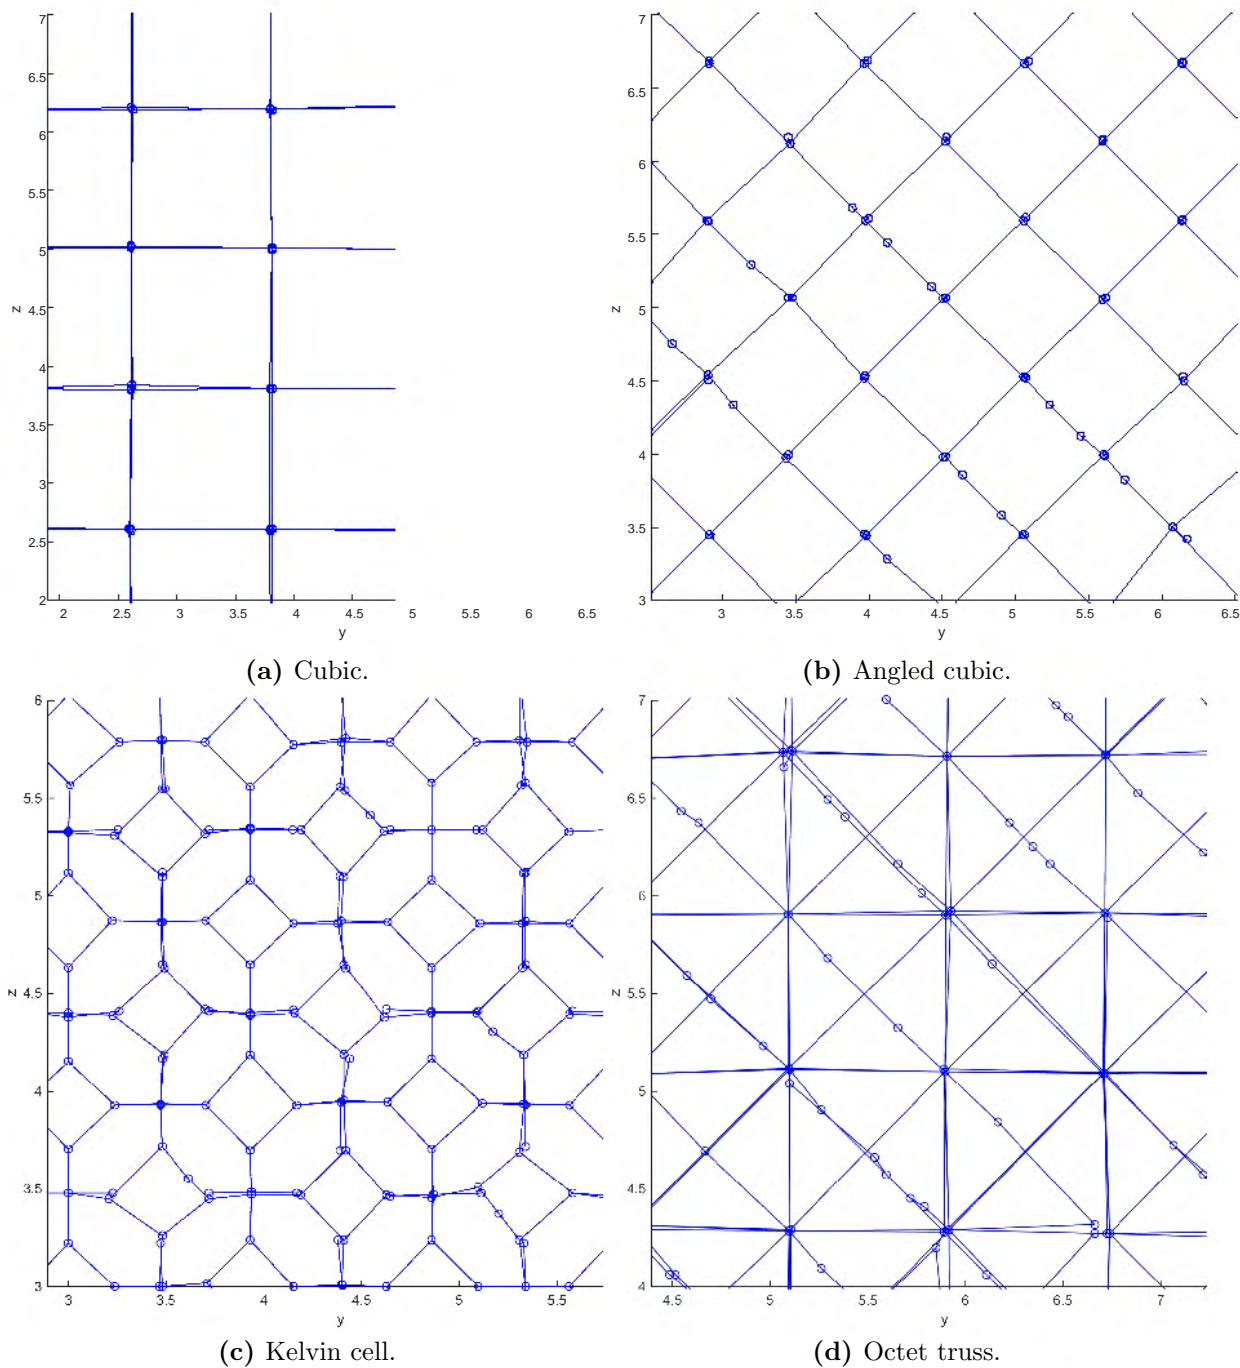


Figure 6.55: Enlarged region of plots of the first layer of cells in a cubic, angled cubic, Kelvin cell and octet truss mesh.




Figure 6.56: Small elements caused by the presence of the extra nodes in the octet truss lattice.

incorrectly angled/joined elements should not cause a major error in the results, almost every strut (at the correct angle) contains extra nodes in the octet truss lattice, and so a significant error is seen in the results.

The allocation of extra nodes in the mesh is a problem stemming from the discretisation code, which seems to occur only for elements in the yz -plane angled at 135° to the y -axis. Although this points to some directional preference in the discretisation code, which is not ideal, it is likely that this will not be an issue when discretising bone specimens. Even if the bone contains beam elements at the correct angle, it is unlikely that there will be a large number and, therefore, any error associated with the presence of the extra nodes should be small. Changing the discretisation code is outside the scope of this work, however it is important to take note of this possible error as an area of improvement for any future BSM codes.

Despite the fact that the apparent modulus and inner modulus results show a constant error of approximately -20% , the modulus ratio of the OT8-t2 lattice shows a good correlation to the baseline value, with all modulus ratios at image sizes greater than 200 voxels falling within the 10% error zone as shown in Figure 6.57. This is similar to the result obtained in the K8-t2 VBM simulation (see Section 6.2.6), albeit for different reasons. The extra nodes in the mesh cause an error in both the apparent modulus and inner modulus, however, by taking the ratio of these two values, the error inherent in both of the individual moduli is “cancelled out”. As such, a good prediction of the modulus ratio is obtained.

6.3.2.2 Variation due to constraint

Tables 6.9 and 6.10 show the baseline results for the apparent modulus and modulus ratio of the OT8-t2 lattice in various configurations. As expected, the octet truss lattice showed some sensitivity to constraint, but not as much as the Kelvin cell lattice. Similar to the Kelvin cell lattice, the modulus ratio decreases to a value below 1 when the top and bottom surfaces of the uncropped lattice are constrained, however when the cropped lattice is constrained, the modulus ratio increases slightly.




Figure 6.57: Modulus ratio results for an uncropped, unconstrained OT8-t2 lattice at varying image sizes.

Table 6.9: Baseline results for the apparent modulus of the uncropped and cropped OT8-t2 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained apparent modulus from unconstrained apparent modulus.

	Unconstrained (MPa)	Constrained (MPa)	Increase (%)
Uncropped	33.45	36.75	9.865
Cropped	34.35	43.16	25.65

Table 6.10: Baseline results for the modulus ratio of the uncropped and cropped OT8-t2 lattices in the unconstrained and constrained conditions, including the percentage increase in constrained modulus ratio from unconstrained modulus ratio.

	Unconstrained	Constrained	Increase (%)
Uncropped	1.081	0.8752	-19.04
Cropped	1.106	1.140	3.074




Figure 6.58: Modulus ratio results for an uncropped, constrained OT8-t2 lattice at varying image sizes.

The uncropped, constrained lattice has a modulus ratio which is significantly different from the other lattice configurations because it is less than 1 (Table 6.10). Although the apparent moduli predicted by the analysis cycle for this lattice show the same -20% offset as the uncropped, unconstrained lattice, the predicted modulus ratio remains excellent, with all representative image results falling within the 10% error zone as shown in Figure 6.58. As such, it can be said that even though the analysis cycle does not provide a very accurate prediction for the apparent modulus of the octet truss lattice, it consistently provides excellent predictions for the modulus ratio, even when a significant change in behaviour is present.

6.3.3 Summary and significance of results

Overall, the analysis cycle provided excellent results for the cubic and octet truss lattices in all their various topologies and test configurations. Similar to the Kelvin cell lattices, the cubic and octet truss lattices produced results with high error at small image sizes and results with low error at large image sizes. Additionally, familiar error surges may be detected in the cubic and octet truss lattices. These error surges may be attributed to over- or under-estimation of beam cross-sectional properties stemming from the rasterisation of the images.

The notable features in the cubic lattice results are summarised as follows:

- **Variation due to strut diameter:** At large image sizes, the cubic lattice showed a significant increase in error. This error stems from the erroneous inclusion of shell elements in the mesh which had zero nodal thickness. These shell elements are removed in a post-processing step, which allows the simulation to run, but produces inaccurate results because the structure is not well-represented. It was concluded that this phenomenon

does not pose a problem to the bone simulations, as there were no zero thickness shell elements included in the bone meshes.

- **Variation due to constraint:** The analysis cycle accurately captured the insensitivity of the cubic lattice to constraint in all configurations. This proves that the lattice is well-represented in the FE mesh and lends confidence to the methodology.
- **Variation due to angle:** The angled cubic lattice was shown to be exceptionally sensitive to constraint, with increases in apparent modulus and modulus ratio values of up to 675%. Despite this massive range of behaviour, the analysis cycle captured the response of the various lattice configurations to within 15%.

The notable features in the octet truss lattice results are summarised as follows:

- **Additional nodes:** It was found that the analysis cycle predictions for apparent and inner modulus were consistently approximately -20% off from the baseline value. This phenomenon is caused by the inclusion of a significant number of additional (and unnecessary) nodes in the mesh, which cause erroneous deformation mechanisms in the lattice and, consequently, decrease the simulated moduli. These additional nodes are only added to struts in the yz -plane, oriented at 135° to the y -axis, therefore this underestimation in modulus is unlikely to occur in bone simulations. Despite the underestimation of apparent and inner moduli, the predictions for modulus ratio were consistently excellent.
- **Variation due to constraint:** The analysis cycle accurately captured the modulus ratios of the octet truss lattice for all levels of constraint.

As with the Kelvin cell lattice, the results from the cubic and octet truss lattice analyses provide a great deal of confidence in the developed methodology. Although there are some phenomena which merit further investigation, such as the directional preference of additional node allocation, **this work has shown that the analysis cycle is capable of accurately capturing the response of a wide variety of structures, with a range of directional dependence and constraint sensitivity.**

Chapter 7

Bone Results

This chapter presents the results generated during the simulation of the bone specimens. Due to the large difference in size and, consequently, behaviour between the bone specimens, the results for the standard ($\phi 10\text{ mm}$) specimens are displayed separately to the results for the large ($\phi 20\text{ mm}$ and $\phi 28\text{ mm}$) specimens. The resource allocation required for the analysis of the bone specimens is presented after the modulus results.

In addition to the above, microstructural indices for the bone specimens are presented in reference to the experimental, BSM and VBM results. Finally, the results from the investigation into the variation of apparent modulus with voxel size, segmentation technique and boundary conditions are presented toward the end of this chapter.

7.1 Standard specimens results

106 specimens from the data set obtained from Hilton [3] were run through the analysis cycle¹. 88 specimens (83%) were discretised and simulated successfully. The remaining 18 specimens could not be discretised by the BSM code and, therefore, simulations of these specimens could not be performed. Figure 7.1 shows the results of the beam-shell simulations compared to the experimental results for the standard specimens. The experimentally determined apparent moduli range between 85.3 MPa and 891 MPa , whereas the simulated apparent moduli range between 9.50 MPa and 343 MPa .

The results of the VBM simulations compared to the experimental results are shown in Figure 7.2. All 106 specimens were discretised and simulated successfully using the VBM, but the failed BSM simulations are indicated in the figure for ease of comparison. The experimentally determined apparent moduli range from 85.3 MPa to 1.12 GPa , whereas the VBM apparent moduli range from 15.1 MPa to 1.03 GPa .

¹The data set obtained from Hilton contained 110 standard specimens, however 4 specimens could not be used for the current work due to missing/corrupt micro-CT scans. The complete bone specimen list may be found in Appendix H




Figure 7.1: Comparison of BSM and experimental results for standard bone specimens. Failures are included on the x -axis and are marked with a red “X”.




Figure 7.2: Comparison of VBM and experimental results for standard bone specimens.

7.2 Large specimen results

The experimental data set contained nine $\phi 20$ mm specimens and eight $\phi 28$ mm specimens, making a total of 17 large specimens. Seven $\phi 20$ mm specimens and four $\phi 28$ mm specimens were successfully discretised and simulated using the BSM. Four of the failed specimens could not be discretised by the beam-shell code, whereas two of the specimens were successfully discretised, but could not be simulated due to memory constraints. Figure 7.3a shows the BSM apparent modulus results, which are skewed due to the presence of an outlier. As such, Figure 7.3b shows an enlarged region² of Figure 7.3a, which excludes the major outlier. The experimental results range between 21.2 MPa and 390 MPa, whereas the BSM results range between 0.310 MPa and 43.0 MPa, excluding the outlier.

The modulus of an inner $\phi 10$ mm core region was also determined for the large specimens, as a measure of the effective modulus of the specimen. The BSM inner modulus results are displayed in Figure 7.4. There is no clear outlier in the inner modulus results as there was in the apparent modulus results, however the comparison of the range of the experimental results to the BSM results for the inner modulus is similar to that of the apparent modulus results. The experimental results range from 70.9 MPa to 725 MPa, whereas the BSM results range between 0.219 MPa and 39.4 MPa.


The results for the BSM modulus ratio, i.e. the ratio of the inner modulus to the apparent modulus, are shown in Figure 7.5. Only three of the BSM ratios are above 1, suggesting that the inner core of most of the specimens is more compliant than the whole, which does not correlate to the experimental results.

As with the standard specimens, there were less failures using the VBM than there were using the BSM. Of the 17 large specimens which were simulated, only one specimen failed to simulate using the VBM, as the memory required for the simulation exceeded the available memory on the computer used for this work. Figure 7.6a shows the apparent modulus results using the VBM. The presence of an erroneous negative results skews the results shown in Figure 7.6a, hence an enlarged region of the graph which excludes the erroneous data point is shown in Figure 7.6b. The experimental results range between 59.4 MPa and 390 MPa, whereas the VBM values range between 12.1 MPa and 157 MPa, excluding the negative value.


The inner modulus results for the large VBM simulations are shown in Figure 7.7. As with BSM simulations, there is no outlier in the VBM inner modulus results, despite the presence of an outlier in the apparent modulus results shown in Figure 7.6. The experimental results range from 60.3 MPa to 725 MPa, whereas the simulation results range from 5.56 MPa to 163 MPa.

The modulus ratio results for the large VBM simulations are shown in Figure 7.8. There are more specimens which have a ratio greater than 1 than there were in the BSM results, however more than half of the VBM ratios fall below 1, which does not correlate to the experimental results.

²The range of the BSM results is approximately one tenth the range of the experimental results. In order to improve readability of the graphs, the scales of the axes are not equivalent.



(a) Full result set.



(b) Enlarged region.

Figure 7.3: Comparison of BSM and experimental results for large bone specimens. Failures are included on the x -axis and are marked with a red “+”.




Figure 7.4: Comparison of BSM and experimental results for the inner core of the large bone specimens. Failures are included on the x -axis and are marked with a red “+”.






Figure 7.5: Comparison of BSM and experimental results for the modulus ratio of the large bone specimens. Failures are included on the x -axis and are marked with a red “+”.



(a) Full result set.



(b) Enlarged region.

Figure 7.6: Comparison of VBM and experimental results for large bone specimens. Failures are included on the x -axis and are marked with a red “+”.




Figure 7.7: Comparison of VBM results for the inner core of the large bone specimens. Failures are included on the x -axis and are marked with a red “+”.




Figure 7.8: Comparison of VBM and experimental results for the modulus ratio of the large bone specimens. Failures are included on the x -axis and are marked with a red “+”.




Figure 7.9: Time taken for discretisation of the standard bone specimens using the BSM and VBM. Results shown in ascending order of BSM time.

7.3 Resource usage

The computational expense of the simulations is quantified by considering two factors, namely the time taken for the analysis and the memory requirements of the analysis. The results from the time and memory usage for both the BSM and VBM simulations are presented in this section.

7.3.1 Time

Figure 7.9 shows the time taken for the discretisation of the standard bone specimens using the BSM and VBM. The BSM is clearly significantly more time consuming than the VBM during discretisation, with the time taken for the BSM discretisation ranging from 2 – 172 minutes and the time taken for the VBM discretisation taking 0.3 – 9 minutes. The average time taken for discretisation using the BSM is 26 minutes, whereas the average time for the VBM is 1 minute.

In contrast, the time taken to simulate a VBM mesh is generally significantly longer than the time taken to simulate a BSM mesh, as is shown in Figure 7.10. The simulation times for the BSM range from 0.2 – 51 minutes, whereas the VBM simulation times range from 1 – 292 minutes. The average time taken for a BSM simulation is 9 minutes; the average time taken for a VBM simulation is 43 minutes.

Figure 7.11 shows the total time taken for the discretisation and simulation of the specimens. Although the time taken for the BSM and VBM analyses is relatively similar for the majority




Figure 7.10: Time taken for simulation of the standard bone specimens using the BSM and VBM. Results shown in ascending order of BSM time.

of the specimens, the BSM analysis took less time than the VBM analysis for 61.4% of the specimens.

Due to the significant difference in size between the standard specimens and the large specimens, the time taken for the discretisation and simulation of the large specimens is considered separately to that of the standard specimens. Figure 7.12 shows the time taken for the discretisation of the large specimens. The average time taken for the BSM discretisation was 142 minutes, whereas the average time taken for the VBM discretisation was 3 minutes.

The time taken for the simulation of the large specimens is shown in Figure 7.13. The average time for the BSM simulations was 14.6 minutes, and the average time for the VBM simulations was 418 minutes, which is a 28-fold increase on the BSM time.

The total time taken for the discretisation and simulation of the large specimens is shown in Figure 7.14. Despite the fact that the VBM has a much shorter discretisation time than the BSM, 82% of the specimens were analysed faster with the BSM than they were with the VBM.

For both the standard specimens and the large specimens, the VBM discretisation was quicker than the BSM discretisation, however the time saved by the BSM during simulation meant that the majority of specimens were analysed faster with the BSM than they were with the VBM.




Figure 7.11: Total time taken for discretisation and simulation of the standard bone specimens using the BSM and VBM. Results shown in ascending order of BSM time.




Figure 7.12: Time taken for discretisation of the large bone specimens using the BSM and VBM.




Figure 7.13: Time taken for simulation of the large bone specimens using the BSM and VBM.




Figure 7.14: Total time taken for discretisation and simulation of the large bone specimens using the BSM and VBM.

7.3.2 Memory

It is important to consider not only the time taken for the analysis cycle, but also the memory allocation required for the simulations. A base memory of $700M^3$ words⁴ was allocated to all the simulations, however it was found that this value was not sufficient in all cases.

The small specimens were successfully simulated with the VBM using the $700M$ word allocation, however during the simulation their stiffness matrices were factorised in out-of-core mode⁵ due to insufficient memory, which severely decreased performance and increased the time taken for the simulations. In contrast, the $700M$ word allocation was sufficient for the small specimens simulated using the BSM, and allowed the stiffness matrices to be factorised in-core.

Both the VBM and BSM suffered from memory issues when the large specimens were simulated. Memory allocations of $2000M$ to $4000M$ words were necessary for the large specimens to run successfully using the VBM, however these increased values were still not sufficient for in-core factorisation of the stiffness matrix. In some cases, the required memory allocation was beyond the specifications of the computer used for the work. The base memory allocation was also not sufficient for some of the large specimens simulated using the BSM, however increasing the memory did not allow for successful simulation of these specimens.

7.4 Microstructural indices

The results for the microstructural indices analysis of the experimental results, BSM results and VBM results⁶ are shown in Figures 7.15-7.19. Each figure represents a particular index, with the subfigures showing how the experimental results, BSM results and VBM results vary with respect to that index. Figures 7.15-7.19 represent the results for bone volume fraction, trabecular thickness, trabecular separation, connectivity and connectivity density, respectively. Various markers are used to illustrate the results in the subfigures: black squares represent the experimental results, black circles represent the BSM results and blue squares represent the VBM results.

In all of the microstructural indices results, the VBM shows the strongest correlation to the indices, whereas the experimental results show a weak correlation or no correlation whatsoever. The BSM results consistently fall between these two extremes.


As mentioned in Section 5.3, the shell-to beam ratio of the specimens was calculated directly from the BSM mesh and used as an additional index. The results for the comparison of experimental, BSM and VBM results to the shell-to-beam ratio are shown in Figure 7.20. Once again, the experimental results show no correlation to this index, however there appears to be some correlation between the simulation results and the shell-to-beam ratio, particularly at smaller ratios.

³ $700M = 700 \times 10^6$


⁴1 word = 2 bytes.

⁵Out-of-core mode involves storing the stiffness matrix on the hard drive, instead of just in memory, i.e. RAM.


⁶Note: Only the specimens which were successfully simulated using both the BSM and VBM are presented in these graphs. Additionally, the outlier in the BSM results (see Figure 7.3a) and the negative result in the VBM results (see Figure 7.6a), are not shown in the microstructural indices graphs.



(a) Experimental.




(b) BSM.




(c) VBM.


Figure 7.15: Comparison of apparent modulus results and bone volume fraction.



(a) Experimental.




(b) BSM.




(c) VBM.


Figure 7.16: Comparison of apparent modulus results and mean trabecular thickness.



(a) Experimental.




(b) BSM.




(c) VBM.


Figure 7.17: Comparison of apparent modulus results and mean trabecular separation.



(a) Experimental.




(b) BSM.




(c) VBM.


Figure 7.18: Comparison of apparent modulus results and connectivity.



(a) Experimental.




(b) BSM.




(c) VBM.


Figure 7.19: Comparison of apparent modulus results and connectivity density.



(a) Experimental.



(b) BSM.



(c) VBM.

Figure 7.20: Comparison of apparent modulus results and shell-to-beam ratio.




Figure 7.21: Comparison of micro-CT scans of specimen 477 at varying voxel sizes.




Figure 7.22: Comparison of micro-CT scans of specimen 518 at varying voxel sizes.

7.5 Voxel size analysis

In order to investigate the variation caused by increasing voxel size, two bone specimens were scanned at three different voxel sizes. Segmented, cropped images from the micro-CT scans are shown for specimens 477 and 518 in Figures 7.21 and 7.22, respectively.

The microstructural indices obtained from the scans of varying voxel sizes are shown in Table 7.1. The apparent modulus results from the voxel size analysis are shown in Table 7.2.

Table 7.1: Microstructural indices for different voxel sizes.

Sample number	Voxel size	BV/TV	Mean Tb.Th		Mean Tb.Sp	
	μm		Pixels	μm	Pixels	μm
477	20	0.339	10.9	218	29.4	588
	40	0.347	5.82	233	14.8	592
	80	0.405	3.73	298	7.10	568
518	20	0.316	11.2	224	16.2	324
	40	0.317	6.32	253	16.6	664
	80	0.365	3.90	312	8.46	677

Table 7.2: Apparent modulus results for different voxel sizes. Failed simulations are indicated with a dash (-).

Sample number	Apparent modulus (MPa)		
	20 μm	40 μm	80 μm
477	-	94.1	105
518	-	84.6	96.1




Figure 7.23: Comparison of BSM and experimental results for standard bone specimens segmented using the *Optimise Threshold* function. Failures are included on the x -axis and are marked with a red “X”.

7.6 Segmentation analysis

The apparent modulus results from the alternative segmentation using the *Optimise Threshold* function are shown in Figure 7.23. There is clearly more scatter in the results produced by the *Optimise Threshold* function than in the *Make Binary* results shown in Figure 7.1. Of the 106 standard specimens, there were 10 discretisation failures, which is less than the 18 failures resulting from the *Make Binary* images.

The modulus ratio results produced using the *Optimise Threshold* segmentation are shown in Figure 7.24. These results are almost identical to the equivalent *Make Binary* results shown in Figure 7.5, with the analysis cycle predicting a modulus ratio below 1 for most of the specimens.

In addition to the standard boundary conditions, a lateral constraint was added to the top and bottom surfaces of the specimens to determine if the sensitivity of the specimens to constraint was affected by the method used for image segmentation. Figure 7.25 shows the variation of the apparent modulus with boundary condition when the *Optimise Threshold* function is used for segmentation. Similarly, Figure 7.26 shows the variation of the modulus ratio results with boundary condition. Neither the variation in apparent modulus, nor the variation in modulus ratio appear to be significantly different to the variation in results observed in the *Make Binary* analysis.




Figure 7.24: Comparison of BSM and experimental modulus ratio results for large bone specimens segmented using the *Optimise Threshold* function. Failures are included on the x -axis and are marked with a red “+”.




Figure 7.25: Comparison of BSM and experimental results for standard bone specimens segmented using the *Optimise Threshold* function with a lateral constraint on the top and bottom surfaces.




Figure 7.26: Comparison of experimental and BSM modulus ratio results for large bone specimens segmented using the *Optimise Threshold* function with a lateral constraint on the top and bottom surfaces.

7.7 Boundary condition analysis

The results from the boundary condition analysis are divided into two sections, for the standard specimens and large specimens, respectively. The apparent modulus is displayed for the standard specimens, whereas the inner modulus is displayed for the large specimens as the effect of the boundary condition on the apparent modulus is captured in the standard specimen results. Additionally, the modulus ratio is displayed for the large specimens, so the effect of the boundary conditions on the interplay between the inner modulus and apparent modulus may be discussed.

7.7.1 Standard specimens

The standard specimen results from the boundary condition analysis are shown in Figure 7.27, with legend terminology as indicated:

- **T** - constraint applied on top surface
- **B** - constraint applied on bottom surface
- **TB** - constraint applied on top and bottom surfaces simultaneously
- **lat** - lateral constraint applied
- **rot** - rotational constraint applied
- **lat/rot** - lateral and rotational constraint applied simultaneously

Table 7.3: Average increase in stiffness experienced by the bone specimens due to the addition of lateral and rotational constraints. Increase is determined relative to the “Standard” boundary condition.

Surface	Constraint	Increase (%)
Bottom	Lateral	4.52
	Rotational	1.44
	Lateral and rotational	5.12
Top	Lateral	4.90
	Rotational	1.45
	Lateral and rotational	5.52
Top and bottom	Lateral	14.9
	Rotational	2.94
	Lateral and rotational	16.1

Table 7.4: Average change in inner stiffness experienced by the large bone specimens due to the addition of lateral and rotational constraints. Change is determined relative to the “Standard” boundary condition. Any boundary condition which caused both an increase and decrease in inner modulus of the specimens is marked with an asterisk (*).

Surface	Constraint	Change (%)
Bottom	Lateral*	4.96
	Rotational	1.95
	Lateral and rotational*	4.85
Top	Lateral	29.2
	Rotational	6.86
	Lateral and rotational	33.1
Top and bottom	Lateral	41.0
	Rotational	9.15
	Lateral and rotational	46.2

Table 7.3 shows the average percentage increase in stiffness of the bone specimens when a particular boundary condition is applied.

It is clear from Figure 7.27 and Table 7.3 that a combined lateral and rotational constraint applied on the top and bottom surface of the specimen has the greatest impact on the stiffness of the specimen, followed closely by a lateral only constraint applied on the top and bottom surfaces of the specimen. The other constraints have a much less significant impact on the results, and all show average increases of 5.5% or less.

7.7.2 Large specimens

The inner modulus results from the boundary condition analysis of the large specimens are shown in Figure 7.28 and the average change in the inner modulus results with varying boundary condition is shown in Table 7.4. This table differs from Table 7.3 in that it documents the absolute change in the results, as opposed to the increase in results. This distinction is important, because some of the inner moduli decreased in magnitude with an increase in constraint, which was not seen in the apparent modulus results shown in Figure 7.27.






Figure 7.27: BSM results for apparent modulus of standard specimens simulated using a variety of boundary conditions. A region of (a) has been enlarged in (b) to better illustrate the influence of the various boundary conditions on the results.



(a) Full result set.



(b) Enlarged region.

Figure 7.28: BSM results for inner modulus of large specimens simulated using a variety of boundary conditions. A region of (a) has been enlarged in (b) to better illustrate the influence of the various boundary conditions on the results.

Table 7.5: Average change in modulus ratio of the large bone specimens due to the addition of lateral and rotational constraints. Change is determined relative to the “Standard” boundary condition. Any boundary condition which caused both an increase and decrease in inner modulus of the specimens is marked with an asterisk (*).

Surface	Constraint	Change (%)
Bottom	Lateral*	7.86
	Rotational*	1.25
	Lateral and rotational*	7.66
Top	Lateral*	10.8
	Rotational*	1.51
	Lateral and rotational*	10.7
Top and bottom	Lateral*	10.4
	Rotational*	2.24
	Lateral and rotational*	9.93

As with the apparent modulus results shown in Figure 7.27, a combined lateral and rotational constraint applied simultaneously on the top and bottom surface of the specimen generally had the greatest impact on the inner modulus of the specimens. In some cases, a lateral and rotational constraint applied only on the top surface had a greater impact than the same constraint applied on the top and bottom surfaces simultaneously. The constraints that had a significant ($> 20\%$) impact on the inner modulus are lateral constraints on the top surface and combined surfaces (i.e. top and bottom) as well as lateral and rotational constraints on the top surface and combined surfaces. The influence of the other boundary conditions was much less significant, showing a maximum change of only 9%.

The modulus ratio results from the boundary condition analysis are shown in Figure 7.29 and the average change in modulus ratio with varying boundary conditions is shown in Table 7.5. All of the boundary conditions caused an alternating increase and decrease in the modulus ratio results, depending on specimen.

It is clear that the addition of constraints has a significant effect on some of the large specimens. In particular, three of the specimens which had a modulus ratio less than 1 using the standard boundary conditions showed a modulus ratio above 1 with increased constraints. Additionally, one of the specimens, the results for which are displayed in Figure 7.29b, showed an increase in modulus ratio from 0 to 0.8–0.9 with the addition of constraints on the bottom surface. Unlike the apparent and inner moduli results, the greatest change in modulus ratio was experienced when a lateral constraint was added to the top surface of the specimen, as shown in Table 7.5. However, there is some variation in the way the various boundary conditions affect each specimen, which makes it more difficult to generalise with this metric.




Figure 7.29: BSM modulus ratio results for large specimens simulated using a variety of boundary conditions.

Chapter 8

Discussion

A variety of simulations were performed in order to investigate the performance of the beam-shell method (BSM) in predicting the apparent and inner (or effective)¹ moduli of trabecular bone. The discussion of the results is divided into sections detailing the performance of the BSM when compared to experimental results (Section 8.1), the performance of the BSM compared to the voxel-based method (VBM) (Section 8.2), the microstructural indices analysis (Section 8.3) and the effect of various factors on the performance of the BSM (Section 8.4). Finally, the differences between the results in the current work and the results produced by Vanderroost, et al [2] are discussed (Section 8.5).

8.1 BSM vs experimental

Simulating the full bone specimens, as opposed to cropped regions from the specimens, allowed the simulation results produced by the BSM to be compared to the experimental results obtained from Hilton [3]. This section investigates the performance of the BSM in relation to the experimental results for the standard bone specimens and large bone specimens separately, so that comparisons between apparent and inner moduli may be highlighted in the corresponding sections.

8.1.1 Standard specimens

Of the 106 standard specimens run through the analysis cycle, 88 were discretised and simulated successfully. Some of the failures, such as specimens 165 and 461 can be explained by poor segmentation (see segmented bone images in Appendix I), however overall the segmentation of the bone specimens is satisfactory, indicating that research is required around making the BSM more robust.

Figure 8.1 shows the standard specimen results obtained using the BSM, including a line of best fit². The BSM results for the specimens with experimental modulus less than 0.2 *GPa* are relatively tightly grouped to the best fit line, however the scatter increases as the experimental

¹In the current work, the effective modulus is taken to be equivalent to the inner modulus, based on the work in [3, 15, 83]. However, the term “inner modulus” is generally used instead of “effective modulus” for clarity.

²The line of best fit was determined using a log-normal analysis. For more information on the process, see Appendix J.

modulus increases. Significant scatter is seen at experimental moduli greater than 0.4 *GPa*, with the maximum deviation from the best fit line being 392%. Despite the scatter, a significant portion (60.9%) of the data points show good agreement³ to the best fit line, as indicated by the 50% error region in Figure 8.1.

It is clear from Figure 8.1 that there is not a 1 : 1 correlation between the BSM results and the experimental results as the line of best fit has a gradient of $m = 0.121$. Although this means that the values of the simulated moduli differ significantly from the experimental moduli, it is not indicative of a loss of accuracy in the BSM. Because the Young's modulus of the bone specimens is unknown, an arbitrary Young's modulus of 1 *GPa* was allocated to the simulations. If the Young's modulus value is increased to 8.30 *GPa*, the results are scaled accordingly and a 1 : 1 correlation is obtained, albeit with more scatter in the results⁴. This back-calculated Young's modulus is well within the range of Young's moduli for trabecular bone reported in the literature (see Section 2.5.2), which indicates that the values predicted by the BSM are realistic. The difference in value of the BSM results compared to the experimental results is, therefore, not a concern and emphasis should instead be placed on the degree to which the results deviate from the trend line.

Although viewing the bone results on a single graph provides information about the overall scatter of the results, this graph represents the amalgamation of three different sets of bone tests, namely the defatted, marrow and confined tests as described in Section 2.6.2. As such, it is important to view the results in these respective subsets, as is shown in Figure 8.2. In these graphs, the results for all the specimens are shown, however the specimens in the subset being considered (i.e. defatted, marrow or confined) are represented with coloured markers, namely red, blue and pink, respectively. Similarly, the line fit for the full data set is provided on the graph as a grey dashed line, while the line fit for the subset being considered is shown as a solid black line.

It is clear from Figure 8.2 that the defatted specimens show the least scatter of the three different types of bone tests. This result makes sense if the micro-CT scans of the different specimen types are considered, as illustrated in Figure 8.3. Because the marrow has been removed from the defatted specimens, there is a clear distinction between bone and non-bone voxels, which leads to good, consistent segmentation as demonstrated in Figures 8.3a and 8.3b. In contrast, the specimens which still contain marrow (i.e. the marrow and confined specimens) have a less clear distinction between bone voxels and non-bone voxels as the grayscale value of the marrow voxels falls somewhere between the two. This leads to the misidentification of marrow voxels as bone voxels in some specimens, which would result in a thicker structure as demonstrated in Figures 8.3c and 8.3d and, consequently, a higher apparent modulus. These higher moduli increase the scatter, as seen in Figure 8.2b and 8.2c. Additionally, it is important to note that the confined specimens were experimentally tested with the aid of an aluminium collar [3]. Although it was found that this collar did not provide the required confinement from the start of the compression test, it is likely that the collar did provide a confinement effect once the specimen had compressed and laterally expanded. As a result, there may be some variation in the experimental results for the confined specimens, which would further increase the scatter.

³The bone results cannot be judged by the same standards as the lattice results, as the bone is far more structurally complex and the Young's modulus is unknown. Additionally, the agreement of the data points to the best fit line varies based on which regression technique is used. Consequently, the bone results were considered to show "good agreement" to the experimental results if they showed an error of less than 50% as measured from the best fit line.

⁴This relationship was confirmed for both the BSM and the VBM.







Figure 8.1: Comparison of BSM and experimental results for standard bone specimens with line fit.



(a) Defatted.



(b) Marrow.



(c) Confined.

Figure 8.2: Comparison of BSM and experimental results for defatted, marrow and confined bone specimens.




Figure 8.3: Comparison of a defatted (specimen 13) and marrow (specimen 26) image slice with corresponding segmentation. Figures (a) and (c) show the grayscale images, while figures (b) and (d) show the segmentation of the image slices using the Otsu method. Clear overthickening of the marrow specimen as a result of improper segmentation is visible in (d).

The overthickening of the marrow-inclusive bone specimens during segmentation can be confirmed by observing the images associated with the six outlying results in the confined graph, as shown in Figure 8.4. If the widths of the trabeculae in the grayscale images are compared to the segmented trabeculae, it is clear that all the specimens experience general overthickening of their trabeculae during segmentation. Additional overthickening may be seen in specific regions of the binarised images. For example, Figures 8.4b, 8.4f, 8.4j and 8.4l show additional overthickening around the circumference of the specimen⁵, while Figure 8.4d shows specific regions of overthickening within the specimen. A combination of the general overthickening and local regions of overthickening is responsible for the scatter seen in the marrow-inclusive specimens. Considering these results, it is likely that there would be less scatter in the BSM results if all the specimens were defatted before scanning or if the quality of the segmentation was better.

Another observation which can be made from Figure 8.3 is that the marrow and confined specimens both show higher apparent moduli, on average, than the defatted specimens. A similar phenomenon was noted in the results from the experimental work [3]. Although the simulation results seem to correlate with the experimental results in this regard, it is important to remember that the marrow is removed from the structure during segmentation and, hence, is not included in the simulation. As such, even though similar phenomena are present in the experimental and simulation results, the experimental increase is due to the marrow lending stiffness to the structure, whereas the simulation increase is due to poor segmentation of the marrow from the bone.

It is important to acknowledge here that some of the variation in the data may stem from a lack of accuracy in the experimental results, rather than the numerical results. Although this is a possibility which must be acknowledged, the work by Hilton [3] has been reviewed by two international examiners and was awarded a distinction due to its excellent quality. As a result, for the duration of the current work, the experimental results by Hilton [3] are viewed as a solid and reliable standard against which to compare the numerical results.

⁵The circumferential overthickening is caused by a combination of machining artifacts, which are discussed in Section 8.4.1, and global segmentation issues as mentioned in Section 2.2.2, where the outer edge of the specimen tends to be overthickened compared to the inner region [14].




Figure 8.4: Comparison of confined specimen images with corresponding segmentation. Figures (a), (c), (e), (g), (i) and (k) show the grayscale images, while figures (b), (d), (f), (h), (j) and (l) show the segmentation of the image slices using the Otsu method.




Figure 8.5: Comparison of BSM and experimental results for the modulus ratio of the large bone specimens with data labels. Outlying result not pictured.

8.1.2 Large specimens

Both the apparent and inner moduli were determined for the large bone specimens, which presented the novel opportunity to determine the ratio of the inner modulus to the apparent modulus and compare this ratio to equivalent experimental results. The modulus ratio is a useful metric to consider, as it removes the influence of Young's modulus from the results, and thus provides a more direct comparison to the experimental results than either the apparent or inner moduli. As such, more emphasis is placed on the modulus ratio results for the large specimens than the apparent or inner moduli.

The modulus ratio results for the large specimen simulations using the BSM are shown in Figure 8.5. According to the experimental results, all the specimens with the exception of specimen 535 should have a modulus ratio greater than 1, however only two of the eleven specimens show the correct relationship between their apparent and inner moduli (i.e. the experimental and simulation ratios are both greater than 1 or the experimental and simulation ratios are both less than 1). The simulated ratio for specimen 534 is 1.32, a decrease of 38% from the experimental value of 2.13. Specimen 510 has a BSM ratio of 1.46, where the experimental ratio is 1.31, a difference of only 12%.

Figure 8.6 shows the results for the apparent modulus and inner modulus for the large bone specimens simulated using the BSM. There were too few data points for a reliable line fit to be performed, however a dashed line has been added to the graph to highlight the results from specimen 510, which was the most accurate data point produced by the BSM. If the line asso-

ciated with specimen 510 is taken as a reference, it seems that the apparent modulus predicted by the BSM (Figure 8.6a) is too large for the majority of the specimens, whereas the inner modulus prediction (Figure 8.6b) is relatively good for most of the specimens. The idea that the apparent modulus is being overestimated is strengthened by the fact that the apparent modulus of specimen 534 is the closest result to the 510 line in Figure 8.6a, and 534 was the only other specimen to show an accurate relationship between apparent and inner moduli. This idea is further strengthened by the known issues caused by global segmentation, where trabeculae toward the outer edge of the specimen are overly thick compared to the trabeculae in the centre of the specimen [14]. These overly thick regions cause an erroneous increase in the apparent modulus, but are not taken into account when the inner modulus is calculated. As a result of the increased apparent modulus, the predicted modulus ratio is lower than what was experimentally measured.

Using the modulus ratio allowed a direct, unbiased comparison to be made between experimental and simulation results. The BSM showed that it is capable of predicting the correct relationship between apparent and inner modulus, however there is clearly still work required in this area, as the correct relationship was only predicted for two of the eleven specimens. Considering the identified issues, it is likely that the accuracy of these results would be significantly increased through the use of a different segmentation technique.

8.2 BSM vs VBM


Although the VBM has been called the computational “gold standard” method for Finite Element simulations of trabecular bone, it is important to establish how accurately the VBM predicted the apparent moduli of the bone specimens used in this study. As such, this section compares the VBM results to the experimental results and discusses the comparison between the BSM results and VBM results. Additionally, the resources required for the BSM and VBM discretisation and simulation are discussed here.

8.2.1 Standard specimens


Figure 8.7 shows the standard specimen results obtained using the VBM, including a trend line calculated as described in Appendix J⁶. Only the results which ran successfully using the BSM are shown in this graph to provide a basis of comparison to the BSM results. Overall, there is much less scatter in the VBM results than was seen in the BSM results, with 69.0% of the specimens lying within the 50% error envelope. This number increases to 80.5% if the gradient of the line fit is increased to 0.28. A noticeable feature in the VBM graph is the presence of six significant outliers. Although specimen 332 was similarly positioned in the BSM results, the significant outliers above the trend line were not seen in the BSM results. The maximum error associated with these outliers is 674% for specimen 19.

If the graphs corresponding to the defatted, marrow and confined specimens shown in Figure 8.8 are considered, it is clear that three of the significant outliers in the VBM results stem from defatted specimens (Figure 8.8a). The maximum outlier in the confined VBM results, specimen 19, corresponds to the maximum outlier in the confined BSM results shown in Figure

⁶The log-normal line fit does not fit the VBM data as well as it fit the BSM data. However, a least-squares line fit was not representative of the VBM data due to the influence of the outliers. Consequently, to ensure consistency between the BSM and VBM graphs, the log-normal line fit was maintained for the VBM data.



(a) Apparent modulus.



(b) Inner modulus.

Figure 8.6: Comparison of BSM and experimental results for large bone specimens.




Figure 8.7: Comparison of VBM and experimental results for standard 10 *mm* bone specimens with line fit. Only specimens which were simulated successfully using the BSM are shown in the graph. Significant outliers have been labelled with their specimen number.

8.2c, and may therefore be attributed to poor image segmentation as shown in Figures 8.4a-8.4b and/or an erroneous experimental result. However, the BSM simulations provided good results for specimens 13, 15, 25 and 502, which makes these outliers in the VBM results particularly interesting.

Image slices for the outlying VBM bone specimens are shown in Figure 8.9. Although there is slight overthickening around the circumference of specimen 13, the other specimens all appear to be well-segmented, with any overthickening effects being concentrated in small regions. So why does the BSM produce good results for these specimens, when the VBM does not? As with the lattices, which were discussed in Section 6.2.6, the answer may be found in the way the meshes are created. The VBM creates a mesh directly from the scan by making each bone voxel a solid element, whereas the BSM uses a thinning algorithm to discretise the structure into beams and shells and then infers the cross-sectional properties of an element based on the group of voxels that make up that particular element. This means that overly thick areas along an element will have less of an effect on the mesh, as the BSM considers the thickness of the whole element and not just one particular region. In contrast, because of its direct approach, the VBM has to take into account these overly thick areas, which will increase the apparent modulus predicted by the method.

The same methodology which allows the BSM to mitigate the effects of small regions of overthickening, can work against it during general overthickening. If most of the specimen is overly thick and only regions of the element are correctly segmented, the BSM will assign cross-sectional properties correlating to the majority of the strut, i.e. the inferred cross-sectional properties will be overly thick. The VBM will take the overly thick regions into account, but will also take the well-segmented regions into account, which will mitigate the effect of the overthickening. As a result, the marrow inclusive specimens, which tend to experience general overthickening, show greater scatter in the BSM results shown in Figure 8.2b than they do in the VBM results shown in Figure 8.8b.

As with the BSM results, the VBM defatted results in Figure 8.8a show a lower trend line gradient than the marrow or confined trend lines in Figures 8.8b and 8.8c, respectively. This confirms that the marrow inclusive specimens (i.e. marrow and confined in the graphs) experience general overthickening during segmentation, which leads to stiffer apparent moduli and a corresponding increase in the gradient of the line of best fit.

Figure 8.10 shows the BSM and VBM results compared to the experimental results. The gradient of the line fit for the VBM data ($m = 0.240$) is approximately twice as large as the BSM line fit gradient ($m = 0.121$). Correspondingly, the back-calculated modulus for the VBM ($E_y = 4.16 \text{ GPa}$) is approximately half the size of the back-calculated BSM modulus ($E_y = 8.30 \text{ GPa}$). This result indicates that the scaling factor of 1.605, which was suggested by Vanderroost, et al [2] is not applicable to this data. Instead, a correction factor of approximately 2 is required. Wang, et al [87] attribute this difference in value (and consequent scaling factor) to an underestimation in apparent modulus by the BSM, however the validation study in the current work has shown that it may, in fact, be due to an overestimation by the VBM, rather than an underestimation by the BSM (see Section 6.2.6). As a result, it is more likely that the increased scaling factor in the current work is due to overthickening in the bone scans which is likely more significant due to the larger voxel size [14] used for the micro-CT scans ($80 \mu\text{m}$ vs $20 \mu\text{m}$ in [2]). As discussed in Section 6.2.6, overthickening is likely to have a more significant effect on the VBM than the BSM, leading to an increase in the apparent modulus produced by the VBM and corresponding increase in the required scaling factor.




Figure 8.8: Comparison of VBM and experimental results for defatted, marrow and confined bone specimens.




Figure 8.9: Comparison of outlying VBM specimen images with corresponding segmentation. Figures (a), (c), (e) and (g) show the grayscale images, while figures (b), (d), (f) and (h) show the segmentation of the image slices using the Otsu method.

In Figure 8.11, the BSM results are compared to the VBM results and the defatted, marrow and confined specimens are indicated using black circles, blue squares and red diamonds, respectively. There appears to be good correlation between the BSM and the VBM results for VBM moduli less than 0.15 GPa . The correlation remains relatively good, albeit with slightly more scatter, between 0.15 GPa and 0.2 GPa , after which the BSM results show an increase which is not seen in the VBM results and, consequently, causes a break in the trend. This increase corresponds to the increase in scatter seen in the BSM results in Figure 8.1, and can therefore be attributed to a loss of accuracy in the BSM simulations. Additionally, there appear to be some outliers in the graph, occurring at high VBM moduli and comparatively low BSM moduli, which are labelled with their specimen number in Figure 8.11a. These outliers correspond to the outliers in the VBM results seen in Figure 8.7, and therefore must be attributed to a loss of accuracy in the VBM simulations. If the VBM outliers are removed, as shown in Figure 8.11b, a strong linear correlation ($R^2 = 0.878$) is obtained between the BSM and VBM results.

A similar comparison between the BSM and VBM was drawn in the paper by Vanderroost, et al [2], however a log-log scale was used to show the comparison. In order to provide an equivalent comparison, the BSM and VBM results in the current work were replotted using a log-log scale and polynomial fit in Figure 8.12a and compared to the graph in the literature, as shown in Figure 8.12b. Viewed in this way, the BSM appears to correlate better with the VBM than in Figures 8.11-8.11b, however this author would argue that the log scale provides a skewed representation of the difference between the results produced by the two methods, particularly because a more complex curve is fitted to the data. Despite this, it is encouraging to see a similar shape and similar correlation in the data generated for the current work (Figure 8.12a) as in the data generated by Vanderroost, et al (Figure 8.12b).

Overall, the VBM has shown a better correlation to the experimental results than the BSM, however there are clear issues present in both modelling methodologies. If an individual bone






Figure 8.10: Comparison of BSM and VBM results to experimental results for standard bone specimens. Equations of line fits are included for both sets of simulation results.




(a) Full result set.




(b) Enlarged region.

Figure 8.11: Comparison of BSM results to VBM results for standard bone specimens.



(a) Results generated by the author.



(b) Results generated by Vanderoost, et al [2].

Figure 8.12: Comparison of BSM results to VBM results using a log-log scale. Figure (a) shows results generated in the current work, whereas Figure (b) shows results generated in the study by Vanderoost, et al [2].

specimen is scanned and simulated using either technique, an accurate result is by no means guaranteed. Hence, it is imperative that results are viewed in the context of other experimental and simulation results to gain an understanding of the accuracy of the simulated data point.

8.2.2 Large specimens

The apparent modulus and inner modulus results produced by the VBM for the large specimens, displayed in Figures 7.7 and 7.7, show a much stronger trend than the equivalent BSM results, shown in Figures 7.3 and 7.4. This may be due to the fact that more of the large specimens were successfully simulated with the VBM, which makes the trend easier to distinguish, however the fact that the VBM was capable of simulating so many more specimens than the BSM in and of itself lends confidence to the discretisation method.

In order to provide a basis of comparison with the BSM, only the specimens which were successfully simulated using the BSM are considered in this analysis. Figure 8.13a and 8.13b show the apparent and inner moduli obtained using the VBM. Even with fewer results displayed, the trend for the VBM data is easily distinguishable, which was not the case with the BSM specimens, particularly the apparent modulus shown in Figure 8.6a. Additionally, specimens 537 and 538 do not show the low apparent or inner moduli seen when the BSM was used.

The modulus ratio results, along with their respective data labels, are shown in Figure 8.14. Even though the apparent and inner moduli predicted by the VBM show strong trends, only five of the eleven⁷ predicted ratios demonstrated the correct relationship between the apparent and inner moduli. Moreover, only two of the predicted VBM ratios, namely specimens 510 and 512, showed less than 20% error when compared to the experimental results.


In order to compare the BSM to the VBM, both sets of modulus ratios are plotted in Figure 8.15. For each specimen, a line is drawn between the BSM result and VBM result to highlight the difference in the results and distinguish the specimens from each other. It is clear that the VBM predicts modulus ratios greater than 1 more often than the BSM, as is demonstrated by specimens 498 and 503 and 512. Although these ratios are greater than 1 for the VBM, the differences between the BSM and VBM ratios for specimens 498 and 503 are small and are due in part to the inherent boundary condition in the VBM mesh⁸ which is not present in the BSM mesh. The modulus ratios produced by the BSM when the boundary conditions are varied are discussed in Section 8.4.3.

Despite the fact that the VBM produces more modulus ratios which are greater than 1, there are multiple cases where the BSM provides a higher, and therefore more accurate, prediction of the modulus ratio than the VBM. Specimens 504, 510, 514 and 534 all produce higher modulus ratios using the BSM than they do using the VBM. The VBM outlier, specimen 535, is more accurately simulated using the BSM, whereas the BSM outlier, specimen 513, is more accurately simulated using the VBM.


Overall, both the BSM and the VBM proved to be capable of predicting the modulus ratio of trabecular bone specimens, with the VBM more regularly predicting the correct relationship between apparent and inner moduli than the BSM. Considering the VBM has accurately predicted the inner modulus of trabecular bone specimens in previous work [15, 83], the incorrect

⁷Note: Specimen 508 produced a negative result, so is not shown in Figure 8.14.

⁸The inherent boundary condition is discussed in Appendix G.



(a) Apparent modulus.



(b) Inner modulus.

Figure 8.13: Comparison of VBM and experimental results for large bone specimens.




Figure 8.14: Comparison of VBM and experimental results for the modulus ratio of the large bone specimens with data labels.

predictions provided by the VBM in this work must be due to the larger voxel size⁹ and issues with segmentation as discussed in Section 8.1.2. These issues both impact the quality of the image which, in turn, affects how well the Finite Element mesh represents the bone structure. Consequently, if these factors are improved upon, the results produced by both discretisation methods will improve.

8.2.3 Resource allocation

The primary reason the BSM was developed was to decrease the time and memory required for Finite Element simulations of trabecular bone specimens [2]. This section compares the time taken and memory used for the BSM and VBM simulations.

8.2.3.1 Time

The results in Section 7.3.1 clearly show that the VBM discretisation is significantly faster than the BSM discretisation. On average, the BSM showed a 26-fold increase in discretisation time over the VBM for the standard specimens and a 47-fold increase for the large specimens. The increased discretisation time is the expected result as the BSM discretisation process is more complex than the VBM discretisation process, and this phenomenon is mentioned explicitly in

⁹The results in previous studies were produced using scans with much smaller voxel sizes than was used in this work. In [83], a voxel size of 22 μm was used; in [15], a voxel size of 37 μm was used; in the current work, a voxel size of 80 μm was used.




Figure 8.15: Comparison of experimental, BSM and VBM results for the modulus ratio of the large bone specimens.

the paper by Vanderroost, et al [2].

It is important to note that one of the mesh refinement processes in the BSM analysis cycle was not included in the VBM analysis cycle, which skews the discretisation time results. This phenomenon is discussed in more detail in the following section. The removal of separable components in the FE mesh, described in Section 3.2.2.2 and Appendix B, is an important step in the analysis cycle as it removes unnecessary, and possibly problematic, free element clusters from the mesh. The removal of free clusters was originally included in the VBM analysis cycle, however it was found that MATLAB did not have the memory required to perform this type of analysis for the large bone specimens. As such, to maintain consistency amongst the VBM results, this section of code was removed from the analysis cycle for both the standard and large bone specimens. This dramatically decreased the discretisation time for the VBM specimens. For example: when the separable components code was included in the VBM discretisation process, the time taken to discretise specimen 15 was 11 hours and 25 minutes; once this section of code was removed, the time taken for the discretisation of the same specimen was 49 seconds¹⁰.

There is an argument to be made that the separable components code could be optimised, however, for the purposes of this analysis the important item to note is that the BSM discreti-

¹⁰The time taken for the identification and removal of separable components is extremely high for the VBM because of the number of nodes on each element as well as the number of elements in the mesh. Each element produced using the VBM has 8 nodes, whereas the elements produced using the BSM have 2 nodes or 3 nodes. As a result, for each element in the VBM mesh the separable components code has to check the elements list 8 times to find all the connected elements, whereas the list only needs to be checked 2 or 3 times in a BSM mesh. Additionally, the code has to loop through significantly more elements in a VBM mesh, than it does in a BSM mesh.

sation time includes the removal of separable components, whereas the VBM discretisation time does not. Considering this fact, it is likely that the increase in discretisation time for the BSM over the VBM is exaggerated in these results.

In contrast to the discretisation time, the BSM simulation time for the specimens was much lower than the VBM simulation time. On average, the BSM showed a 5-fold decrease in simulation time over the VBM for the standard specimens and a 29-fold decrease for the large specimens.

Despite the fact that the BSM discretisation took much longer than the VBM discretisation, the time saved during the BSM simulation was significant. This meant that the time taken for the combined BSM discretisation and simulation was lower than the equivalent VBM time for 61% of the standard specimens and 82% of the large specimens. Considering that the BSM discretisation time includes the removal of separable components, where the VBM does not, it is clear that the BSM is significantly more time-efficient than the VBM. The time-saving effect of using the BSM is compounded with an increase in the size of the bone specimen, as there is a significant increase in the time taken for the VBM simulation.

If multiple simulations are to be performed on a single bone specimen, there can be no argument that the BSM is the better choice where time is concerned. Even with a single simulation, the BSM is likely to be quicker than the VBM, especially as specimen size increases.

8.2.3.2 Memory

Both the BSM and VBM suffered from memory issues, albeit in different ways. The 700M word base memory allocation was sufficient for the successful simulation of the small VBM specimens, but was insufficient for the factorisation of the stiffness matrix during the simulation. The stiffness matrix factorisation was forced to occur in out-of-core mode, which severely decreased the performance and, consequently, increased the time taken for the simulations. The large VBM specimens required memory allocations which were 185 – 471% bigger than the base memory allocation. Even with the additional memory, the simulation for specimen 508 could not be performed as it required a memory allocation beyond the specifications of the computer used for the work.

The 700M word allocation was more than sufficient for the small BSM simulations, which allowed the stiffness matrix to be factorised in-core and significantly reduced the time taken for the simulations. However, the large BSM simulations also suffered from a memory issue. Because of the massive number of section definitions in the BSM meshes, a large amount of memory is required to import the mesh information into LS-DYNA. For some of the large specimens, the memory required to import the mesh is too high, which prevents the simulation from occurring. This is different to the VBM memory issue, where the mesh is imported into LS-DYNA, but there is insufficient memory to run the simulation. Despite increasing the memory allocation for the large specimens, this BSM memory issue is responsible for 2 of the 6 large specimen failures.

Overall, the BSM specimens required and used much less memory for the simulation process than the VBM specimens. However, the memory issues discovered while simulating the large specimens, using both the BSM and VBM, provide a barrier to successful FEA of trabecular bone which necessitates further research.

8.3 Microstructural indices


The results for the microstructural indices analysis of the experimental results, BSM results and VBM results are shown in Figures 7.15-7.19. This comparison of experimental, BSM and VBM apparent moduli to microstructural indices is, to the knowledge of the author, the first of its kind in the literature¹¹. The comparison between the apparent moduli microstructural indices for the experimental, BSM and VBM results is significant. When considering the experimental results, there appears to be a weak connection between the bone volume fraction (Figure 7.15a) and the apparent modulus, as well as the trabecular separation (Figure 7.17a) and the apparent modulus, however there is no discernible link between the other microstructural indices and the experimental results as shown in Figures 7.16a, 7.18a and 7.19a.

In contrast, the VBM results show a strong connection to the bone volume fraction (Figure 7.15c), trabecular thickness (Figure 7.16c) and trabecular separation (Figure 7.17c), which matches reports in the literature [67]. Similar to the VBM results, the BSM results show a correlation to the bone volume fraction (Figure 7.15b), trabecular separation (Figure 7.17b) and, to a lesser degree, the trabecular thickness (Figure 7.16b). The BSM correlation to the trabecular separation and trabecular thickness is not as strong as the VBM correlation, but a trend is clearly distinguishable from the data. Neither the BSM nor the VBM showed any discernible link between the apparent modulus results and the connectivity or connectivity density, which correlates with reports in the literature [63].


Figure 8.16 shows a comparison of the experimental, BSM and VBM results to the bone volume fraction and includes a least squares curve fit for each data set. A straight line was fitted to both the experimental and VBM results, however an exponential curve provided the best fit to the BSM data. There is too much scatter in the experimental data for any useful inferences to be made regarding the connection between experimental results and bone volume fraction. The scatter in the BSM and VBM graphs is significantly less than the scatter in the experimental results graph, with both data sets showing a correlation coefficient greater than 77%. The decreased scatter in the simulation results is expected as both modelling methods are based directly on the identification of bone voxels in the TIFF stack, which is a similar process to the calculation of the bone volume fraction. The scatter should be particularly small in the VBM results, as is shown in Figure 8.16c as an element is created for each bone voxel in the TIFF stack.

Assuming that there is a linear relationship (weak or strong) between the bone volume fraction and the apparent modulus, as is suggested by the results produced by the experimental results and VBM results, it is strange that the BSM results should have a non-linear correlation to the volume fraction. If a linear fit is considered for the data instead, the results suggest that the BSM tends to overestimate the apparent modulus of specimens with higher volume fractions. This result correlates to the results seen in Figure 8.11b, where the BSM predicted higher apparent moduli than the VBM for VBM moduli greater than 0.4 *GPa*, causing a break from the trend seen in the results from 0 *GPa* to 0.4 *GPa*. Additionally, the five clear outliers in the comparison of VBM apparent modulus to bone volume fraction (Figure 8.16c) correlate to the outliers in the VBM results shown in Figure 8.7. This serves as further proof that these high moduli are caused by a loss of accuracy in the VBM.


¹¹Wang, et al [87] presented a comparison of BSM and VBM apparent moduli to microstructural indices, but did not include a similar comparison for experimental results. The BSM used in [87] is not the same method used in the current work.



(a) Experimental.



(b) BSM.



(c) VBM.

Figure 8.16: Comparison of apparent modulus results and bone volume fraction with fitted curve. The five outlying results are excluded from the line of best fit analysis for the VBM specimens.

Based on this data, it seems that bone volume fraction can be used to obtain an estimate of the simulated apparent moduli in this dataset, however the volume fraction cannot be used to determine the experimental apparent modulus. Additionally, considering there is a tight correlation between the simulation results (both BSM and VBM) and the volume fraction which is not seen in the experimental results, the effectiveness of a modulus prediction based on microstructural indices must be questioned.

Figure 8.17 shows a comparison of the experimental, BSM and VBM results to the trabecular separation and includes an exponential least squares curve fit for each data set. Similar to the results shown in Figure 8.16, the experimental results show the most scatter of the three result sets. Even though a weak trend may be seen in the data, there is too much scatter to allow for an accurate prediction of apparent modulus to be made based on the trabecular separation. The BSM results show less scatter than the experimental results, however there is still too much scatter for an accurate prediction of the apparent modulus to be made based on the trabecular separation. The VBM results show the strongest link to the trabecular separation, however the usefulness of this link is questionable considering the experimental results do not show the same link.

Another interesting result stemming from the analysis of the microstructural indices is the comparison between the BSM results and the shell-to-beam (S2B) ratio. Because the S2B ratio is calculated directly from the BSM mesh, a strong correlation was expected between the S2B ratio and the BSM results. Although this holds true for small S2B ratios, the trend is broken at higher S2B ratios as shown in Figure 7.20b. In order to investigate this phenomenon further, BSM results were plotted against the experimental results and the specimens with S2B ratios above 10 were highlighted, as shown in Figure 8.18.


Figure 8.18 clearly shows that all of the significant scatter in the BSM results stems from the specimens with high (≥ 10) shell-to-beam ratios. Because the value of the ratio is calculated directly from the volume of shells and beams in the discretised mesh, it is unclear whether the specimens are inherently shell-dominant which causes the high shell-to-beam ratio, or if the BSM discretisation erroneously includes a higher proportion of shells than is required. Either way, these results show that the BSM is likely to produce less accurate results at high S2B ratios.

8.4 Factors affecting simulation results


Although there are many factors which influence the accuracy of a simulation, three factors were chosen for investigation in this work, namely voxel size, segmentation technique and boundary conditions. The aim of the investigation was not to determine which factors make the simulation most accurate, but rather to build a general understanding of the extent to which each of these factors could impact the simulation results. As a result, the influence of these factors on the results is discussed with the intention of establishing and recommending research areas for future work.

8.4.1 Voxel size


The results from the analysis of the scans at varying voxel size, displayed in Section 7.5, clearly show that the choice of voxel size can have a significant impact on the quality of the images



(a) Experimental.



(b) BSM.



(c) VBM.

Figure 8.17: Comparison of apparent modulus results and trabecular separation with fitted curve. The five outlying results are excluded from the line of best fit analysis for the VBM specimens.




Figure 8.18: Comparison of BSM and experimental results for standard bone specimens with high S2B ratios highlighted.

obtained from the scanning process. There is a clear visual difference between the images with smaller voxel sizes (i.e. $20\ \mu\text{m}$ and $40\ \mu\text{m}$) and those with larger voxel sizes ($80\ \mu\text{m}$), as shown in Figures 7.21 and 7.22. The $20\ \mu\text{m}$ and $40\ \mu\text{m}$ images also show distinct differences in clarity, however visually the difference between these scans is not as significant as the difference between the $40\ \mu\text{m}$ and $80\ \mu\text{m}$ scans and it appears as if most of the structural complexity captured in the $20\ \mu\text{m}$ scan is captured in the $40\ \mu\text{m}$ scan. This observation may mean that the $80\ \mu\text{m}$ voxel size used for the micro-CT scans in [3] was insufficient. However, by using a larger voxel size, a large number of specimens could be scanned, allowing for a thorough comparison of experimental and numerical techniques in the current work, which was advantageous.

If the images in Figure 7.22 are carefully considered, subtle structural differences between the images may be identified. This structural variation is due to the varying voxel size used to scan the image, but may also be attributed to slight differences in the chosen regions of interest (ROIs), which is influenced by the varying voxel size. In order to designate the ROI, an initial image slice was chosen¹² for each of the voxel sizes of a specimen, based on structural features in different regions of the image. However, because of the difference in voxel size, not all of the regional features in the $80\ \mu\text{m}$ image slice will be visible in the $20\ \mu\text{m}$ image slice, e.g. a structural feature at $50\ \mu\text{m}$ will be included in the first slice of the $80\ \mu\text{m}$ scan, but will only be included in the third slice of a $20\ \mu\text{m}$ scan. Consequently, initial slices were chosen based on best overall visual agreement and, hence, there may be a slight variation in the ROI chosen for the various voxel sizes. This variation is unlikely to have a significant effect on the results, as it is extremely small (likely $\leq 80\ \mu\text{m}$), meaning that the chosen ROIs are, overall, representative of the same volume of material.

An interesting phenomenon was observed when full images of the bone specimen were compared at varying voxel sizes, which is demonstrated in Figure 8.19. In the $80\ \mu\text{m}$ scan shown in Figure 8.19c, it appears as if the bone is solid around its circumference. However, if the $20\ \mu\text{m}$ and $40\ \mu\text{m}$ images are considered (Figure 8.19a and 8.19b, respectively), it may be seen that this thicker circumference stems from bone shards which must have broken off during machining of the specimens. These shards are small, but plentiful, which makes it difficult to distinguish

¹²Once the initial image slice was chosen, the ROI was numerically determined using in-house code developed for this purpose, so as to ensure consistency.




Figure 8.19: Comparison of micro-CT scans of specimen 518 at varying voxel size.

them from the bulk of the bone at large voxel sizes. Consequently, in the 80 μm scan, these machining artifacts manifest as a solid ring around the circumference of the specimen.

The discovery of the machining artifacts in the images with small voxel sizes explains some of the poor segmentation that was discussed in Section 8.1.1. During segmentation, these machining artifacts are identified as bone voxels, which is technically correct, however because of the larger voxel size, these bone voxels are erroneously attached to the bulk of the bone specimen, thereby creating a thick ring at the circumference of the specimen which links to the circumferential overthickening discussed in relation to Figures 8.4b, 8.4f and 8.4l. Physically, these bone shards are not connected to the bulk of the specimen, which means that experimentally they do not add stiffness to the bone, however when the specimen is simulated, these shards will be included in the mesh and add stiffness to the bone. In the large bone specimens¹³, this circumferential thickening will increase the value of the apparent modulus, but will not affect the value of the inner modulus, thereby decreasing the simulated modulus ratios, as discussed in Sections 8.1.2 and 8.2.2.

The discovery of the machining artifacts in the large voxel size scans is important as it may inform future research. In order to obtain an accurate simulation result, these artifacts must be removed from the scan before discretisation and simulation, using numerical or experimental techniques. This can be achieved either by numerical cropping of the micro-CT scan (which is discussed further in Section 8.5), or defatting of the specimens before micro-CT scanning.

The microstructural indices measured from the scans clearly show a bigger difference between the 40 μm and 80 μm scans than between the 20 μm and 40 μm scans, which reinforces the observation that a voxel size of 80 μm may be insufficient. Table 8.1 shows the percentage difference between the 40 μm and 80 μm indices¹⁴ when compared to the 20 μm indices. There is a clear increase in the bone volume fraction as the voxel size increases. Although the increase in volume fraction is negligible from 20 μm to 40 μm , the increase experienced from 20 μm to 80 μm is significant (> 15%).

The trabecular thickness and spacing results are equally interesting. For specimen 477, the trabecular thickness increases from 20 μm to 40 μm and increases more significantly from 40

¹³Examples of circumferential overthickening in large bone specimens may be seen in the images of specimens 529, 530, 532 and 538 in Appendix I.

¹⁴To calculate the percentage difference in trabecular thickness and spacing, the values were scaled up based on their size, i.e. a 40 μm voxel is twice the size of a 20 μm voxel, therefore the percentage difference (d) was calculated as:
$$d = \frac{2(TbTh_{40}) - (TbTh_{20})}{(TbTh_{20})}$$

Table 8.1: Percentage difference in microstructural indices when compared to 20 μm indices.

Sample number	Voxel size (μm)	BV/TV (%)	Mean Tb.Th (%)	Mean Tb.Sp (%)
477	40	2.36	7.20	0.510
	80	19.5	37.2	-3.37
518	40	0.316	12.5	104.7
	80	15.5	-30.6	108.6

Table 8.2: Increase in apparent modulus results from 40 μm to 80 μm .

Sample number	Apparent modulus (MPa)		Increase (%)
	40 μm	80 μm	
477	94.06	105.0	11.63
518	84.57	96.11	13.65

μm to 80 μm . Specimen 518 also shows an increase between the trabecular thickness at 40 μm and 20 μm , but shows a significant decrease in trabecular thickness at 80 μm . This decrease in trabecular thickness is particularly unusual because there is an increase in the bone volume fraction. Essentially, the combination of these two factors means that the trabeculae in the 80 μm image are thinner than the trabeculae in the 20 μm image, but there are more trabeculae present in the 80 μm image than in the 20 μm image, possibly due to small trabeculae created as a result of poor marrow rasterisation.

Similar to the trabecular thickness, the trabecular separation also generally shows an increase with an increase in voxel size, except for the scan of specimen 477 at 80 μm , which shows a decrease. The results for the trabecular separation for specimen 518 are clear outliers in Table 8.1. The 20 μm image for this specimen shows a noticeable amount of noise, as shown in Figure 8.20, so it is postulated that the trabecular separation determined for the 20 μm image is erroneous, rather than the 40 μm or 80 μm values, which show good agreement with each other. It is clear that although the bone volume fraction shows a steady increase with increasing voxel size, the trabecular thickness and trabecular separation are not so predictable.

Finally, the comparison of the apparent modulus results in Table 8.2 shows that there is an increase in the apparent modulus as the voxel size increases. Specimens 477 and 518 show increases of 11.6% and 13.7%, respectively between the 40 μm and 80 μm results¹⁵. This result clearly shows that the apparent modulus predictions at 80 μm are too large, which contradicts the findings by Chen, et al [40] and supports the observation made in Section 8.1.2 based on the apparent and inner modulus results (Figure 8.6) for the large specimens. Consequently, if the voxel size is decreased, the apparent modulus produced by the BSM will decrease¹⁶ and the modulus ratio will increase, causing more correct predictions of the modulus ratio than were seen in Section 8.1.2.

¹⁵Although meshes were successfully constructed for the 20 μm images for both specimens, the memory required to import the mesh information into LS-DYNA was too high, which caused the simulation to fail.

¹⁶Decreasing the voxel size will increase the inner modulus of the specimen to a lesser degree than the apparent modulus due to the improved segmentation seen at smaller voxel sizes [14]. Improved segmentation will not only allow the bone to be better represented in the FE mesh, but will also minimise circumferential overthickening which will decrease the apparent modulus, but not the inner modulus.




Figure 8.20: Noise in $20\ \mu\text{m}$ scan of specimen 518. Although this noise will not impact the bone volume fraction, the sphere-filling algorithm used to determine the trabecular separation will be significantly affected, causing a corresponding decrease in the trabecular separation reported for this scan.

Although only two specimens were used in this analysis, it is clear from this small subset that the voxel size is a key component in both a morphometric and numerical analysis of trabecular bone.

8.4.2 Segmentation technique

The difference in the defatted, marrow and confined specimen results, discussed in Sections 8.1.1 and 8.2.1, clearly show that the quality of the segmentation can have a noticeable effect on the results of a simulation, regardless of whether the BSM or VBM is used. Figure 8.21 shows the comparison of the apparent modulus results obtained using the *Optimise Threshold* segmentation function to the results obtained using the *Make Binary* segmentation function which was used throughout this work. A large portion of the *Optimise Threshold* results do not differ significantly ($< 10\%$) from the *Make Binary* results, however there is a significant portion of the specimens which show an increase or decrease greater than 10% as shown in Figure 8.22 and approximately 13% of the specimens show an increase in modulus which exceeds 20% . This difference in value of the apparent moduli for the two segmentation methods generally agrees with the work by Hara, et al [45] and highlights the importance of accurate segmentation in the image processing phase.

Figure 8.23 shows the comparison between the modulus ratio results obtained using the *Make Binary* function and the modulus ratio results obtained using the *Optimise Threshold* function. Clearly, the effect of using a different segmentation technique on the relationship between the apparent modulus and inner modulus is negligible. Considering these results and the results seen in Figure 8.21, it seems that the choice of a different segmentation technique may increase the scatter in the results and possibly change the value of the back-calculated Young's modulus, but not change the way the apparent modulus and inner modulus relate to each other.




Figure 8.21: Comparison of BSM and experimental apparent modulus results for standard bone specimens segmented using the *Make Binary* and *Optimise Threshold* functions.




Figure 8.22: Percentage increase in apparent modulus results using the *Optimise Threshold* function versus the *Make Binary* function.




Figure 8.23: Comparison of BSM and experimental modulus ratio results for large bone specimens segmented using the *Make Binary* and *Optimise Threshold* functions.

The *Optimise Threshold* function was used for another set of bone simulations, where a lateral constraint was added on the top and bottom surfaces of the specimen. The effect of the boundary condition on the *Optimise Threshold* results (Figure 7.25) was extremely similar to the effect of the boundary condition on the *Make Binary* results (Figure 7.29).

Another result worth noting in the segmentation analysis is the decreased number of failed simulations using the *Optimise Threshold* function (10 failures) when compared to the *Make Binary* function (18 failures). The *Make Binary* function calculates a threshold value for each image slice in the TIFF stack, whereas the *Optimise Threshold* function calculates a threshold value for the entire TIFF stack. It is likely that the calculation of a threshold value for each image slice results in discontinuities in the structure, which the BSM is incapable of handling. This indicates that the segmentation technique can not only affect the magnitude of the results obtained using the BSM, but also the rate of success of the simulations.

8.4.3 Boundary conditions

The boundary condition results displayed in Section 7.7 clearly show that the addition of constraints to the mesh can have a large impact on the apparent modulus and inner modulus produced by the simulation. The influence of the change in constraint will be discussed in this section with reference to the apparent modulus, inner modulus and modulus ratio, respectively.




Figure 8.24: Increase in stiffness caused by “TB - lat/rot” boundary condition compared to the “Standard” boundary condition results.

8.4.3.1 Influence on apparent modulus

The results from the boundary condition analysis presented in Section 7.7 show that the addition of a rotational constraint had a negligible effect ($< 3\%$) on the apparent modulus of the bone specimens, regardless of the surface/s on which the constraint was applied. Similarly, the addition of a lateral constraint on the top or bottom surface only increased the apparent modulus by approximately 5%. However, the addition of a lateral constraint on the top and bottom surfaces simultaneously produced a significant increase of, on average, 15%. If rotational constraints on the top and bottom surfaces were added to these lateral constraints, the bone specimens experienced an increase in stiffness of approximately 16%.

The application of lateral and rotational constraints simultaneously to the top and bottom surfaces (hereafter referred to as the “TB - lat/rot” boundary condition) produced an average increase in stiffness of 16%, however it is important to consider the full spectrum of the results and not just the average value to determine how much variation there is in the sensitivity of the bone specimens to constraint. To this end, Figure 8.24 shows a histogram of the percentage increase in stiffness between the “Standard” boundary condition results and the “TB - lat/rot” boundary condition results. The maximum increase caused by the boundary condition is 50.4%, but most of the specimens experience an increase in the region of 5 – 25%.

Specimens 119, 124, 332, 365, 435, 440 and 466 all show an increase in stiffness which is greater than 25%. With the exception of specimen 332, which is an outlier in both the BSM and VBM results, all the results fall relatively close to the trend line of the BSM results as shown in Figure 8.25 and, therefore, are considered to be good data points. Interestingly, all of these specimens belong to either the marrow or confined data sets, which means that the specimens are all marrow-inclusive. This result seems to indicate that simulating marrow-inclusive specimens




Figure 8.25: Comparison of BSM and experimental results for standard 10 *mm* bone specimens with line fit.

will not only increase the scatter in the results, but may also cause the specimen meshes to be more sensitive to constraint. However, it may be that these particular specimens are inherently more sensitive to constraint and, therefore, the BSM results are simply capturing this inherent behaviour. Unfortunately, there is no definitive way to confirm this because corresponding experimental data is only available for one boundary condition set-up, however this phenomenon should be investigated in future work.

It is assumed that the experimental boundary condition falls somewhere between the minimum “Standard” boundary condition and the maximum “TB - lat/rot” boundary condition due to friction between the top and bottom surfaces of the bone specimens and the plates of the testing equipment. Considering that the maximum simulated boundary condition increases the apparent stiffness of the majority of the specimens by 5 – 25%, the error associated with assigning boundary conditions to the mesh which do not exactly replicate the experimental boundary conditions should be in the same range or less for the majority of specimens. Although determining the exact error associated with the application of non-realistic boundary conditions is outside the scope of this work, simulating multiple boundary conditions to establish a range of error provides a useful estimate of the error associated with these non-realistic boundary conditions. This error must be taken into account when discussing the correlation of simulation results to experimental results.

8.4.3.2 Influence on inner modulus and modulus ratio

The impact of the different boundary conditions on the inner modulus of the specimens is very different from the impact on the apparent modulus, as shown in Figure 7.28 and Table 7.4. There are three major differences:

1. The addition of a constraint consistently increased the apparent modulus of the specimens, regardless of the type of constraint or the surface on which it was applied. In the large specimen results, the addition of particular constraints caused a decrease in the inner modulus of some of the specimens.
2. The change in the apparent modulus results with the varying boundary conditions was

not exactly the same for each specimen, but there was some consistency, particularly regarding which boundary condition caused the smallest/largest change. The influence of the boundary conditions on inner modulus was not as consistent.

3. The largest average increase in the apparent modulus results was 16%, whereas the largest average change in the inner modulus results was 46%.

In the apparent modulus results, the standard boundary condition consistently produced the lowest apparent modulus of all of the boundary conditions. Although this is true for the inner modulus of some of the specimens, in the majority of cases, the addition of a lateral constraint or a combined lateral and rotational constraint on the bottom surface of the specimen caused a decrease in the inner modulus as shown in Figure 7.28, meaning that the standard boundary condition no longer produced the lowest value for inner modulus. This result may seem strange initially, however a similar phenomenon was seen in the lattice results discussed in Section 6.3.1.4, where the addition of lateral constraints on the top and bottom surface of the lattice caused a decrease in the inner modulus of the lattice. **This is just one example of how having a full set of validation simulations can lend confidence to the results produced for the bone specimens.**

In general, as with the apparent modulus results, the inner modulus of a specimen was more sensitive to a lateral constraint than a rotational constraint, with all rotational constraints causing an average change less than 10% as shown in Table 7.4. The lateral constraint on the bottom surface only caused a 5% change in results, however the addition of a lateral constraint on the top or combined surfaces caused changes of 29% and 41%, respectively, which is not only significant in the realm of the inner modulus results, but is a much more significant change than was seen in any of the apparent modulus results (Table 7.3). Overall, a constraint on the bottom surface of the specimen had a negligible impact (< 5%) on the inner modulus, whereas a constraint on the top surface had a significant impact (up to 33%) on the inner modulus. Constraints applied simultaneously on the top and bottom surfaces of the specimen had the greatest impact, showing changes up to 47%.

It is important to consider how the changes in apparent and inner modulus caused by the varying boundary conditions affect the modulus ratio as this is the only metric which does not rely on the Young's modulus of the specimen. There was an expectation that the increased constraint on the specimen would cause a significant increase in the modulus ratios, which would allow them to correlate better with the experimental results [3], however this was not the case as shown in Table 7.5. The modulus ratios experienced a far less significant change due to the varying boundary conditions than the inner moduli, with the largest average change in results being 11% for a lateral constraint on the top surface of the specimen.

The addition of a lateral constraint on the top surface caused a large enough increase in three of the modulus ratio results to push them above a value of 1, i.e. the stiffness relationship of the specimens changed from $E_a > E_i$ to being $E_i > E_a$ as shown in Figure 7.29a. This increases the number of specimens showing the correct stiffness relationship from 2 out of 11, to 5 out of 11 for the BSM, which is equivalent to the number of specimens which showed the correct stiffness relationship using the VBM¹⁷. This is an important result as it shows that **the BSM is equally as capable of predicting the modulus ratio of trabecular bone specimens as the VBM, provided that appropriate boundary conditions are applied.**

¹⁷Interestingly, the specimens which show the correct stiffness relationship are slightly different for the BSM and the VBM. The BSM and VBM both show the correct relationship for specimens 503, 510, 512 and 534, however the BSM shows the correct relationship for specimen 504, where the VBM does not and, similarly, the VBM shows the correct relationship for specimen 498, where the BSM does not.

Overall, the addition of constraints to the mesh had a more significant impact on the inner modulus than it did on the apparent modulus. When a constraint was added, the apparent modulus consistently increased, whereas the inner modulus increased or decreased, depending on the specimen. The variation in the apparent modulus and inner modulus results meant that the overall change in modulus ratio was relatively small ($< 11\%$), however this small change could have a significant effect on the stiffness relationship of the bone by increasing the modulus ratio to a value above 1.

8.5 Difference in results

It is clear that the BSM results in this work show some phenomena which were not seen in the BSM results generated by Vanderoost, et al [2]. Specifically, the BSM scaling factor of 1.605 in the Vanderoost study was found to be approximately 2 in the current work. Furthermore, there appears to be more scatter in the BSM vs VBM comparison generated in this work than there was in the work by Vanderoost (although this scatter is, at least visually, vastly reduced by the use of a log scale). The following section will discuss some of the possible causes for the variations between datasets.

The first factor that must be taken into account is the type of bone used in the studies. Vanderoost, et al [2] studied human trabecular bone from a variety of different anatomical sites, whereas the current work considered bovine bone from one anatomical location, the femur. Initially, this difference may seem insignificant given that the validation study performed in the current work showed that the BSM is capable of capturing the response of a wide variety of structures, however the structural differences in human versus bovine trabecular bone may still have an affect on the results. The human specimens used in the work by Vanderoost, et al [2] had an average bone volume fraction of 0.145, whereas the bovine specimens used in the current work had an average bone volume fraction of 0.347 indicating higher trabecular thickness. It is likely that the bovine specimens, which have high trabecular thickness and were scanned at large voxel sizes, show more significant overthickening in the micro-CT scans than the human specimens, which have lower trabecular thickness and were scanned at smaller voxel size. If the bovine scans show more significant overthickening than the human scans, the VBM, which creates a mesh directly from the image, would overestimate¹⁸ the apparent properties of the bovine specimens to a greater degree than in the human specimens, leading to a greater difference between BSM and VBM results and an increase in the scaling factor.

The second factor which must be considered is the shape of the samples. Square samples were used in the work by Vanderoost [2], whereas this work made use of cylindrical samples. Because of the variation in shape, changes had to be made to BSM code, which limited the 2D skeletonisation to the top and bottom surfaces of the specimen¹⁹ in the current work. Vanderoost, et al [2] report that the 2D skeletonisation on the boundaries of the specimen increases the accuracy of the simulation, therefore the removal of some of the 2D skeletonisation to facilitate the simulation of cylindrical specimens may be causing a decrease in the accuracy of the results.

¹⁸In the validation study it was found that overthickening in images causes the VBM to overestimate the apparent modulus of the structure, meaning that the BSM provides a more accurate result. See Section 6.2.6 for more details.

¹⁹Note: this edited code was used to analyse the cylindrical bone samples, but was also used to analyse the rectangular lattices used in the validation study.




Figure 8.26: Comparison between micro-CT scans used in the Vanderoost study [2] and the current work. No machining artifacts are visible on the scan from the Vanderoost study [2], however small pieces of bone are clearly visible around the edge of the machined specimens used in the current work.

The shape of the sample brings to light another variation between the Vanderoost study and the current work. Judging by the micro-CT scans used in the Vanderoost study (shown in Figure 8.26a), it is likely that these scans represent regions of interest which have been removed from a larger image with the aid of image processing software. The cropping of the scans is not explicitly mentioned in the Vanderoost study, however this is a common technique used in image processing of trabecular bone and the lack of visible machining artifacts on scans used by Vanderoost (an example of which is shown in Figure 8.26b) makes it likely that this cropping process was followed. The inclusion of the machining artifacts manifests as increased thickness around the circumference of the specimen in lower resolution scans, as was discussed in Section 8.4.1, and will cause an increase in the apparent modulus predicted by the BSM. Despite this, it was important for the full specimens to be used in the current work so that a comparison could be drawn between the simulation results and equivalent experimental results - particularly because such a comparison was missing from the Vanderoost study.

The cropping of the specimen is not the only image-related difference between the current study and the Vanderoost study. In the current work, a voxel size of $80 \mu m$ was used during micro-CT scanning of the specimens, whereas in the Vanderoost study, a voxel size of $30 \mu m$ was used. It has been shown in the current work as well as in literature (see Section 2.2.1) that the voxel size of a micro-CT scan can have a large impact on the results. Consequently, it makes sense that the Vanderoost study would produce less scatter as it was based on images with a smaller voxel size. Additionally, as mentioned previously in this section, the voxel size seems to effect the scaling factor between the VBM and BSM results, as it determines how much the thickness of the structure is overestimated during segmentation.

Clearly, there are a multitude of reasons which could explain the difference between the results in the Vanderoost study and the results presented in the current work. It is likely that a combination of these factors affected not only the correlation of the BSM to the VBM, but also the correlation of both simulation methodologies to experimental results.

Chapter 9

Conclusions

This work has taken important steps in investigating the capabilities of the beam-shell method in predicting the mechanical properties of trabecular bone. Specifically, it has addressed the limitations in the literature, discussed in Section 1.2, by:

1. Thoroughly validating the developed BSM methodology before applying the methodology to bone.
2. Comparing results produced by the BSM and VBM to equivalent experimental results.
3. Investigating the effective properties of the bone, using both the BSM and the VBM.

Although somewhat similar work has been performed by Wang, et al [87] using a BSM created by that research group, the validation study performed for the current work included a greater variety of structures and image sizes and, overall, a larger number of validation simulations were performed in the current work than in [88]. Similarly, the current work simulated just over double the number of bone specimens as were simulated in the work by Wang, et al [87]. Finally, the current work is the first study which draws a comparison between the effective modulus obtained experimentally and numerically, using both the BSM and VBM.

A general conclusion is provided at the outset of the chapter, followed by a summary of the work and conclusions drawn from various aspects of the work. Closing remarks are included at the end of the chapter to emphasise the most important findings.

9.1 General conclusion

This work has shown that the apparent modulus of trabecular bone may be accurately predicted using a beam-shell Finite Element model. Currently, the scatter produced by the BSM for apparent modulus is significant, but improved segmentation and the removal of machining artifacts before scanning has been shown to significantly decrease the scatter, making the BSM a viable alternative to the computational “gold-standard”, the VBM.

In addition to the excellent prediction of apparent modulus, the BSM has also showed promise in the prediction of the effective modulus of the bone, where the effective modulus is defined as the modulus of an inner core region of the bone. The BSM predicted the correct relationship between apparent and effective moduli as regularly as the VBM, given that appropriate

boundary conditions were applied. This correlation provides further confidence in the BSM as an alternative to the VBM.

The conclusions drawn above confirm the hypothesis provided in Section 1.4, however further research in this field is necessary to refine the results produced by the BSM.

9.2 Validation

In this work, an automated analysis cycle was developed which allowed the apparent modulus of a structure to be determined from a micro-CT scan using a beam-shell Finite Element simulation. The analysis cycle was validated using over 3000 artificially generated cubic, Kelvin cell and octet truss lattices in various configurations. The analysis cycle generally performed well, although there were issues which stemmed from the generation of the lattice images and post-processing of the results. Additionally, a directional preference was identified in the BSM which resulted in the generation of multiple extra nodes in the mesh. This phenomena was judged to be unlikely to occur in bone, as it only occurs at a very specific angle. Overall, the analysis cycle accurately captured the behaviour of a wide variety of structures displaying a significant range of directional dependence and constraint sensitivity.

The most interesting result to stem from the validation study is the fact that the VBM provided a less accurate prediction of the apparent modulus than the BSM for Kelvin cell lattices. Rasterisation of the lattices results in struts which are, at times, overly thick. Because the VBM mesh is directly based on the images, the VBM tends to predict a higher apparent modulus than the baseline value. In contrast, the BSM mitigates the overthickening of the struts during skeletonisation. Consequently, it stands to reason that the BSM provides a better prediction of apparent modulus than the VBM in cases where there is overthickening of the trabeculae.

The validation study was an integral part of this work, yet is not something that is often reported in the literature. The validation study carried out in this work provided valuable insight into the efficacy and accuracy of the developed analysis cycle and highlighted behaviour trends which were mirrored in the bone specimens. Consequently, the validation study was an essential part of this work, and similar validation studies are recommended for future work on trabecular bone.

9.3 Performance of BSM

As mentioned in Section 5.4, the performance of the beam-shell method was determined by considering three factors: the success rate of the simulations, the accuracy of the apparent and effective moduli predicted using the developed analysis cycle and the resource requirements of the discretisation and simulation processes. Each of these three factors is addressed here, with particular emphasis being placed on the accuracy of the results.

9.3.1 Success rate

106 standard bone specimens ($\phi 10\text{ mm}$) and 17 large bone specimens ($\phi 20\text{ mm}$ and $\phi 28\text{ mm}$, respectively) were run through the developed analysis cycle, making for a total of 123 specimens. Of these specimens, 18% could not be discretised by the BSM code, whereas 2% could

not be simulated due to memory constraints. This makes for a total of 20% of specimens which failed during the BSM analysis.

In contrast, all of the specimens were successfully discretised using the VBM, and only 1% of the specimens failed to simulate due to memory constraints. It is clear from these results that the BSM is less robust than the VBM in the area of discretisation. However, the BSM failures can be significantly reduced by scanning bone samples at smaller voxel sizes and choosing a segmentation technique which applies a single threshold to the TIFF stack, as opposed to applying a unique threshold to each image in the stack. By increasing the success rate of the BSM through these methods, the performance of the BSM is improved, making it a viable choice over the VBM for trabecular bone simulations.

9.3.2 Accuracy of results

The accuracy of the results was judged in two key measures, namely the apparent and effective moduli. Although the BSM has been used to determine the apparent modulus of bone specimens in previous work [2], this is the first time to the knowledge of the author that it has been used to determine the effective modulus using an embedding approach.

9.3.2.1 Apparent modulus

The apparent moduli generated by the BSM show a clear correlation to the experimental results, however the scatter in the data is significant, particularly at high moduli, and therefore an accurate result cannot be guaranteed for a BSM simulation of an arbitrary bone specimen. Despite the scatter in the results, a large portion of the data correlates well with experimental results, which shows that the BSM is capable of accurately predicting the apparent modulus. Various areas of improvement have been identified in this work which are capable of reducing the scatter in the results:

- **Voxel size.** This work has shown that apparent modulus increases as voxel size increases, due to partial volume effects and subsequent overthickening of the trabeculae. If bone specimens are scanned at smaller voxel sizes, the overthickening of the trabeculae can be limited, causing a decrease in overall scatter.
- **Defatting.** It was found that removing the marrow from the bone specimens before micro-CT scanning resulted in scans with better contrast between bone and non-bone voxels and improved segmentation. Additionally, machining artifacts around the edge of the specimen, which artificially increase the bone volume fraction and apparent modulus of the specimen, are likely to be removed during defatting. Consequently, defatted specimens are more accurately represented in the segmented image, which decreases the scatter in the results.
- **Segmentation technique.** A known issue associated with global segmentation is the overthickening of trabeculae around the edge of the specimen, relative to central trabeculae [14]. This overthickening causes an artificial increase in the bone volume fraction and apparent modulus of the sample. A more effective segmentation technique would limit this local overthickening, leading to a better representation of the bone structure and, consequently, less scatter in the results.

The apparent moduli generated by the BSM were also compared to apparent moduli generated from VBM simulations of the bone specimens. Because the Young's modulus of the bone specimens is unknown, an arbitrary Young's modulus of 1 *GPa* was assigned to the BSM and VBM simulations, making it impossible for the simulation data from either the BSM or the VBM to exactly predict the experimental value of the apparent stiffness, i.e. provide a 1 : 1 correlation with the experimental results. Back-calculation of Young's modulus was performed on both the BSM and VBM data to determine values for Young's modulus, however the simulation results were not adjusted based on these values as the goal was to provide a true, unedited comparison between the simulation results.

It was observed that there was a significant difference in the values predicted by the BSM and those predicted by the VBM, with corresponding back-calculated moduli for the simulation techniques being 8.30 *GPa* and 4.16 *GPa*, respectively. The difference between BSM and VBM results is mentioned in the work on the BSM [2], where a correction factor of 1.605 was applied to the BSM results to allow for correlation with the VBM results. It was found that this correction factor did not hold for the results generated in the current work, which required a correction factor of approximately 2. This difference may be due to a number of factors which vary between the current work and the Vanderroost study (see Section 8.5), with voxel size being considered a key factor.

The comparison of the BSM and VBM results highlighted some of the challenges associated with the simulation of trabecular bone. The uncertainty surrounding the Young's modulus of trabecular bone is a significant difficulty in any trabecular bone simulations. Because an arbitrary modulus (or range of moduli) is chosen for the simulation of a bone specimen, the simulated apparent modulus does not correlate with the experimental apparent modulus. Consequently, the confidence in a simulation result stems from its position relative to other data points, both experimental and numerical. For example, the poor simulation results produced by the BSM (scatter) and VBM (outliers) can only be identified because of their deviation from the trend line, which was determined by considering 123 simulation results and their corresponding experimental results. While there is debate around the Young's modulus of trabecular bone, it is essential for studies incorporating FEA of trabecular bone to discuss the accuracy of an individual simulation in relation to a collection of other data points, so as to provide confidence in the generated results.

Considering the results produced by the BSM and VBM, it is tempting to conclude that the VBM produces more accurate results than the BSM. However, this work has shown that this may not be the case. Although the VBM results generally show a stronger correlation to the experimental results than the BSM moduli, the results from the validation study showed that the VBM produced significantly less accurate results than the BSM for images containing struts which were overly thick. Considering that the segmentation of the bone specimens creates images which can be generally or locally overthickened, it may very well be that the BSM provides a more accurate prediction of apparent modulus than the VBM. This observation further strengthens the argument for using the BSM as an alternative to the VBM, particularly if the scatter associated with the results can be reduced.

9.3.2.2 Effective modulus

The novel aspect of this work involved determining whether the BSM was capable of capturing the effective modulus of trabecular bone specimens, where the effective modulus is defined as the modulus of an inner core region of the specimen. The VBM has been used previously [15] to

predict the effective modulus in this way, therefore the BSM must be able to provide a similarly accurate prediction if it is to be used as an alternative to the VBM. In this work, the apparent modulus and inner modulus were determined for each large bone specimen using both the BSM and the VBM, after which the modulus ratio was calculated by dividing the inner modulus by the apparent modulus. Instead of using the inner modulus (or effective modulus) as a basis for comparison, the ratio of the inner modulus and apparent modulus was used as this metric does not include the influence of Young's modulus. Using the modulus ratio allowed the simulation results and experimental results to be directly compared to each other, with no back-calculation of Young's modulus required.

Eleven large specimens were successfully simulated using the BSM and modulus ratios which showed the correct relationship between apparent and effective moduli were produced for 45% of the specimens, once appropriate boundary conditions had been applied. The VBM, similarly, produced the correct relationship between apparent and effective moduli for 45% of the specimens. Generally, the simulated modulus ratios produced by both the BSM and the VBM were too low, meaning that either the calculated apparent modulus was too high, or the effective modulus was too low. Considering the VBM has provided accurate predictions of effective moduli in previous work [15], the incorrect relationships produced for some of the specimens in the current work must stem from the images used for the analysis. When considering the images, it was found that the apparent modulus of the specimens was too high, because of:

- **Issues with segmentation**, which cause the trabeculae around the edge of the specimen to overthicken relative to the centre. The increased circumferential thickness increases the apparent modulus of the structure, but does not increase the effective modulus, as this value is measured using an inner core region in the specimen.
- **Machining artifacts** manifested in the form of bone fragments at the circumference of the specimen in the micro-CT scan. These bone fragments are erroneously attached to the bulk of the mesh during discretisation, as the 80 μm scans used for the majority of this work are too large to capture the separation. These machining artifacts result in further circumferential overthickening, which increases the apparent modulus but leaves the effective modulus unchanged.

The identified issues are identical to the issues described in Section 9.3.2.1. Consequently, by applying techniques to decrease the scatter in the apparent modulus results produced by the BSM, the modulus ratio results will be improved and accurate predictions of the relationship between apparent and effective moduli will be increased.

Although there is clearly work still to be done in this area, particularly in increasing the correlation of the simulation results (generated by both the BSM and the VBM) with the experimental results, this work has shown that the BSM is equally as capable of predicting the modulus ratio of trabecular bone specimens as the computational "gold-standard", the VBM. Combined with the results obtained from the analysis of apparent modulus, this result solidifies the view that the BSM is a viable alternative to the VBM.

9.3.3 Resource usage

The results show that the BSM discretisation took significantly longer than the VBM discretisation, however the VBM simulations took significantly longer than the BSM simulations. Despite the fact that one of the mesh refinement processes in the BSM discretisation could not be included in the VBM discretisation, 61% of the standard specimens were analysed (i.e. discretised

and simulated) quicker using the BSM than they were using the VBM. The difference in total analysis time increased as the specimen size increased, with 82% of the large specimens being analysed quicker with the BSM than the VBM. If multiple simulations were to be run on a single specimen, the BSM would further outperform the VBM.

The BSM was shown to be preferable to the VBM not only in time taken for the analysis, but also in the memory requirements for the discretisation and simulation processes. As mentioned previously, one of the mesh refinement processes had to be removed from the VBM code because the memory requirements for the refinement of the large specimens exceeded the available memory in MATLAB. Additionally, the VBM required the allocation of additional memory to its simulations, which was not necessary for the BSM specimens. Both the BSM and VBM experienced simulation failures due to memory issues, however overall the BSM was less computationally expensive than the VBM for both the discretisation and simulation processes.

As expected, the BSM outperformed the VBM in the areas of memory usage and time taken for the analysis of the bone specimens.

9.4 Microstructural indices

Overall, this work has shown that the experimental results show, at best, a weak correlation to the microstructural indices, whereas the simulation results can show very strong correlations to the indices. These results bring to light two contrasting questions, which will be stated and discussed in the following section.

The first question which stems from this work is: has the correlation between the microstructural indices and apparent modulus been overstated? The microstructural indices analysis performed in this work clearly shows that the simulation results generated using both the BSM and VBM correlate well to the bone volume fraction and trabecular separation. However, the simulation results, particularly the VBM, contain an inherent bias towards the microstructural indices as the Finite Element mesh is constructed directly from the TIFF stack, just as the indices are measured directly from the TIFF stack. In contrast, the experimental results, which do not contain the bias inherent in the simulation results, show little to no correlation to the indices. Although microstructural indices are widely used in the literature, much of the work is based on results generated using VBM simulations, which the current work has shown to be biased towards the microstructural indices. Hence, the reported correlation between microstructural indices and apparent modulus must be questioned.

If it is assumed that there is a link between the microstructural indices and apparent modulus, the second question stemming from this work is: are the BSM and VBM accurately predicting the apparent modulus? It stands to reason that if the BSM and VBM were accurately predicting the apparent modulus, they would show a similar correlation to the indices as the experimental results, however this is clearly not the case.

The two questions stemming from the microstructural indices analysis are in contrast to one another. If the correlation between the apparent modulus and the microstructural indices has been overstated (question 1), then there is no reason that the experimental and simulation results should show similar correlations to the microstructural indices (question 2). Consequently, the difference between experimental-index correlation and simulation-index correlation does not call into question the accuracy of the numerical methodology. In contrast, if there is a correlation

between the apparent modulus and the microstructural indices (question 1), then there should be little difference in the correlation of the experimental and simulation results to the indices, meaning the difference in behaviour of the experimental and simulation results must call into question the accuracy of the simulations (question 2). It seems unlikely that the microstructural indices, which are scalar quantities, can capture the complexity of the bone structure as well as a detailed FE simulation. Consequently, the results presented in this work imply that the link between the microstructural indices and the apparent modulus of trabecular bone has been overstated.

One non-conventional index was calculated for the specimens, namely the shell-to-beam ratio. It was found that all of the scatter in the BSM results stems from specimens with exceptionally high shell-to-beam ratios, i.e. specimens which had a much higher volume of shells in the mesh than beams. These results indicate that the BSM is producing too many shells in the mesh, or that there is a problem with the implementation of the shells. This phenomenon should be investigated in future work but, in the meantime, provides a useful method to identify poor simulation results generated by the BSM. This identification process may even negate the necessity for verification using additional data points, which is one of challenges associated with trabecular bone simulations as discussed in Section 9.3.2.1.

9.5 Factors affecting performance

The analysis of various factors affecting the performance of trabecular bone simulations highlighted important information, which could be used to understand and explain the results obtained for the apparent modulus and effective modulus. These points are reiterated in the following sections, along with other interesting conclusions drawn from the work.

9.5.1 Voxel size

The size of the voxels used for the micro-CT scan was shown to be a significant factor in both the microstructural indices and apparent moduli results. Notably, the bone volume fraction was shown to increase with increasing voxel size. Trabecular thickness and separation also showed variation with voxel size, although neither consistently increased or decreased so a pattern could not be established. Both the microstructural indices analysis and a visual comparison seem to imply that the difference between the 20 μm and 40 μm scans is not as significant as the difference between the 40 μm and 80 μm scans, which leads the author to believe that the 80 μm scan may be an insufficient voxel size for trabecular bone simulations.

The presence of machining artifacts in the bone specimens, which are clearly seen in the 20 μm images, serve to overthicken the circumference of the bone specimen in the lower resolution images. This phenomenon is the worst in the 80 μm images, though the 40 μm images are also affected, albeit to a lesser degree. The presence of these machining artifacts provides motivation for cropping a region of interest out of the micro-CT scans, however this methodology makes it impossible to compare the simulation results to experimental results, which is an important step for providing confidence in simulation results. Alternatively, the machining artifacts could be removed prior to scanning by defatting the specimens. This would maintain the structural integrity of the scanned specimen and allow for the simulation results to be compared to the experimental results.

Overall, the image voxel size can produce significant variation in simulation results. Although one might say that scans should be performed at the smallest possible voxel size, this is not practical because as voxel size decreases:

- The time taken for each scan increases.
- The cost of each scan increases.
- The computational resources required to analyse the data increases.
- The size of the specimen which can be scanned decreases.

A thorough investigation needs to be performed which links voxel size to simulation accuracy, so that researchers are aware of the effect of the voxel size on trabecular bone simulations and are able to make informed decisions about voxel size based on their time, budget and computational constraints.

9.5.2 Segmentation technique

This work has shown that the quality of the image segmentation has a significant effect on the results produced by the simulation. Defatted specimens seem to provide a better basis for segmentation than marrow-inclusive specimens, as there is better differentiation between bone and non-bone voxels. This leads to more accurate segmentation and, consequently, more accurate simulated apparent moduli.

The use of a different segmentation technique changed the BSM results by 0 – 20% for the majority of specimens and increased the scatter. In contrast, the modulus ratio results remained mostly unaffected, leading to the conclusion that the alternative segmentation technique did not change the relationship between the apparent modulus and effective modulus. Additionally, the alternative segmentation did not cause a significant change in the behaviour of the bone when simulated under various boundary conditions, leading to the conclusion that the overall structure of the bone was not severely impacted by the change in segmentation technique.

Interestingly, the use of a single threshold per TIFF stack caused less failures in the BSM than the use of an individual threshold applied to each image in the stack. This phenomenon should be taken into account in any future work involving trabecular bone simulation using the BSM.

Considering these results, it can be concluded that the overall structural behaviour of the bone did not change significantly under the alternative segmentation used in this work. However, the quality of the segmentation, which is influenced particularly by the inclusion of marrow in the specimens, may significantly alter the results.

9.5.3 Boundary conditions

Ten different boundary conditions were investigated for the bone simulations, because:

1. It is difficult to determine what type/s of constraint were experienced by the bone specimens in the experimental test set-up.

2. The experimental boundary conditions are likely to be complex and, therefore, difficult to replicate in a simulation.

The goal of the boundary condition analysis was not to determine which constraint, or combination of constraints, was the closest to the experimental boundary condition, but rather to demonstrate the effect that boundary conditions can have on results obtained from the BSM.

The results from the boundary condition analysis show that, on average, the addition of a rotational constraint produces a negligible change in the apparent modulus results and a slightly larger change in the inner modulus and modulus ratio results. Overall, the lateral constraints had a much larger impact on the apparent modulus, inner modulus and modulus ratio.

The apparent modulus results did not show much difference based on the placement of the constraint, i.e. whether the constraint was applied on the top or bottom surface of the specimen. However, the inner modulus showed a much larger variation when a constraint was applied on the top surface versus when it was applied on the bottom surface. As a result of the sensitivity of the inner modulus to top surface constraints, the modulus ratio showed a slightly larger variation when the constraint was applied on the top surface. However, the difference between the top and bottom surface variation is so small that it may be said that the modulus ratio, like the apparent modulus, is insensitive to constraint placement.

The application of lateral and rotational constraints on the top and bottom surface of the specimen caused the greatest variation in the apparent modulus and inner modulus results, respectively. In contrast, a lateral constraint on the top surface of the specimen caused the greatest change in modulus ratio results.

Although the largest average variation in the modulus ratio results was only 11%, this small change was enough to increase some of the modulus ratios to values above 1. This shift brings the BSM modulus ratio results on par with the VBM modulus ratio results, in terms of the number of specimens for which the relationship between the apparent modulus and inner modulus was correctly predicted.

Overall, the boundary conditions were shown to be a significant factor in the inner modulus and modulus ratio of the bone specimens, but a less important factor in the apparent modulus. These results show that any study taking the effective modulus to be equal to the inner modulus should carefully consider the boundary conditions applied to the specimen during simulation.

9.6 Closing remarks

In this work, a wide variety of simulations were performed to analyse a range of metrics, most importantly, the performance of the beam-shell method in capturing the effective modulus of trabecular bone. Because of the multitude of tests, some closing remarks are included in point form in this section to summarise the most important findings.

- A validation study using structures with known mechanical properties is essential to numerical trabecular bone studies. The validation study provides confidence in simulation results and can explain some of the behaviour of trabecular bone (if cellular solids are used for validation).
- The BSM is capable of accurately predicting the apparent modulus of trabecular bone,

however the scatter in the data is still significant. Techniques have been identified to reduce the scatter, which will increase the accuracy of the BSM.

- The BSM was equally as capable of predicting the modulus ratio as the VBM. Although the simulation results showed variation from the experimental results, these results may be improved through the removal of machining artifacts and application of improved segmentation techniques.
- The BSM mitigates the effects of overthickening in an image, which is not done by the VBM. Consequently, the BSM has the potential to simulate poorly segmented (overly thick) images more accurately than the VBM.
- Until the Young's modulus of trabecular bone specimens can be determined reliably and accurately, any trabecular bone simulation should be accompanied by a dataset which includes simulation and corresponding experimental data points, to lend confidence to the result.

Overall, both the BSM and VBM have inherent advantages and disadvantages. This work has specifically highlighted the potential of the BSM in numerical studies involving trabecular bone and has shown that the BSM may be considered a viable alternative to the VBM. Consequently, the choice of modelling methodology is non-trivial and researchers must take care to choose the correct methodology for the given problem.

Chapter 10

Recommendations

This work has taken important steps in the validation of a beam-shell Finite Element method, and has provided valuable insight into the capabilities of the method regarding the prediction of mechanical properties of trabecular bone. Additionally, this work has brought to light some areas of research which should be further explored to gain a better understanding of the nuances of simulating trabecular bone.

A large portion of this work focused on validating the chosen Finite Element method using artificially generated cellular solids with known structures. Although this was a useful endeavour, the generated lattices only consisted of rod-like trabeculae, meaning that the generation of shell elements in the mesh was not explored during the validation process. In order to further validate the methodology, it would be useful to artificially generate closed cell foams, so as to assess the ability of the beam-shell code in capturing the response of these plate-rod structures in addition to the rod-rod structures assessed in this work.

The validation process would not have been possible without the image rasterisation code, which constructed TIFF stacks of the lattices from node and element data. Although the rasterisation code provided valuable insight into the discretisation and simulation process, it is likely that this code could be improved and optimised to be more efficient. Additionally, the quality of the rasterisation could be better quantified by rasterising simple structures and performing a series of small translations and rotations, assessing the accuracy of the representation after each manipulation.

The voxel size was shown to have a significant effect on the results produced by the BSM, however in-depth research into the effect of voxel size on the results produced by both the BSM and VBM is needed. This research should specify the minimum voxel size required for accurate simulation of trabecular bone, but should also take into account the time and expense associated with micro-CT scanning samples.

A system which lays out the steps to follow during the image processing of trabecular bone micro-CT scans is sorely needed to bring some uniformity and repeatability to the literature. This system should address pre-processing steps such as filtering or sharpening, and should specify which segmentation method is best for the analysis of trabecular bone. Moreover, this research should investigate if there is a single best segmentation method for trabecular bone, or if the best segmentation method changes with image resolution and/or the inclusion of marrow in the specimens.

Probably the biggest barrier to computational work on trabecular bone currently, is the lack of knowledge of the Young's modulus. The numerical techniques used to assign Young's modulus to simulations have clear issues, and the uncertainty associated with the Young's modulus brings uncertainty to the simulation results, which cannot be resolved without more information. A technique is required to determine the Young's modulus, which does not rely too heavily on images of the bone (which are influenced by voxel size and segmentation technique), but also takes into account the possible variation of the modulus throughout the structure. It is unclear what such a method would entail, however this author suspects a combination of experimental and numerical techniques would be best.

This work has shown that the BSM is a viable alternative to the VBM. However, a key aspect of the BSM that should be investigated is why some of the bone specimens fail during the discretisation of the specimen. If more specimens could be simulated using the BSM, it would significantly reduce the performance gap between the BSM and the VBM. Additionally, a significant portion of the large specimens could not be simulated using the BSM because of the large number of section definitions in the mesh. If these section definitions could be minimised while maintaining the accuracy of the simulation, the performance gap between the BSM and VBM could be further reduced.

Arguably the most important finding of the current work is the fact that the BSM is equally as capable of predicting the modulus ratio of trabecular bone specimens as the VBM. Only 17 large specimens were available for analysis of modulus ratios, which is a limitation of this work. However, this work has showed significant potential and it is recommended that further research is performed in this area. Future work should include the use of experimental and numerical results, similar to this work.

If the supplementary research is performed and discussed improvements are made to the BSM, this method will become a powerful tool in the analysis of trabecular bone.

References

- [1] J.-Y. Rho, L. Kuhn-Spearing, and P. Zioupos, “Mechanical properties and the herarchical structure of bone,” *Medical Engineering and Physics*, vol. 20, no. 2, pp. 92–102, 1998.
- [2] J. Vanderroost, S. V. Jaecques, G. Van der Perre, S. Boonen, J. D’hooge, W. Lauriks, and G. H. van Lenthe, “Fast and accurate specimen-specific simulation of trabecular bone elastic modulus using novel beam-shell finite element models,” *Journal of Biomechanics*, vol. 44, no. 8, pp. 1566–1572, 2011. [Online]. Available: <http://dx.doi.org/10.1016/j.jbiomech.2011.02.082>
- [3] K. Hilton, “The effect of boundary conditions and architecture on the response of cancellous bone,” MSc Dissertation, University of Cape Town, 2018.
- [4] P. R. Townsend, R. M. Rose, R. E. Miegel, and E. L. Radin, “The Distribution and Anisotropy of the Stiffness of Cancellous Bone in the Human Patella,” Massachusetts Institute of Technology, Tech. Rep., 1974. [Online]. Available: <https://apps.dtic.mil/docs/citations/ADA000271>
- [5] J. W. Pugh, R. M. Rose, and E. L. Radin, “On the Elastic and Viscoelastic Properties of Trabecular Bone: Dependence on Structure,” Massachusetts Institute of Technology, Tech. Rep., 1972. [Online]. Available: <https://apps.dtic.mil/docs/citations/AD0752980>
- [6] Ł. Cyganik, M. Binkowski, G. Kokot, T. Rusin, P. Popik, F. Bolechała, R. Nowak, Z. Wróbel, and A. John, “Prediction of Young’s modulus of trabeculae in microscale using macro-scale’s relationships between bone density and mechanical properties,” *Journal of the Mechanical Behavior of Biomedical Materials*, vol. 36, pp. 120–134, 2014.
- [7] F. A. Sabet, A. Raeisi Najafi, E. Hamed, and I. Jasiuk, “Modelling of bone fracture and strength at different length scales: a review,” *Interface Focus*, vol. 6, no. 1, p. 20150055, 2016. [Online]. Available: <http://rsfs.royalsocietypublishing.org/lookup/doi/10.1098/rsfs.2015.0055>
- [8] M. Ding, A. Odgaard, and I. Hvid, “Accuracy of cancellous bone volume fraction measured by micro-CT scanning,” *Journal of Biomechanics*, vol. 32, no. 3, pp. 323–326, 1999.
- [9] B. van Rietbergen, H. Weinans, R. Huiskes, and A. Odgaard, “A new method to determine trabecular bone elastic properties and loading using micromechanical finite-elements models,” *J. Biomech.*, vol. 28, no. 1, pp. 69–81, 1995.
- [10] J. H. Keyak, J. M. Meagher, H. B. Skinner, and C. D. Mote, “Automated three-dimensional finite element modeling of bone: a new method,” *Journal of Biomedical Engineering*, vol. 12, pp. 389–397, 1990.
- [11] D. H. Pahr and P. K. Zysset, “A comparison of enhanced continuum FE with micro FE models of human vertebral bodies,” *Journal of Biomechanics*, vol. 42, no. 4, pp. 455–462, 2009.

- [12] X. S. Liu, E. Shane, D. J. McMahon, and X. E. Guo, "Individual trabecula segmentation (ITS)-based morphological analysis of microscale images of human tibial trabecular bone at limited spatial resolution," *Journal of Bone and Mineral Research*, vol. 26, no. 9, pp. 2184–2193, 2011.
- [13] M. L. Bouxsein, S. K. Boyd, B. A. Christiansen, R. E. Guldberg, K. J. Jepsen, and R. Müller, "Guidelines for assessment of bone microstructure in rodents using micro-computed tomography," *Journal of Bone and Mineral Research*, vol. 25, no. 7, pp. 1468–1486, 2010.
- [14] J. H. Waarsing, J. S. Day, and H. Weinans, "An improved segmentation method for in vivo μ CT imaging," *Journal of Bone and Mineral Research*, vol. 19, no. 10, pp. 1640–1650, 2004.
- [15] K. Daszkiewicz, G. Maquer, and P. K. Zysset, "The effective elastic properties of human trabecular bone may be approximated using micro-finite element analyses of embedded volume elements," *Biomechanics and Modeling in Mechanobiology*, vol. 16, no. 3, pp. 731–742, 2017.
- [16] S. Nawathe, B. P. Nguyen, N. Barzanian, H. Akhlaghpour, M. L. Bouxsein, and T. M. Keaveny, "Cortical and trabecular load sharing in the human femoral neck," *Journal of Biomechanics*, vol. 48, no. 5, pp. 816–822, 2015. [Online]. Available: <http://dx.doi.org/10.1016/j.jbiomech.2014.12.022>
- [17] 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*, vol. 21, no. 2, pp. 307–314, 2006. [Online]. Available: <http://doi.wiley.com/10.1359/JBMR.051027>
- [18] L. J. Gibson and M. F. Ashby, *Cellular solids: structure and properties-Second edition*, 2nd ed. Cambridge: Cambridge University Press, 1997.
- [19] J. J. Kim, J. Nam, and I. G. Jang, "Computational study of estimating 3D trabecular bone microstructure for the volume of interest from CT scan data," *International Journal for Numerical Methods in Biomedical Engineering*, vol. 34, no. 4, pp. 1–13, 2018.
- [20] X. S. Liu, P. Sajda, P. K. Saha, F. W. Wehrli, G. Bevil, T. M. Keaveny, and X. E. Guo, "Complete volumetric decomposition of individual trabecular plates and rods and its morphological correlations with anisotropic elastic moduli in human trabecular bone," *Journal of Bone and Mineral Research*, vol. 23, no. 2, pp. 223–235, 2008. [Online]. Available: <http://doi.wiley.com/10.1359/jbmr.071009>
- [21] T. Hildebrand and P. Rüegsegger, "Quantification of bone microarchitecture with the structure model index," *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 1, no. 1, pp. 15–23, 1997. [Online]. Available: <http://www.tandfonline.com/doi/abs/10.1080/01495739708936692>
- [22] "Prophylaxis and treatment of osteoporosis," in *Consensus Development Conference*, vol. 1, 1991, pp. 114–117. [Online]. Available: <https://lirias.kuleuven.be/handle/123456789/173726>
- [23] J. Y. Rho, R. B. Ashman, and C. H. Turner, "Young's modulus of trabecular and cortical bone material: Ultrasonic and microtensile measurements," *Journal of Biomechanics*, vol. 26, no. 2, pp. 111–119, 1993.

- [24] H. H. Bayraktar, E. F. Morgan, G. L. Niebur, G. E. Morris, E. K. Wong, and T. M. Keaveny, "Comparison of the elastic and yield properties of human femoral trabecular and cortical bone tissue," *Journal of Biomechanics*, vol. 37, no. 1, pp. 27–35, 2004.
- [25] B. Helgason, E. Perilli, E. Schileo, F. Taddei, S. Brynjólfsson, and M. Viceconti, "Mathematical relationships between bone density and mechanical properties: A literature review," *Clinical Biomechanics*, vol. 23, no. 2, pp. 135–146, 2008.
- [26] S. Eberle, M. Göttlinger, and P. Augat, "Individual density-elasticity relationships improve accuracy of subject-specific finite element models of human femurs," *Journal of Biomechanics*, vol. 46, no. 13, pp. 2152–2157, 2013. [Online]. Available: <http://dx.doi.org/10.1016/j.jbiomech.2013.06.035>
- [27] J. Hong, Y. Park, S. Lee, G. Khang, and Y. Kim, "Elastic moduli and Poisson's ratios of microscopic human femoral trabeculae," in *Proceedings of the 11th Mediterranean Conference on Medical and Biomedical Engineering and Computing*, 2007, pp. 274–277.
- [28] W. Pistoia, B. Van Rietbergen, E. M. Lochmüller, C. A. Lill, F. Eckstein, and P. Rügsegger, "Estimation of distal radius failure load with micro-finite element analysis models based on three-dimensional peripheral quantitative computed tomography images," *Bone*, vol. 30, no. 6, pp. 842–848, 2002.
- [29] J. J. Kim, J. Nam, and I. G. Jang, "Computational study of estimating 3D trabecular bone microstructure for the volume of interest from CT scan data," *Int J Numer Meth Biomed Engng.*, vol. 34, pp. 1–13, 2018. [Online]. Available: <https://doi.org/10.1002/cnm.2950>
- [30] Z. Yosibash, N. Trabelsi, and C. Milgrom, "Reliable simulations of the human proximal femur by high-order finite element analysis validated by experimental observations," *Journal of Biomechanics*, vol. 40, no. 16, pp. 3688–3699, 2007.
- [31] Z. Yosibash, R. Padan, L. Joskowicz, and C. Milgrom, "A CT-based high-order finite element analysis of the human proximal femur compared to in-vitro experiments," *Journal of Biomechanical Engineering*, vol. 129, no. 3, pp. 297–309, 2007.
- [32] A. C. Vale, M. F. Pereira, A. Maurício, P. Amaral, L. G. Rosa, A. Lopes, A. Rodrigues, J. Caetano-Lopes, B. Vidal, J. Monteiro, J. E. Fonseca, H. Canhão, and M. F. Vaz, "Micro-computed tomography and compressive characterization of trabecular bone," *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, vol. 438, pp. 199–205, 2013. [Online]. Available: <http://dx.doi.org/10.1016/j.colsurfa.2013.01.057>
- [33] J. D. Boerckel, D. E. Mason, A. M. McDermott, and E. Alsberg, "Microcomputed tomography: Approaches and applications in bioengineering," *Stem Cell Research and Therapy*, vol. 5, no. 6, pp. 1–12, 2014.
- [34] Y. Wu, S. Adeeb, and M. R. Doschak, "Using micro-CT derived bone microarchitecture to analyze bone stiffness - a case study on osteoporosis rat bone," *Frontiers in Endocrinology*, vol. 6, no. MAY, pp. 1–7, 2015.
- [35] G. L. Niebur, M. J. Feldstein, J. C. Yuen, T. J. Chen, and T. M. Keaveny, "High-resolution finite element models with tissue strength asymmetry accurately predict failure of trabecular bone," *Journal of Biomechanics*, vol. 33, no. 12, pp. 1575–1583, 2000.
- [36] G. Bevil, S. K. Eswaran, F. Farahmand, and T. M. Keaveny, "The influence of boundary conditions and loading mode on high-resolution finite element-computed trabecular tissue properties," *Bone*, vol. 44, no. 4, pp. 573–578, 2009. [Online]. Available: <http://dx.doi.org/10.1016/j.bone.2008.11.015>

- [37] D. Ulrich, B. Van Rietbergen, H. Weinans, and P. Rügsegger, "Finite element analysis of trabecular bone structure: A comparison of image-based meshing techniques," *Journal of Biomechanics*, vol. 31, no. 12, pp. 1187–1192, 1998.
- [38] N. M. Harrison and P. E. McHugh, "Comparison of trabecular bone behavior in core and whole bone samples using high-resolution modeling of a vertebral body," *Biomechanics and Modeling in Mechanobiology*, vol. 9, no. 4, pp. 469–480, 2010.
- [39] A. J. Ladd and J. H. Kinney, "Numerical errors and uncertainties in finite-element modeling of trabecular bone," *Journal of Biomechanics*, vol. 31, no. 10, pp. 941–945, 1998.
- [40] Y. Chen, M. Pani, F. Taddei, C. Mazzà, X. Li, and M. Viceconti, "Large-Scale Finite Element Analysis of Human Cancellous Bone Tissue Micro Computer Tomography Data: A Convergence Study," *Journal of Biomechanical Engineering*, vol. 136, no. 10, p. 101013, 2014. [Online]. Available: <http://biomechanical.asmedigitalcollection.asme.org/article.aspx?doi=10.1115/1.4028106>{%}5Cn<http://www.ncbi.nlm.nih.gov/pubmed/25070476>
- [41] I. H. Parkinson, A. Badieli, and N. L. Fazzalari, "Variation in segmentation of bone from micro-CT imaging: Implications for quantitative morphometric analysis," *Australasian Physical and Engineering Sciences in Medicine*, vol. 31, no. 2, pp. 160–164, 2008.
- [42] N. Otsu, "A Threshold Selection Method from Gray-Level Histograms," *IEEE Transactions on Systems*, vol. C, no. 1, pp. 62–66, 1979.
- [43] K. Rovaris, P. M. Queiroz, K. D. F. Vasconcelos, L. d. S. Corpas, B. M. da Silveira, and D. Q. Freitas, "Segmentation methods for micro CT images: A comparative study using human bone samples," *Brazilian Dental Journal*, vol. 29, no. 2, pp. 150–153, 2018.
- [44] S. Tassani and G. K. Matsopoulos, "The micro-structure of bone trabecular fracture: An inter-site study," *Bone*, vol. 60, pp. 78–86, 2014.
- [45] T. Hara, E. Tanck, J. Homminga, and R. Huiskes, "The influence of microcomputed tomography threshold variations on the assessment of structural and mechanical trabecular bone properties," *Bone*, vol. 31, no. 1, pp. 107–109, 2002. [Online]. Available: [http://dx.doi.org/10.1016/S8756-3282\(02\)00782-2](http://dx.doi.org/10.1016/S8756-3282(02)00782-2)
- [46] T. Hildebrand, A. Laib, R. Muller, J. Dequeker, and P. Ruegsegger, "Direct Three-Dimensional Morphometric Analysis of Human Cancellous Bone: Microstructural Data from Spine, Femur, Iliac Crest, and Calcaneus," *Journal of Bone and Mineral Research*, vol. 14, no. 7, pp. 1167–1174, 1999.
- [47] T. Uchiyama, T. Tanizawa, H. Muramatsu, N. Endo, H. E. Takahashi, and T. Hara, "Three-dimensional microstructural analysis of human trabecular bone in relation to its mechanical properties," *Bone*, vol. 25, no. 4, pp. 487–491, 1999.
- [48] M. R. A. Kadir, A. Syahrom, and A. Öchsner, "Finite element analysis of idealised unit cell cancellous structure based on morphological indices of cancellous bone," *Medical and Biological Engineering and Computing*, vol. 48, no. 5, pp. 497–505, 2010.
- [49] D. Ulrich, B. Van Rietbergen, A. Laib, and P. Ruegsegger, "The ability of three-dimensional structural indices to reflect mechanical aspects of trabecular bone," *Bone*, vol. 25, no. 1, pp. 55–60, 1999.
- [50] R. Müller, H. Van Campenhout, B. Van Damme, G. Van Der Perre, J. Dequeker, T. Hildebrand, and P. Rügsegger, "Morphometric analysis of human bone

- biopsies: A quantitative structural comparison of histological sections and micro-computed tomography,” *Bone*, vol. 23, no. 1, pp. 59–66, 1998. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S8756328298000684>
- [51] M. Stauber, A. Nazarian, and R. Müller, “Limitations of global morphometry in predicting trabecular bone failure,” *Journal of Bone and Mineral Research*, vol. 29, no. 1, pp. 134–141, 2014.
- [52] A. Nazarian, M. Stauber, D. Zurakowski, B. D. Snyder, and R. Müller, “The interaction of microstructure and volume fraction in predicting failure in cancellous bone,” *Bone*, vol. 39, no. 6, pp. 1196–1202, 2006.
- [53] M. E. A. Doube, “UKPMC Funders Group,” *UKPMC Funders Group*, vol. 47, no. 6, pp. 1076–1079, 2010.
- [54] J. Schindelin, I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzich, S. Preibisch, C. Rueden, S. Saalfeld, B. Schmid, J.-Y. Tinevez, D. J. White, V. Hartenstein, K. Eliceiri, P. Tomanack, and A. Cardona, “Fiji: an open-source platform for biological-image analysis,” *Nature Methods*, vol. 9, no. 7, pp. 676–682, 2012.
- [55] C. A. Schneider, W. S. Rasband, and K. W. Eliceiri, “NIH Image to ImageJ: 25 years of image analysis,” *Nature Methods*, vol. 9, no. 7, pp. 671–675, 2012. [Online]. Available: <http://www.nature.com/doi/10.1038/nmeth.2089>
- [56] R. Dougherty and K.-H. Kunzelmann, “Computing Local Thickness of 3D Structures with ImageJ,” *Microscopy and Microanalysis*, vol. 13, no. S02, pp. 1678–1679, 2007. [Online]. Available: <http://www.journals.cambridge.org/abstract/S1431927607074430>
- [57] T. Hildebrand and P. Rueggsegger, “A new method for the model-independent assessment of thickness in three-dimensional images,” *Journal of Microscopy*, vol. 185, no. 1, pp. 67–75, 1997. [Online]. Available: <http://doi.wiley.com/10.1046/j.1365-2818.1997.1340694.x>
- [58] I. Arganda-Carreras, R. Fernández-González, A. Muñoz-Barrutia, and C. Ortiz-De-Solorzano, “3D reconstruction of histological sections: Application to mammary gland tissue,” *Microscopy Research and Technique*, vol. 73, no. 11, pp. 1019–1029, 2010.
- [59] A. Odgaard, “Three-dimensional methods for quantification of cancellous bone architecture,” *Bone*, vol. 20, no. 4, pp. 315–328, 1997. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S8756328297000070>
- [60] T. P. Harrigan and R. W. Mann, “Characterization of microstructural anisotropy in orthotropic materials using a second rank tensor,” *Journal of Materials Science*, vol. 19, no. 3, pp. 761–767, 1984. [Online]. Available: <https://link.springer.com/article/10.1007/BF00540446>
- [61] A. Odgaard and H. J. G. Gundersen, “Quantification of connectivity in cancellous bone, with special emphasis on 3D reconstructions,” *Bone*, vol. 14, no. 2, pp. 173–182, 1993.
- [62] J. Toriwaki and T. Yonekura, “Euler number and connectivity indexes of a three dimensional digital picture,” *Forma-Tokyo-*, vol. 17, no. 3, pp. 183–209, 2002.
- [63] M. Ding, A. Odgaard, C. C. Danielsen, and I. Hvid, “Mutual associations among microstructural, physical and mechanical properties of human cancellous bone,” *The Journal of Bone and Joint Surgery*, vol. 84, no. 6, pp. 900–907, 2002. [Online]. Available: <http://www.bjj.boneandjoint.org.uk/cgi/doi/10.1302/0301-620X.84B6.11994>

- [64] K. Moore, “A Numerical Assessment of Architectural Parameters for Anisotropic Behaviour in Idealised Trabecular Structures,” Ph.D. dissertation, University of Cape Town, 2018.
- [65] P. L. Salmon, C. Ohlsson, S. J. Shefelbine, and M. Doube, “Structure model index does not measure rods and plates in trabecular bone,” *Frontiers in Endocrinology*, vol. 6, no. OCT, pp. 1–10, 2015.
- [66] E. Perilli, M. Baleani, C. Öhman, R. Fognani, F. Baruffaldi, and M. Viceconti, “Dependence of mechanical compressive strength on local variations in microarchitecture in cancellous bone of proximal human femur,” *Journal of Biomechanics*, vol. 41, no. 2, pp. 438–446, 2008.
- [67] W. Q. Cui, Y. Y. Won, M. H. Baek, and K. K. Kim, “Contribution of Three-Dimensional Trabecular Bone Microstructure of the Proximal Femur to its Mechanical Properties as Assessed by Micro-Finite Element Analysis,” *Key Engineering Materials*, vol. 321-323, no. March, pp. 278–281, 2006. [Online]. Available: <http://www.scientific.net/KEM.321-323.278>
- [68] C. V. Lawrence, “Effective stiffness models of regular honeycombs and implications for trabecular bone models,” Final Year Project, University of Cape Town, 2016.
- [69] L. J. Gibson, “Biomechanics of cellular solids,” *Journal of Biomechanics*, vol. 38, no. 3, pp. 377–399, 2005.
- [70] L. Kelvin, “On Homogeneous Division of Space,” *Proceedings of the Royal Society of London*, vol. 55, no. 331-335, pp. 1–16, 1894.
- [71] R. B. Fuller, *Inventions: The Patented Works of R. Buckminster Fuller*. New York: St. Martin’s Press, 1983.
- [72] L. Dong, V. Deshpande, and H. Wadley, “Mechanical response of Ti-6Al-4V octet-truss lattice structures,” *International Journal of Solids and Structures*, vol. 60, pp. 107–124, 2015. [Online]. Available: <http://dx.doi.org/10.1016/j.ijsolstr.2015.02.020>
- [73] I. Goda, M. Assidi, and J. F. Ganghoffer, “A 3D elastic micropolar model of vertebral trabecular bone from lattice homogenization of the bone microstructure,” *Biomechanics and Modeling in Mechanobiology*, vol. 13, no. 1, pp. 53–83, 2014.
- [74] I. Diamant, R. Shahar, Y. Masharawi, and A. Gefen, “A method for patient-specific evaluation of vertebral cancellous bone strength: In vitro validation,” *Clinical Biomechanics*, vol. 22, no. 3, pp. 282–291, 2007.
- [75] O. C. Yeh and T. M. Keaveny, “Biomechanical effects of intraspecimen variations in trabecular architecture: A three-dimensional finite element study,” *Bone*, vol. 25, no. 2, pp. 223–228, 1999.
- [76] M. Lengsfeld, J. Schmitt, P. Alter, J. Kaminsky, and R. Leppek, “Comparison of geometry-based and CT voxel-based finite element modelling and experimental validation,” *Medical Engineering & Physics*, vol. 20, pp. 515–522, 1998.
- [77] J. Panyasantisuk, D. H. Pahr, T. Gross, and P. K. Zysset, “Comparison of Mixed and Kinematic Uniform Boundary Conditions in Homogenized Elasticity of Femoral Trabecular Bone Using Microfinite Element Analyses,” *Journal of Biomechanical Engineering*, vol. 137, no. 1, p. 011002, 2015. [Online]. Available: <http://biomechanical.asmedigitalcollection.asme.org/article.aspx?doi=10.1115/1.4028968>

- [78] R. Müller and P. Rügsegger, “Three-dimensional finite element modelling of non-invasively assessed trabecular bone structures,” *Medical Engineering and Physics*, vol. 17, no. 2, pp. 126–133, 1995.
- [79] G. Bevill, S. K. Easley, and T. M. Keaveny, “Side-artifact errors in yield strength and elastic modulus for human trabecular bone and their dependence on bone volume fraction and anatomic site,” *Journal of Biomechanics*, vol. 40, no. 15, pp. 3381–3388, 2007.
- [80] R. Hambli, “3D finite element simulation of human proximal femoral fracture under quasi-static load,” *Advances in biomechanics and applications*, vol. 1, no. 1, pp. 1–14, 2014. [Online]. Available: <http://koreascience.or.kr/journal/view.jsp?kj=TPTPOC&py=2014&vnc=vIn1&sp=1>
- [81] X.-X. Wen, H.-L. Yu, Y.-B. Yan, C.-L. Zong, H.-J. Ding, X.-Y. Ma, T.-S. Wang, and W. Lei, “Influence of the shape of the micro-finite element model on the mechanical properties calculated from micro-finite element analysis,” *Experimental and Therapeutic Medicine*, vol. 14, pp. 1744–1748, 2017.
- [82] Y. Chen, E. DallAra, E. Sales, K. Manda, R. Wallace, P. Pankaj, and M. Viceconti, “Micro-CT based finite element models of cancellous bone predict accurately displacement once the boundary condition is well replicated: A validation study,” *Journal of the Mechanical Behavior of Biomedical Materials*, vol. 65, no. July 2016, pp. 644–651, 2017. [Online]. Available: <http://dx.doi.org/10.1016/j.jmbbm.2016.09.014>
- [83] K. Ün, G. Bevill, and T. M. Keaveny, “The effects of side-artifacts on the elastic modulus of trabecular bone,” *Journal of Biomechanics*, vol. 39, no. 11, pp. 1955–1963, 2006.
- [84] D. H. Pahr and P. K. Zysset, “Influence of boundary conditions on computed apparent elastic properties of cancellous bone,” *Biomechanics and Modeling in Mechanobiology*, vol. 7, no. 6, pp. 463–476, 2008.
- [85] D. H. Pahr, E. D. Ara, P. Varga, and P. K. Zysset, “Stiffness and Strength of Human Vertebrae Predicted By Novel Homogenized and,” *Journal of Biomechanics*, pp. 3–4, 2011.
- [86] F. J. Hou, S. M. Lang, S. J. Hoshaw, D. A. Reimann, and D. P. Fyhrie, “Human vertebral body apparent and hard tissue stiffness,” *Journal of Biomechanics*, vol. 31, no. 11, pp. 1009–1015, 1998.
- [87] J. Wang, B. Zhou, X. S. Liu, A. J. Fields, A. Sanyal, X. Shi, M. Adams, T. M. Keaveny, and X. E. Guo, “Trabecular plates and rods determine elastic modulus and yield strength of human trabecular bone,” *Bone*, vol. 72, no. March 2016, pp. 71–80, 2015.
- [88] H. Wang, X. Sherry Liu, B. Zhou, J. Wang, B. Ji, Y. Huang, K.-C. Hwang, and X. Edward Guo, “Accuracy of Individual Trabecula Segmentation Based Plate and Rod Finite Element Models in Idealized Trabecular Bone Microstructure,” *Journal of Biomechanical Engineering*, vol. 135, no. 4, p. 044502, 2013. [Online]. Available: <http://biomechanical.asmedigitalcollection.asme.org/article.aspx?doi=10.1115/1.4023983>
- [89] P. Saha, B. Chaudhuri, and D. Dutta Majumder, “A new shape preserving parallel thinning algorithm for 3D digital images,” *Pattern Recognition*, vol. 30, no. 12, pp. 1939–1955, 1997.
- [90] P. K. Saha, B. R. Gomberg, and F. W. Wehrli, “Three-dimensional digital topological characterization of cancellous bone architecture,” *International Journal of Imaging Systems and Technology*, vol. 11, no. 1, pp. 81–90, 2000.

- [91] L. Pothuaud, P. Porion, E. Lespessailles, C. L. Benhamou, and P. Levitz, "A new method for three-dimensional skeleton graph analysis of porous media: Application to trabecular bone microarchitecture," *Journal of Microscopy*, vol. 199, no. 2, pp. 149–161, 2000.
- [92] G. L. Niebur, M. J. Feldstein, J. C. Yuen, T. J. Chen, and T. M. Keaveny, "High-resolution finite element models with tissue strength asymmetry accurately predict failure of trabecular bone," *Journal of Biomechanics*, vol. 33, no. 12, pp. 1575–1583, 2000.
- [93] Y. Chevalier, D. Pahr, H. Allmer, M. Charlebois, and P. Zysset, "Validation of a voxel-based FE method for prediction of the uniaxial apparent modulus of human trabecular bone using macroscopic mechanical tests and nanoindentation," *Journal of Biomechanics*, vol. 40, no. 15, pp. 3333–3340, 2007.
- [94] C. Huet, "Application of variational concepts to size effects in elastic heterogeneous bodies," *Journal of the Mechanics and Physics of Solids*, vol. 38, no. 6, pp. 813–841, 1990.
- [95] W. M. Harris and W. K. Chiu, "Determining the representative volume element size for three-dimensional microstructural material characterization. Part 1: Predictive models," *Journal of Power Sources*, vol. 282, pp. 552–561, 2015. [Online]. Available: <http://dx.doi.org/10.1016/j.jpowsour.2015.02.035>
- [96] J. L. Grenestedt, "Effective elastic behavior of some models for perfect cellular solids," *International Journal of Solids and Structures*, vol. 36, no. 10, pp. 1471–1501, 1999. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0020768398000481>
- [97] F. Hualin and Y. Wei, "An equivalent continuum method of lattice structures," *Acta Mechanica Sinica*, vol. 19, no. 2, pp. 103–113, 2006. [Online]. Available: <http://dx.doi.org/10.1007/s10338-006-0612-x>
- [98] A. Vigliotti and D. Pasini, "Stiffness and strength of tridimensional periodic lattices," *Computer Methods in Applied Mechanics and Engineering*, vol. 229–232, pp. 27–43, 2012. [Online]. Available: <http://dx.doi.org/10.1016/j.cma.2012.03.018>
- [99] R. K. Burla, A. V. Kumar, and B. V. Sankar, "Implicit boundary method for determination of effective properties of composite microstructures," *International Journal of Solids and Structures*, vol. 46, no. 11–12, pp. 2514–2526, 2009. [Online]. Available: <http://dx.doi.org/10.1016/j.ijsolstr.2009.02.003>
- [100] R. Hedayati, M. Sadighi, M. Mohammadi-Aghdam, and A. A. Zadpoor, "Mechanics of additively manufactured porous biomaterials based on the rhombicuboctahedron unit cell," *Journal of the Mechanical Behavior of Biomedical Materials*, vol. 53, pp. 272–294, 2016. [Online]. Available: <http://dx.doi.org/10.1016/j.jmbbm.2015.07.013>
- [101] A. Yoo and I. Jasiuk, "Couple-stress moduli of a trabecular bone idealized as a 3D periodic cellular network," *Journal of Biomechanics*, vol. 39, no. 12, pp. 2241–2252, 2006.
- [102] M. Ostoja-Starzewski, "Material spatial randomness: From statistical to representative volume element," *Probabilistic Engineering Mechanics*, vol. 21, no. 2, pp. 112–132, 2006.
- [103] P. Kowalczyk, "Elastic properties of cancellous bone derived from finite element models of parameterized microstructure cells," *Journal of Biomechanics*, vol. 36, no. 7, pp. 961–972, 2003.
- [104] S. C. Cowin and M. M. Mehrabadi, "Identification of the elastic symmetry of bone and other materials," *Journal of Biomechanics*, vol. 22, pp. 503–515, 1989.

- [105] B. Van Rietbergen, A. Odgaard, J. Kabel, and R. Huiskes, “Direct mechanics assessment of elastic symmetries and properties of trabecular bone architecture,” *Journal of Biomechanics*, vol. 29, no. 12, pp. 1653–1657, 1996.
- [106] R. E. Newnham, *Properties of materials*, 1st ed. New York: Oxford University Press, 2005.
- [107] A. I. Hussein, D. T. Louzeiro, G. U. Unnikrishnan, and E. F. Morgan, “Differences in Trabecular Microarchitecture and Simplified Boundary Conditions Limit the Accuracy of Quantitative Computed Tomography-Based Finite Element Models of Vertebral Failure,” *Journal of Biomechanical Engineering*, vol. 140, no. 2, p. 021004, 2018. [Online]. Available: <http://biomechanical.asmedigitalcollection.asme.org/article.aspx?doi=10.1115/1.4038609>
- [108] S. P. Friedrichsdorf, V. E. Arana-Chavez, P. M. Cattaneo, R. Spin-Neto, and G. C. Dominguez, “Effect of the software binning and averaging data during microcomputed tomography image acquisition,” *Scientific Reports*, vol. 9, no. 1, pp. 1–8, 2019. [Online]. Available: <http://dx.doi.org/10.1038/s41598-019-46530-z>
- [109] E. F. Moore, “The shortest path through a maze,” in *Proceedings of the International Symposium on the Theory of Switching*. Harvard University Press, 1959, pp. 285–292.
- [110] L. Kahane, *Regression Basics*. 2455 Teller Road, Thousand Oaks California 91320 United States: SAGE Publications, Inc., 2008. [Online]. Available: <http://methods.sagepub.com/book/regression-basics>
- [111] C. Forbes, M. Evans, N. Hastings, and B. Peacock, *Statistical Distributions: Fourth Edition*, 2010.
- [112] E. L. Crow and K. Shimizu, Eds., *Lognormal Distributions: Theory and Applications*. New York: Marcel Dekker, Inc., 1988.

Appendix A

Hardware and software

Information about the computing hardware and software used in this work is presented in Tables A.1, A.2 and A.3.

Table A.1: Hardware used for discretisation.

Name	Version number
Operating system	Windows 7
System type	64-bit
RAM	8.00 <i>GB</i>
Processor	Intel Core i5 CPU @ 3.20 <i>GHz</i>

Table A.2: Hardware used for simulation.

Name	Version number
Operating system	Windows 10 Enterprise
System type	64-bit
RAM	16.0 <i>GB</i>
Processor	Intel Core i5 CPU @ 3.30 <i>GHz</i>

Table A.3: Software used in this work.

Name	Version number
MATLAB	R2015b
LS-DYNA	R7.1.0
ImageJ	1.52p
BoneJ	1.4.2
Python	3.6.2

Appendix B

Removal of free clusters

Trabecular bone has an extremely complex structure, a fact which is evident in the bone samples used in this work. Because of this complex structure, as well as the imperfect nature of the mesh discretisation (due to voxel size, thinning algorithm used, etc.), it is possible that after discretisation there will exist one or more elements which are not connected to the bulk of the structure. These free elements may take the form of a single floating element or a cluster of elements which are locally connected to each other, but not connected to the bulk. These clusters unnecessarily increase the size of the mesh, which increases simulation time, and may lead to erroneous results during Finite Element Analysis as they are ill-constrained. As such, it is important that these free components be identified and removed before simulation.

Although the identification of individual floating elements is relatively easily done, the identification of free element clusters is slightly more complex as a distinction needs to be made between those elements which are connected locally to each other and those which are connected to the bulk. It is convenient to develop code which works for both cases (free elements and free clusters) alike.

At its core, a Finite Element mesh is an example of a graph, meaning that there are a number of established methods from graph theory which can be drawn upon. In-house MATLAB code was developed, which makes use of a breadth-first traversal [109] to identify and remove free cluster in the mesh. The approach followed is described below:

1. All the nodes connected to node 1 (i.e. share an element with node 1) are identified and are added to a storage vector called *tempNodesNew*. The nodes which are added to *tempNodesNew* are removed from the node list used for the search.
2. The next node in *tempNodesNew* is targeted, and all nodes connected to this node are identified, added to the end of *tempNodesNew* and removed from the node list.
3. Step 2 is repeated for all the nodes in *tempNodesNew*, i.e. until all connected nodes have been identified.
4. If *tempNodesNew* contains all the nodes in the mesh, then clearly there are no free clusters in the mesh. However, if *tempNodesNew* contains fewer nodes than the total, this collection of nodes is identified as a cluster.

If the mesh contains clusters, it means that the bulk of the mesh will also be identified as a cluster, so there needs to be a check to determine whether a given cluster is a free cluster or a valid addition to the mesh. As such, if a cluster is identified, the code checks that the cluster is properly bounded, i.e. contains at least one node which lies on the top surface and one node which lies on the bottom surface, meaning that the cluster will contribute to the calculation of the mechanical properties of the sample during simulation. If the cluster is properly bounded, these nodes are discarded, however if a cluster is not bounded properly, it is identified as a free cluster and stored for later reference. A new starting node is identified from the remaining nodes in the node list (step 1 above) and the process starts again. Once all the nodes have been identified and assigned either to the bulk of the mesh or to a free cluster, the free nodes and elements are deleted from the mesh, along with their corresponding geometric properties.

Appendix C

Rasterisation

In-house code was developed for the rasterisation of beams and shells, respectively. In both of these methods, it is important to understand the distinction between mathematical concepts and rasterised objects. To define its orientation, a beam may be described as a mathematical line with a corresponding 3D line equation. Similarly, a shell may be defined as a collection of mathematical planes, with corresponding plane equations. Defining the beams and shells using mathematical techniques allows for calculations to be performed using the voxel coordinates, including calculations of perpendicular and parallel distances which are essential for accurate rasterisation. As a result, in this appendix, the rasterised objects will be referred to as “beams” and “shells”, whereas the mathematical concepts will be referred to as “lines” and “planes”.

C.1 Beam rasterisation

The reconstruction of the beams may be classified as 3D thick line rasterisation. Although there are some standard algorithms for rasterisation of lines (Bresenham’s Algorithm, Murphy’s Modified Bresenham Algorithm), extending these algorithms to 3D (i.e. using voxels rather than pixels) and non-integer points proved to be difficult and, consequently, a different approach was necessary.

The algorithm developed to rasterise the beams is based on the fact that for any rasterised voxel:

- The perpendicular distance from the voxel centre to the mathematical line must be less than or equal to the radius of the beam ($d_{\perp} \leq r$).
- The parallel distance from the voxel centre to the midpoint of the mathematical line must be less than or equal to half of the total line length ($d_{\parallel} \leq \frac{l}{2}$).

This concept is demonstrated in Figures C.1-C.2, which show examples of voxels which meet the perpendicular and parallel distance criteria, respectively. These figures are based on a circular beam, perpendicular to the xy -plane for ease of illustration, but in reality, the beam could be oriented in any direction.




Figure C.1: Perpendicular distance criteria for beam rasterisation. The distance from the centre of the circle to voxel A is larger than the radius of the circle, therefore voxel A will not be rasterised. The distance from the centre of the circle to voxel B is smaller than the radius of the circle, therefore voxel B will be rasterised.




Figure C.2: Parallel distance criteria for beam rasterisation. The distance from the midpoint (shown as a dashed line) of the beam to voxel A is larger than half of the total beam length, therefore voxel A will not be rasterised. The distance from the midpoint of the beam to voxel B is smaller than the half of the total beam length, therefore voxel B will be rasterised.




Figure C.3: Bounding mathematical planes used in the thick shell rasterization. Here the coloured planes represent the bounding planes, otherwise viewed as the faces of the shell.

The algorithm for the beam rasterisation may be summarised as follows:

1. The coordinate values for the starting point of the beam are rounded to the nearest integer values. The voxel associated with the resulting coordinates is chosen as the starting voxel and the voxel coordinates are added to the array, *voxelCoords*.
2. For each of the 26 voxels neighbouring the starting voxel, the parallel and perpendicular distances from the voxel to the line are calculated. If the calculated distances meet the criteria specified earlier in this section, the voxel coordinates are added to *voxelCoords*.
3. Once all 26 neighbouring voxels have been checked, the process is repeated for the next voxel in *voxelCoords*.
4. The process is terminated when all the voxels in *voxelCoords* have been analysed.

C.2 Shell rasterisation

A more complex technique was required for the rasterisation of the shells because the thickness of the shells varies across the element. As such, in-house code was developed which determines whether or not to rasterise a voxel based on its position relative to the 5 mathematical planes which bound the shell. These bounding planes, which are shown in Figure C.3, represent the 5 faces of the shell and, therefore, take into account the varying thickness across the element.

The equation of a plane may be defined as:

$$ax + by + cz + d = 0. \quad (\text{C.1})$$

Any point (x, y, z) which lies on the plane will satisfy this equation. If a point which does not lie on the plane is substituted into the plane formula¹, the result will be either positive or negative, depending on the placement of the point relative to the plane. For example, consider a mathematical plane with normal, $\underline{n} = (0, 0, 1)$:

$$0x + 0y + 1z - 1 = 0 \quad \rightarrow \quad f(z) = z - 1 = 0 \quad (\text{C.2})$$

If a point above the plane, $p_1 = (0, 0, 1.5)$, is substituted into the plane formula, a positive result is obtained:

$$f(1.5) = 1.5 - 1 = 0.5 > 0. \quad (\text{C.3})$$

However, if a point below the plane, $p_2 = (0, 0, 0.5)$, is substituted into the plane formula, a negative result is obtained:

$$f(0.5) = 0.5 - 1 = -0.5 < 0. \quad (\text{C.4})$$

Although this phenomenon is demonstrated using a simple example, the same technique may be applied to any mathematical plane. In order to apply this technique to the rasterisation of the shells, it was noted that any rasterised voxel must satisfy a condition for each of the bounding planes. For example, if all the voxels with $0 \leq z \leq 1$ were to be rasterised, 2 bounding planes could be defined, namely the plane in Equation C.2 and an additional plane:

$$0x + 0y + 1z - 0 = 0 \quad \rightarrow \quad g(z) = z = 0. \quad (\text{C.5})$$

If p_2 is substituted into Equation C.5, a positive result is obtained:

$$g(0.5) = 0.5 > 0. \quad (\text{C.6})$$

Therefore, any voxel which produces a negative result for plane $f(z)$ and a positive result for plane $g(z)$ satisfies the bounding criteria and should be rasterised. If a voxel only satisfies one of the two criteria, it does not lie in the bounded region and should not be rasterised. This technique may be extended to include any number of bounding planes, which is the approach followed for the rasterisation of the shell elements.

The algorithm for the shell rasterisation may be summarised as follows:

1. A mathematical plane is defined for each face of the shell. This is done by finding three points on each plane and using the coordinates to calculate the equation of the plane using the cross product and substitution.
2. A point in the centre of the shell is identified. The coordinates associated with this point are substituted into each plane equation to determine whether a positive or negative result is produced. This result is stored for each plane.
3. The coordinates of the starting point are rounded to the nearest integer values. This voxel is defined as the starting voxel and is added to the array *voxelCoords*.
4. For each of the 26 voxels neighbouring the starting voxel, the coordinates are substituted into the plane equations and compared to the benchmark described in (2). If the point produces the same combination of positive and negative values as is listed in the benchmark, it means that the point lies within the shell and the voxel coordinates are added to *voxelCoords*.
5. Once all 26 neighbouring voxels have been checked, the process is repeated for the next voxel in *voxelCoords*.
6. The process is terminated when all the voxels in *voxelCoords* have been analysed.

¹The plane formula here refers to the left hand side of the equation, i.e. $f(x, y, z) = ax + by + cz + d$.

C.3 Optimisation

Although the basic algorithms for the beam and shell rasterisation fulfil their purpose, the efficiency of the methods leaves much to be desired. The algorithms save time by only checking the neighbours of rasterised voxels (as opposed to checking all voxels in a given region), but using this method, a single voxel may still be analysed multiple times. In order to optimise the algorithms, additional code was written to ensure that each voxel is only analysed once, i.e. the distance calculations and plane substitutions are only performed for voxels which have not previously been analysed.


To this end, a 3D matrix is created at the beginning of the rasterisation process such that each entry in the matrix represents a voxel in a chosen volume around the beam or shell, i.e. the volume matrix runs from $(x_{min} \rightarrow x_{max}, y_{min} \rightarrow y_{max}, z_{min} \rightarrow z_{max})$. Each time a voxel is analysed, the corresponding entry in the bounding matrix is changed to 1, regardless of whether it is rasterised or not. When a neighbouring voxel is analysed, the code checks the volume matrix before performing the distance calculations. If an entry is specified as 1, it has already been analysed and, therefore, the code skips this voxel and moves on to the next without redoing the distance or plane calculations. Using this method, each voxel is only analysed once, vastly decreasing the processing time of the algorithm.

Appendix D


Validation results

D.1 Cubic lattices

D.1.1 C6-t1 - unconstrained



(a) Uncropped.




(b) Cropped.

Figure D.1: Error plots for apparent modulus of unconstrained C6-t1 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.2: Error plots for inner modulus of unconstrained C6-t1 lattices at varying image sizes.



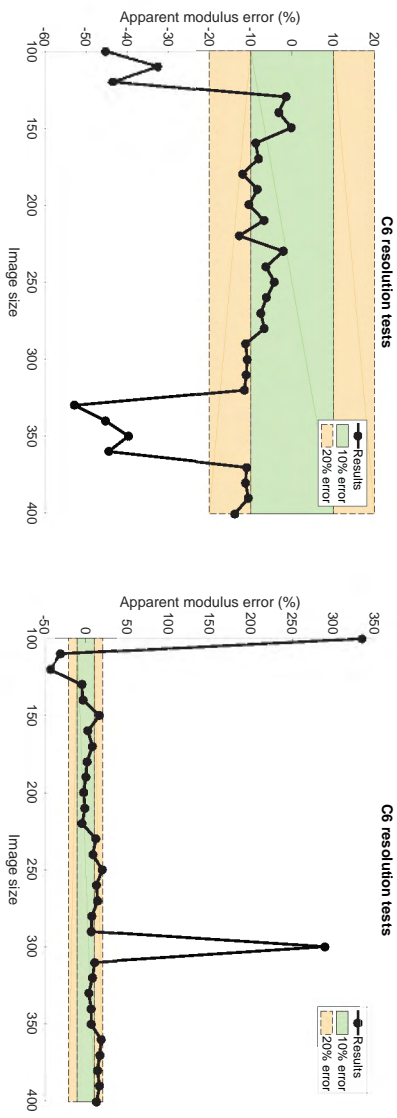
(a) Uncropped.



(b) Cropped.

Figure D.3: Modulus ratio results of unconstrained C6-t1 lattices at varying image sizes.

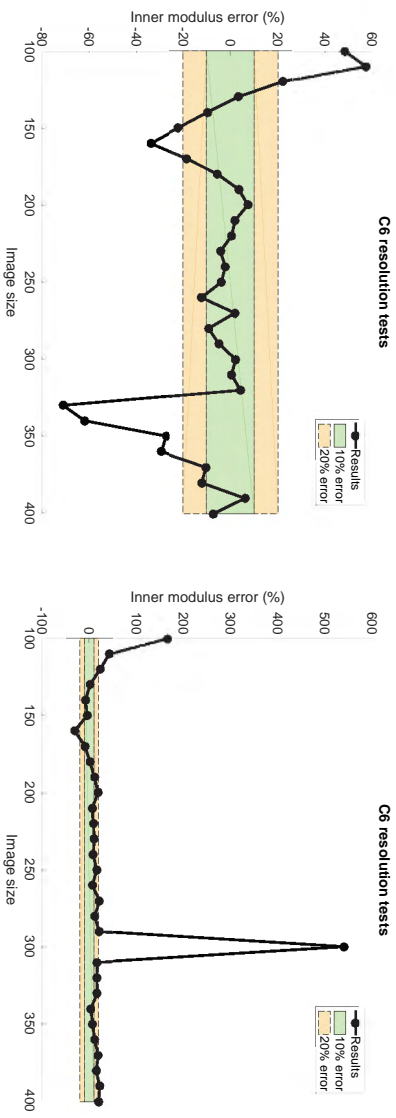
D.1.2 C6-t1 - constrained



(a) Uncropped.

(b) Cropped.


Figure D.4: Error plots for apparent modulus of constrained C6-t1 lattices at varying image sizes.



(a) Uncropped.

(b) Cropped.

Figure D.5: Error plots for inner modulus of constrained C6-t1 lattices at varying image sizes.




(a) Uncropped.


(b) Cropped.

Figure D.6: Modulus ratio results of constrained C6-t1 lattices at varying image sizes.

D.1.3 C6-t2 - unconstrained




(a) Uncropped.




(b) Cropped.

Figure D.7: Error plots for apparent modulus of unconstrained C6-t2 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.8: Error plots for inner modulus of unconstrained C6-t2 lattices at varying image sizes.



(a) Uncropped.



(b) Cropped.

Figure D.9: Modulus ratio results of unconstrained C6-t2 lattices at varying image sizes.

D.1.4 C6-t2 - constrained

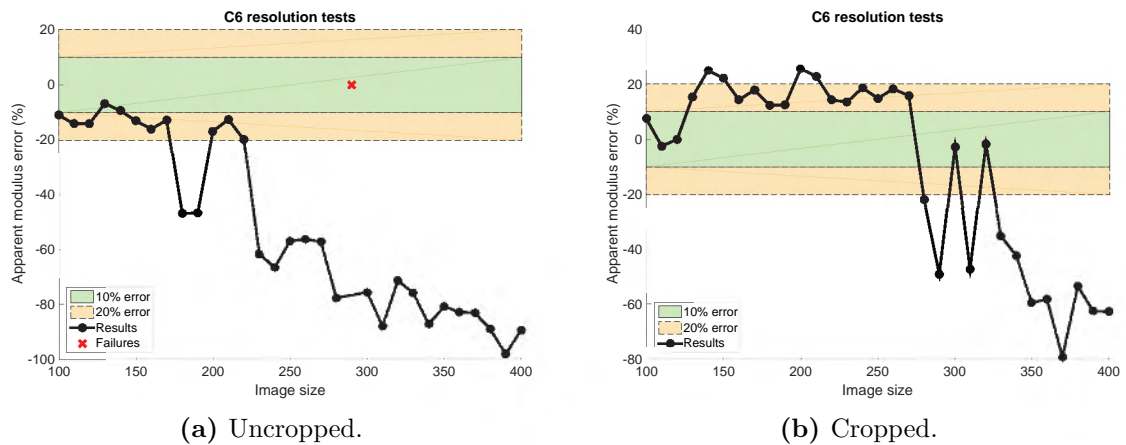


Figure D.10: Error plots for apparent modulus of constrained C6-t2 lattices at varying image sizes.

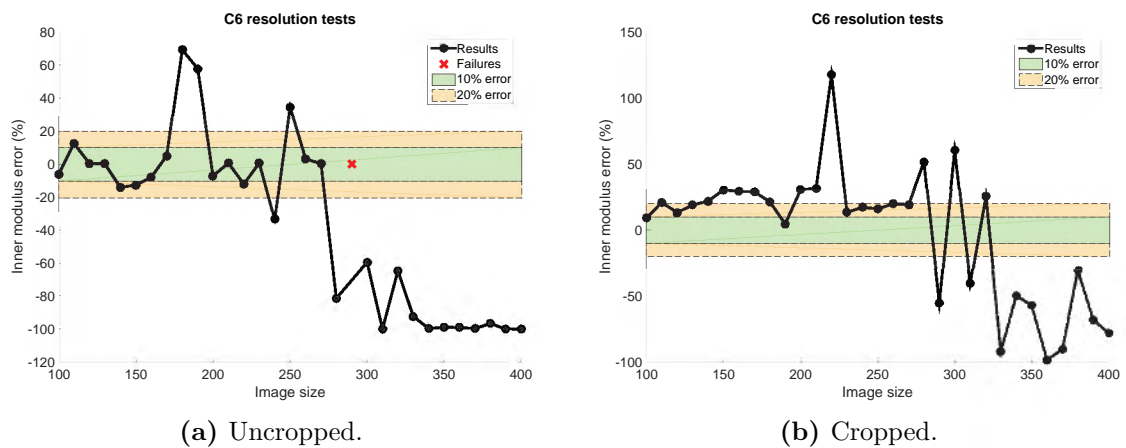


Figure D.11: Error plots for inner modulus of constrained C6-t2 lattices at varying image sizes.

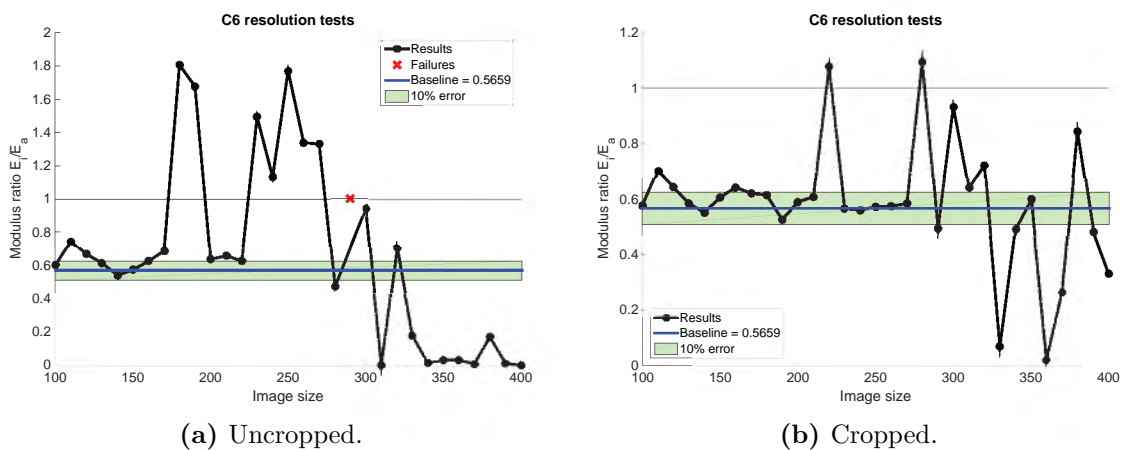


Figure D.12: Modulus ratio results of constrained C6-t2 lattices at varying image sizes.

D.1.5 C6-t3 - unconstrained

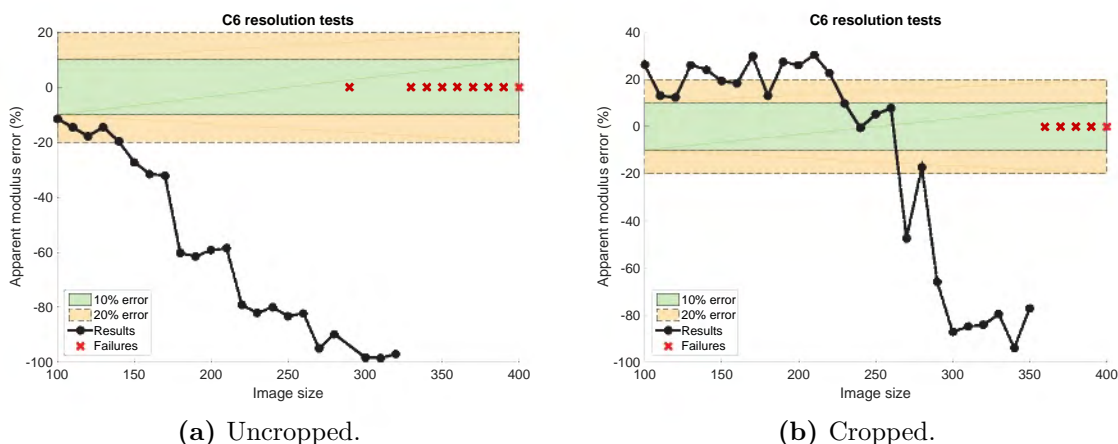


Figure D.13: Error plots for apparent modulus of unconstrained C6-t3 lattices at varying image sizes.

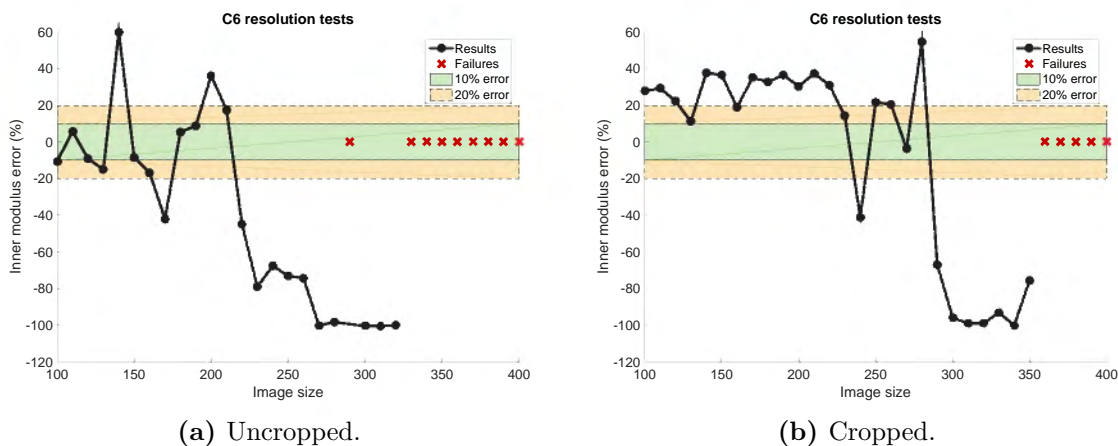


Figure D.14: Error plots for inner modulus of unconstrained C6-t3 lattices at varying image sizes.

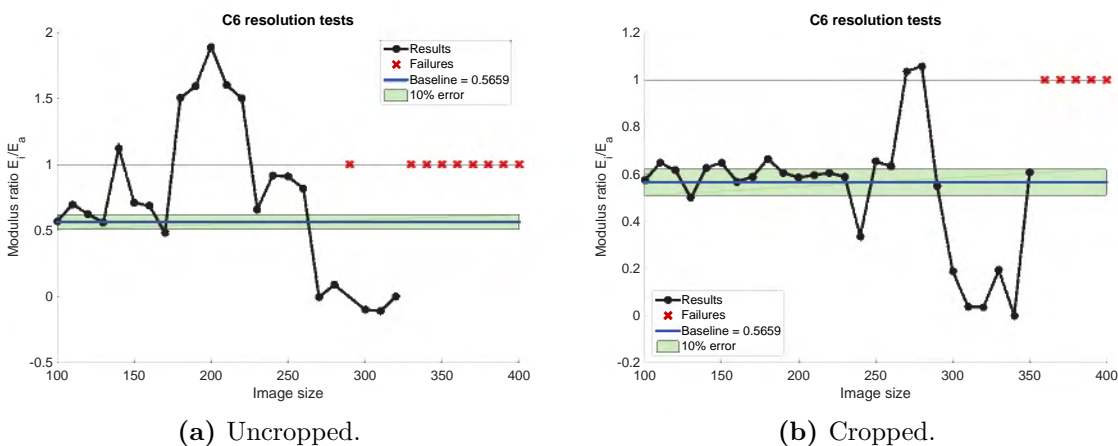




Figure D.15: Modulus ratio results of unconstrained C6-t3 lattices at varying image sizes.

D.1.6 C6-t3 - constrained




(a) Uncropped.




(b) Cropped.

Figure D.16: Error plots for apparent modulus of constrained C6-t3 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.17: Error plots for inner modulus of constrained C6-t3 lattices at varying image sizes.



(a) Uncropped.



(b) Cropped.

Figure D.18: Modulus ratio results of constrained C6-t3 lattices at varying image sizes.

D.1.7 C8-t1 - unconstrained

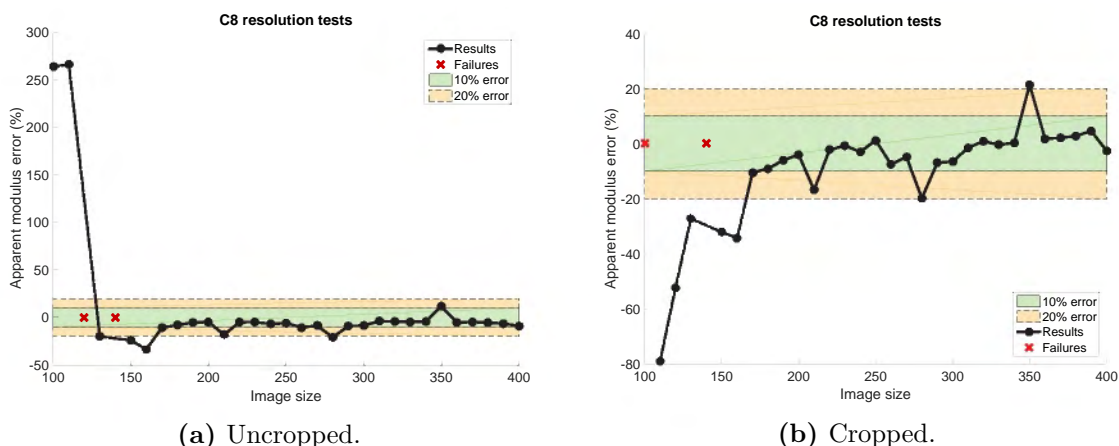


Figure D.19: Error plots for apparent modulus of unconstrained C8-t1 lattices at varying image sizes.

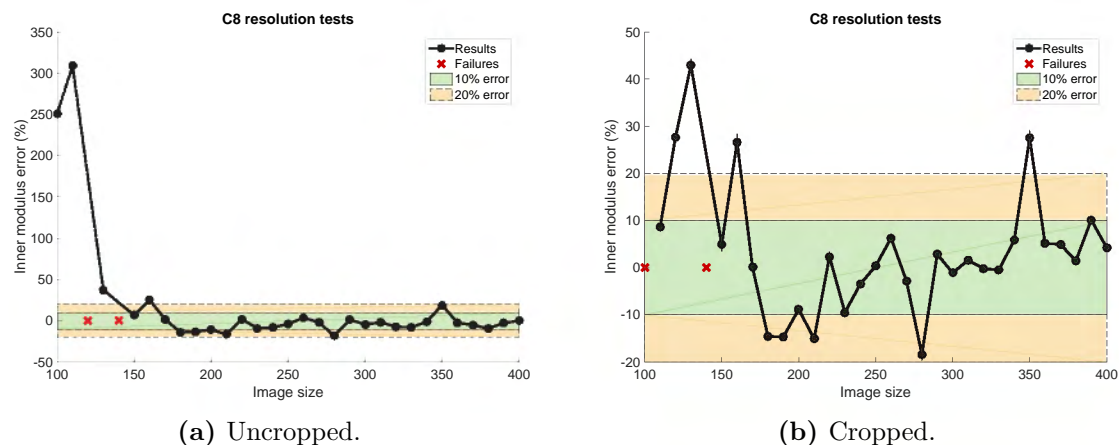


Figure D.20: Error plots for inner modulus of unconstrained C8-t1 lattices at varying image sizes.

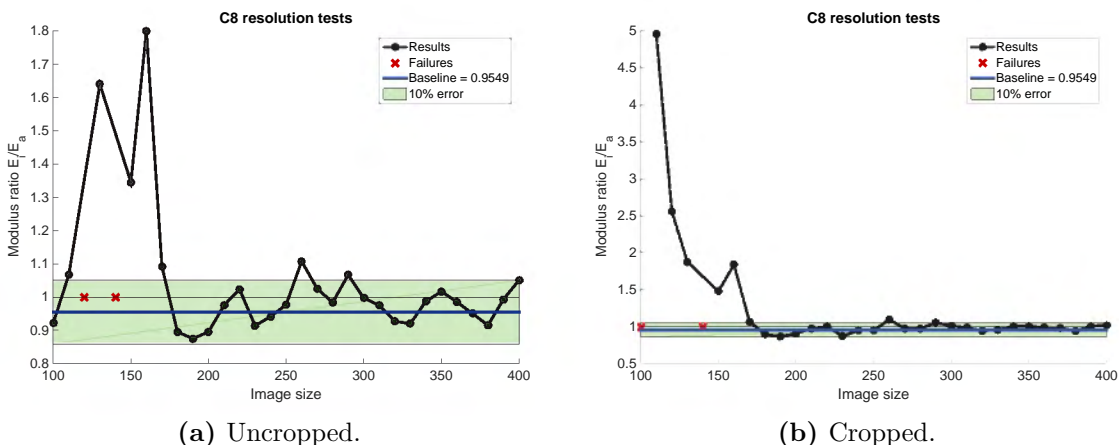


Figure D.21: Modulus ratio results of unconstrained C8-t1 lattices at varying image sizes.

D.1.8 C8-t1 - constrained




Figure D.22: Error plots for apparent modulus of constrained C8-t1 lattices at varying image sizes.

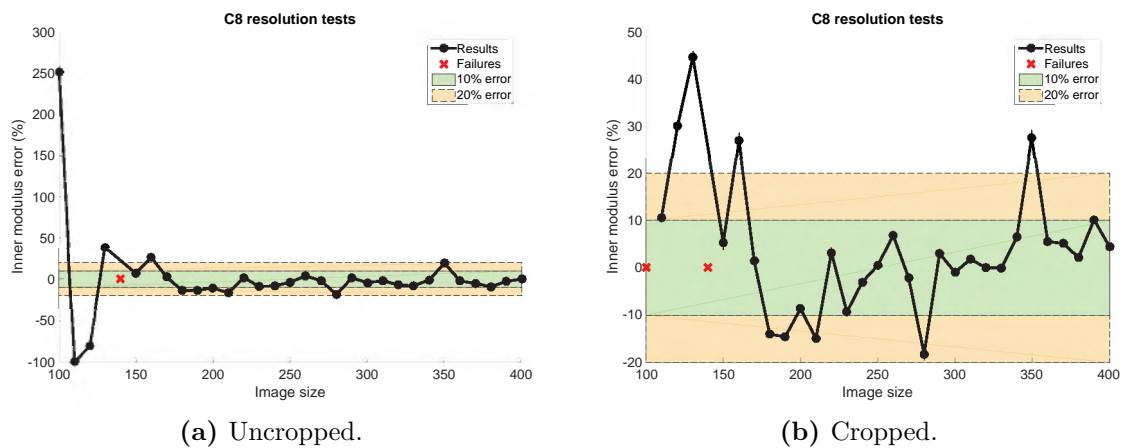


Figure D.23: Error plots for inner modulus of constrained C8-t1 lattices at varying image sizes.

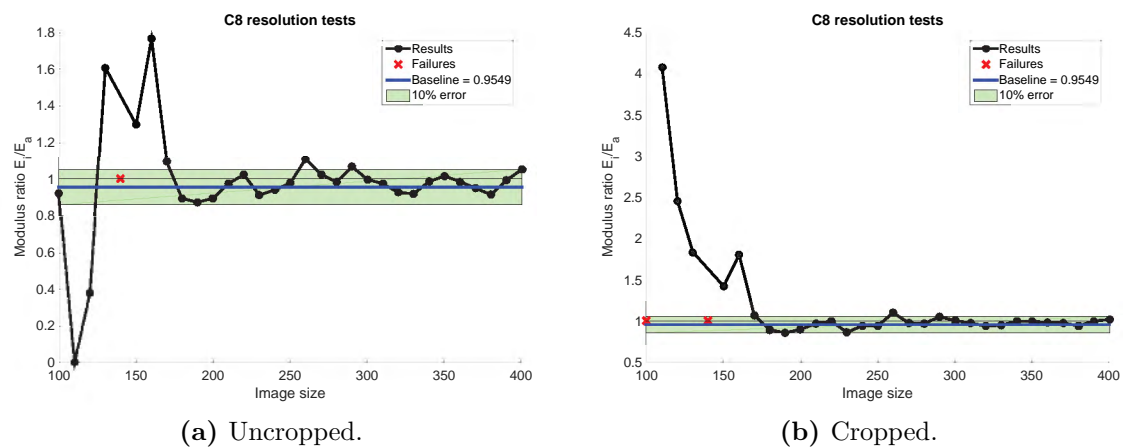


Figure D.24: Modulus ratio results of constrained C8-t1 lattices at varying image sizes.

D.1.9 C8-t2 - unconstrained

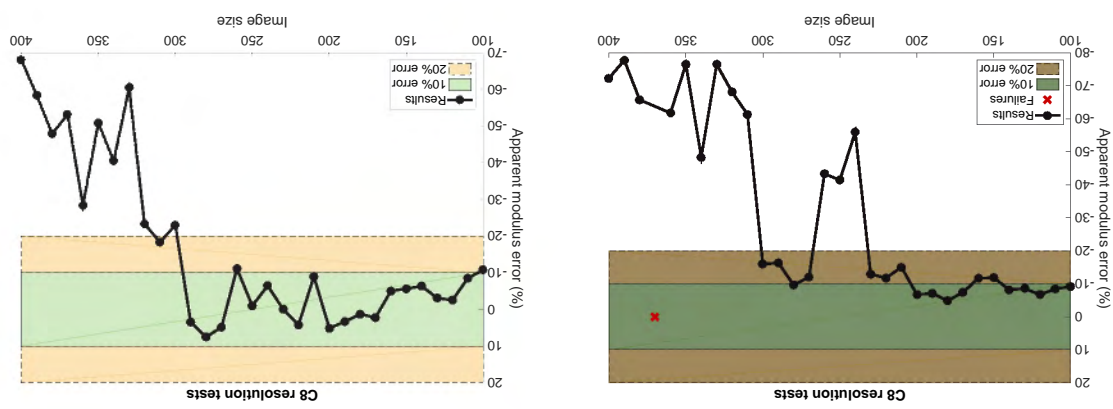


Figure D.25: Error plots for apparent modulus of unconstrained C8-t2 lattices at varying image sizes.

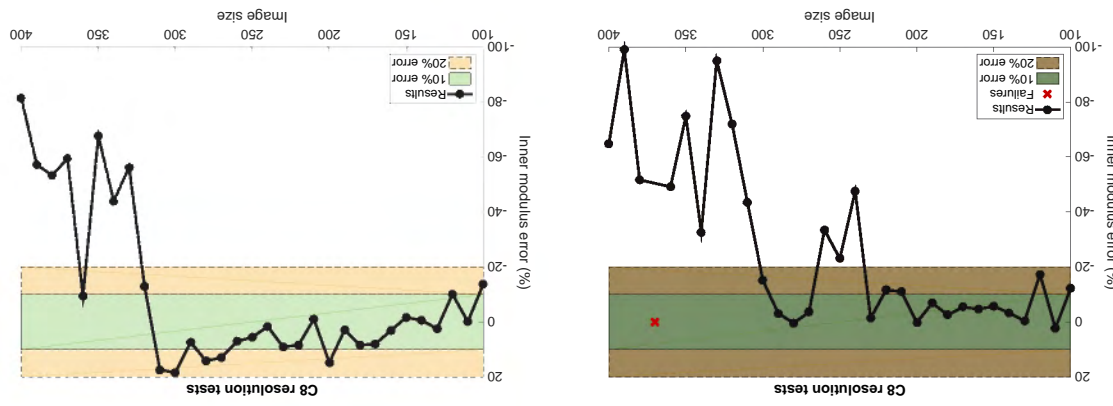


Figure D.26: Error plots for inner modulus of unconstrained C8-t2 lattices at varying image sizes.

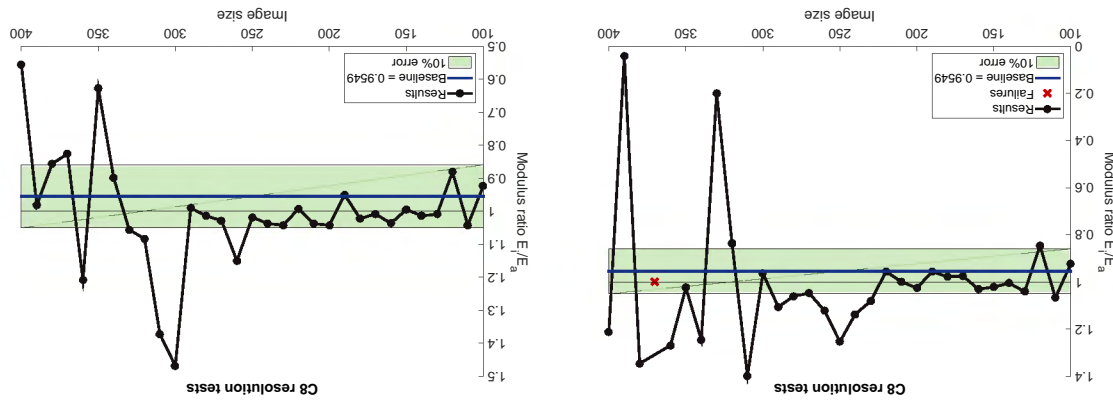


Figure D.27: Modulus ratio of unconstrained C8-t2 lattices at varying image sizes.

D.1.10 C8-t2 - constrained

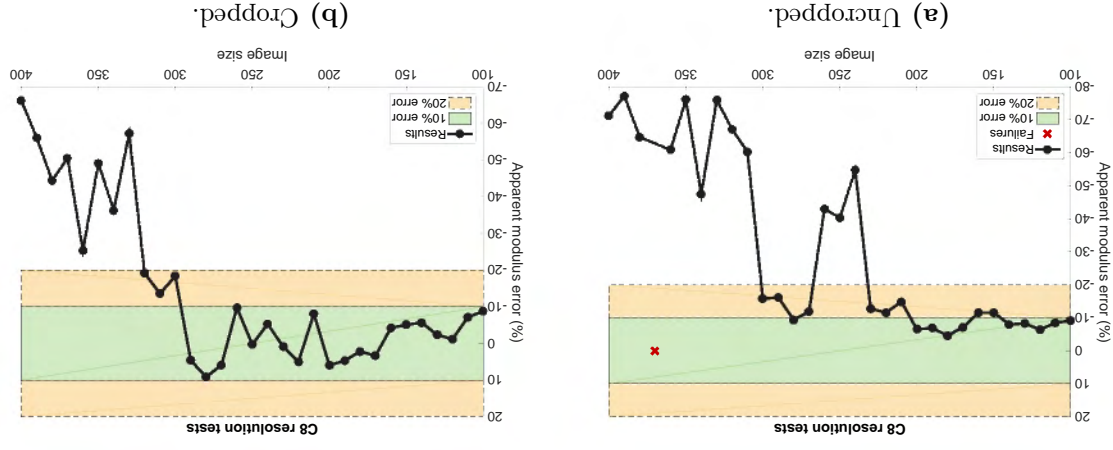


Figure D.28: Error plots for apparent modulus of constrained C8-t2 lattices at varying image sizes.

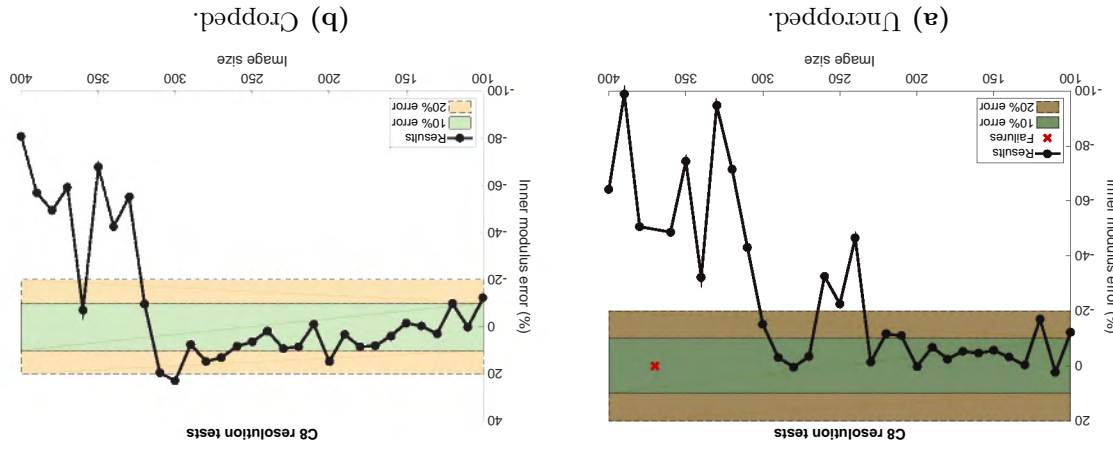


Figure D.29: Error plots for inner modulus of constrained C8-t2 lattices at varying image sizes.

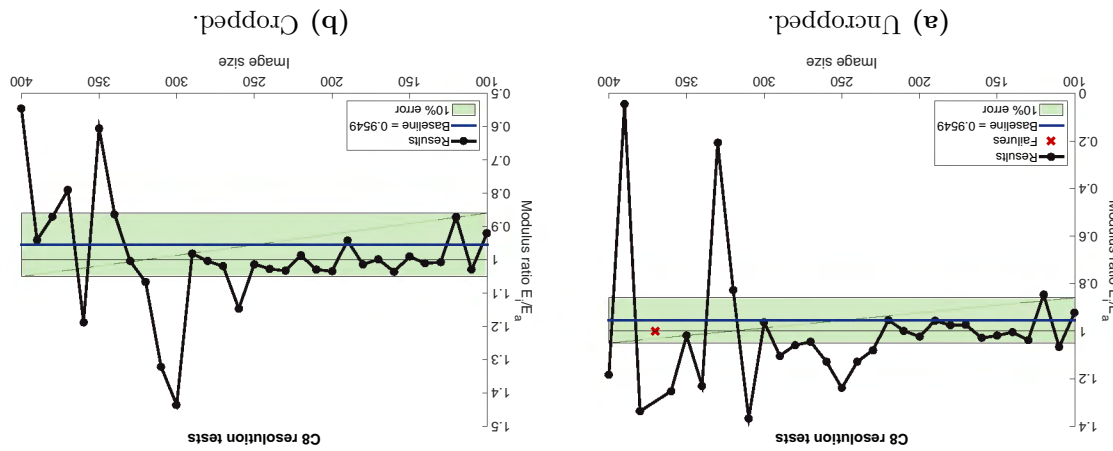


Figure D.30: Modulus ratio results of constrained C8-t2 lattices at varying image sizes.

D.1.11 C8-t3 - unconstrained

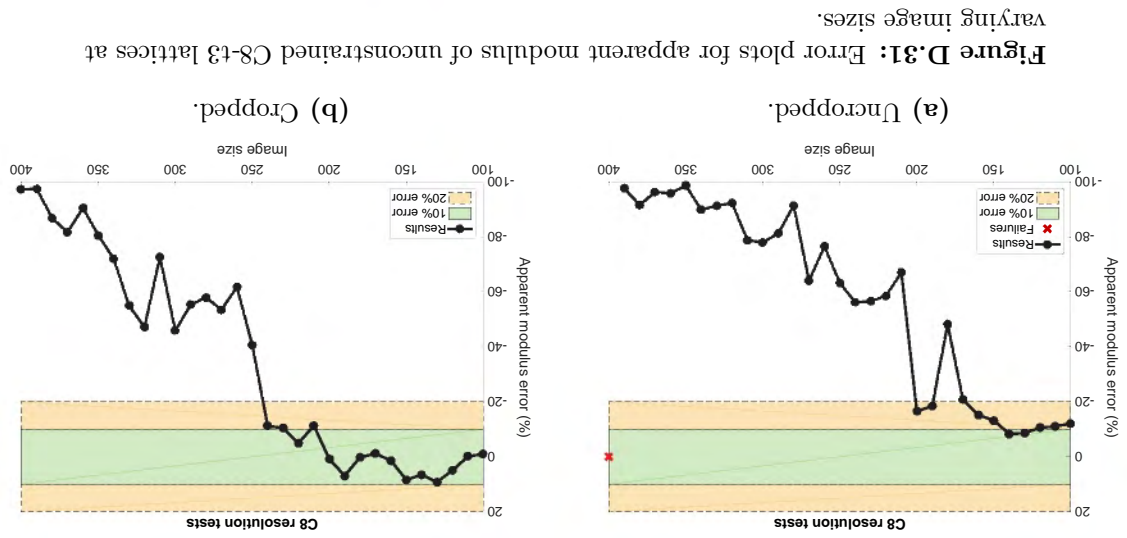


Figure D.31: Error plots for apparent modulus of unconstrained C8-t3 lattices at varying image sizes.

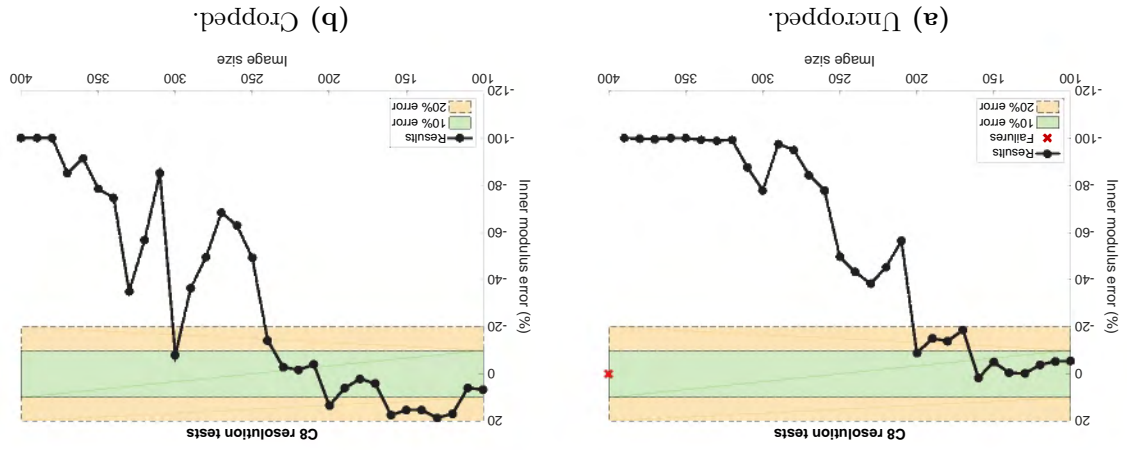


Figure D.32: Error plots for inner modulus of unconstrained C8-t3 lattices at varying image sizes.

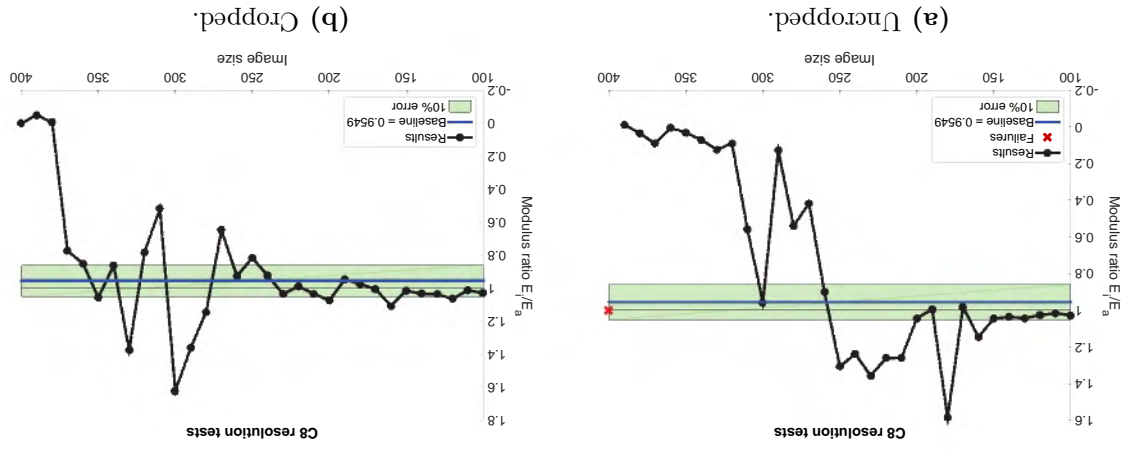


Figure D.33: Modulus ratio results of unconstrained C8-t3 lattices at varying image sizes.




Figure D.34: Error plots for apparent modulus of constrained C8-t3 lattices at varying image sizes.

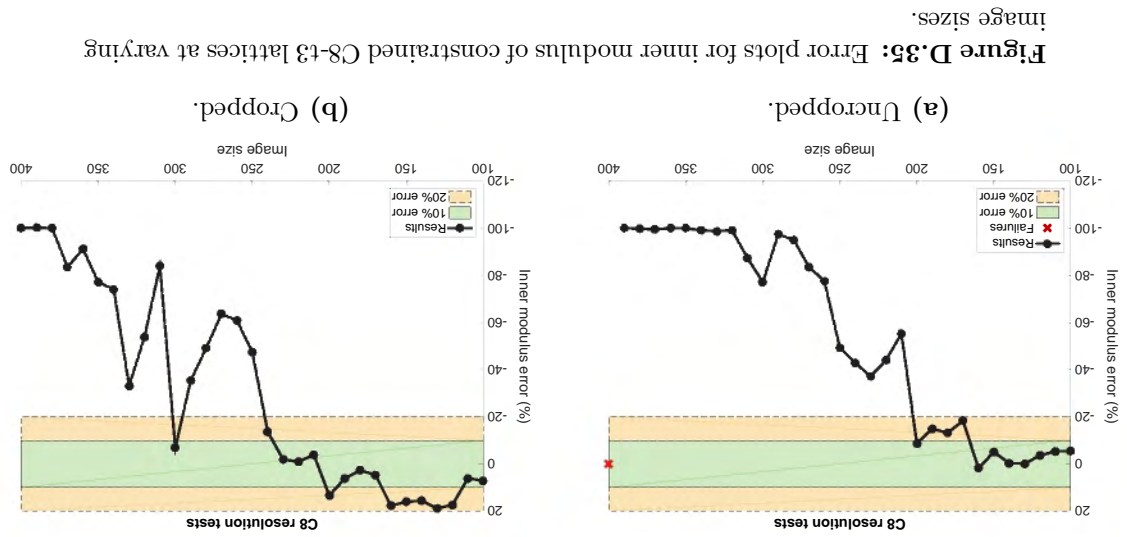


Figure D.35: Error plots for inner modulus of constrained C8-t3 lattices at varying image sizes.

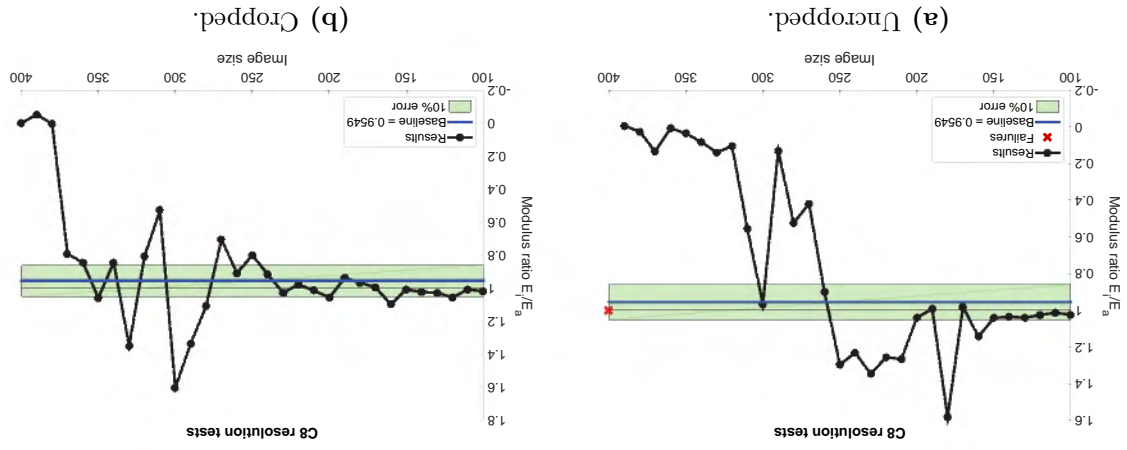


Figure D.36: Modulus ratio of constrained C8-t3 lattices at varying image sizes.

D.1.13 C10-t1 - unconstrained

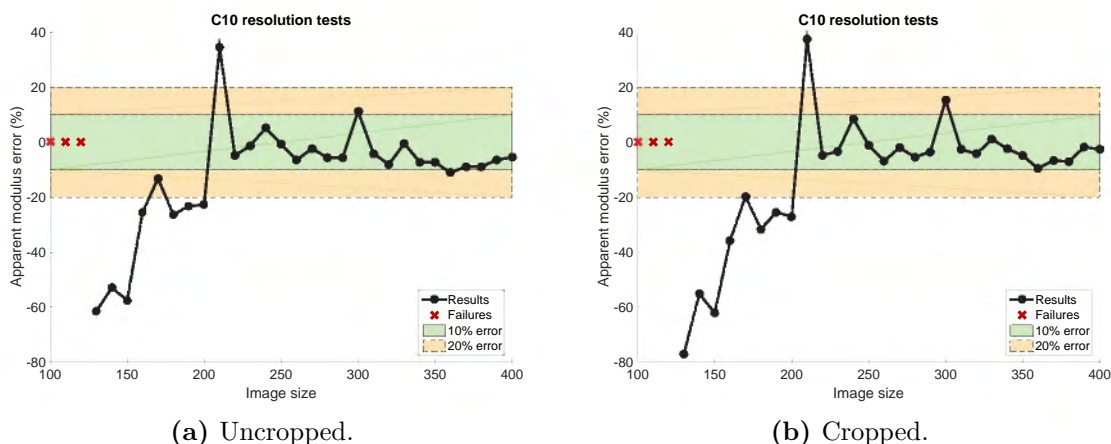


Figure D.37: Error plots for apparent modulus of unconstrained C10-t1 lattices at varying image sizes.

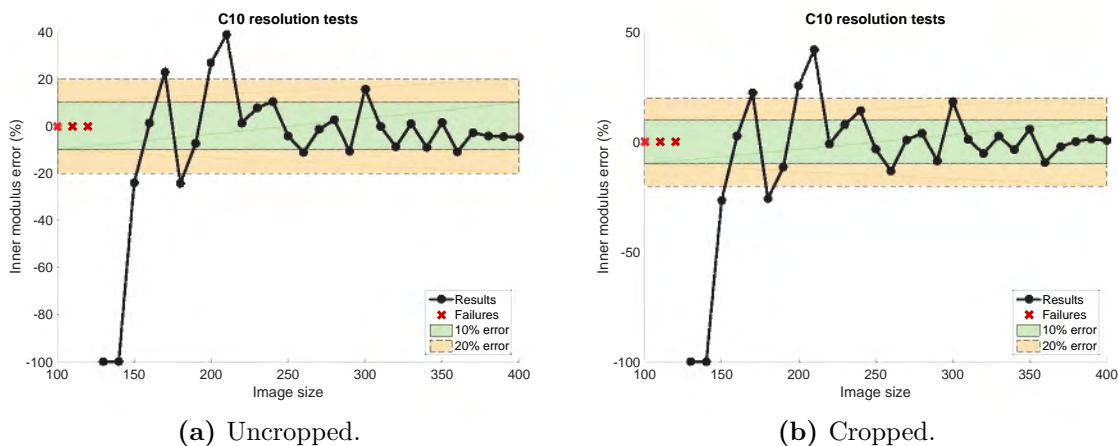


Figure D.38: Error plots for inner modulus of unconstrained C10-t1 lattices at varying image sizes.

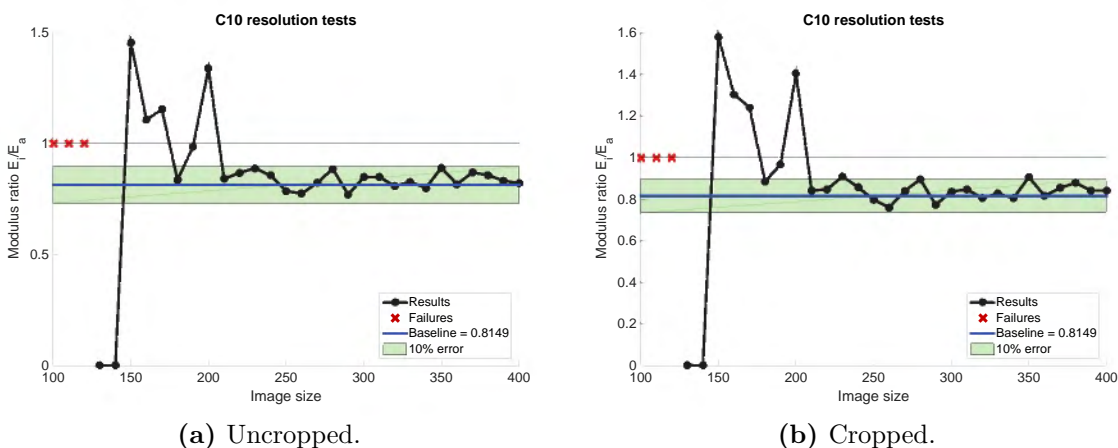


Figure D.39: Modulus ratio results of unconstrained C10-t1 lattices at varying image sizes.

D.1.14 C10-t1 - constrained

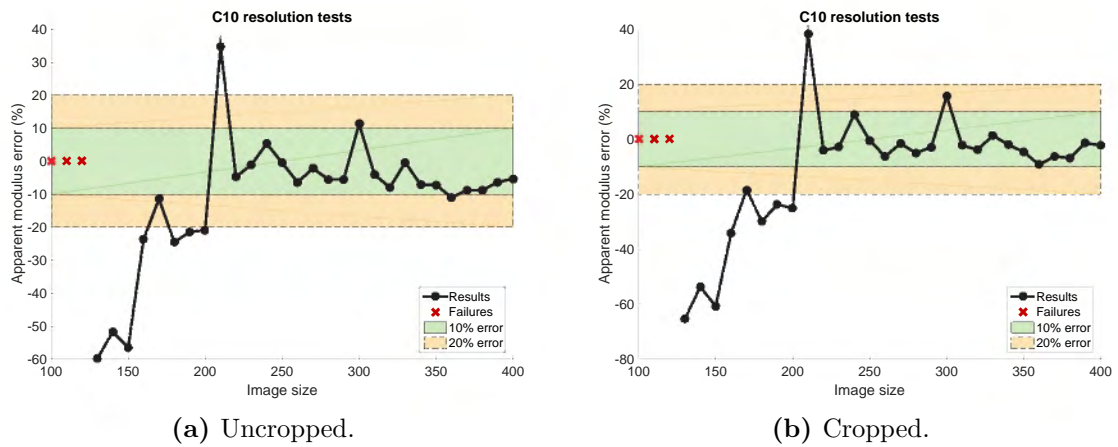


Figure D.40: Error plots for apparent modulus of constrained C10-t1 lattices at varying image sizes.




Figure D.41: Error plots for inner modulus of constrained C10-t1 lattices at varying image sizes.

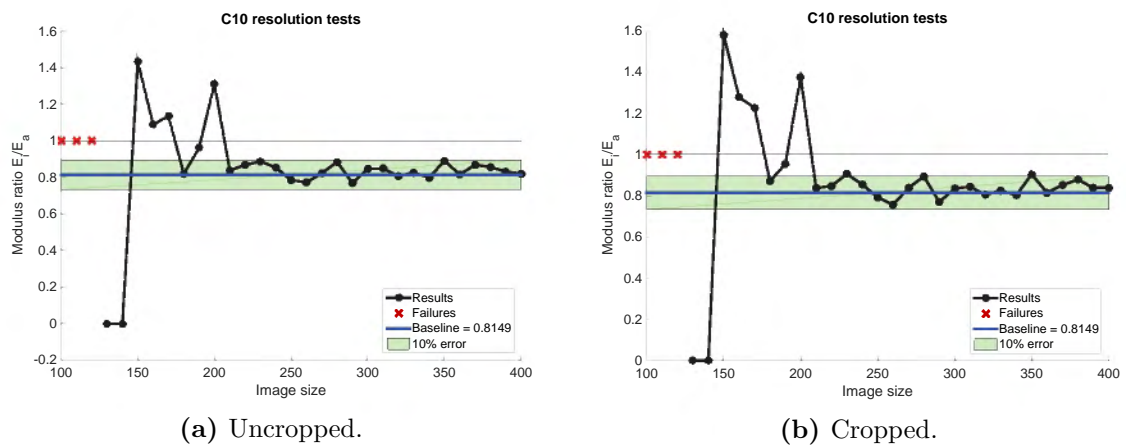


Figure D.42: Modulus ratio results of constrained C10-t1 lattices at varying image sizes.

D.1.15 C10-t2 - unconstrained




Figure D.43: Error plots for apparent modulus of unconstrained C10-t2 lattices at varying image sizes.




Figure D.44: Error plots for inner modulus of unconstrained C10-t2 lattices at varying image sizes.




Figure D.45: Modulus ratio results of unconstrained C10-t2 lattices at varying image sizes.

D.1.16 C10-t2 - constrained




Figure D.46: Error plots for apparent modulus of constrained C10-t2 lattices at varying image sizes.

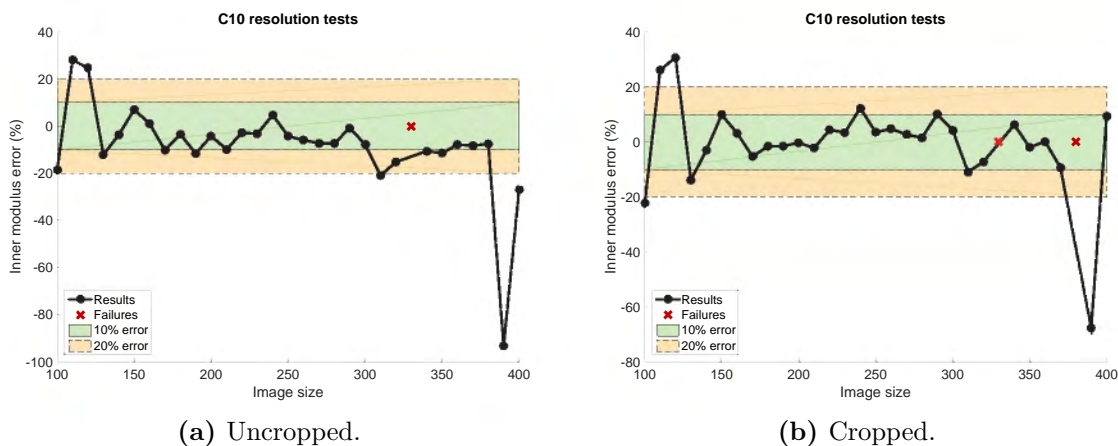


Figure D.47: Error plots for inner modulus of constrained C10-t2 lattices at varying image sizes.




Figure D.48: Modulus ratio results of constrained C10-t2 lattices at varying image sizes.

D.1.17 C10-t3 - unconstrained

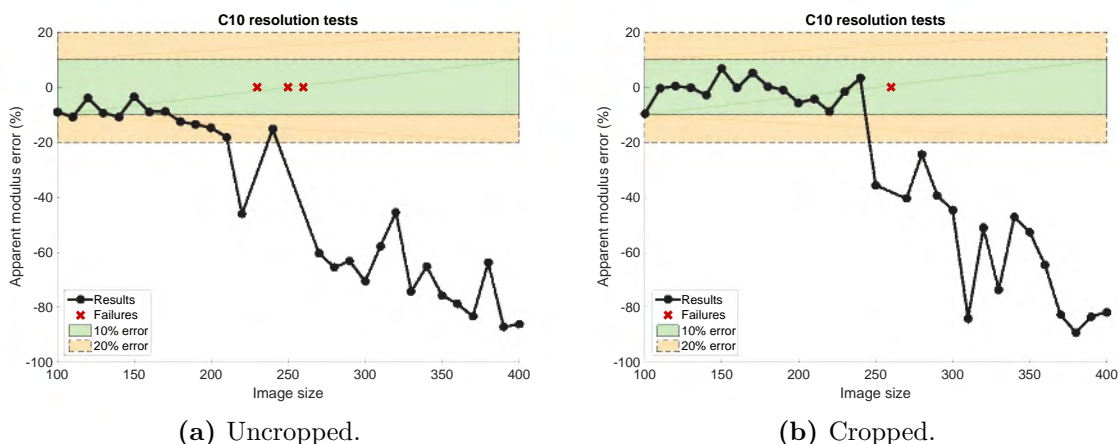


Figure D.49: Error plots for apparent modulus of unconstrained C10-t3 lattices at varying image sizes.




Figure D.50: Error plots for inner modulus of unconstrained C10-t3 lattices at varying image sizes.




Figure D.51: Modulus ratio results of unconstrained C10-t3 lattices at varying image sizes.

D.1.18 C10-t3 - constrained

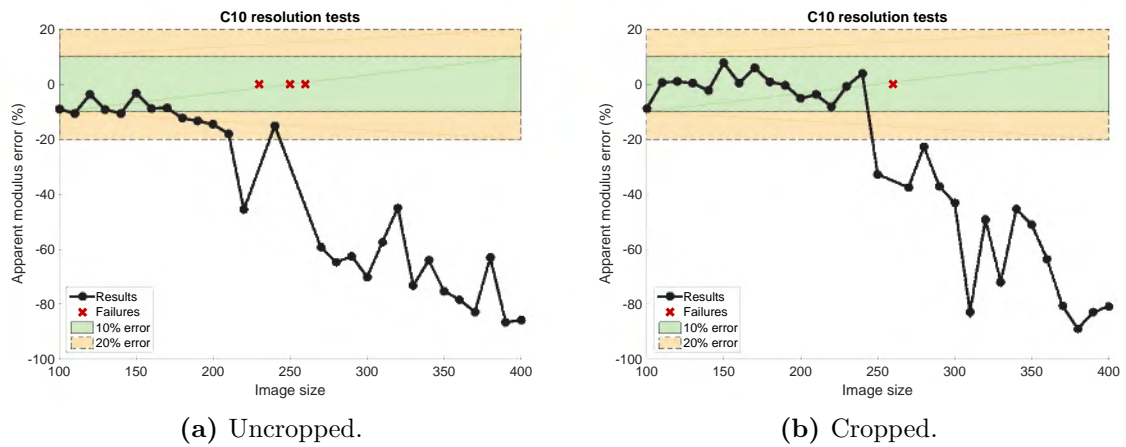


Figure D.52: Error plots for apparent modulus of constrained C10-t3 lattices at varying image sizes.

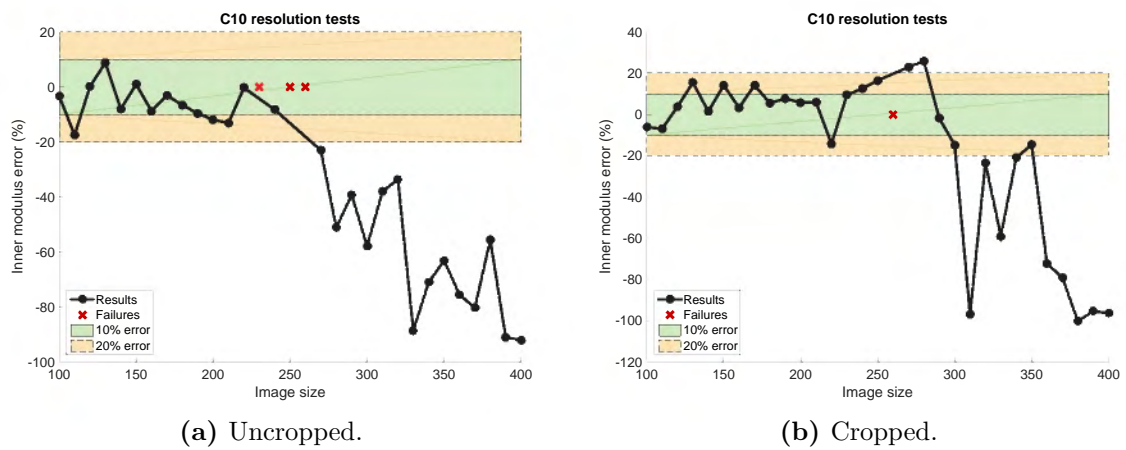


Figure D.53: Error plots for inner modulus of constrained C10-t3 lattices at varying image sizes.




Figure D.54: Modulus ratio results of constrained C10-t3 lattices at varying image sizes.

D.2 Kelvin cell lattices

D.2.1 K6-t1 - unconstrained

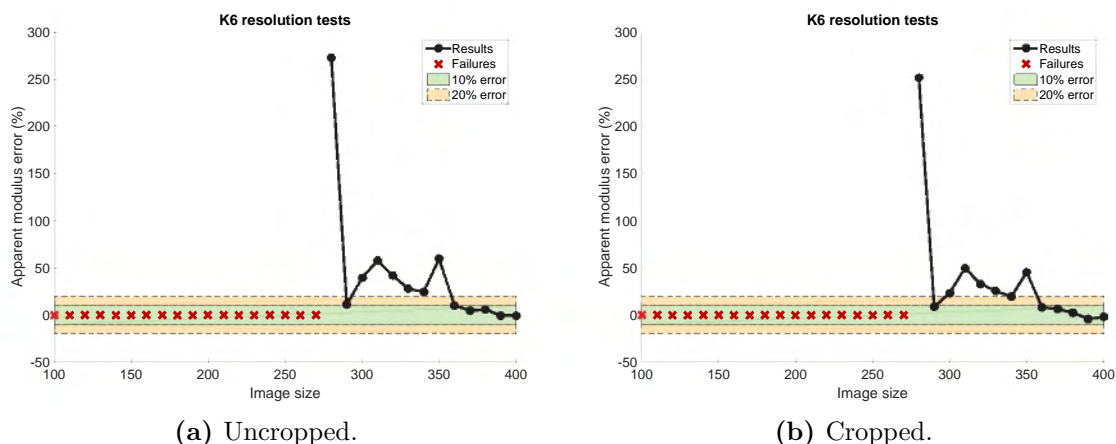


Figure D.55: Error plots for apparent modulus of unconstrained K6-t1 lattices at varying image sizes.

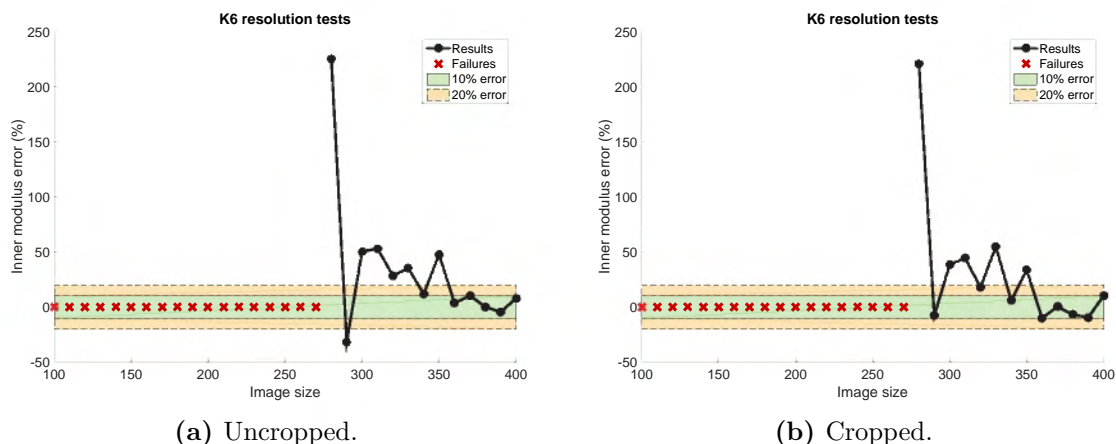


Figure D.56: Error plots for inner modulus of unconstrained K6-t1 lattices at varying image sizes.




Figure D.57: Modulus ratio results of unconstrained K6-t1 lattices at varying image sizes.

D.2.2 K6-t1 - constrained

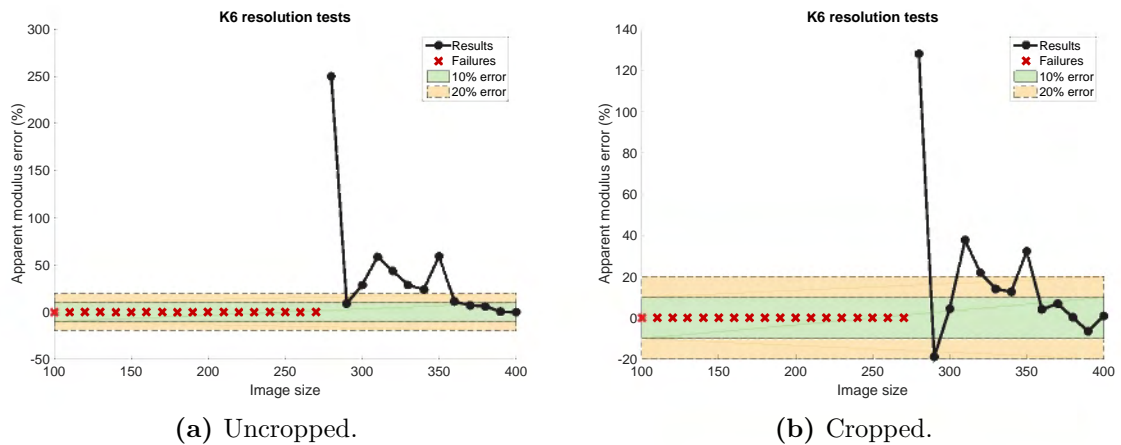


Figure D.58: Error plots for apparent modulus of constrained K6-t1 lattices at varying image sizes.

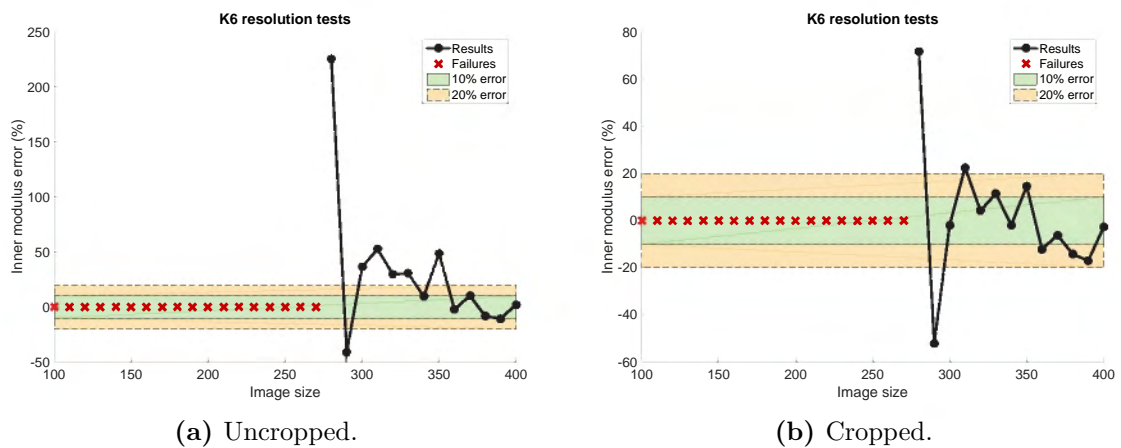


Figure D.59: Error plots for inner modulus of constrained K6-t1 lattices at varying image sizes.




Figure D.60: Modulus ratio results of constrained K6-t1 lattices at varying image sizes.

D.2.3 K6-t2 - unconstrained

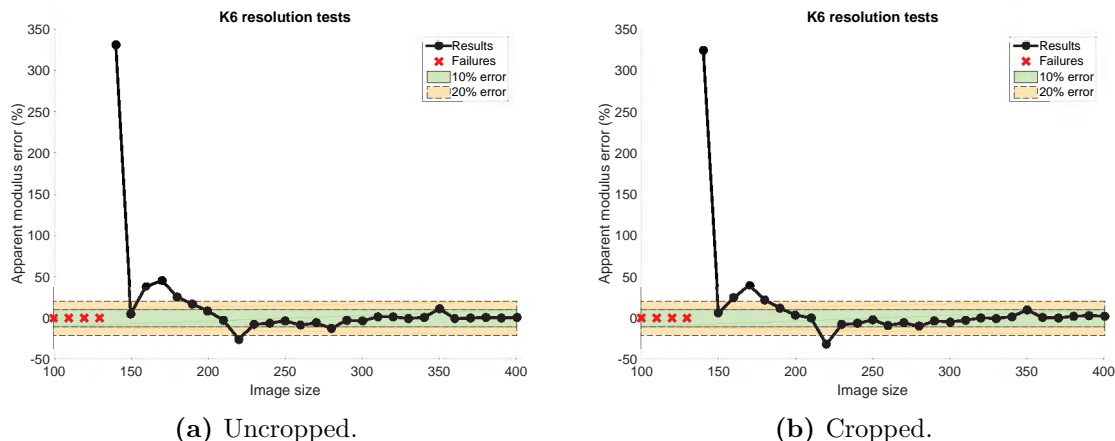


Figure D.61: Error plots for apparent modulus of unconstrained K6-t2 lattices at varying image sizes.

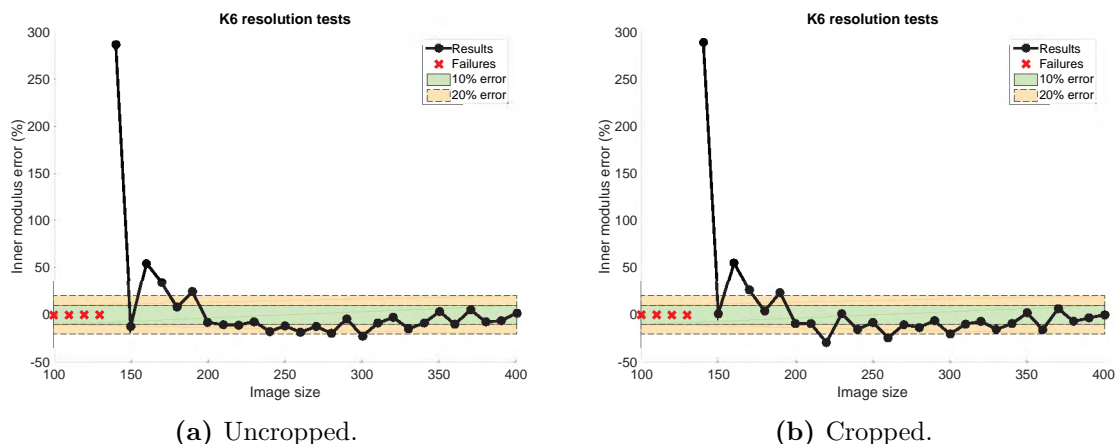


Figure D.62: Error plots for inner modulus of unconstrained K6-t2 lattices at varying image sizes.




Figure D.63: Modulus ratio results of unconstrained K6-t1 lattices at varying image sizes.

D.2.4 K6-t2 - constrained

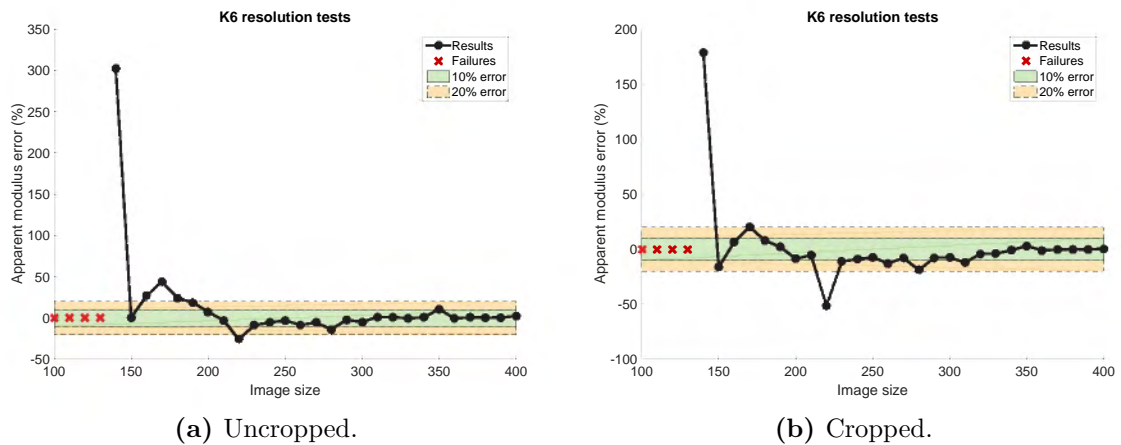


Figure D.64: Error plots for apparent modulus of constrained K6-t2 lattices at varying image sizes.

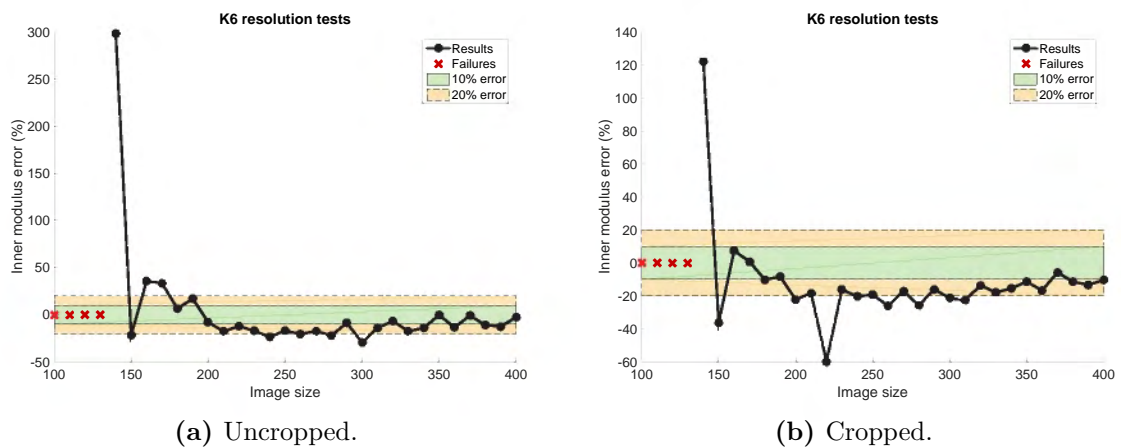


Figure D.65: Error plots for inner modulus of constrained K6-t2 lattices at varying image sizes.

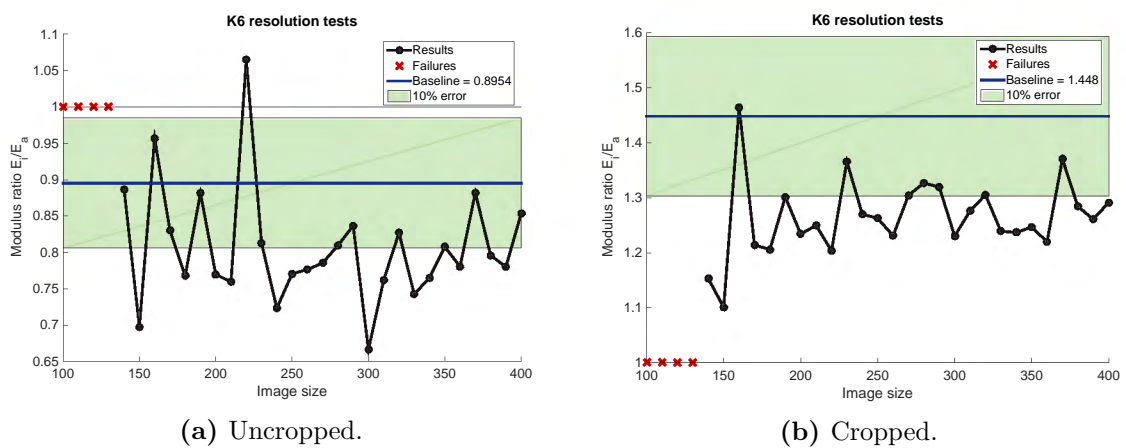
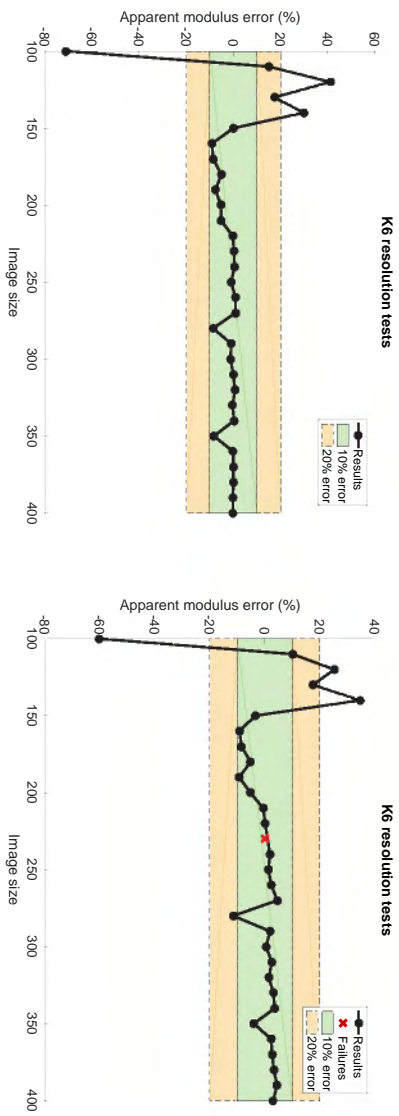


Figure D.66: Modulus ratio results of constrained K6-t2 lattices at varying image sizes.

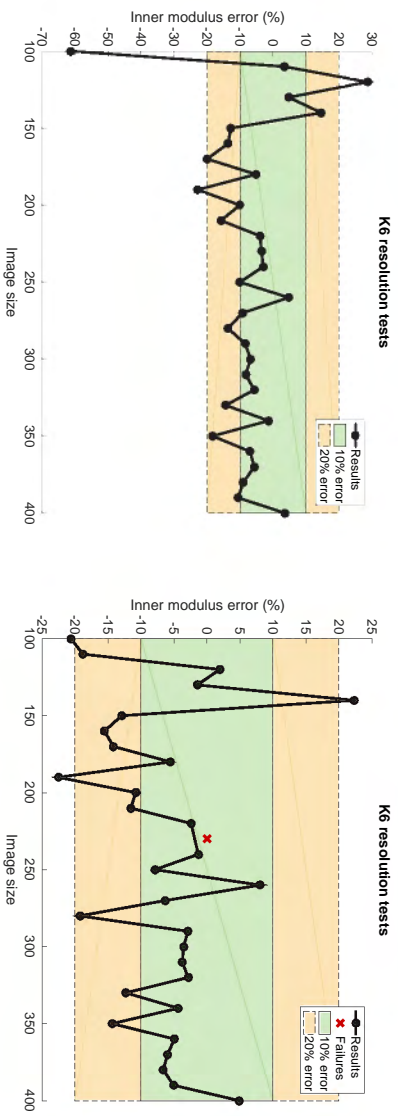
D.2.5 K6-t3 - unconstrained



(a) Uncropped.

(b) Cropped.

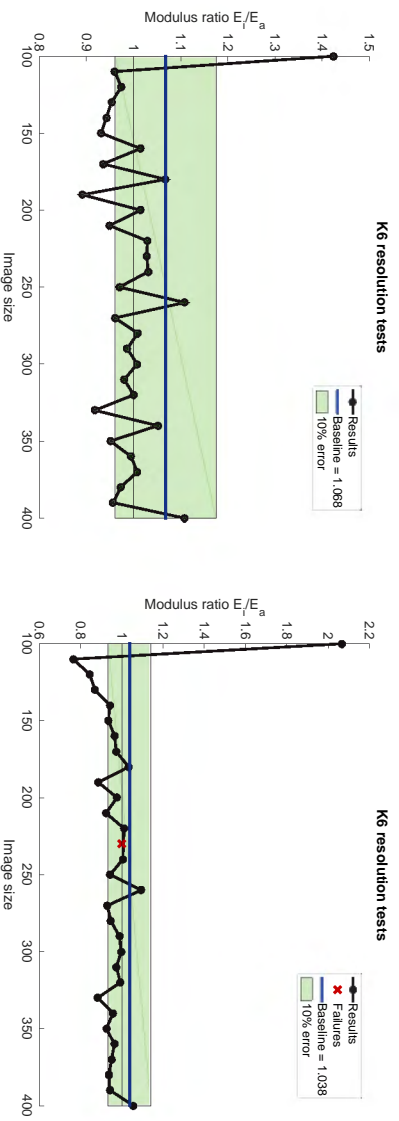
Figure D.67: Error plots for apparent modulus of unconstrained K6-t3 lattices at varying image sizes.



(a) Uncropped.

(b) Cropped.

Figure D.68: Error plots for inner modulus of unconstrained K6-t3 lattices at varying image sizes.



(a) Uncropped.

(b) Cropped.

Figure D.69: Modulus ratio results of unconstrained K6-t3 lattices at varying image sizes.

D.2.6 K6-t3 - constrained




Figure D.70: Error plots for apparent modulus of constrained K6-t3 lattices at varying image sizes.

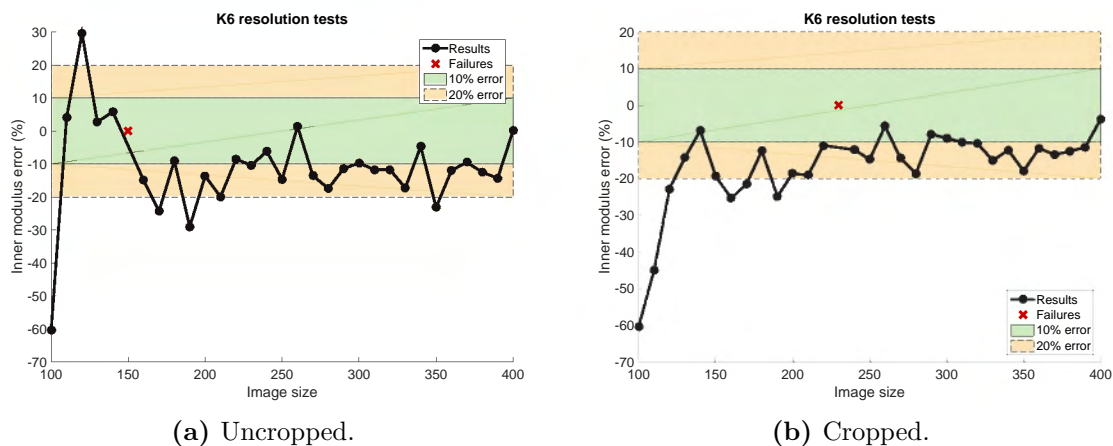


Figure D.71: Error plots for inner modulus of constrained K6-t3 lattices at varying image sizes.




Figure D.72: Modulus ratio results of constrained K6-t3 lattices at varying image sizes.

D.2.7 K8-t1 - unconstrained




Figure D.73: Error plots for apparent modulus of unconstrained K8-t1 lattices at varying image sizes.




Figure D.74: Error plots for inner modulus of unconstrained K8-t1 lattices at varying image sizes.

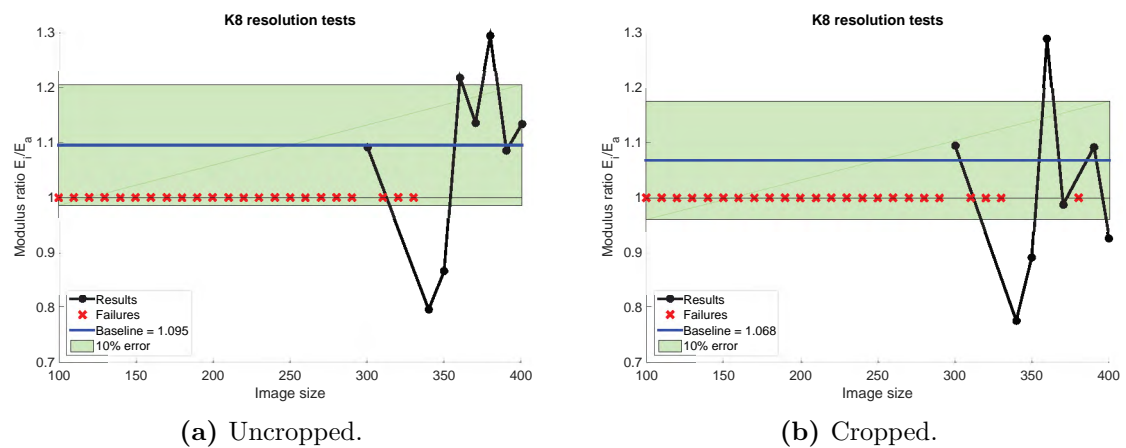


Figure D.75: Modulus ratio results of unconstrained K8-t1 lattices at varying image sizes.

D.2.8 K8-t1 - constrained




Figure D.76: Error plots for apparent modulus of constrained K8-t1 lattices at varying image sizes.

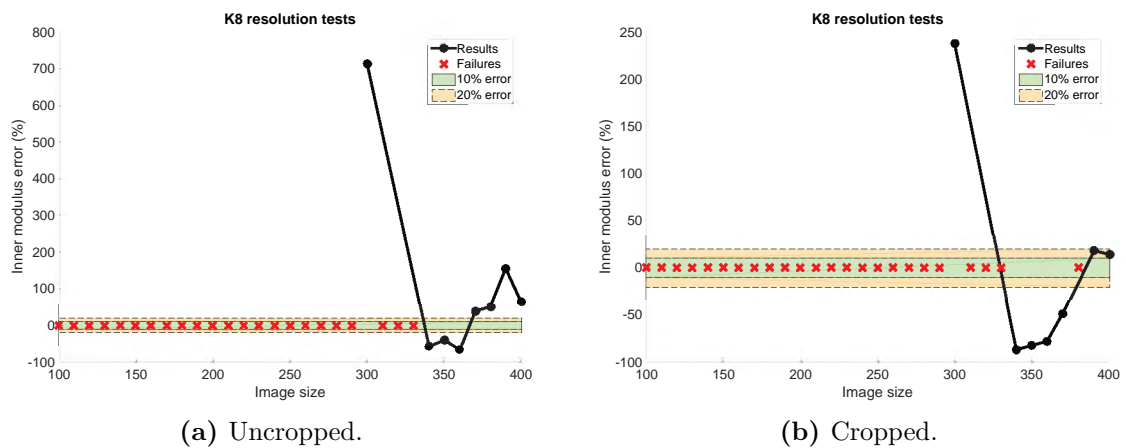


Figure D.77: Error plots for inner modulus of constrained K8-t1 lattices at varying image sizes.

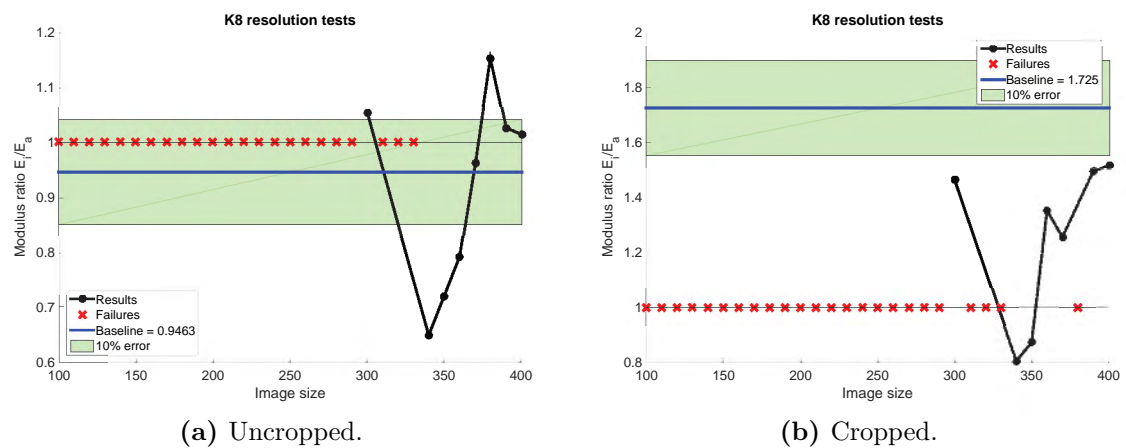


Figure D.78: Modulus ratio results of constrained K8-t1 lattices at varying image sizes.

D.2.9 K8-t2 - unconstrained




Figure D.79: Error plots for apparent modulus of unconstrained K8-t2 lattices at varying image sizes.




Figure D.80: Error plots for inner modulus of unconstrained K8-t2 lattices at varying image sizes.




Figure D.81: Modulus ratio results of unconstrained K8-t1 lattices at varying image sizes.

D.2.10 K8-t2 - constrained




Figure D.82: Error plots for apparent modulus of constrained K8-t2 lattices at varying image sizes.




Figure D.83: Error plots for inner modulus of constrained K8-t2 lattices at varying image sizes.




Figure D.84: Modulus ratio results of constrained K8-t2 lattices at varying image sizes.

D.2.11 K8-t3 - unconstrained

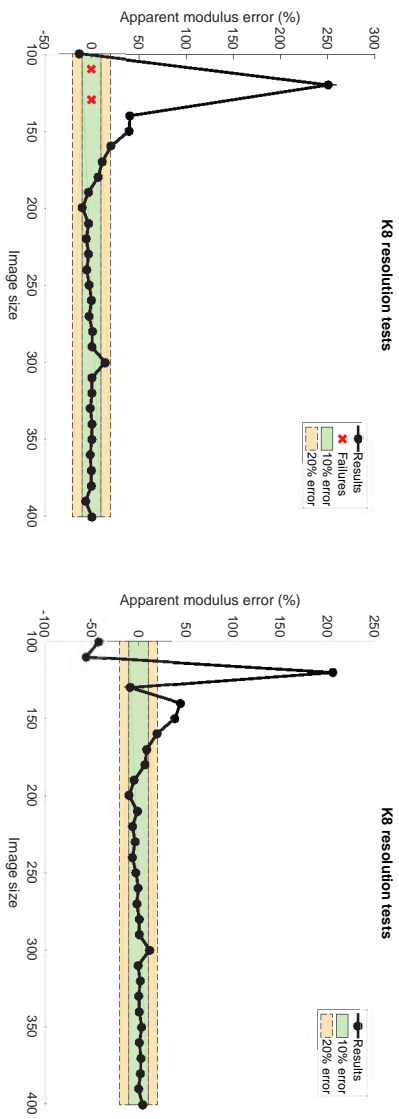


Figure D.85: Error plots for apparent modulus of unconstrained K8-t3 lattices at varying image sizes.

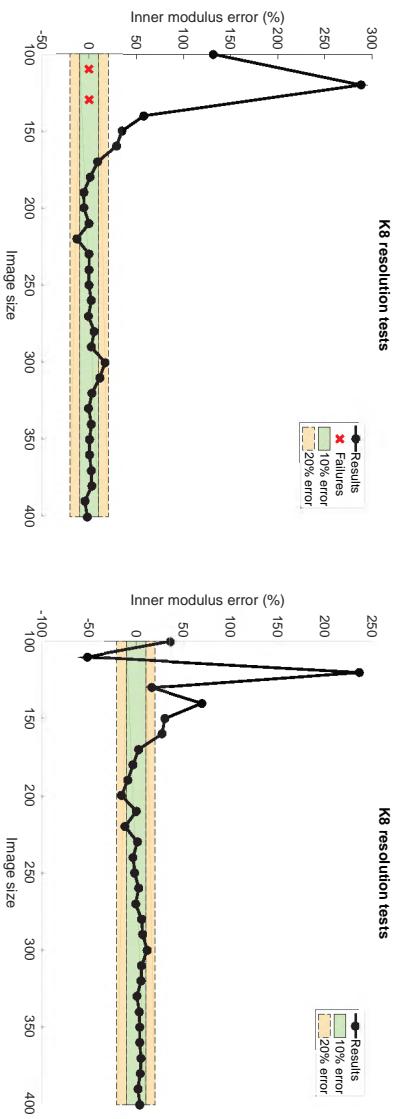


Figure D.86: Error plots for inner modulus of unconstrained K8-t3 lattices at varying image sizes.

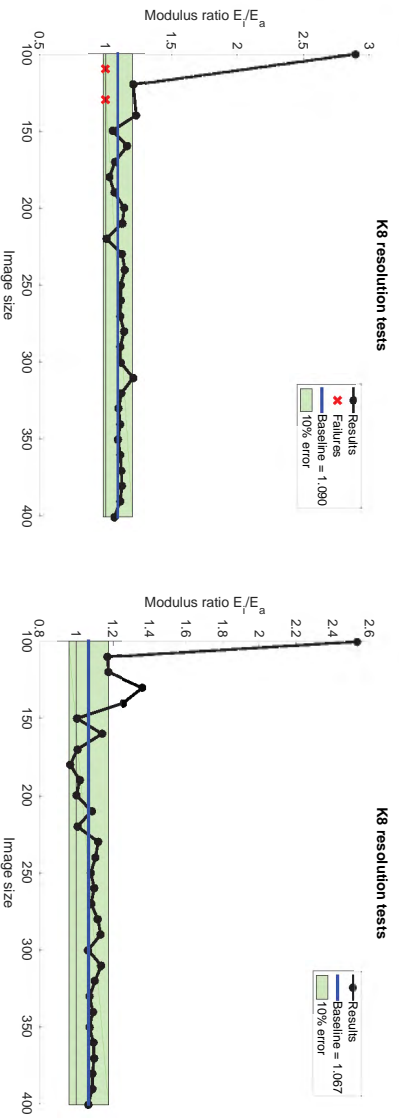


Figure D.87: Modulus ratio results of unconstrained K8-t3 lattices at varying image sizes.

D.2.12 K8-t3 - constrained

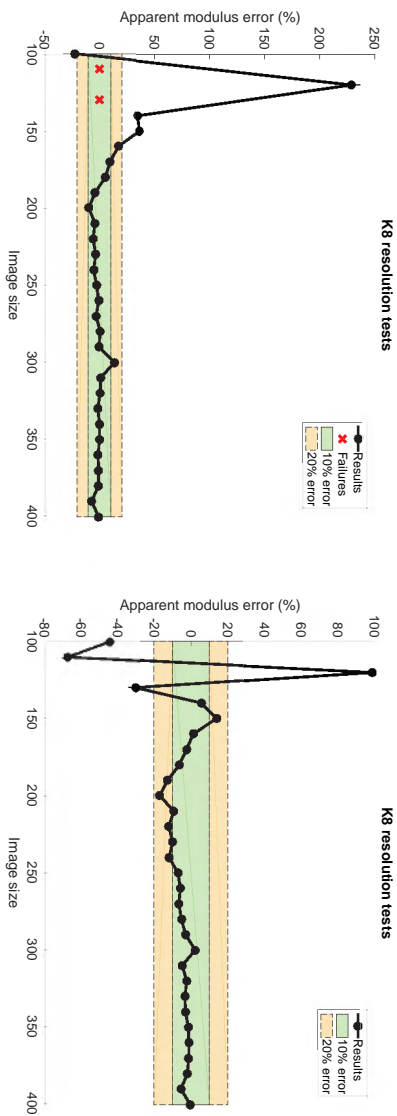


Figure D.88: Error plots for apparent modulus of constrained K8-t3 lattices at varying image sizes.

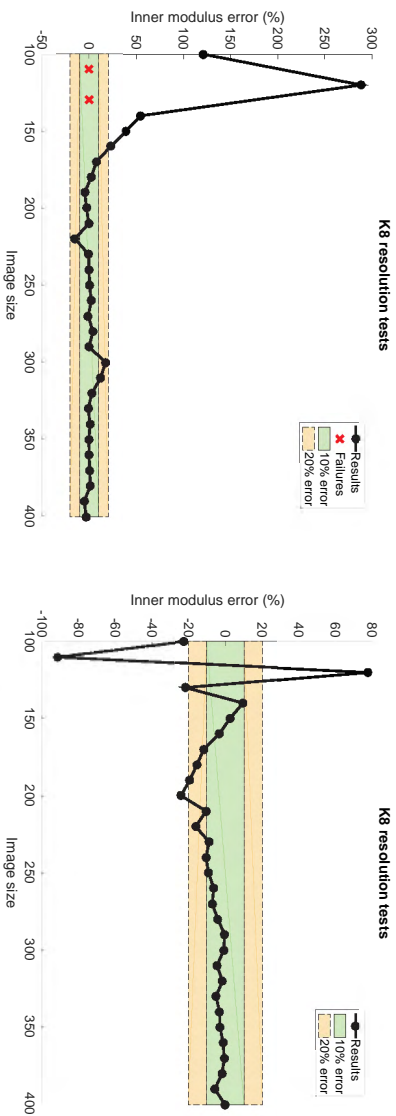


Figure D.89: Error plots for inner modulus of constrained K8-t3 lattices at varying image sizes.

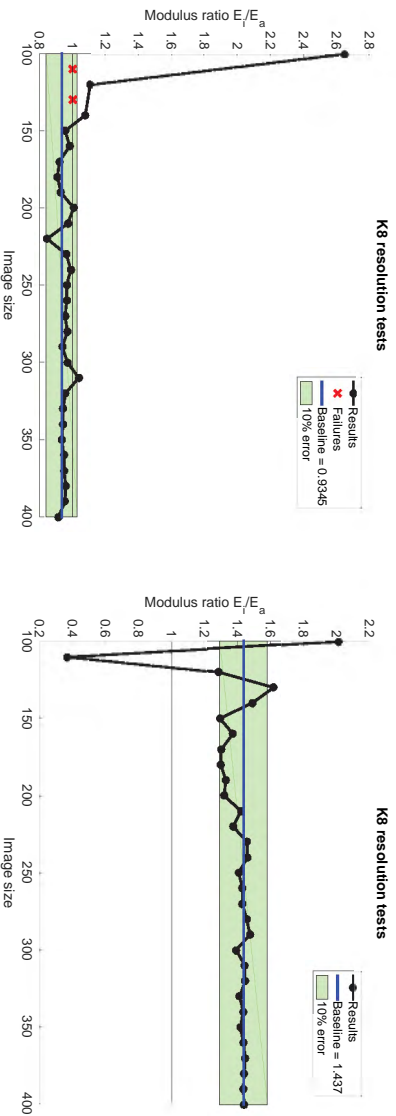


Figure D.90: Modulus ratio results of constrained K8-t3 lattices at varying image sizes.

D.2.13 K10-t1 - unconstrained

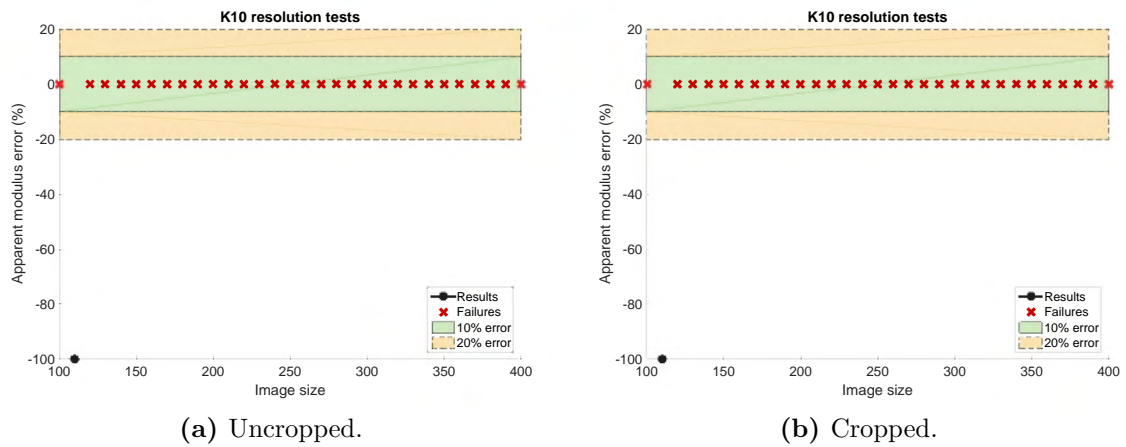


Figure D.91: Error plots for apparent modulus of unconstrained K10-t1 lattices at varying image sizes.

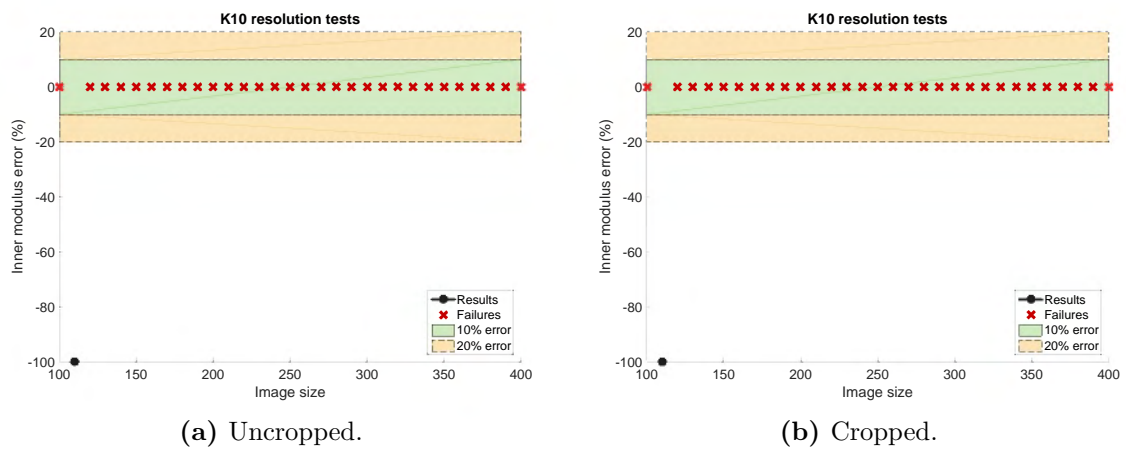


Figure D.92: Error plots for inner modulus of unconstrained K10-t1 lattices at varying image sizes.

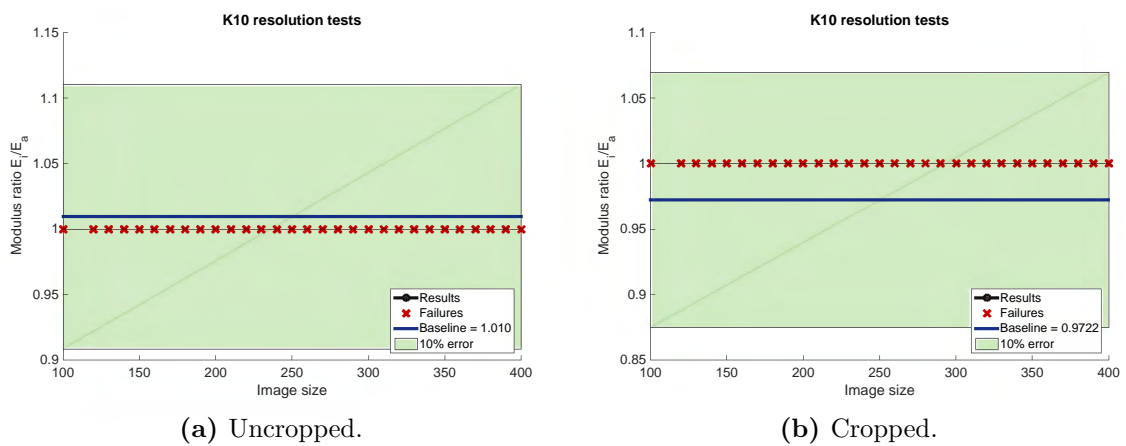




Figure D.93: Modulus ratio results of unconstrained K10-t1 lattices at varying image sizes.

D.2.14 K10-t1 - constrained



(a) Uncropped.




(b) Cropped.

Figure D.94: Error plots for apparent modulus of constrained K10-t1 lattices at varying image sizes.



(a) Uncropped.




(b) Cropped.

Figure D.95: Error plots for inner modulus of constrained K10-t1 lattices at varying image sizes.



(a) Uncropped.



(b) Cropped.

Figure D.96: Modulus ratio results of constrained K10-t1 lattices at varying image sizes.

D.2.15 K10-t2 - unconstrained

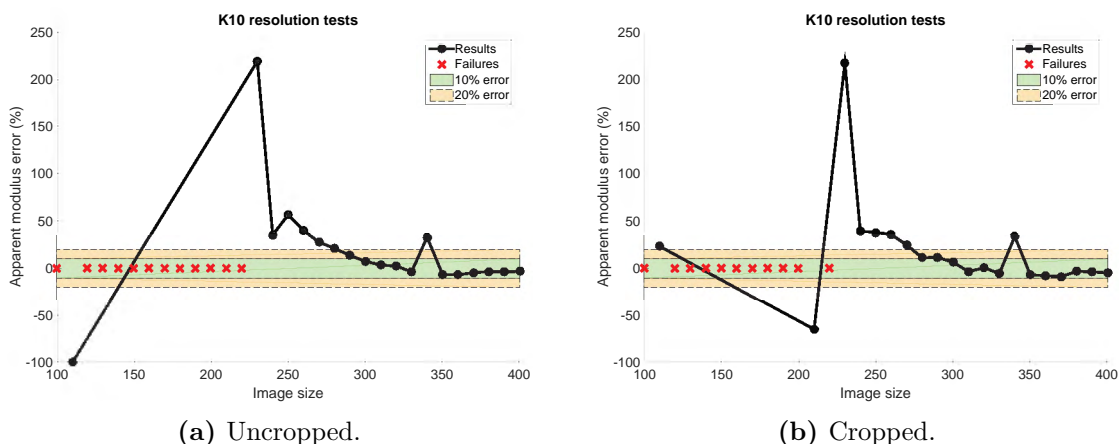


Figure D.97: Error plots for apparent modulus of unconstrained K10-t2 lattices at varying image sizes.

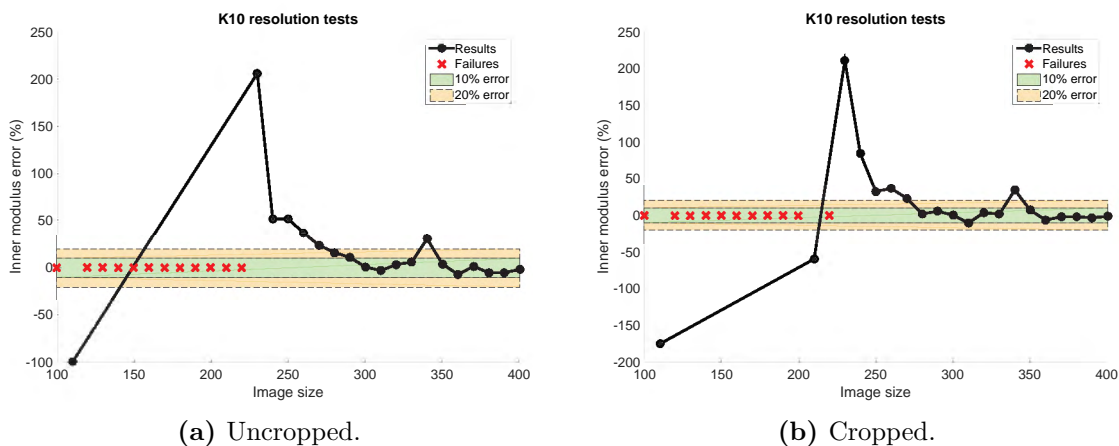


Figure D.98: Error plots for inner modulus of unconstrained K10-t2 lattices at varying image sizes.

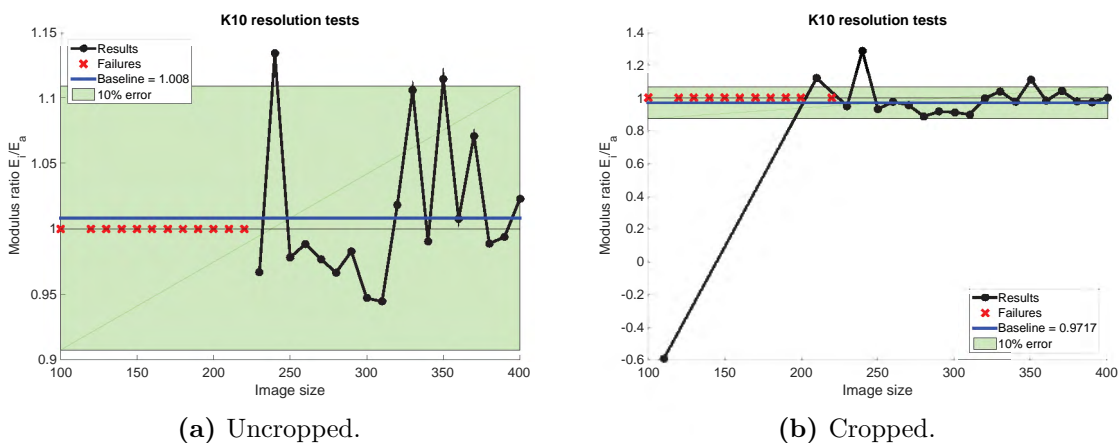


Figure D.99: Modulus ratio results of unconstrained K10-t1 lattices at varying image sizes.

D.2.16 K10-t2 - constrained

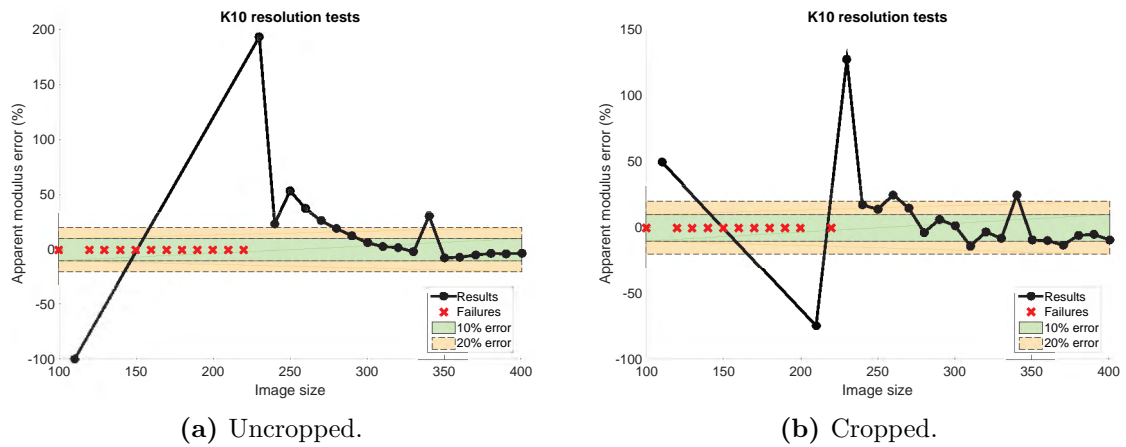


Figure D.100: Error plots for apparent modulus of constrained K10-t2 lattices at varying image sizes.




Figure D.101: Error plots for inner modulus of constrained K10-t2 lattices at varying image sizes.

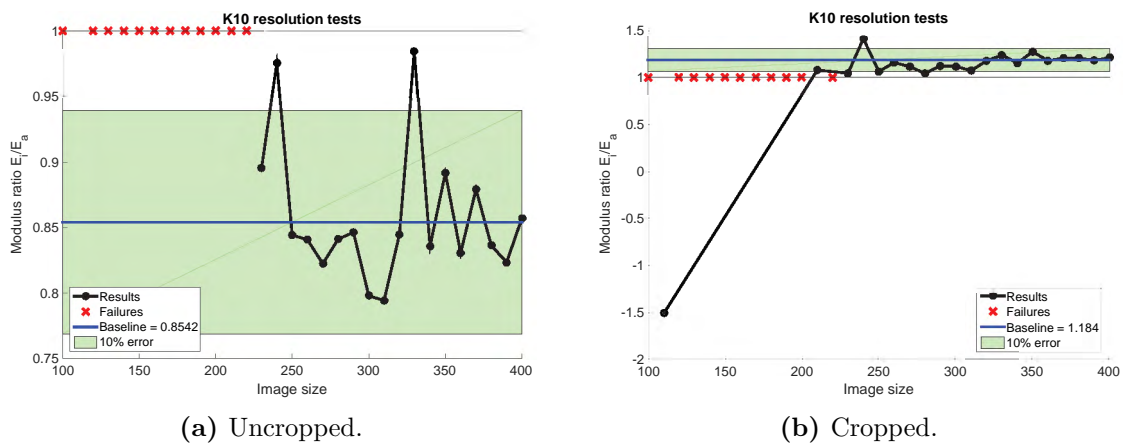




Figure D.102: Modulus ratio results of constrained K10-t2 lattices at varying image sizes.

D.2.17 K10-t3 - unconstrained




(a) Uncropped.




(b) Cropped.

Figure D.103: Error plots for apparent modulus of unconstrained K10-t3 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.104: Error plots for inner modulus of unconstrained K10-t3 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.105: Modulus ratio results of unconstrained K10-t3 lattices at varying image sizes.

D.2.18 K10-t3 - constrained




(a) Uncropped.




(b) Cropped.

Figure D.106: Error plots for apparent modulus of constrained K10-t3 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.107: Error plots for inner modulus of constrained K10-t3 lattices at varying image sizes.



(a) Uncropped.



(b) Cropped.

Figure D.108: Modulus ratio results of constrained K10-t3 lattices at varying image sizes.

D.3 Octet truss lattices

D.3.1 OT6-t1 - unconstrained

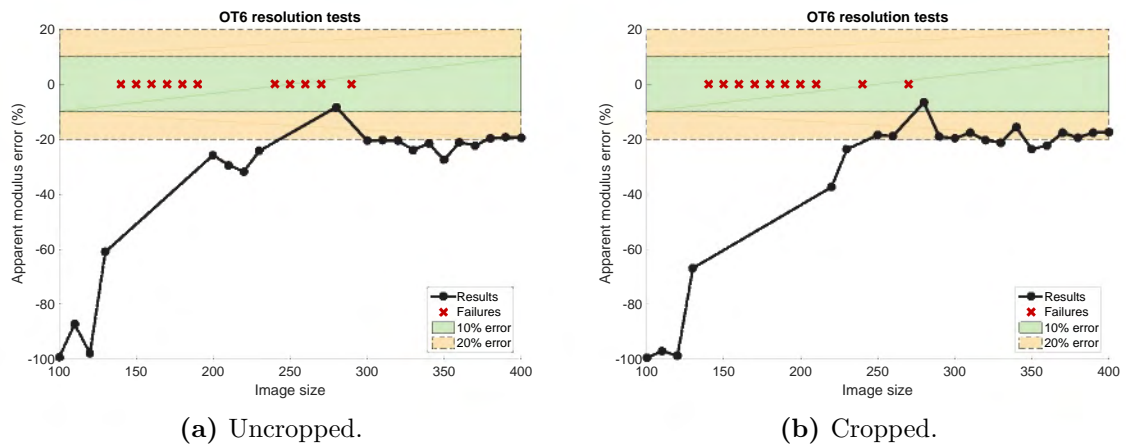


Figure D.109: Error plots for apparent modulus of unconstrained OT6-t1 lattices at varying image sizes.

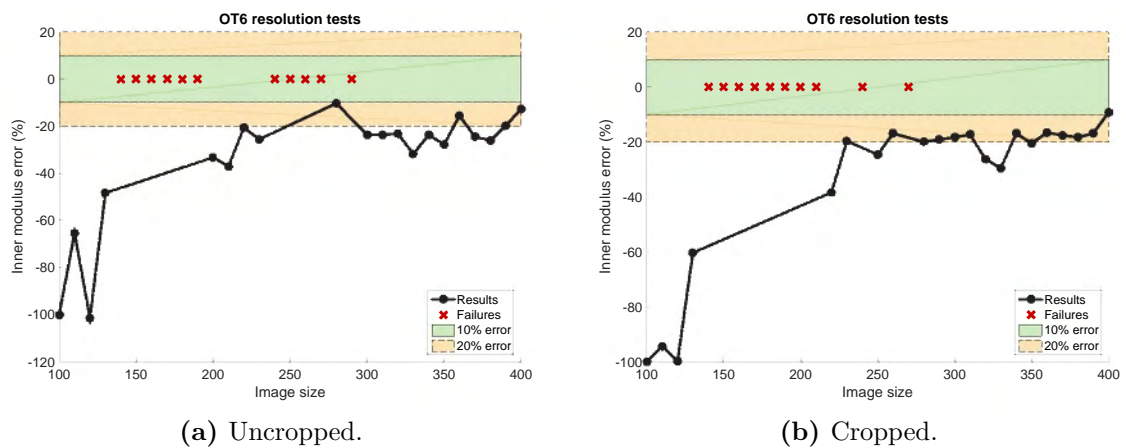


Figure D.110: Error plots for inner modulus of unconstrained OT6-t1 lattices at varying image sizes.

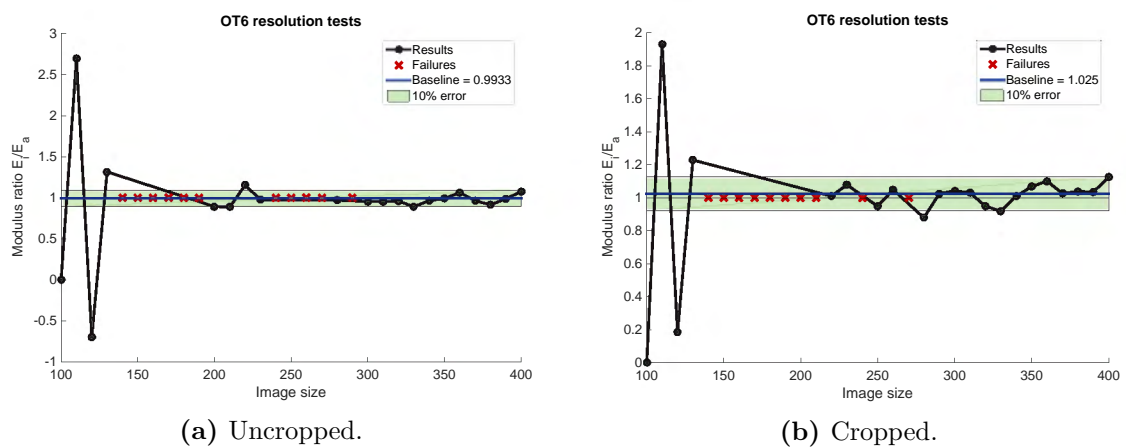


Figure D.111: Modulus ratio results of unconstrained OT6-t1 lattices at varying image sizes.

D.3.2 OT6-t1 - constrained

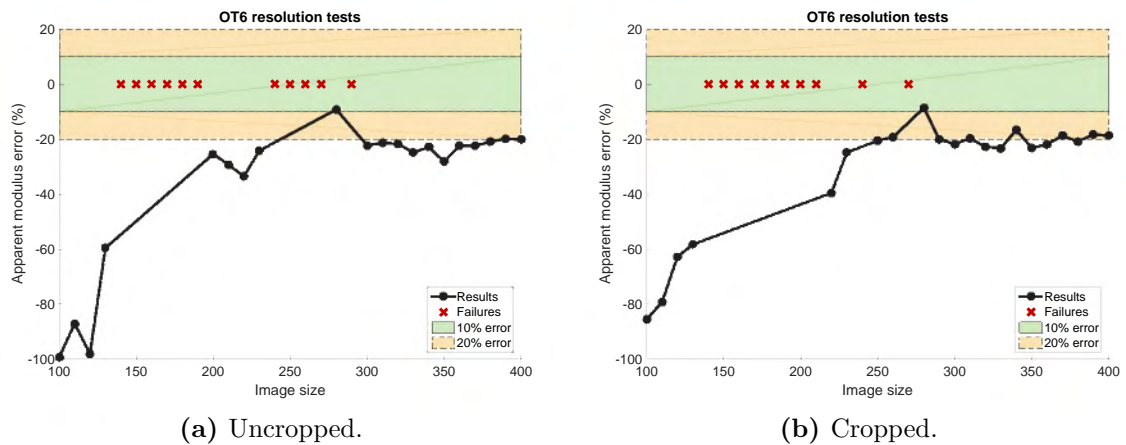


Figure D.112: Error plots for apparent modulus of constrained OT6-t1 lattices at varying image sizes.

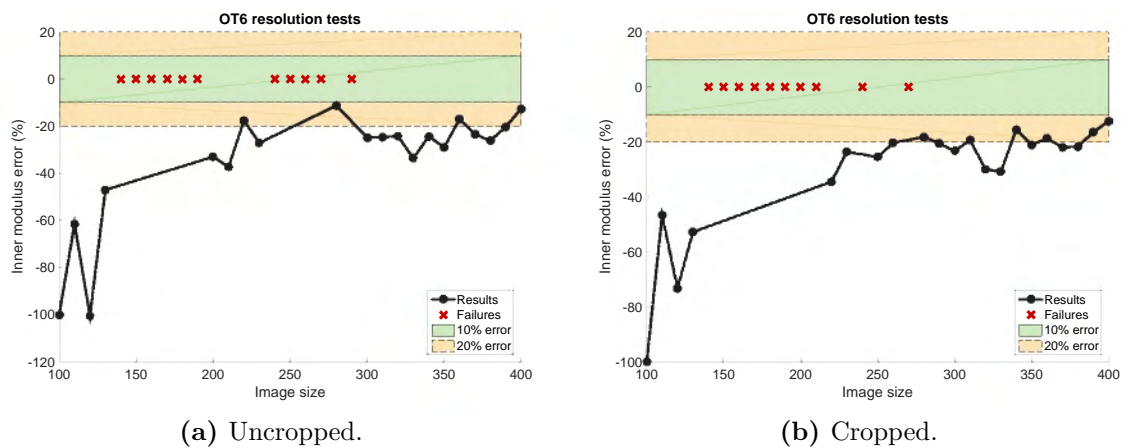


Figure D.113: Error plots for inner modulus of constrained OT6-t1 lattices at varying image sizes.

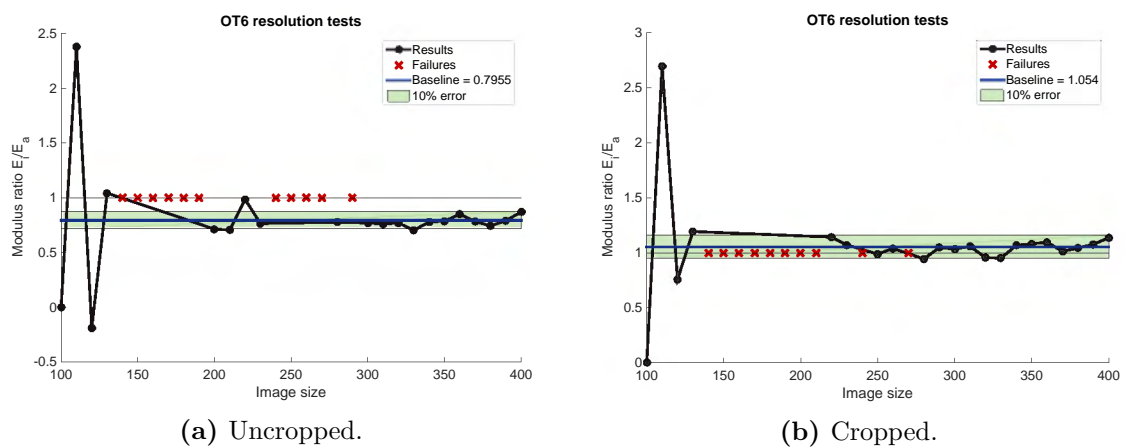




Figure D.114: Modulus ratio results of constrained OT6-t1 lattices at varying image sizes.

D.3.3 OT6-t2 - unconstrained




(a) Uncropped.




(b) Cropped.

Figure D.115: Error plots for apparent modulus of unconstrained OT6-t2 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.116: Error plots for inner modulus of unconstrained OT6-t2 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.117: Modulus ratio results of unconstrained OT6-t2 lattices at varying image sizes.

D.3.4 OT6-t2 - constrained




(a) Uncropped.




(b) Cropped.

Figure D.118: Error plots for apparent modulus of constrained OT6-t2 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.119: Error plots for inner modulus of constrained OT6-t2 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.120: Modulus ratio results of constrained OT6-t2 lattices at varying image sizes.

D.3.5 OT6-t3 - unconstrained




(a) Uncropped.




(b) Cropped.

Figure D.121: Error plots for apparent modulus of unconstrained OT6-t3 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.122: Error plots for inner modulus of unconstrained OT6-t3 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.123: Modulus ratio results of unconstrained OT6-t3 lattices at varying image sizes.

D.3.6 OT6-t3 - constrained




(a) Uncropped.




(b) Cropped.

Figure D.124: Error plots for apparent modulus of constrained OT6-t3 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.125: Error plots for inner modulus of constrained OT6-t3 lattices at varying image sizes.



(a) Uncropped.



(b) Cropped.

Figure D.126: Modulus ratio results of constrained OT6-t3 lattices at varying image sizes.

D.3.7 OT8-t1 - unconstrained




Figure D.127: Error plots for apparent modulus of unconstrained OT8-t1 lattices at varying image sizes.




Figure D.128: Error plots for inner modulus of unconstrained OT8-t1 lattices at varying image sizes.

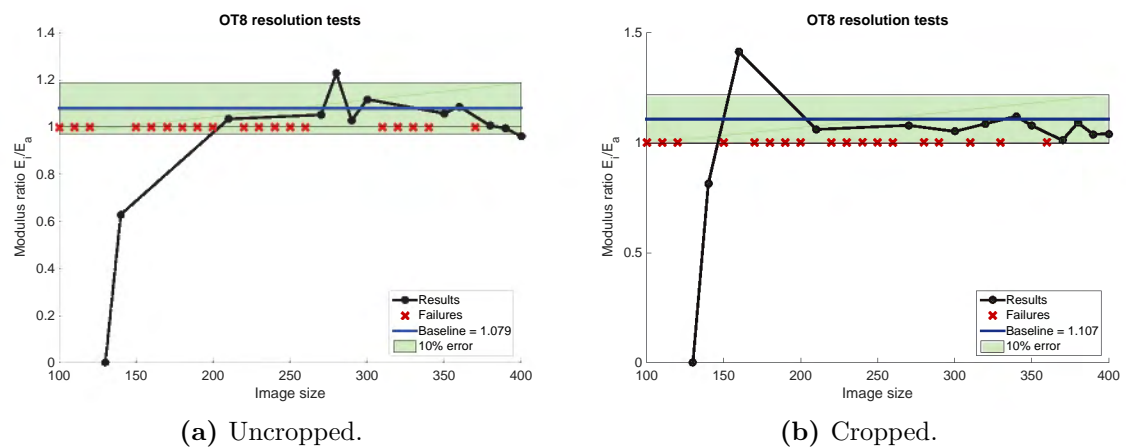


Figure D.129: Modulus ratio results of unconstrained OT8-t1 lattices at varying image sizes.

D.3.8 OT8-t1 - constrained

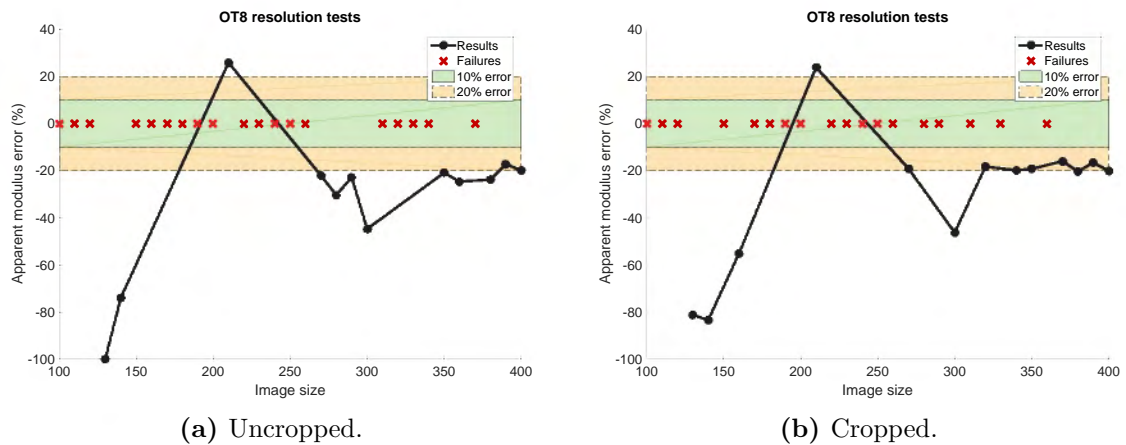


Figure D.130: Error plots for apparent modulus of constrained OT8-t1 lattices at varying image sizes.

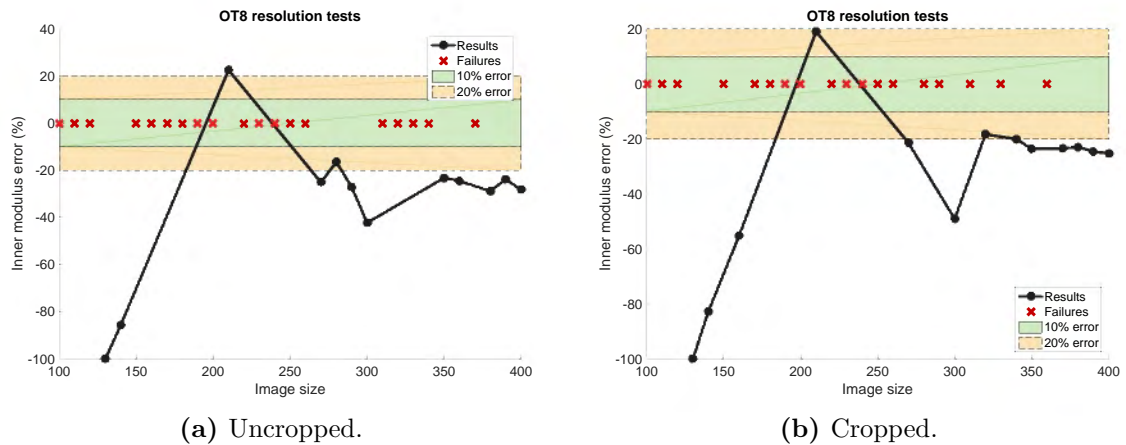


Figure D.131: Error plots for inner modulus of constrained OT8-t1 lattices at varying image sizes.

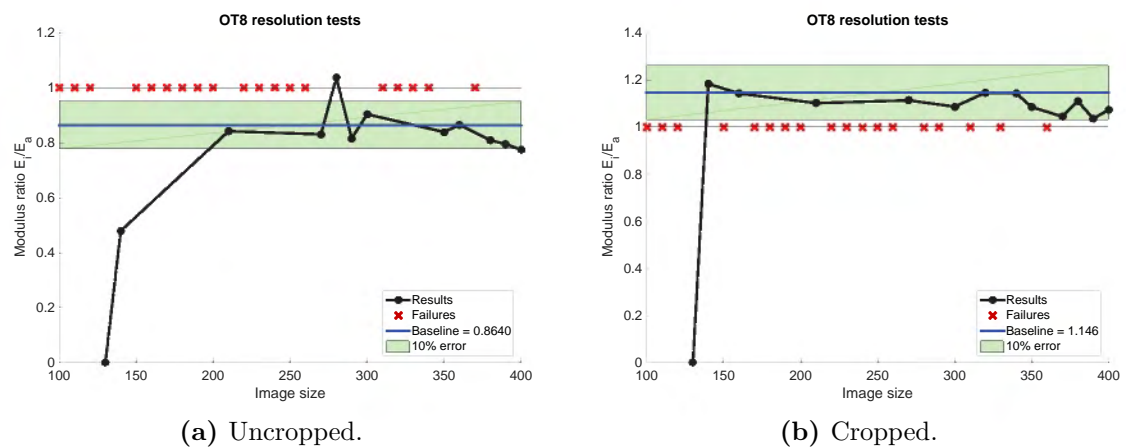


Figure D.132: Modulus ratio results of constrained OT8-t1 lattices at varying image sizes.

D.3.9 OT8-t2 - unconstrained

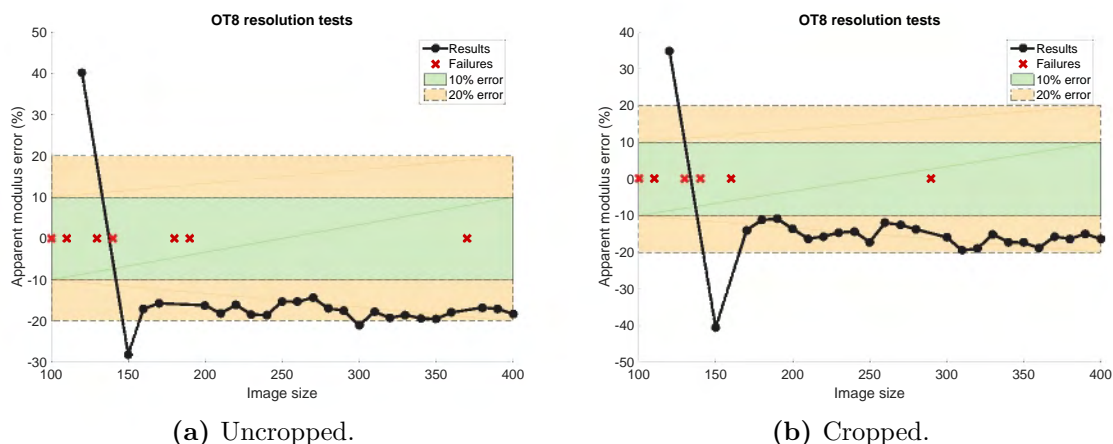


Figure D.133: Error plots for apparent modulus of unconstrained OT8-t2 lattices at varying image sizes.

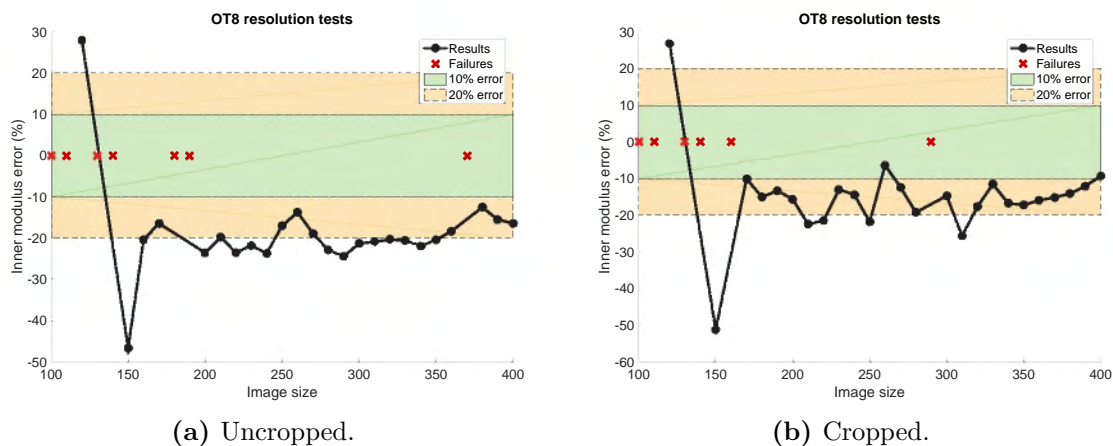


Figure D.134: Error plots for inner modulus of unconstrained OT8-t2 lattices at varying image sizes.






Figure D.135: Modulus ratio results of unconstrained OT8-t2 lattices at varying image sizes.

D.3.10 OT8-t2 - constrained




(a) Uncropped.




(b) Cropped.

Figure D.136: Error plots for apparent modulus of constrained OT8-t2 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.137: Error plots for inner modulus of constrained OT8-t2 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.138: Modulus ratio results of constrained OT8-t2 lattices at varying image sizes.

D.3.11 OT8-t3 - unconstrained




(a) Uncropped.




(b) Cropped.

Figure D.139: Error plots for apparent modulus of unconstrained OT8-t3 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.140: Error plots for inner modulus of unconstrained OT8-t3 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.141: Modulus ratio results of unconstrained OT8-t3 lattices at varying image sizes.

D.3.12 OT8-t3 - constrained




(a) Uncropped.




(b) Cropped.

Figure D.142: Error plots for apparent modulus of constrained OT8-t3 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.143: Error plots for inner modulus of constrained OT8-t3 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.144: Modulus ratio results of constrained OT8-t3 lattices at varying image sizes.

D.3.13 OT10-t1 - unconstrained




(a) Uncropped.




(b) Cropped.

Figure D.145: Error plots for apparent modulus of unconstrained OT10-t1 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.146: Error plots for inner modulus of unconstrained OT10-t1 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.147: Modulus ratio results of unconstrained OT10-t1 lattices at varying image sizes.

D.3.14 OT10-t1 - constrained




(a) Uncropped.




(b) Cropped.

Figure D.148: Error plots for apparent modulus of constrained OT10-t1 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.149: Error plots for inner modulus of constrained OT10-t1 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.150: Modulus ratio results of constrained OT10-t1 lattices at varying image sizes.

D.3.15 OT10-t2 - unconstrained




(a) Uncropped.




(b) Cropped.

Figure D.151: Error plots for apparent modulus of unconstrained OT10-t2 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.152: Error plots for inner modulus of unconstrained OT10-t2 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.153: Modulus ratio results of unconstrained OT10-t2 lattices at varying image sizes.

D.3.16 OT10-t2 - constrained




(a) Uncropped.



(b) Cropped.

Figure D.154: Error plots for apparent modulus of constrained OT10-t2 lattices at varying image sizes.




(a) Uncropped.



(b) Cropped.

Figure D.155: Error plots for inner modulus of constrained OT10-t2 lattices at varying image sizes.



(a) Uncropped.



(b) Cropped.

Figure D.156: Modulus ratio results of constrained OT10-t2 lattices at varying image sizes.

D.3.17 OT10-t3 - unconstrained




Figure D.157: Error plots for apparent modulus of unconstrained OT10-t3 lattices at varying image sizes.




Figure D.158: Error plots for inner modulus of unconstrained OT10-t3 lattices at varying image sizes.

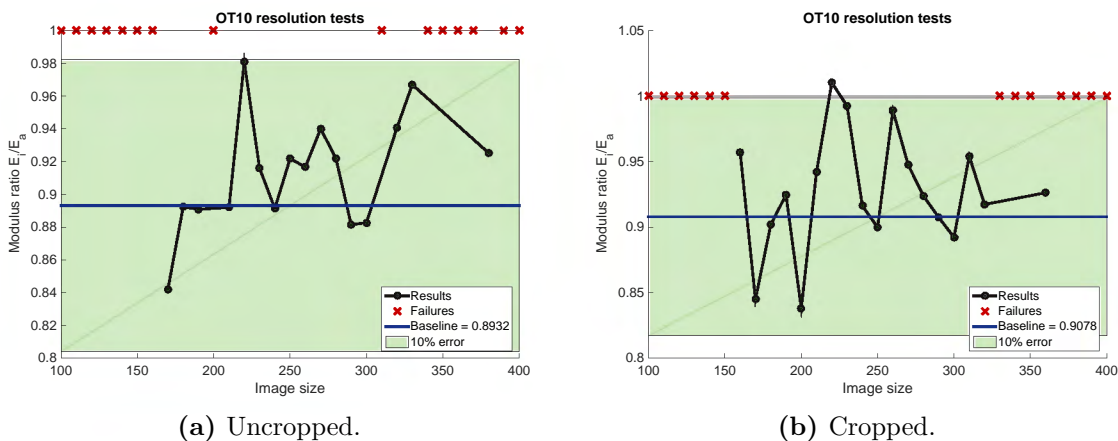




Figure D.159: Modulus ratio results of unconstrained OT10-t3 lattices at varying image sizes.

D.3.18 OT10-t3 - constrained




(a) Uncropped.




(b) Cropped.

Figure D.160: Error plots for apparent modulus of constrained OT10-t3 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.161: Error plots for inner modulus of constrained OT10-t3 lattices at varying image sizes.



(a) Uncropped.



(b) Cropped.

Figure D.162: Modulus ratio results of constrained OT10-t3 lattices at varying image sizes.

D.4 Additional analysis: angled cubic

D.4.1 C8-t1 - angled unconstrained

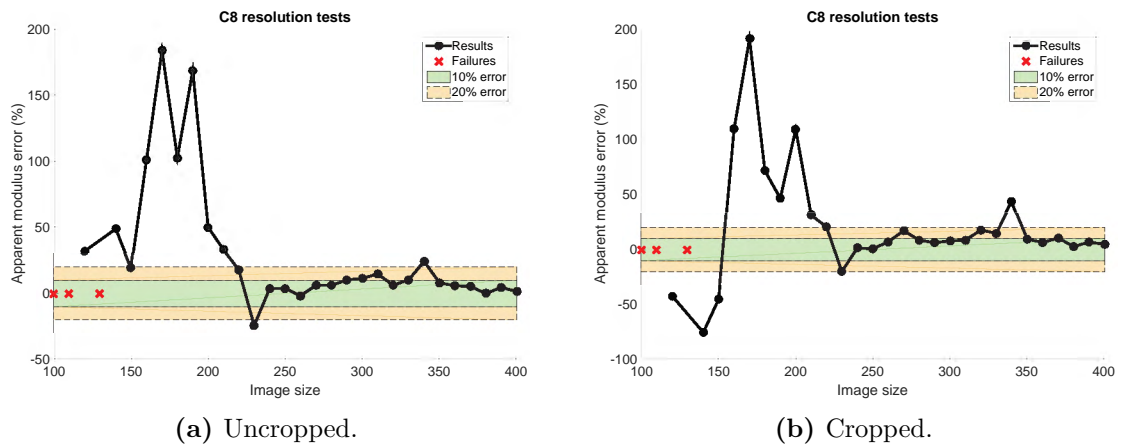


Figure D.163: Error plots for apparent modulus of unconstrained angled C8-t1 lattices at varying image sizes.




Figure D.164: Error plots for inner modulus of unconstrained angled C8-t1 lattices at varying image sizes.






Figure D.165: Modulus ratio results of unconstrained angled C8-t1 lattices at varying image sizes.

D.4.2 C8-t1 - angled constrained




(a) Uncropped.




(b) Cropped.

Figure D.166: Error plots for apparent modulus of constrained angled C8-t1 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.167: Error plots for inner modulus of constrained angled C8-t1 lattices at varying image sizes.



(a) Uncropped.



(b) Cropped.

Figure D.168: Modulus ratio results of constrained angled C8-t1 lattices at varying image sizes.

D.4.3 C8-t2 - angled unconstrained




Figure D.169: Error plots for apparent modulus of unconstrained C8-t2 lattices at varying image sizes.




Figure D.170: Error plots for inner modulus of unconstrained C8-t2 lattices at varying image sizes.

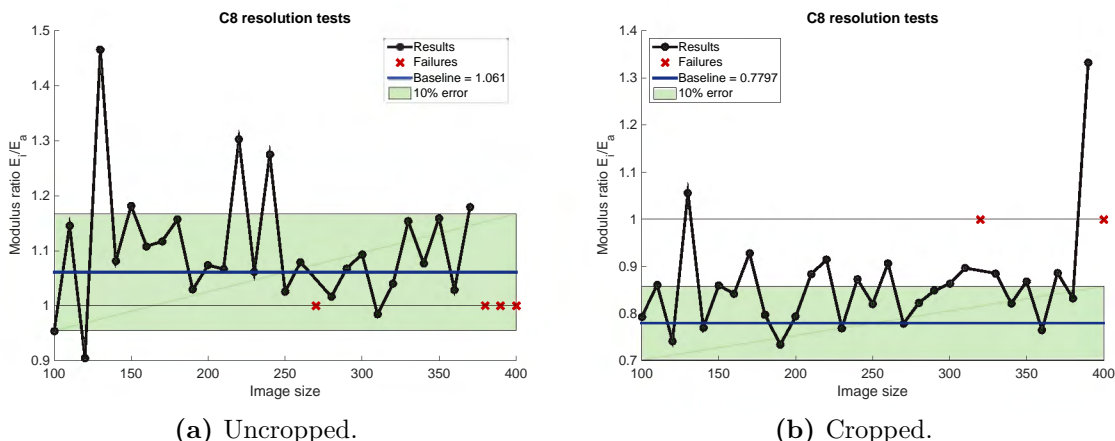




Figure D.171: Modulus ratio results of unconstrained C8-t2 lattices at varying image sizes.

D.4.4 C8-t2 - angled constrained




(a) Uncropped.




(b) Cropped.

Figure D.172: Error plots for apparent modulus of constrained C8-t2 lattices at varying image sizes.




(a) Uncropped.




(b) Cropped.

Figure D.173: Error plots for inner modulus of constrained C8-t2 lattices at varying image sizes.



(a) Uncropped.



(b) Cropped.

Figure D.174: Modulus ratio results of constrained C8-t2 lattices at varying image sizes.

D.5 Additional analysis: VBM

D.5.1 K8-t3 - cropped and cropped constrained

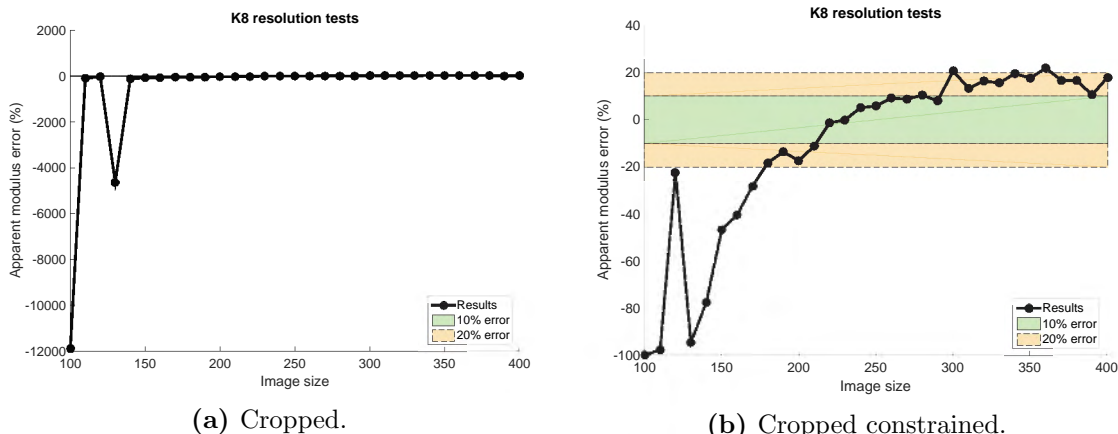


Figure D.175: Error plots for apparent modulus of cropped unconstrained and cropped constrained K8-t3 lattices at varying image sizes obtained using the VBM.

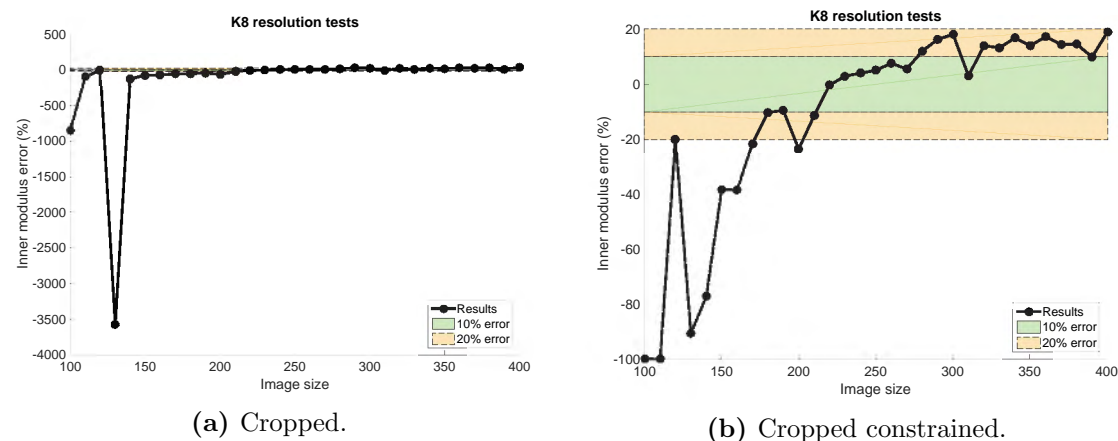


Figure D.176: Error plots for inner modulus of cropped unconstrained and cropped constrained K8-t3 lattices at varying image sizes obtained using the VBM.

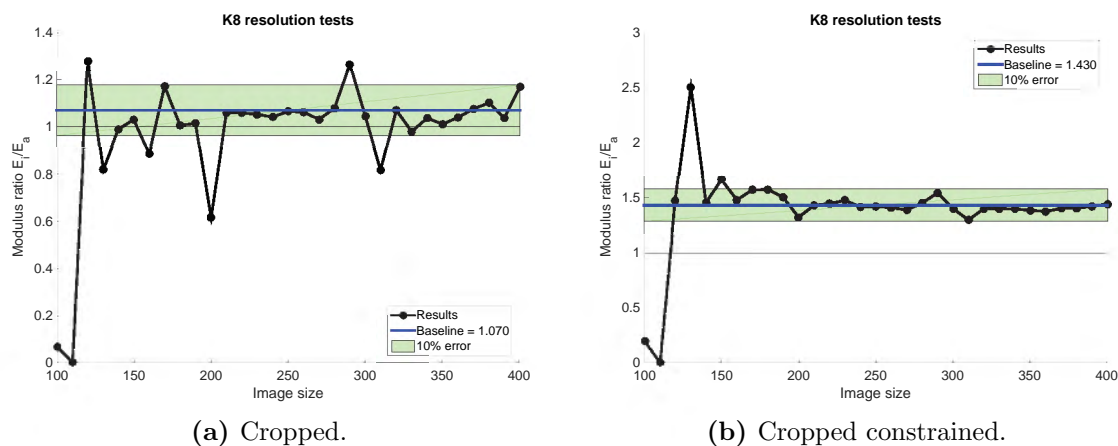


Figure D.177: Modulus ratio results of cropped unconstrained and cropped constrained K8-t3 lattices at varying image sizes obtained using the VBM.

Appendix E

Additional validation results




Figure E.1: Enlarged view of graph of apparent modulus error and average beam cross-sectional area error for an uncropped, unconstrained K6-t3 lattice at varying image sizes.




Figure E.2: Enlarged view of the error plots for the moments of inertia in the primary and secondary directions for an uncropped, unconstrained K6-t3 lattice at varying image sizes.




Figure E.3: Enlarged view of graph of apparent modulus error and average beam cross-sectional area error for an uncropped, unconstrained K10-t3 lattice at varying image sizes.




Figure E.4: Enlarged view of the error plots for the moments of inertia in the primary and secondary directions for an uncropped, unconstrained K10-t3 lattice at varying image sizes.

Appendix F

Rasterised lattices

F.1 Cubic

F.1.1 Size 6

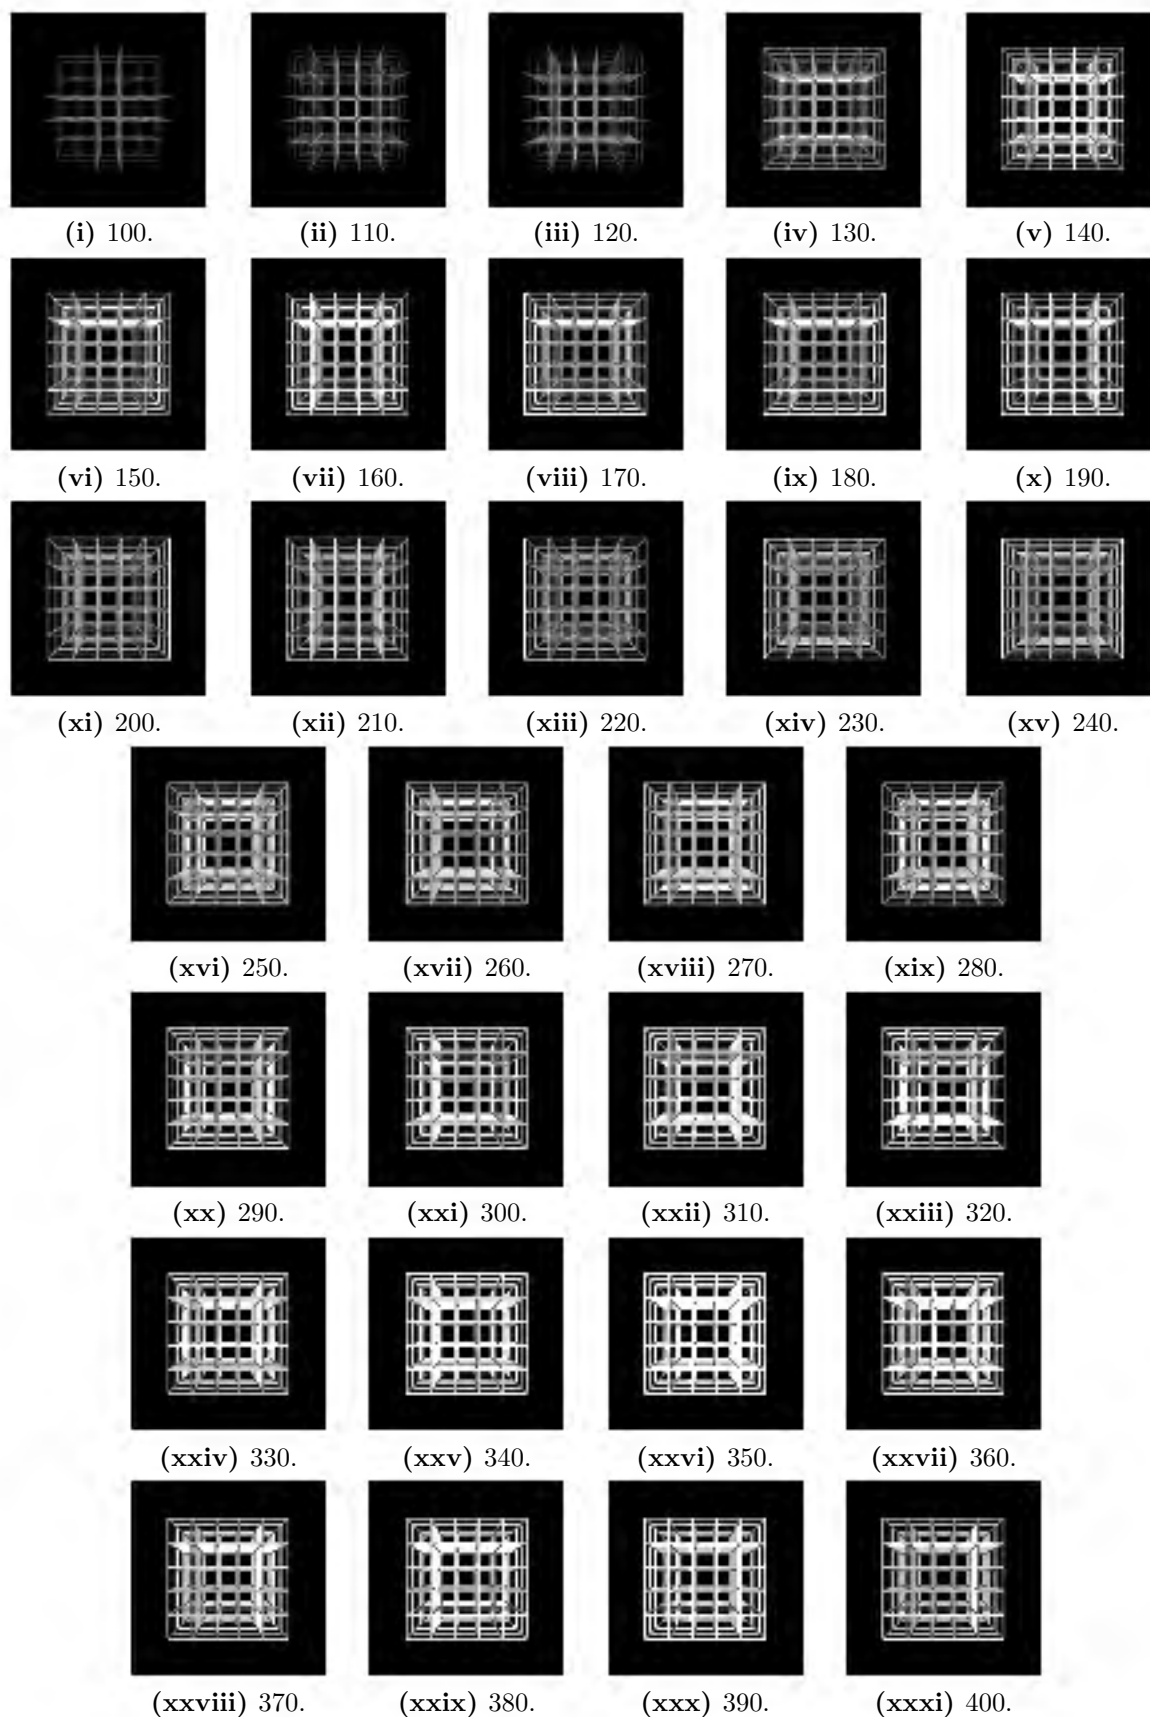


Figure F.1: Rasterisation of C6-t1 lattice at varying image sizes.

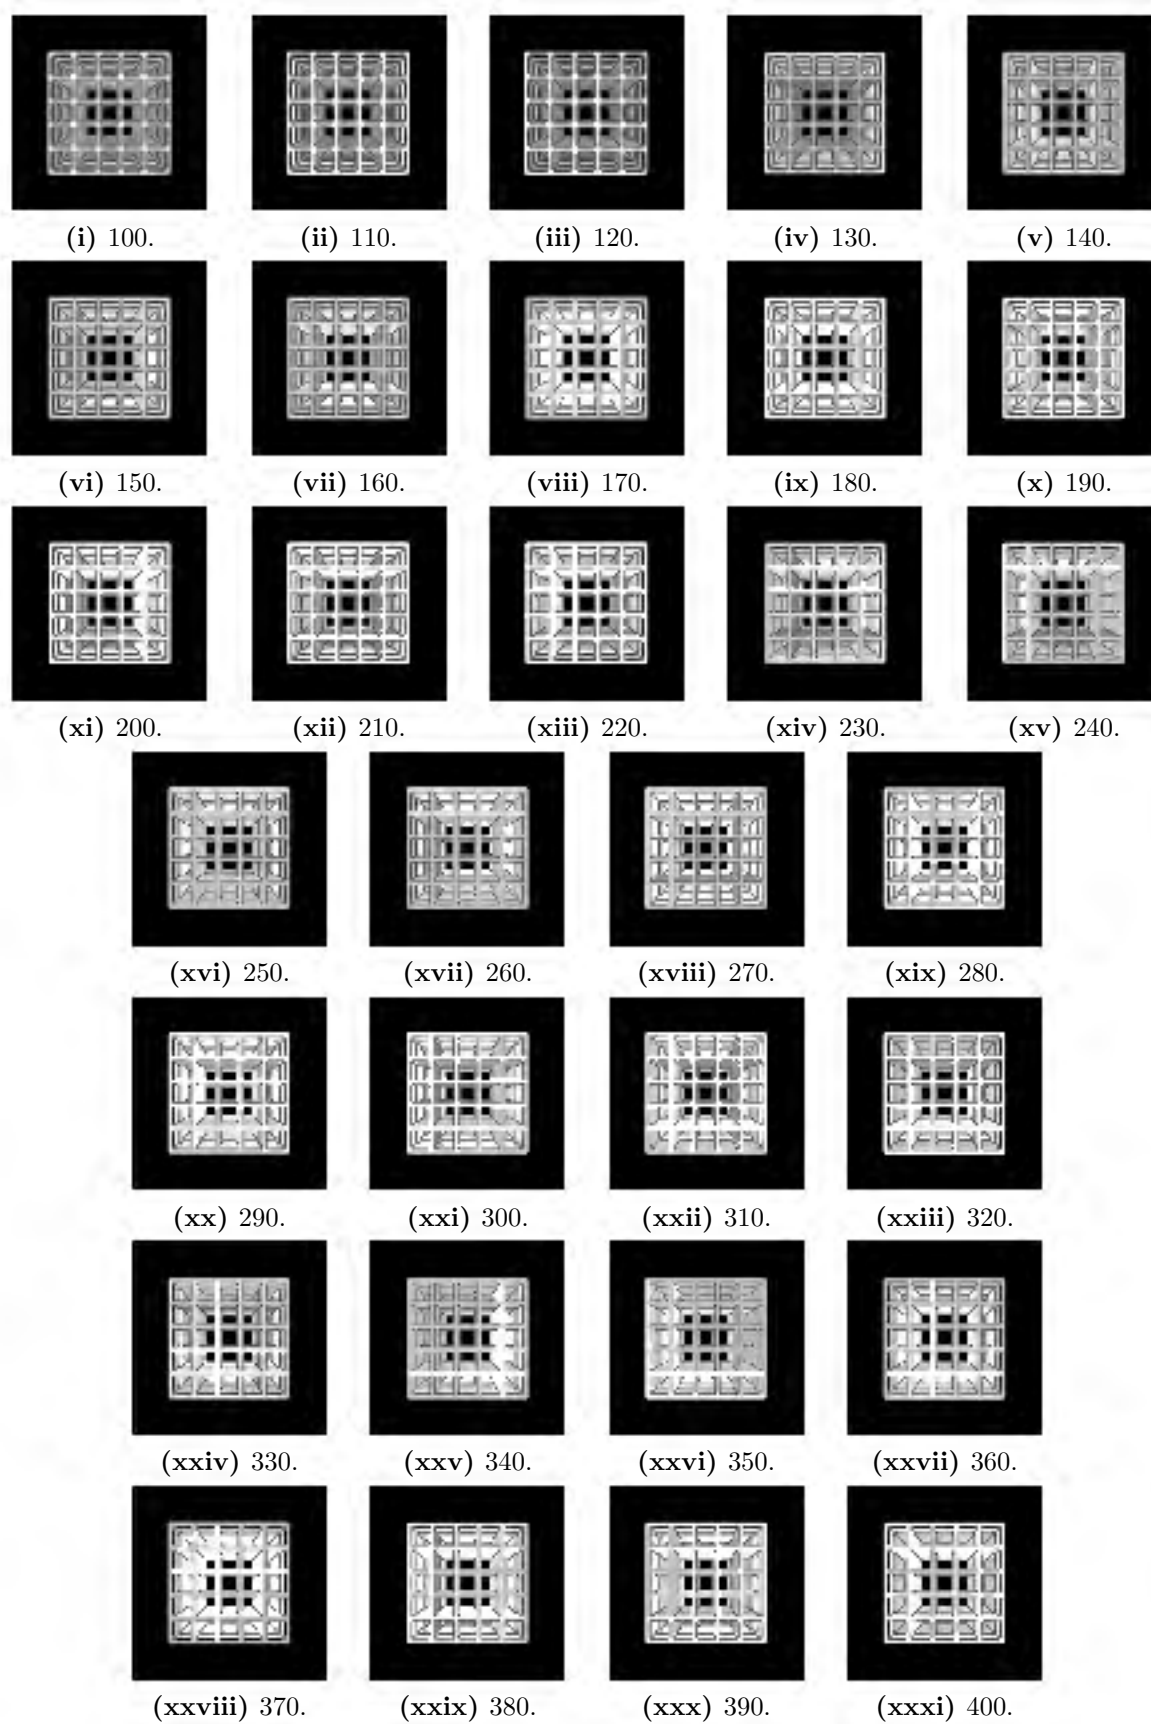


Figure F.2: Rasterisation of C6-t2 lattice at varying image sizes.

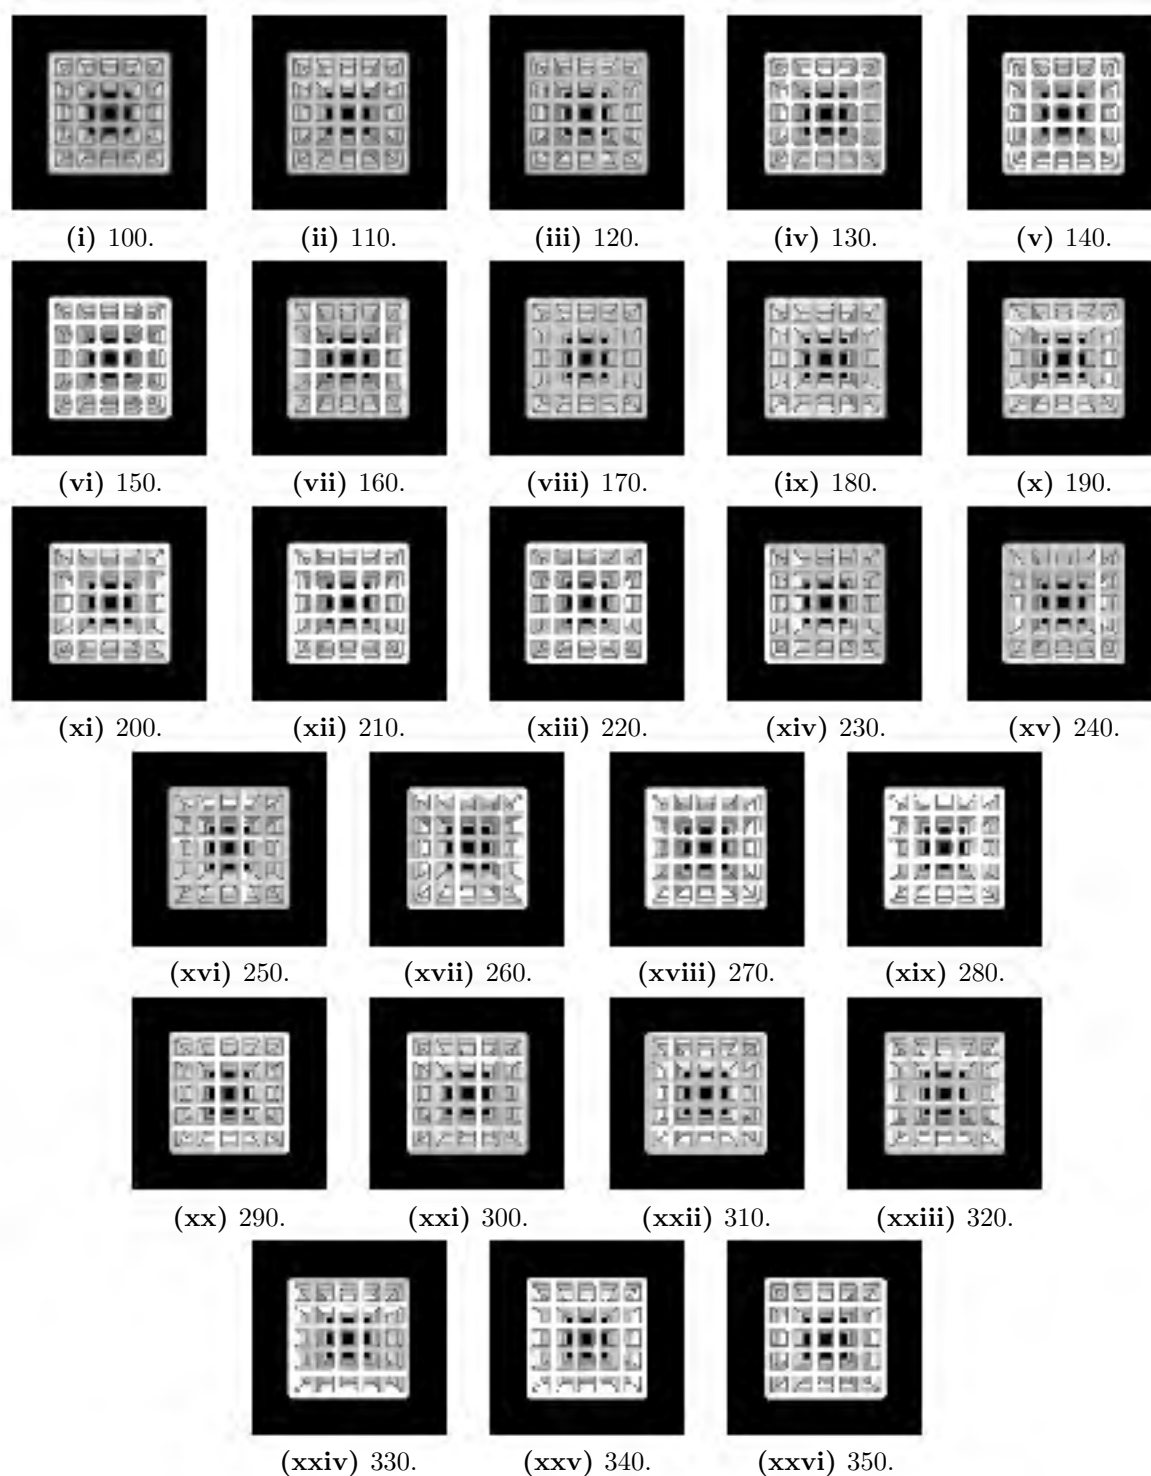


Figure F.3: Rasterisation of C6-t3 lattice at varying image sizes.

F.1.2 Size 8

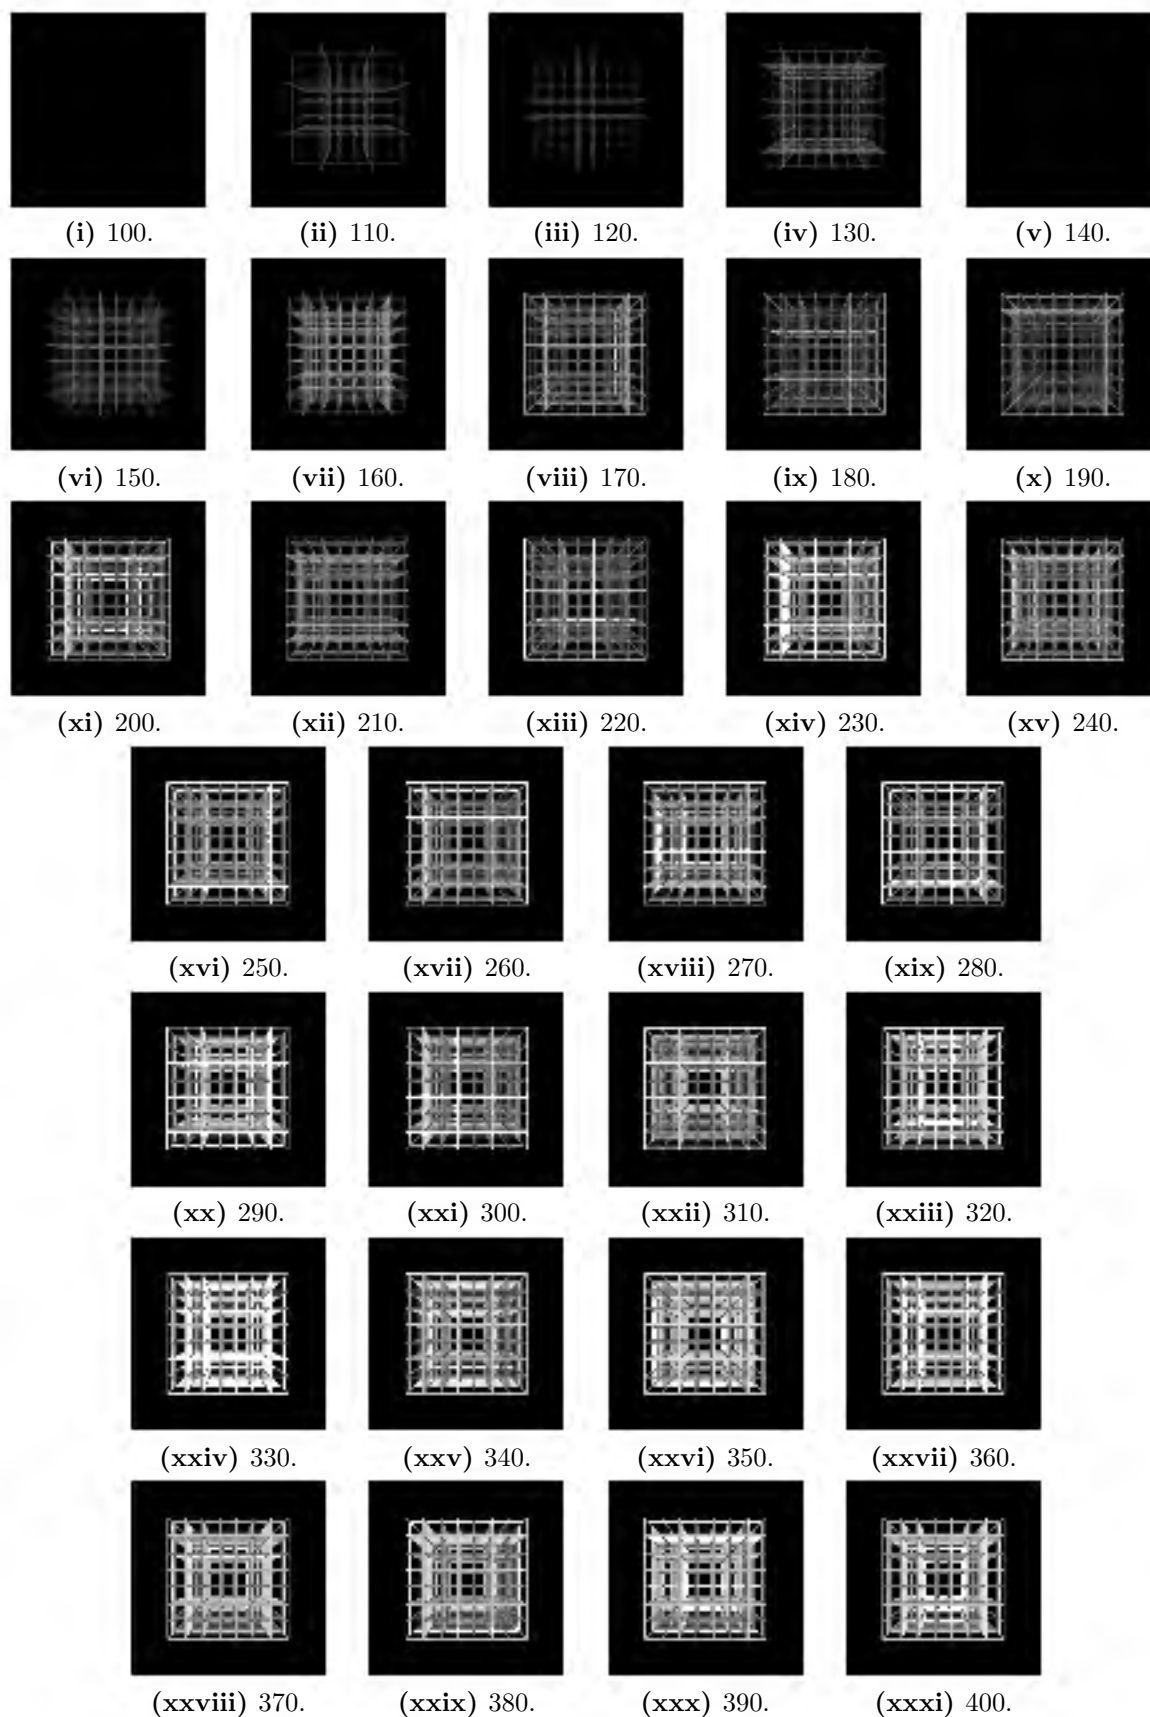


Figure F.4: Rasterisation of C8-t1 lattice at varying image sizes.

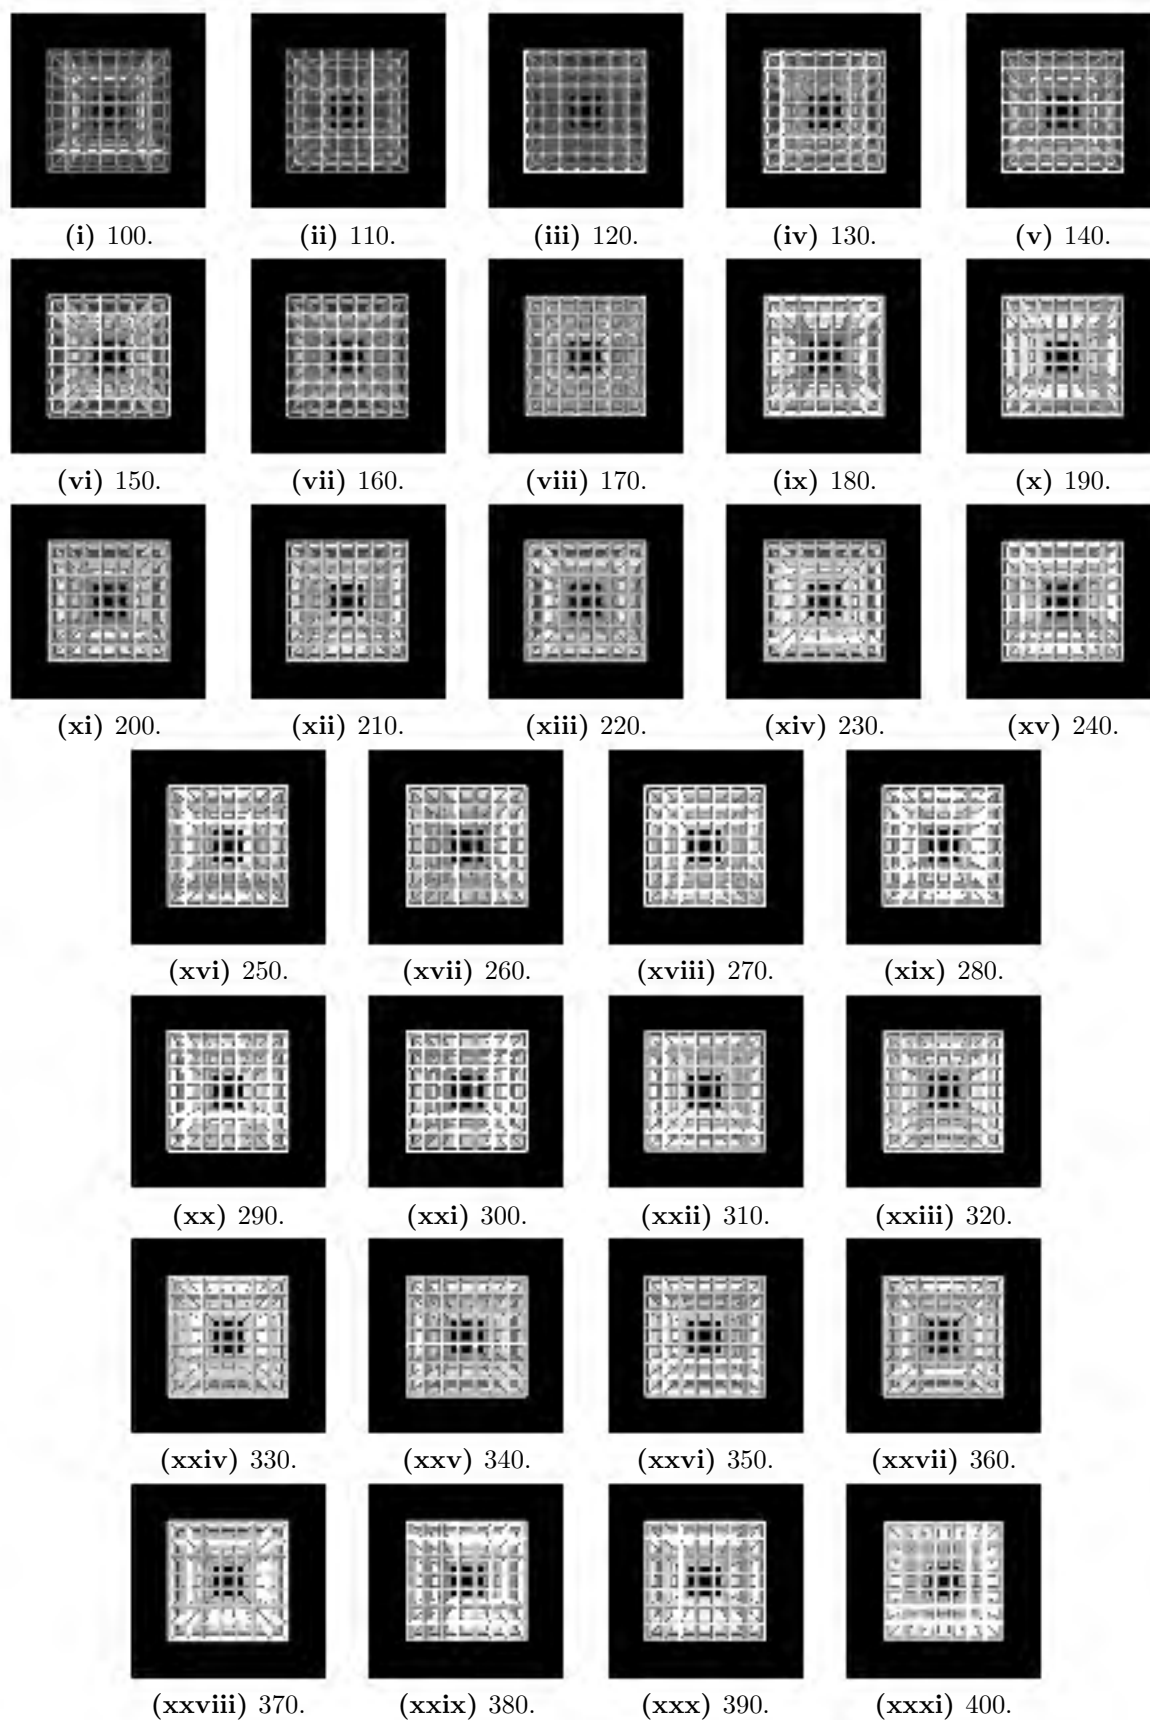


Figure F.5: Rasterisation of C8-t2 lattice at varying image sizes.

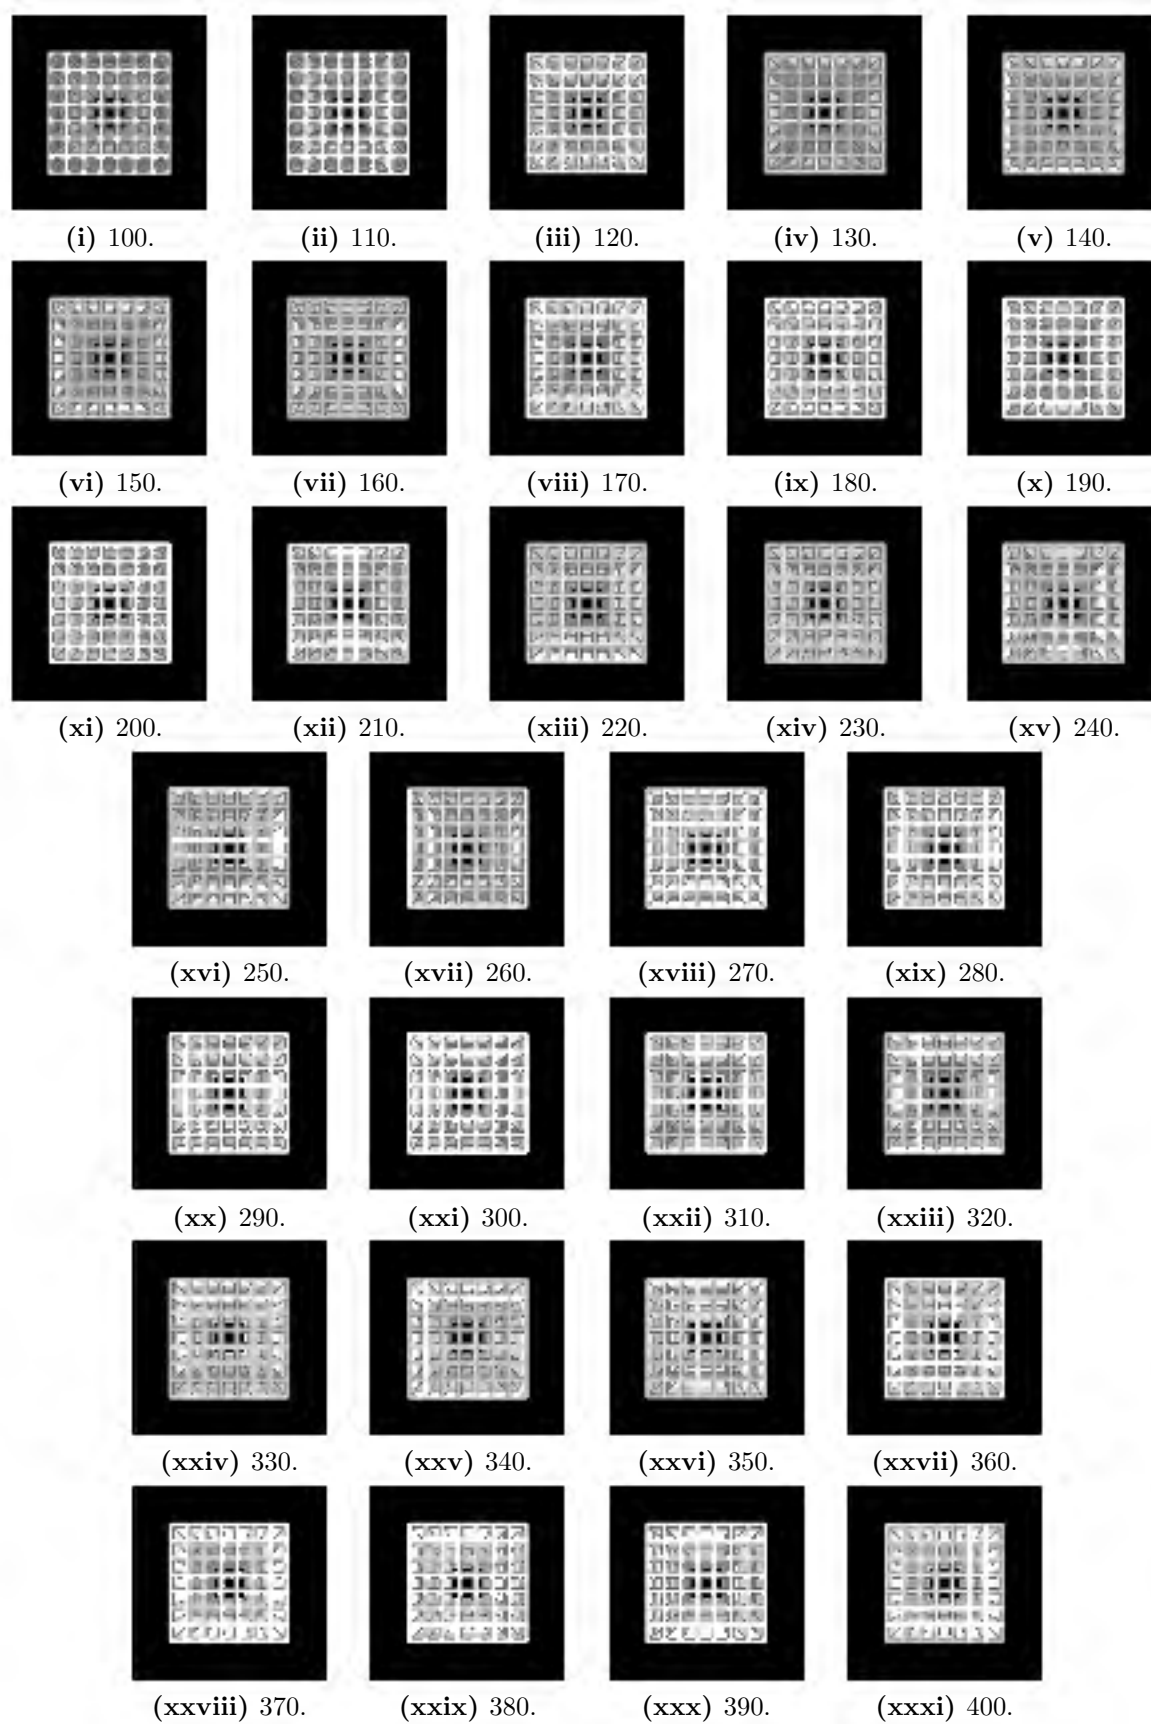


Figure F.6: Rasterisation of C8-t3 lattice at varying image sizes.

F.1.3 Size 10

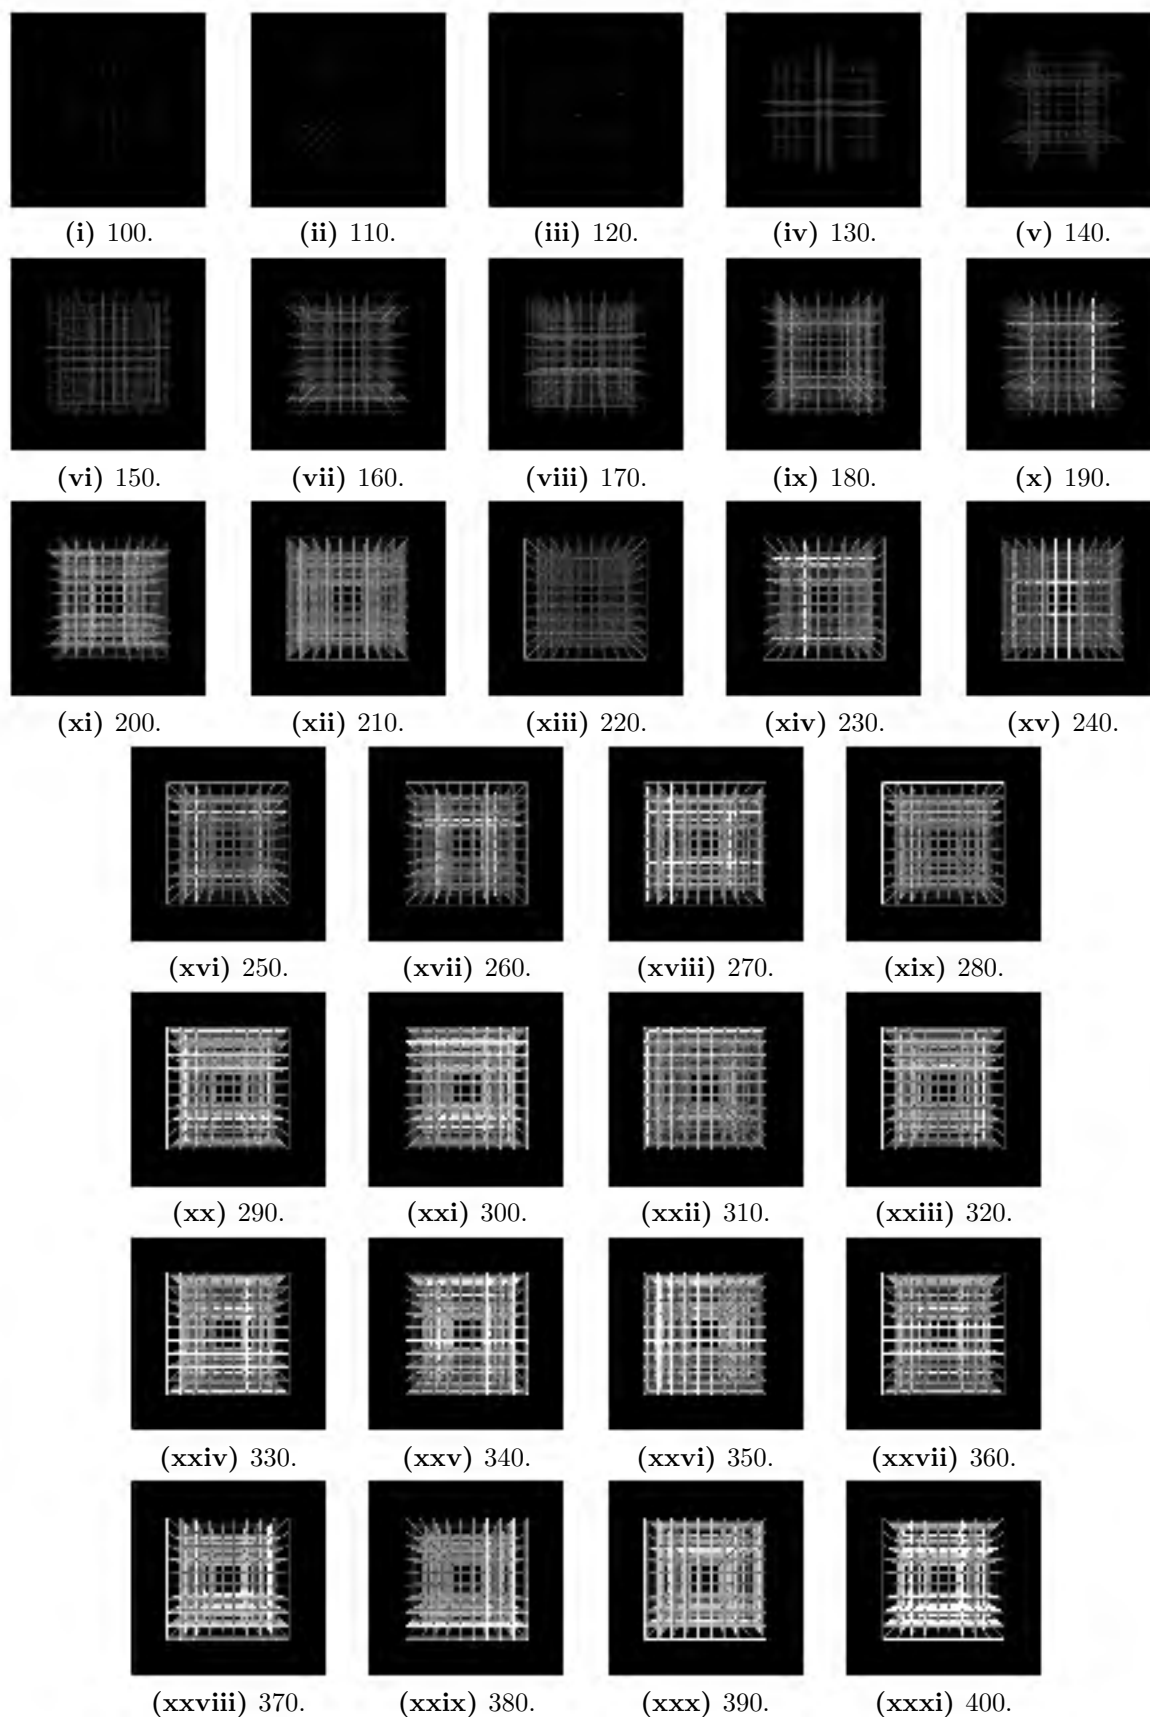


Figure F.7: Rasterisation of C10-t1 lattice at varying image sizes.

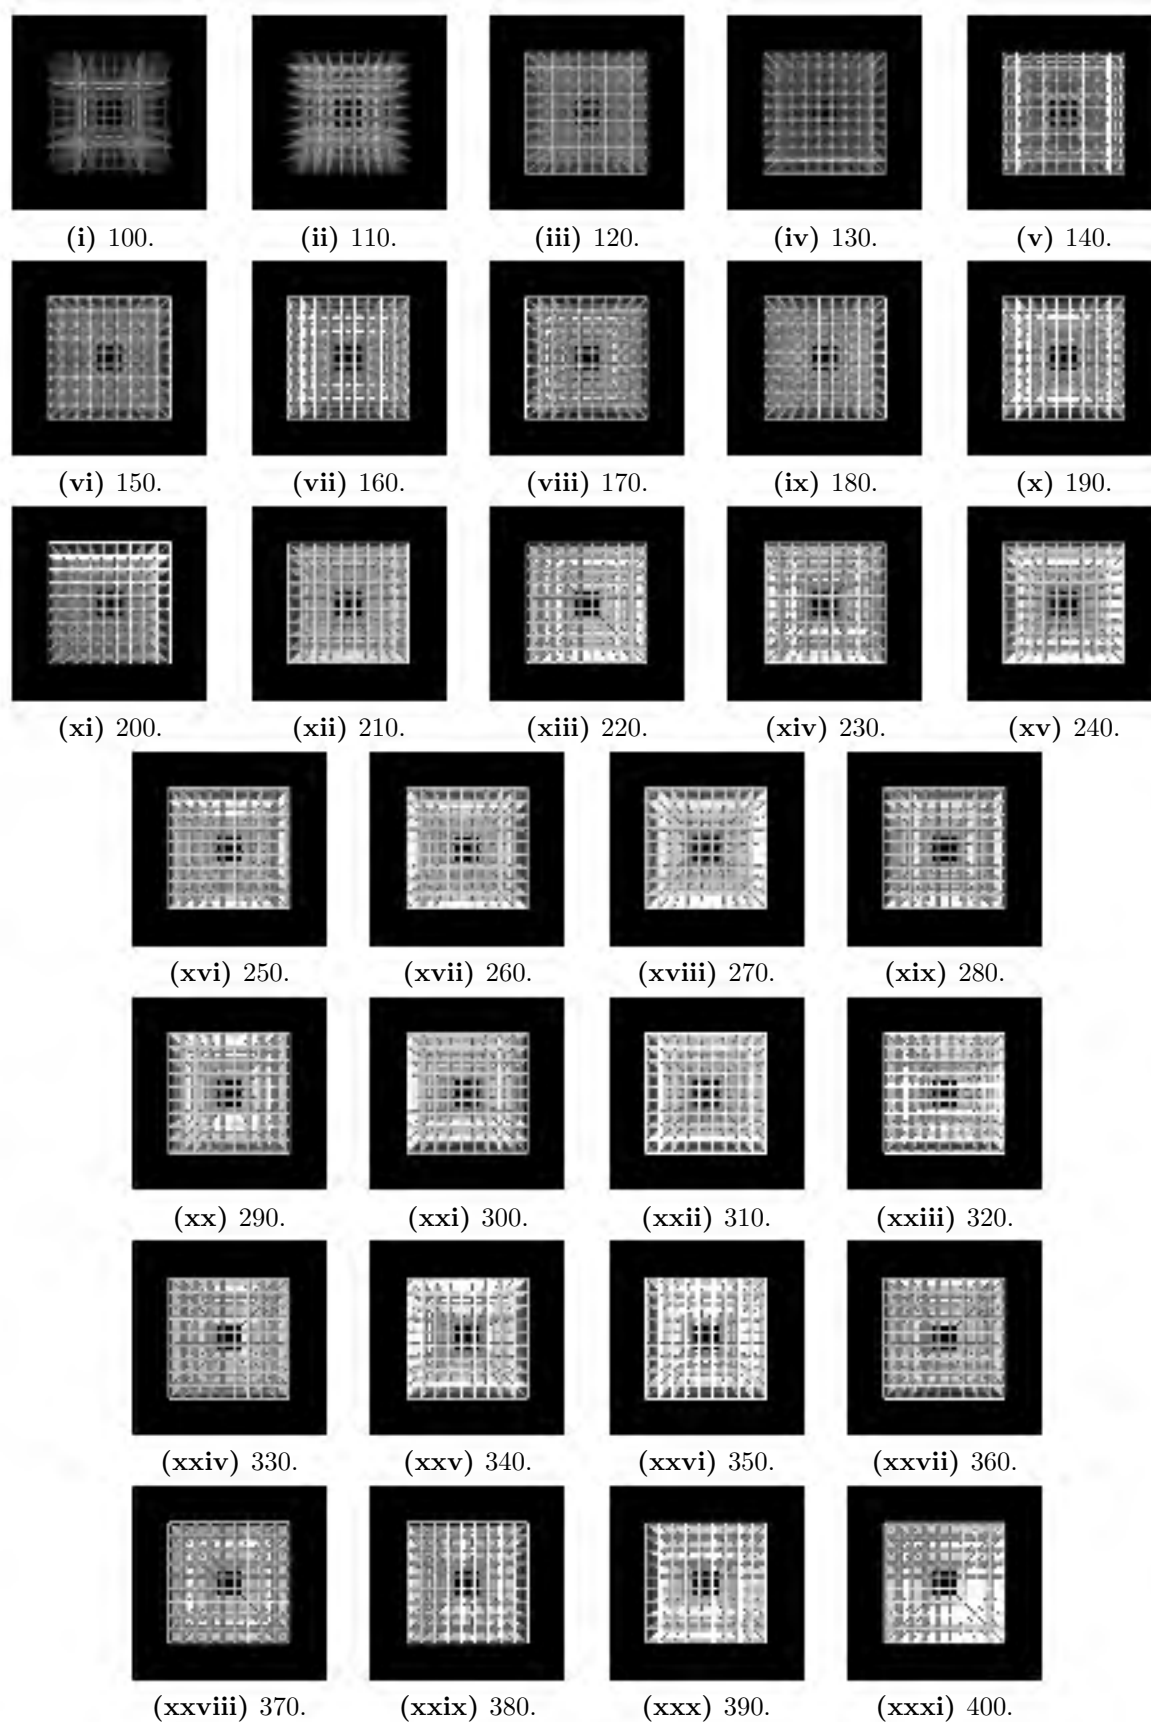


Figure F.8: Rasterisation of C10-t2 lattice at varying image sizes.

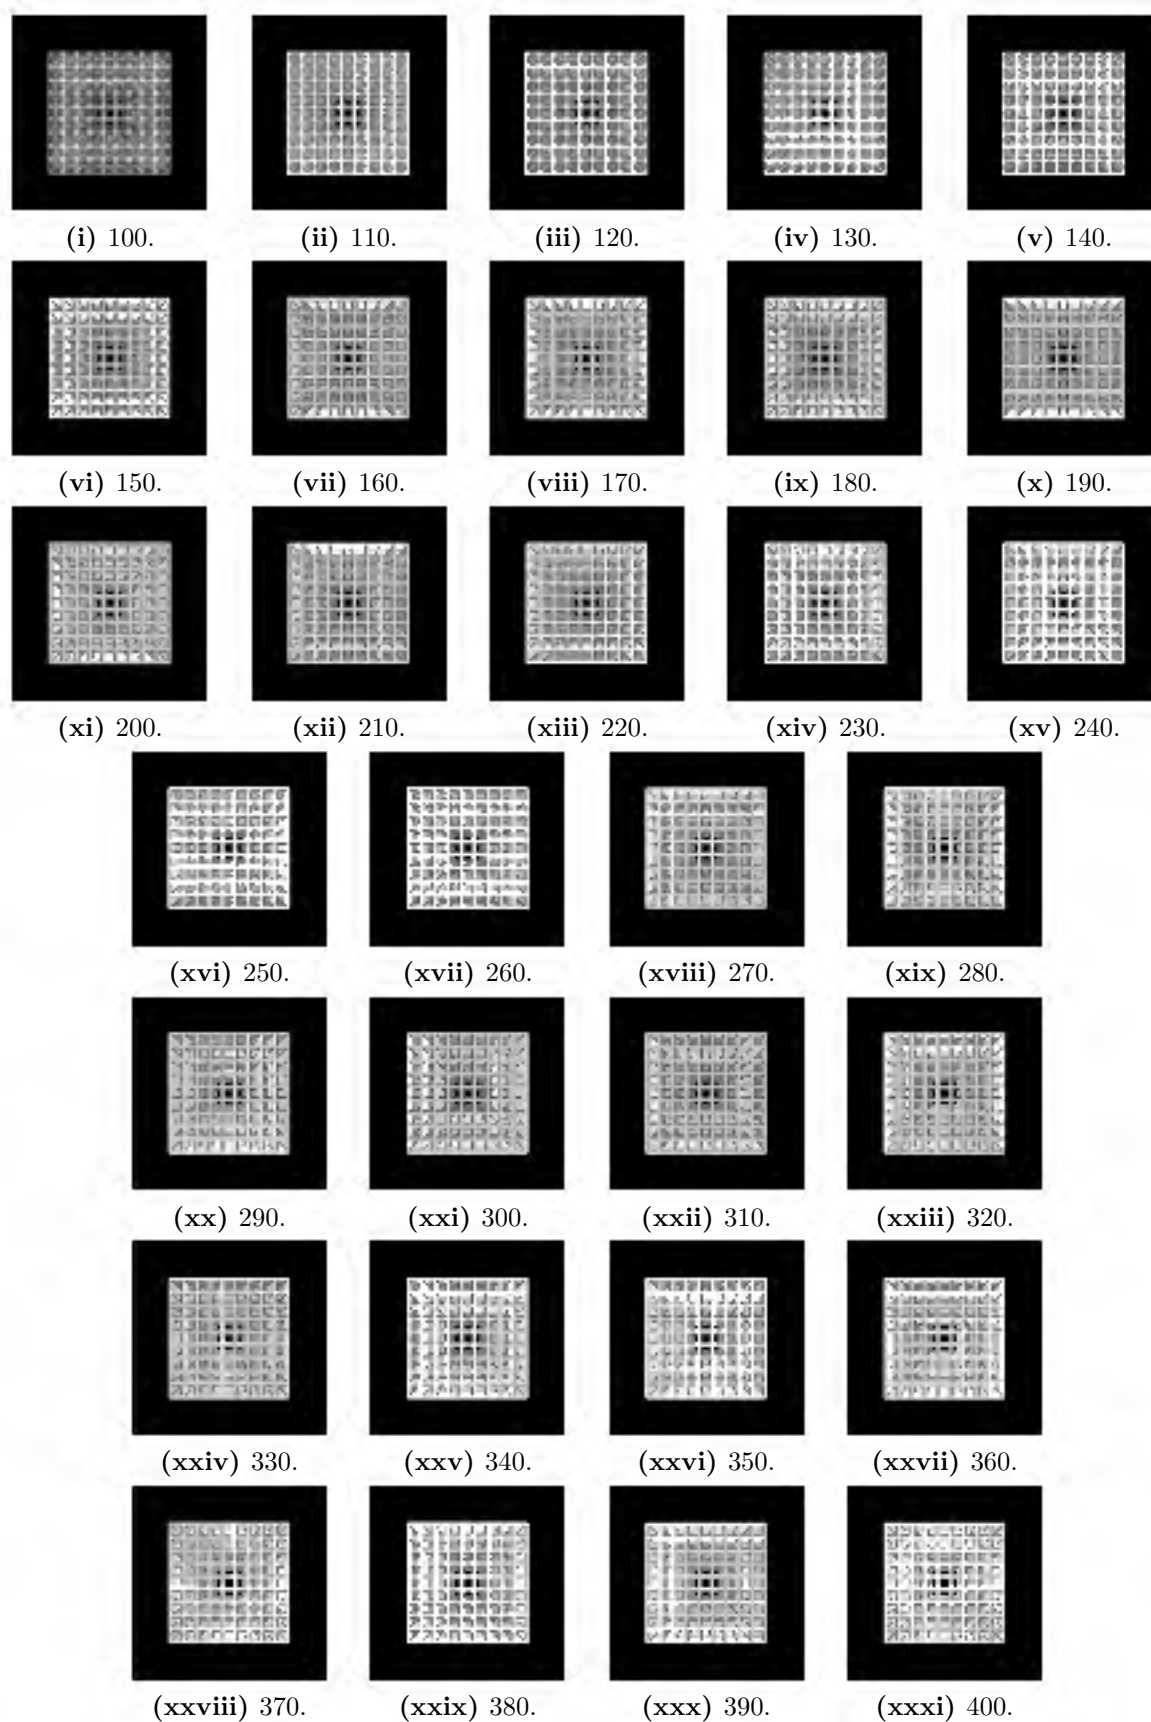


Figure F.9: Rasterisation of C10-t3 lattice at varying image sizes.

F.2 Kelvin cell

F.2.1 Size 6

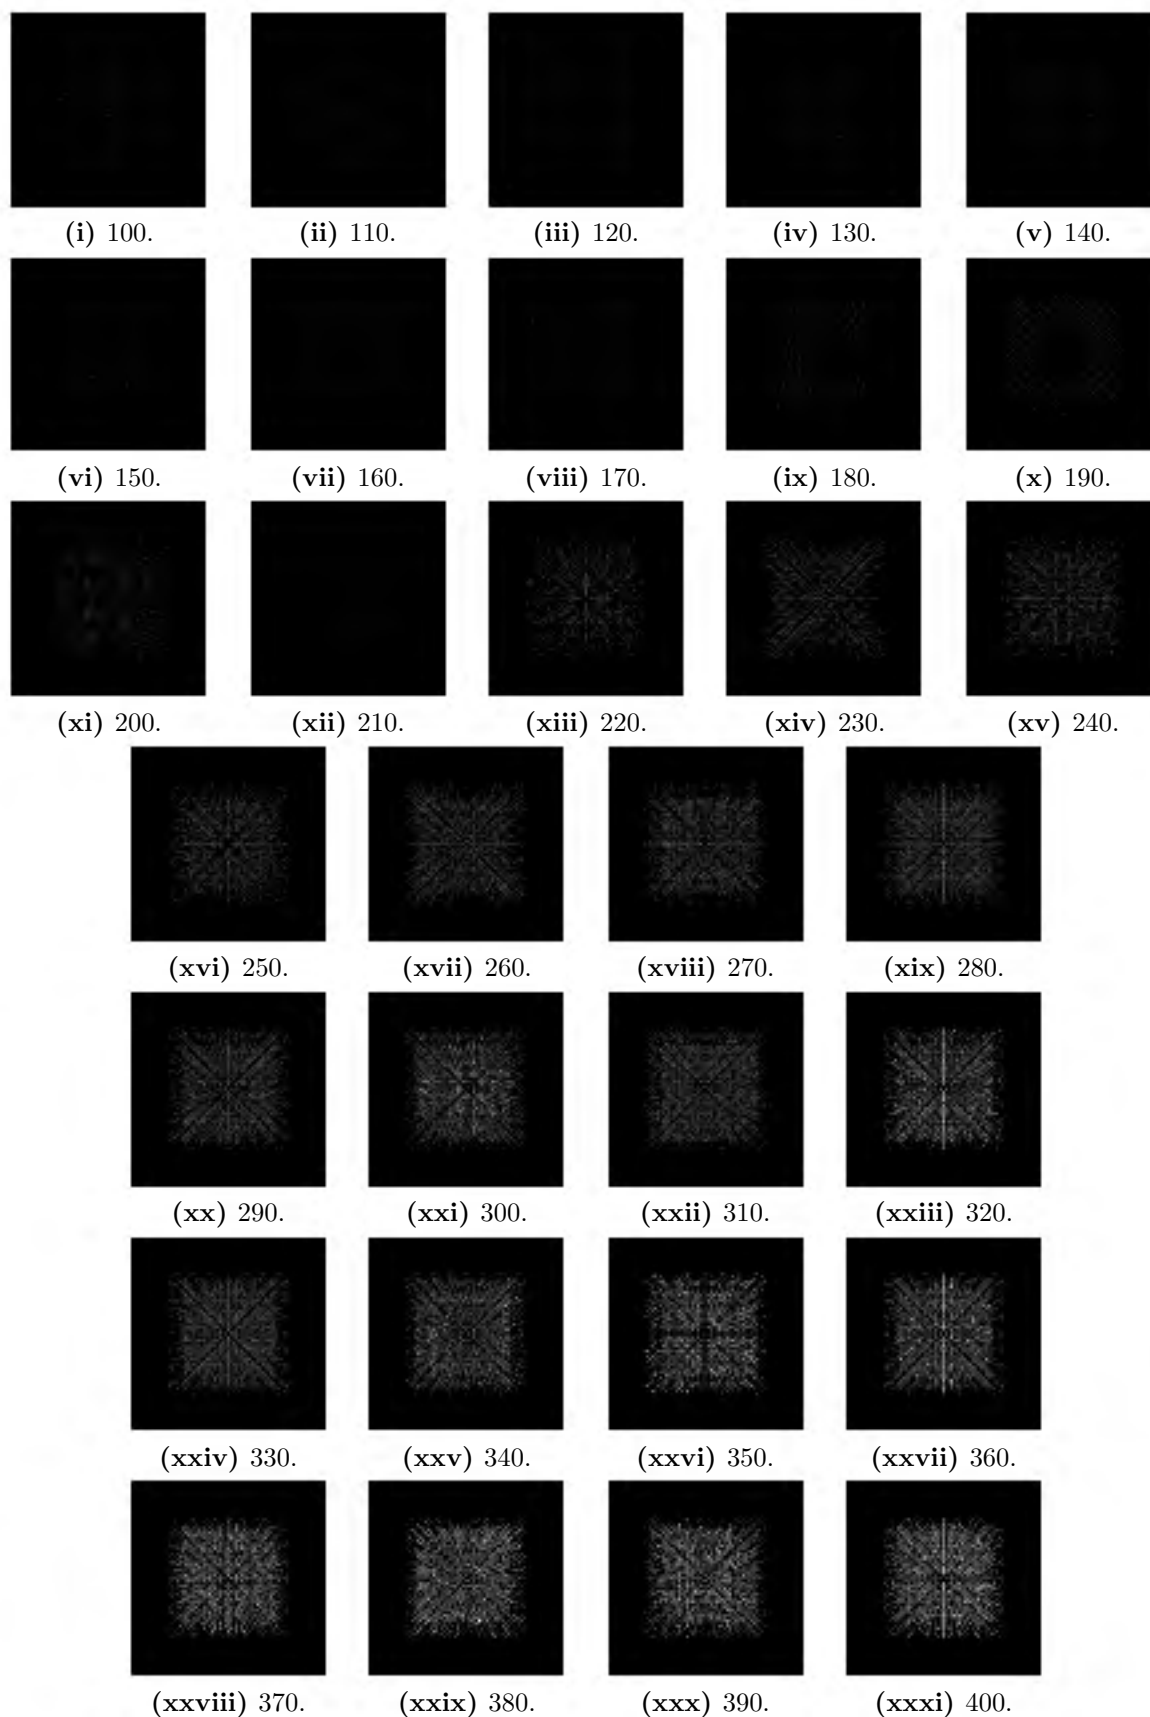


Figure F.10: Rasterisation of K6-t1 lattice at varying image sizes.

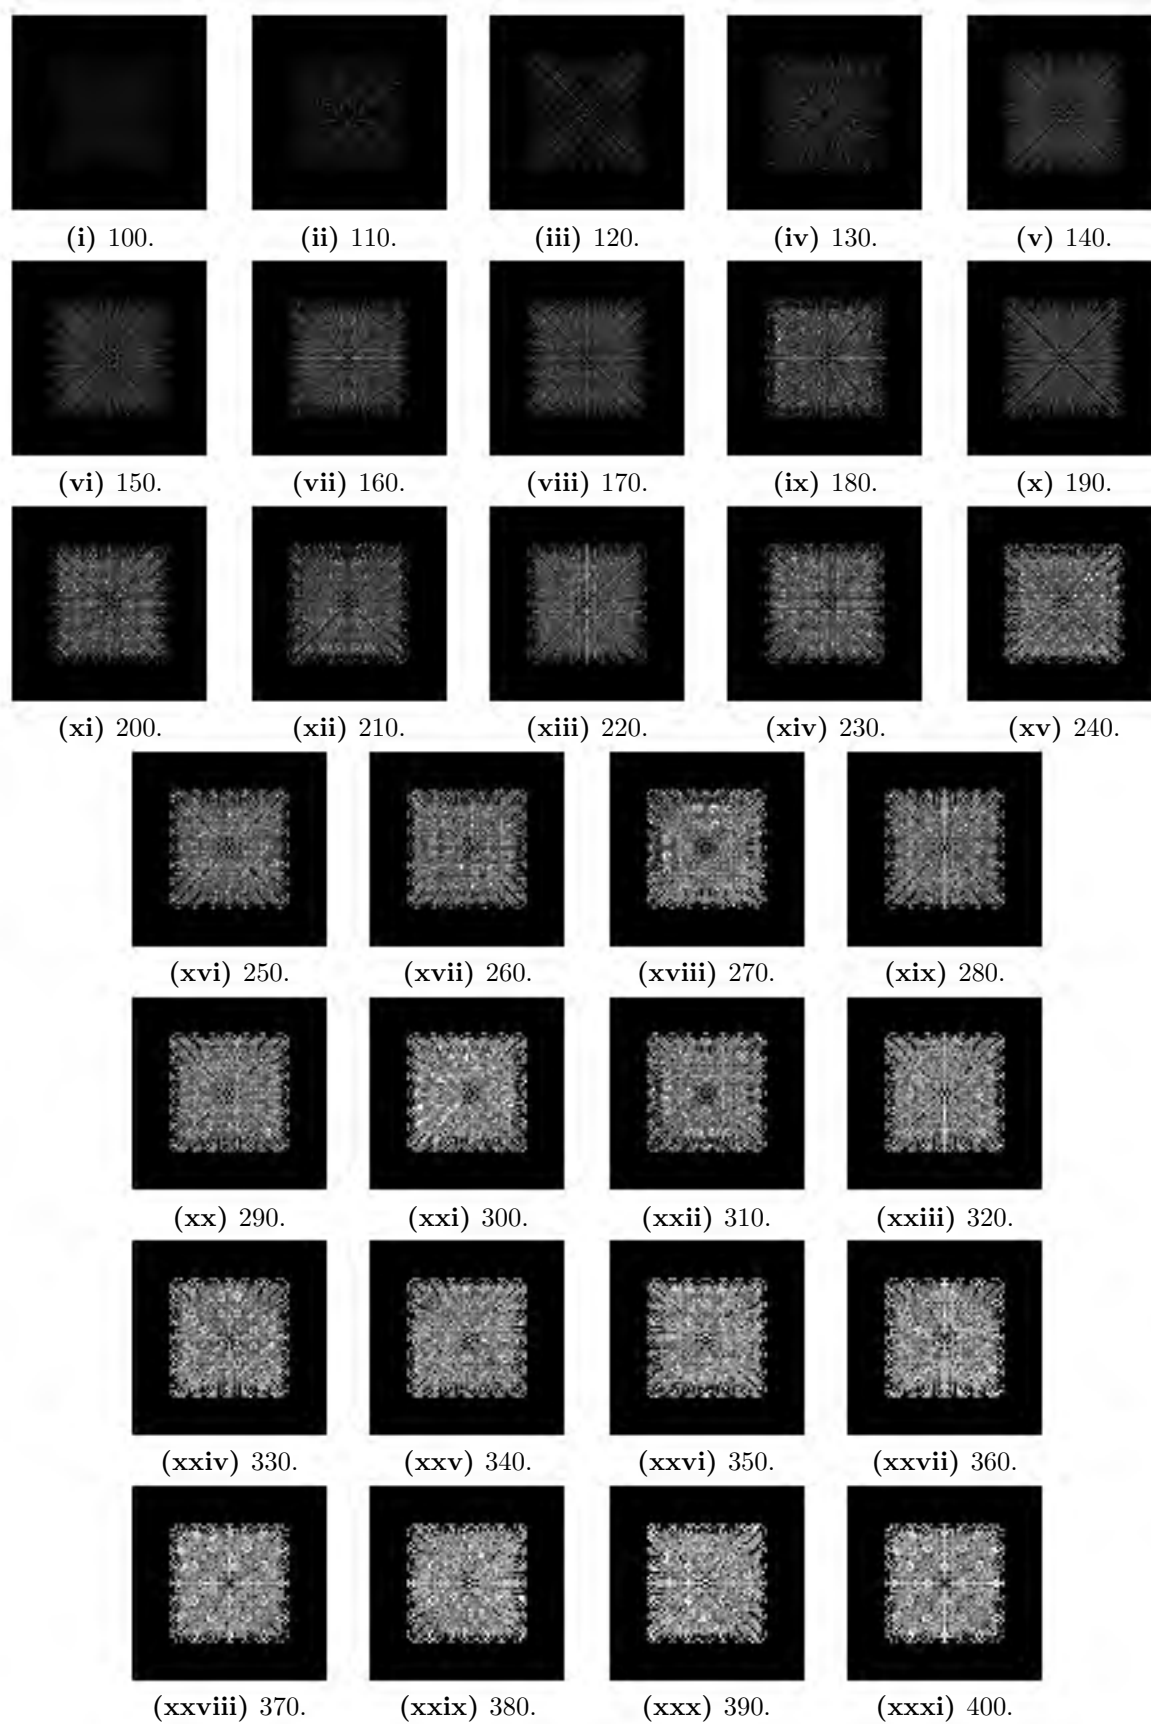


Figure F.11: Rasterisation of K6-t2 lattice at varying image sizes.

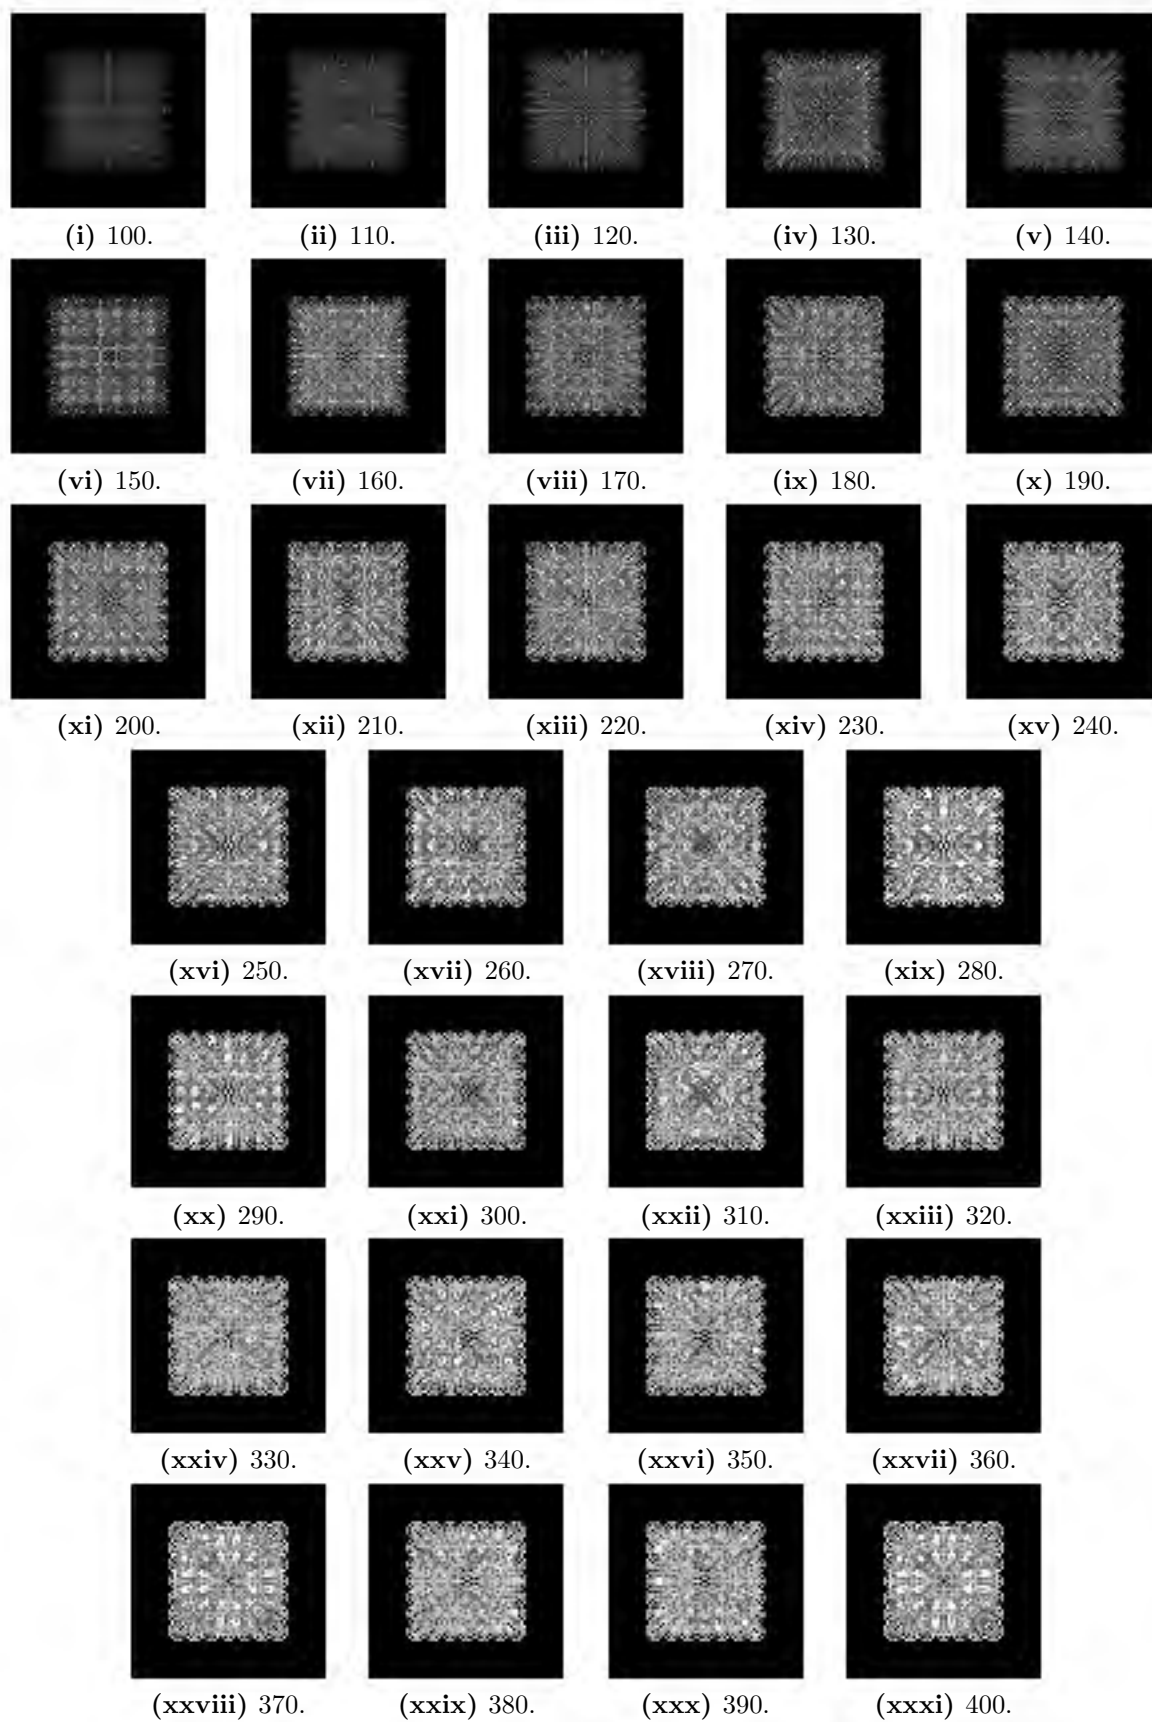


Figure F.12: Rasterisation of K6-t3 lattice at varying image sizes.

F.2.2 Size 8

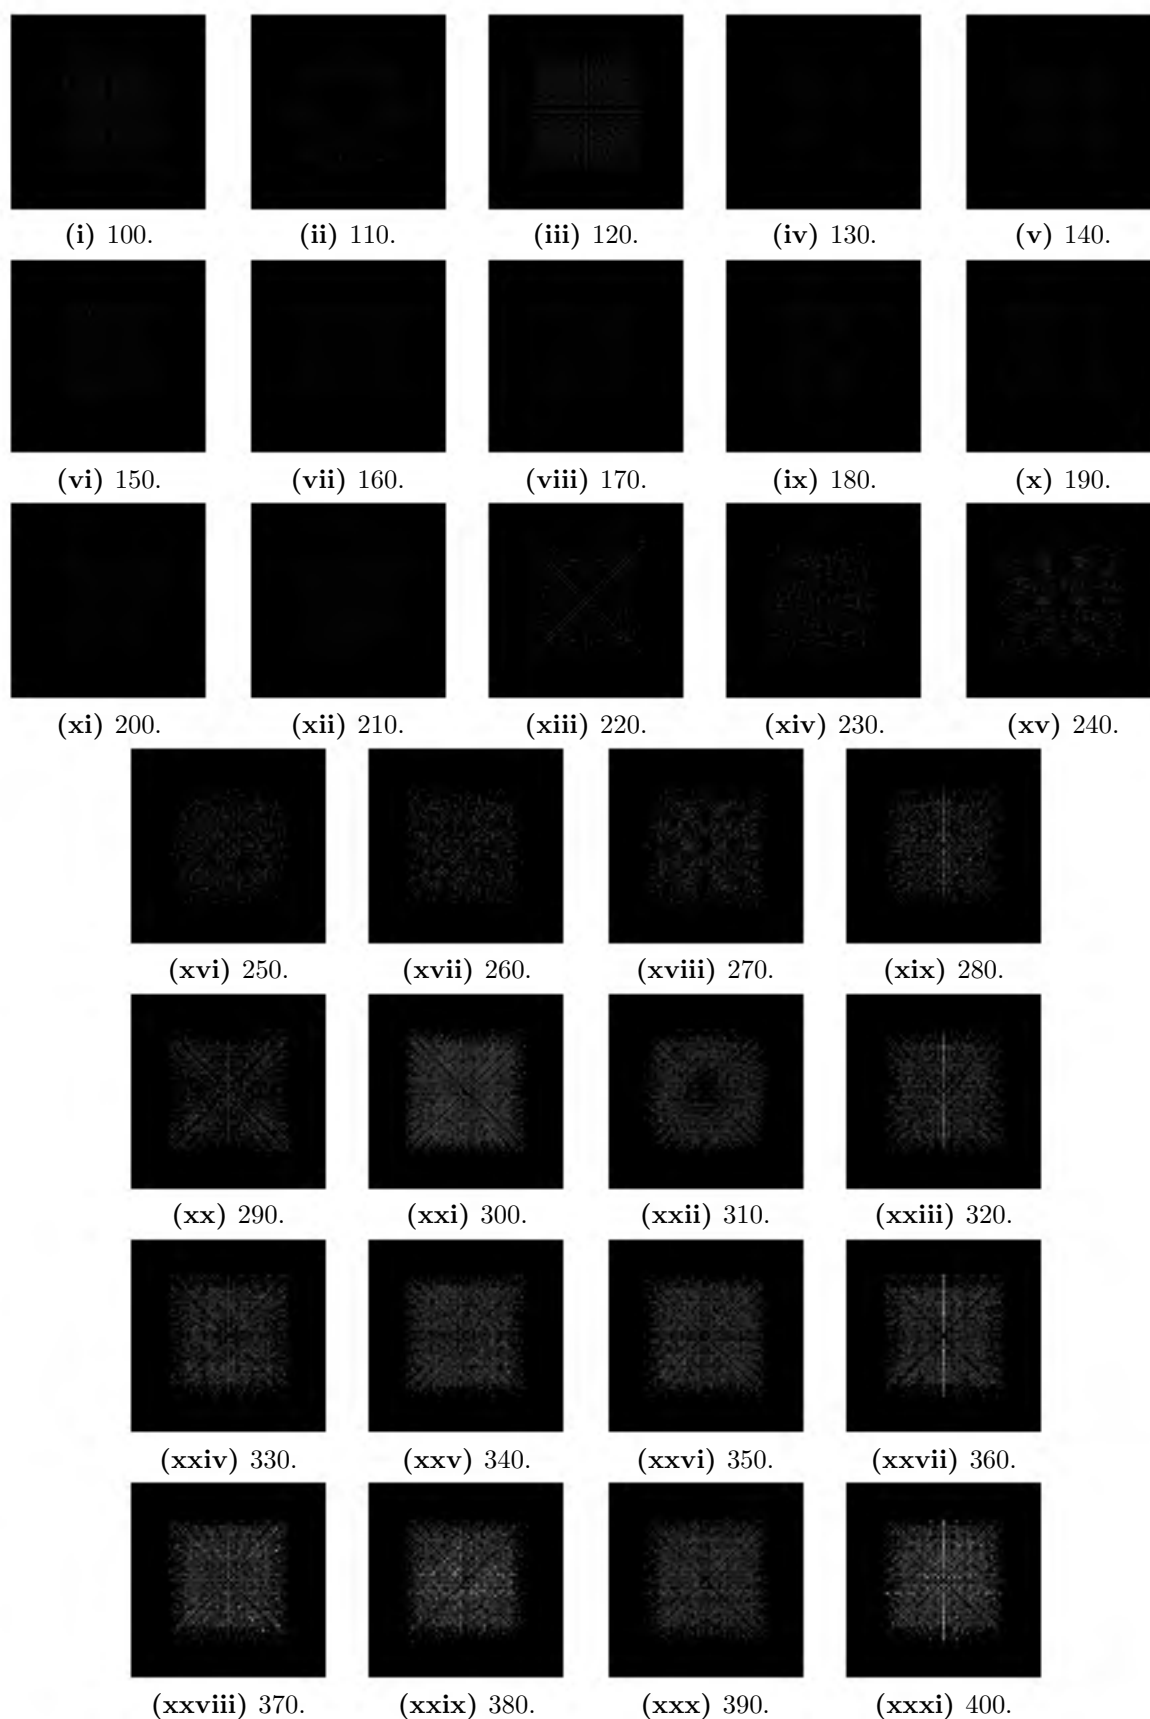


Figure F.13: Rasterisation of K8-t1 lattice at varying image sizes.




Figure F.14: Rasterisation of K8-t2 lattice at varying image sizes.

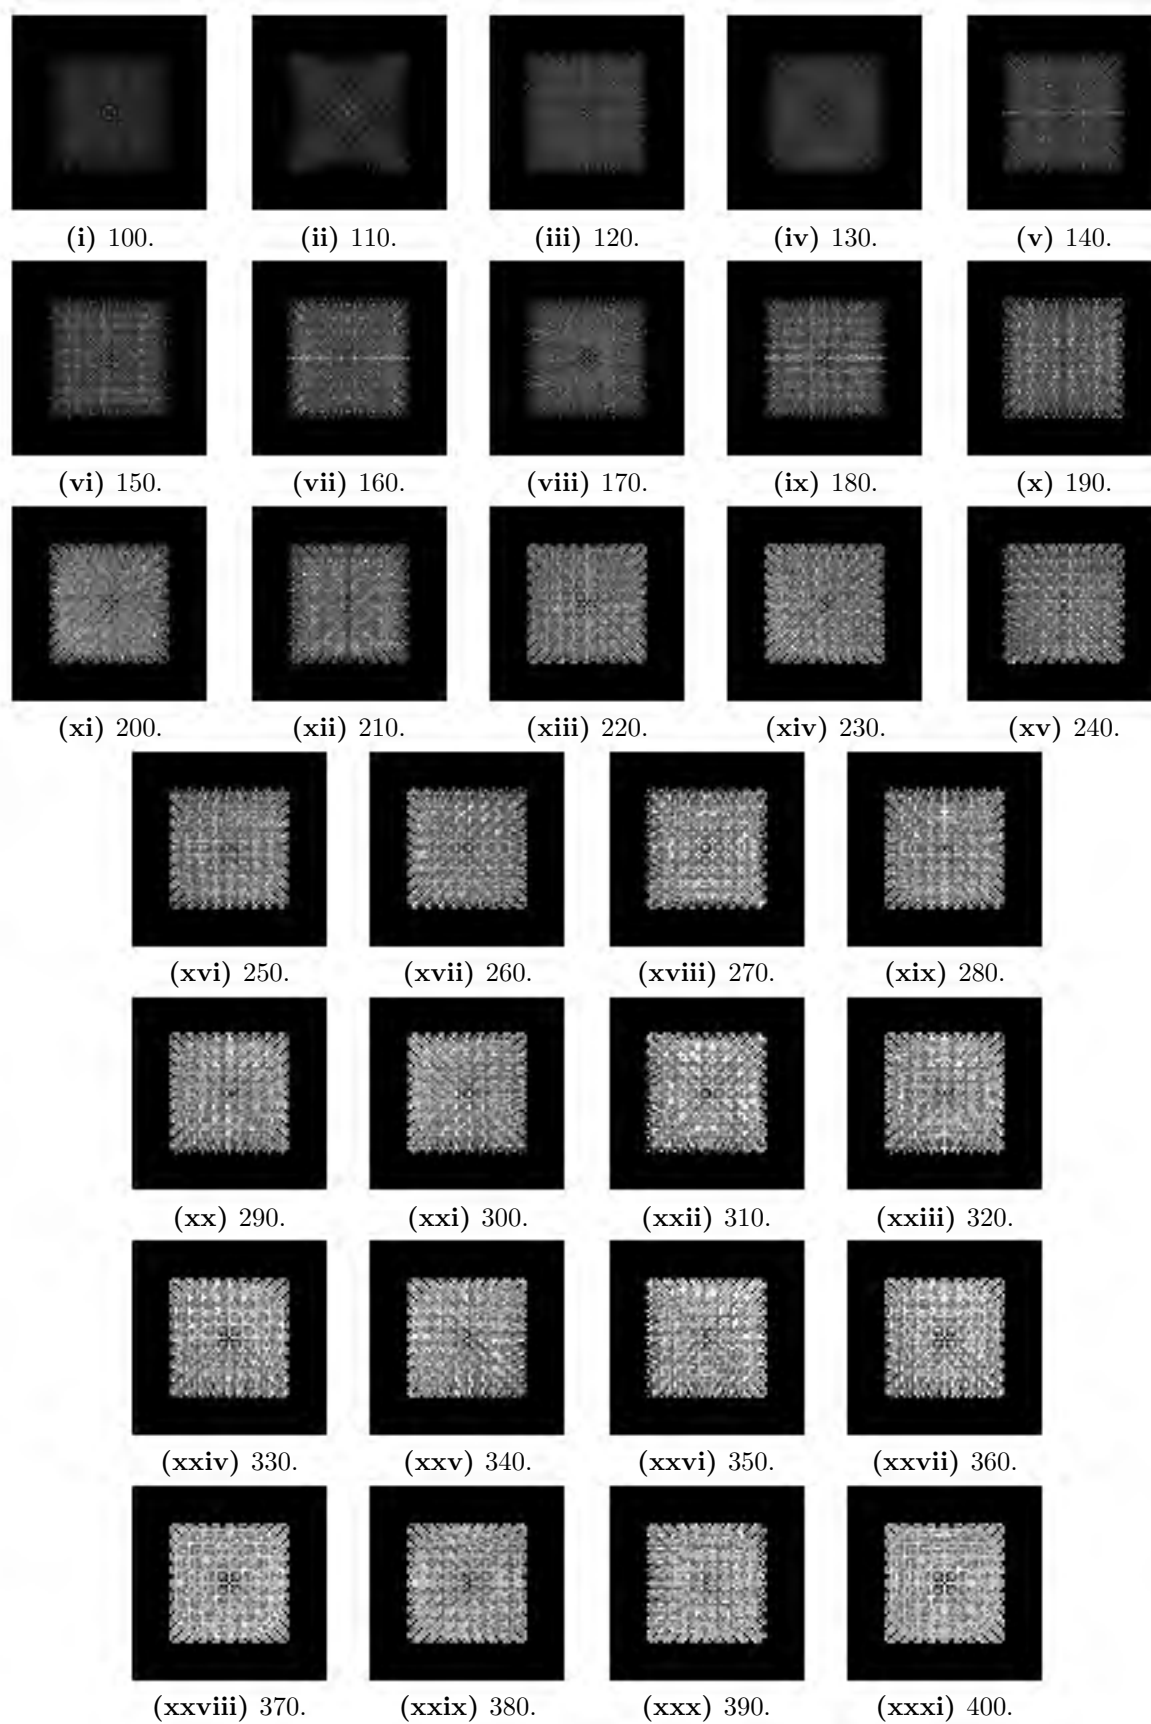


Figure F.15: Rasterisation of K8-t3 lattice at varying image sizes.

F.2.3 Size 10



Figure F.16: Rasterisation of K10-t1 lattice at varying image sizes.

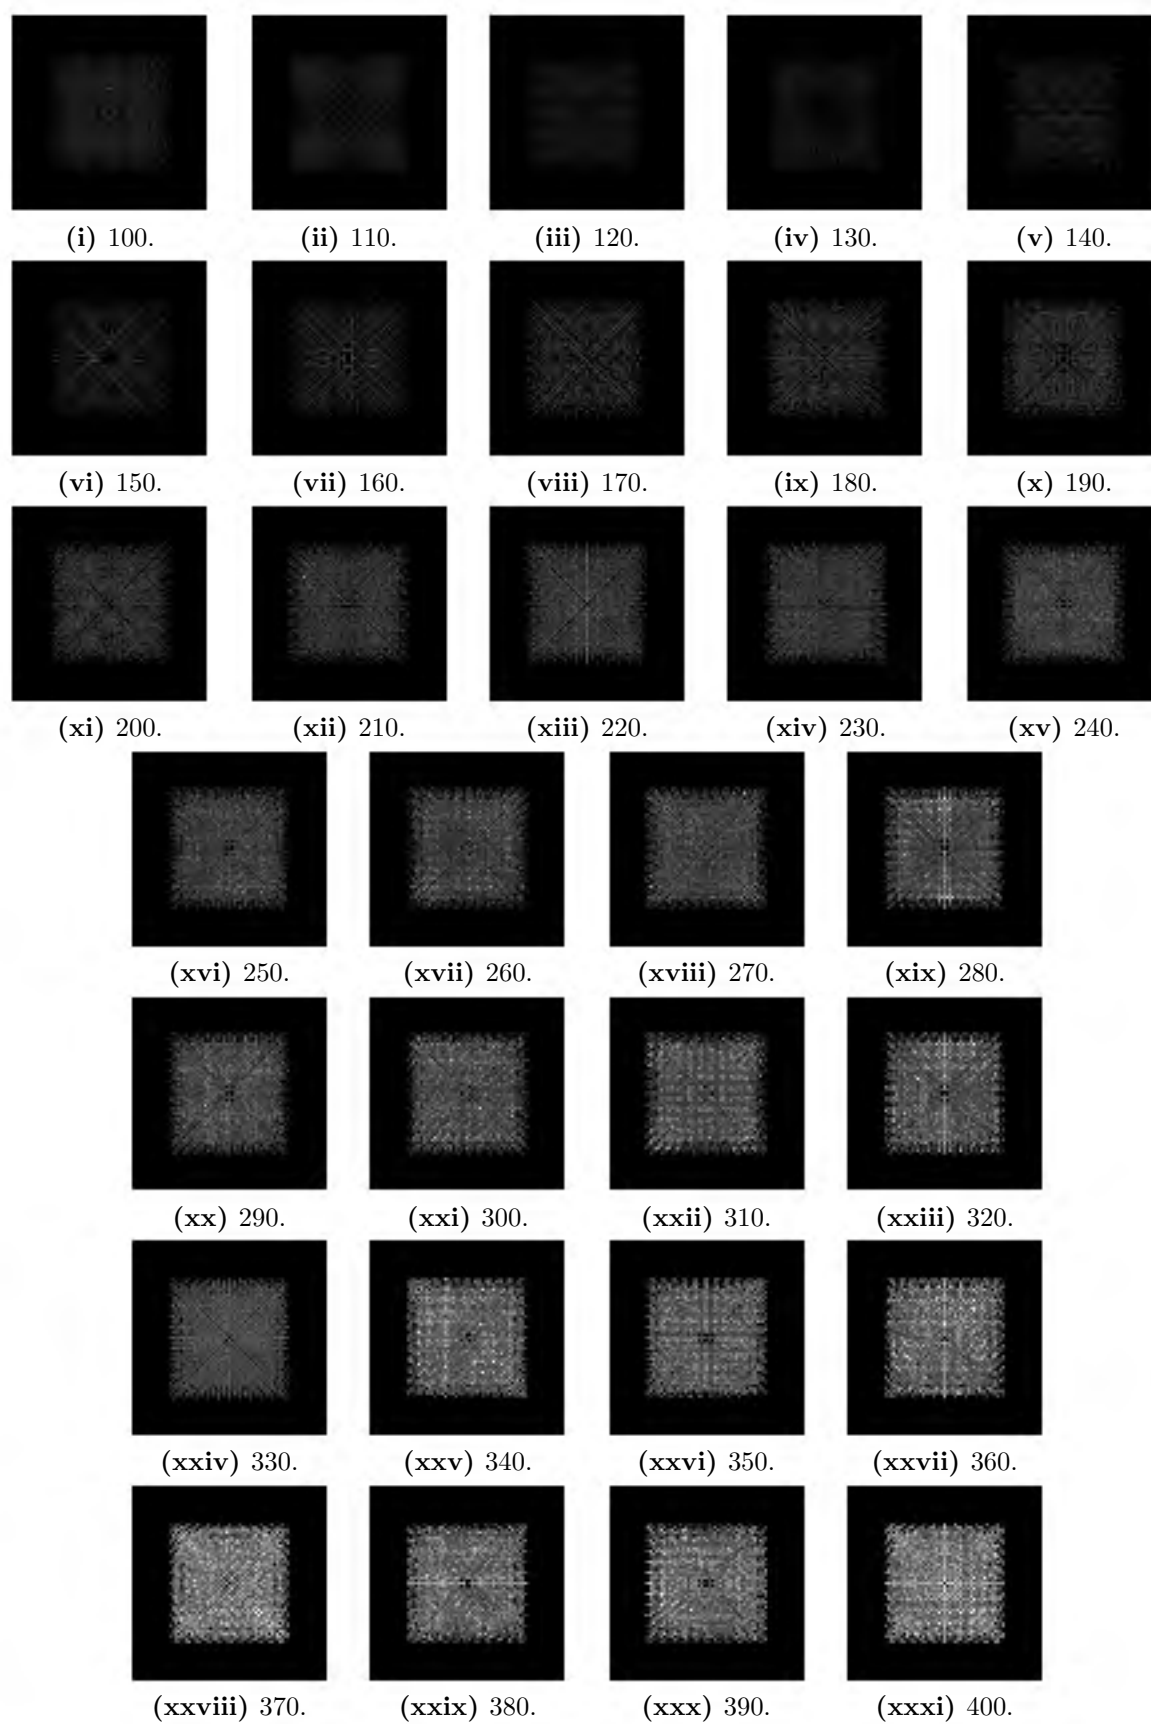


Figure F.17: Rasterisation of K10-t2 lattice at varying image sizes.

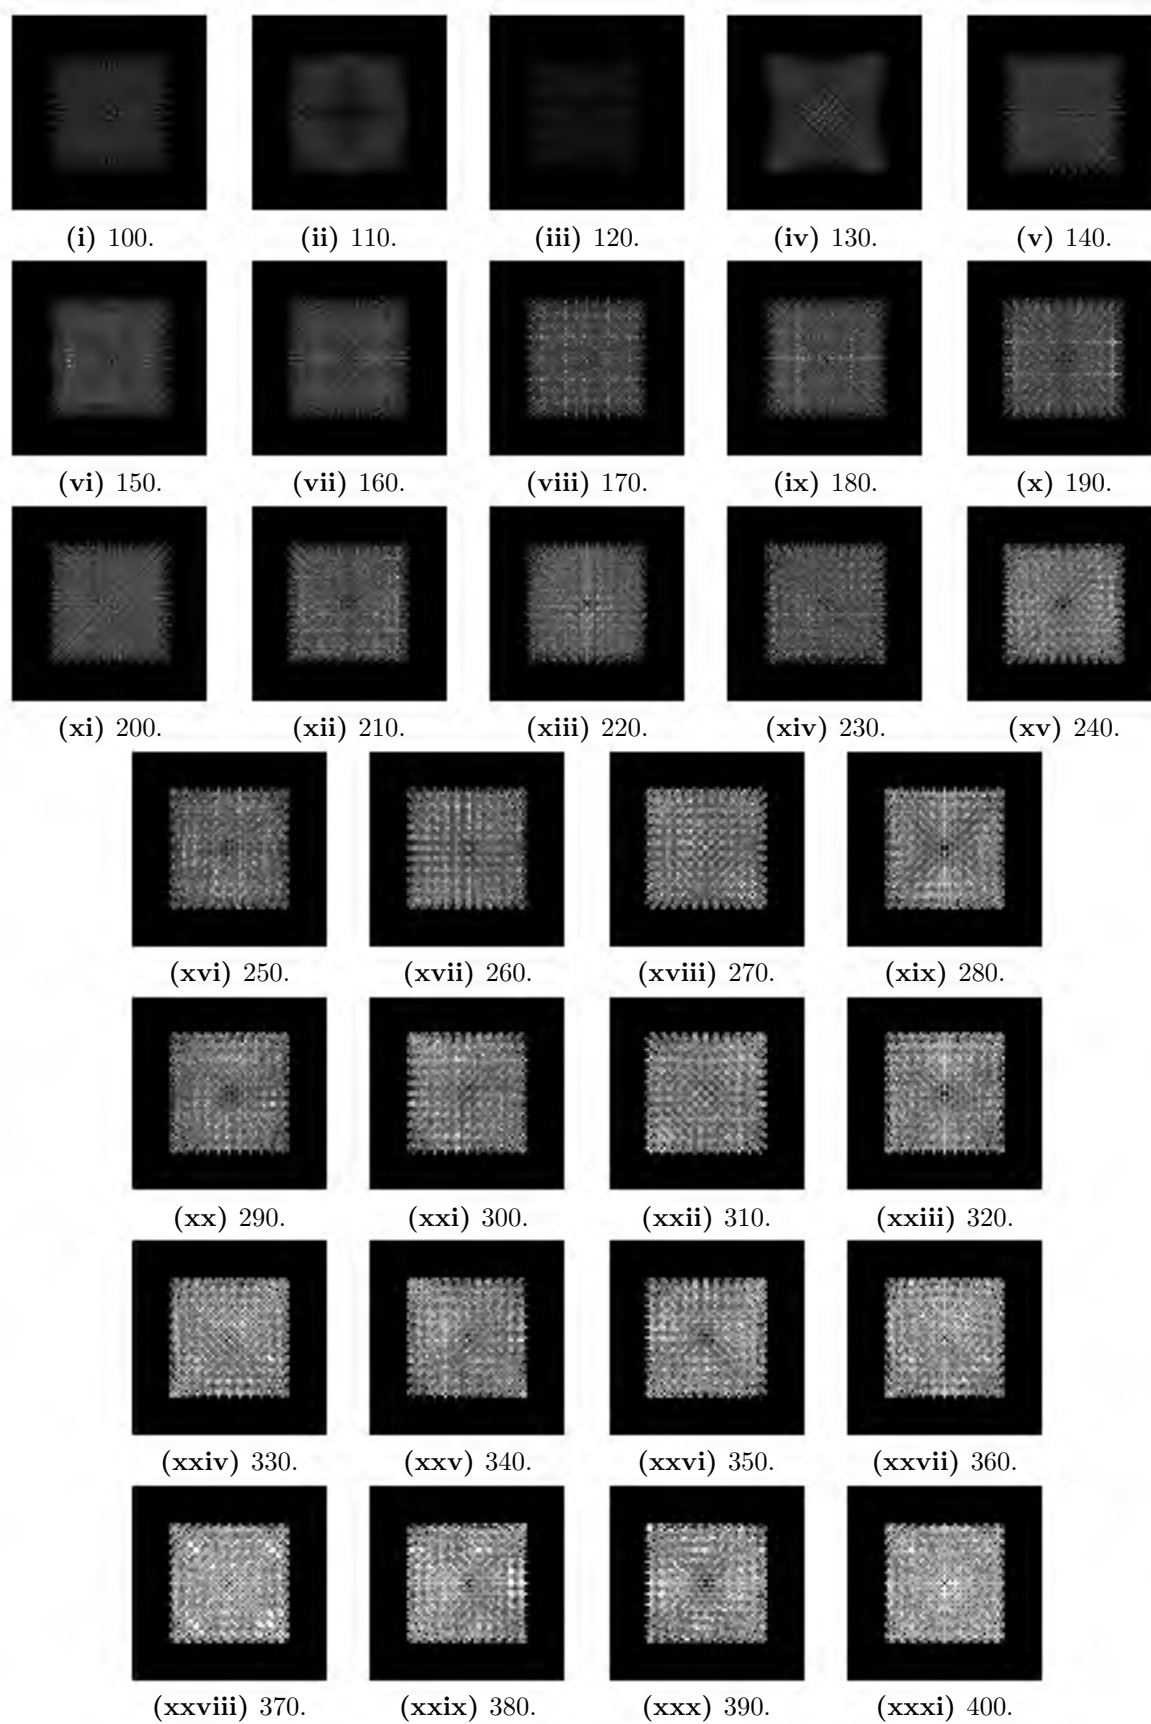


Figure F.18: Rasterisation of K10-t3 lattice at varying image sizes.

F.3 Octet truss

F.3.1 Size 6

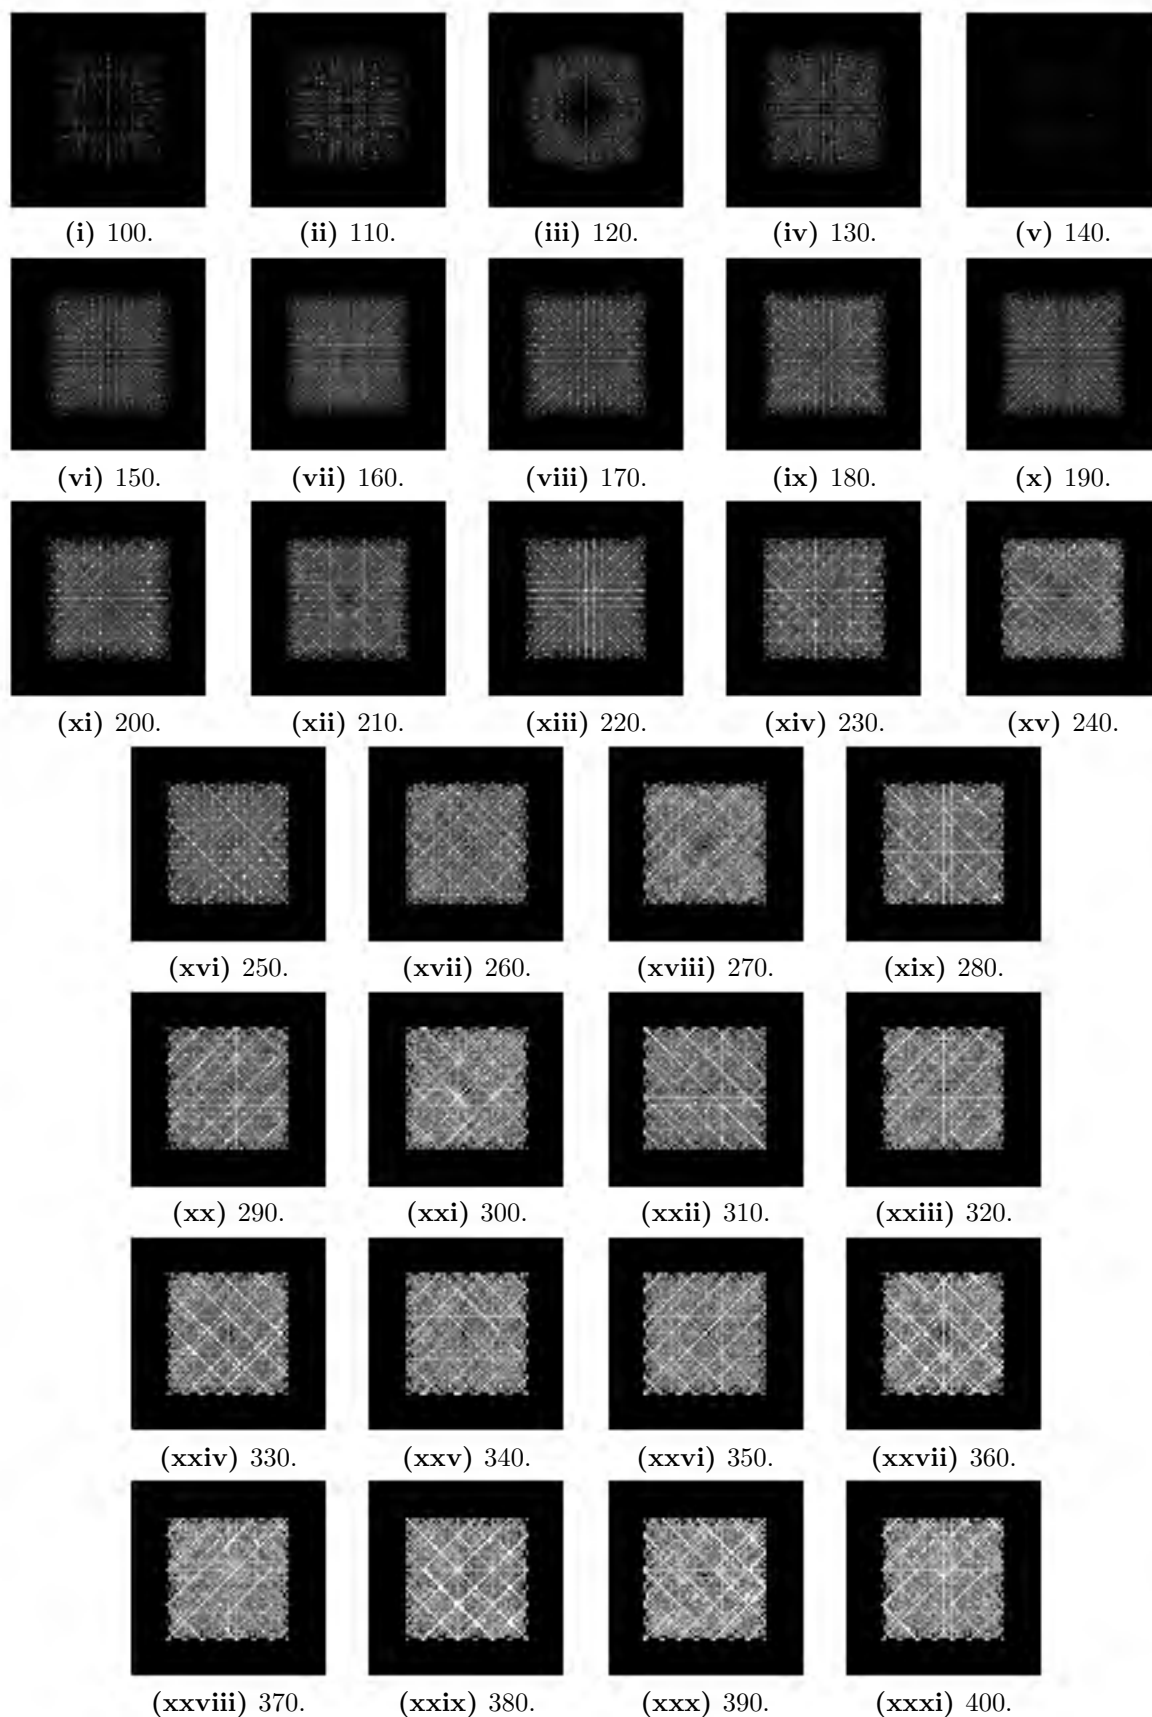


Figure F.19: Rasterisation of OT6-t1 lattice at varying image sizes.

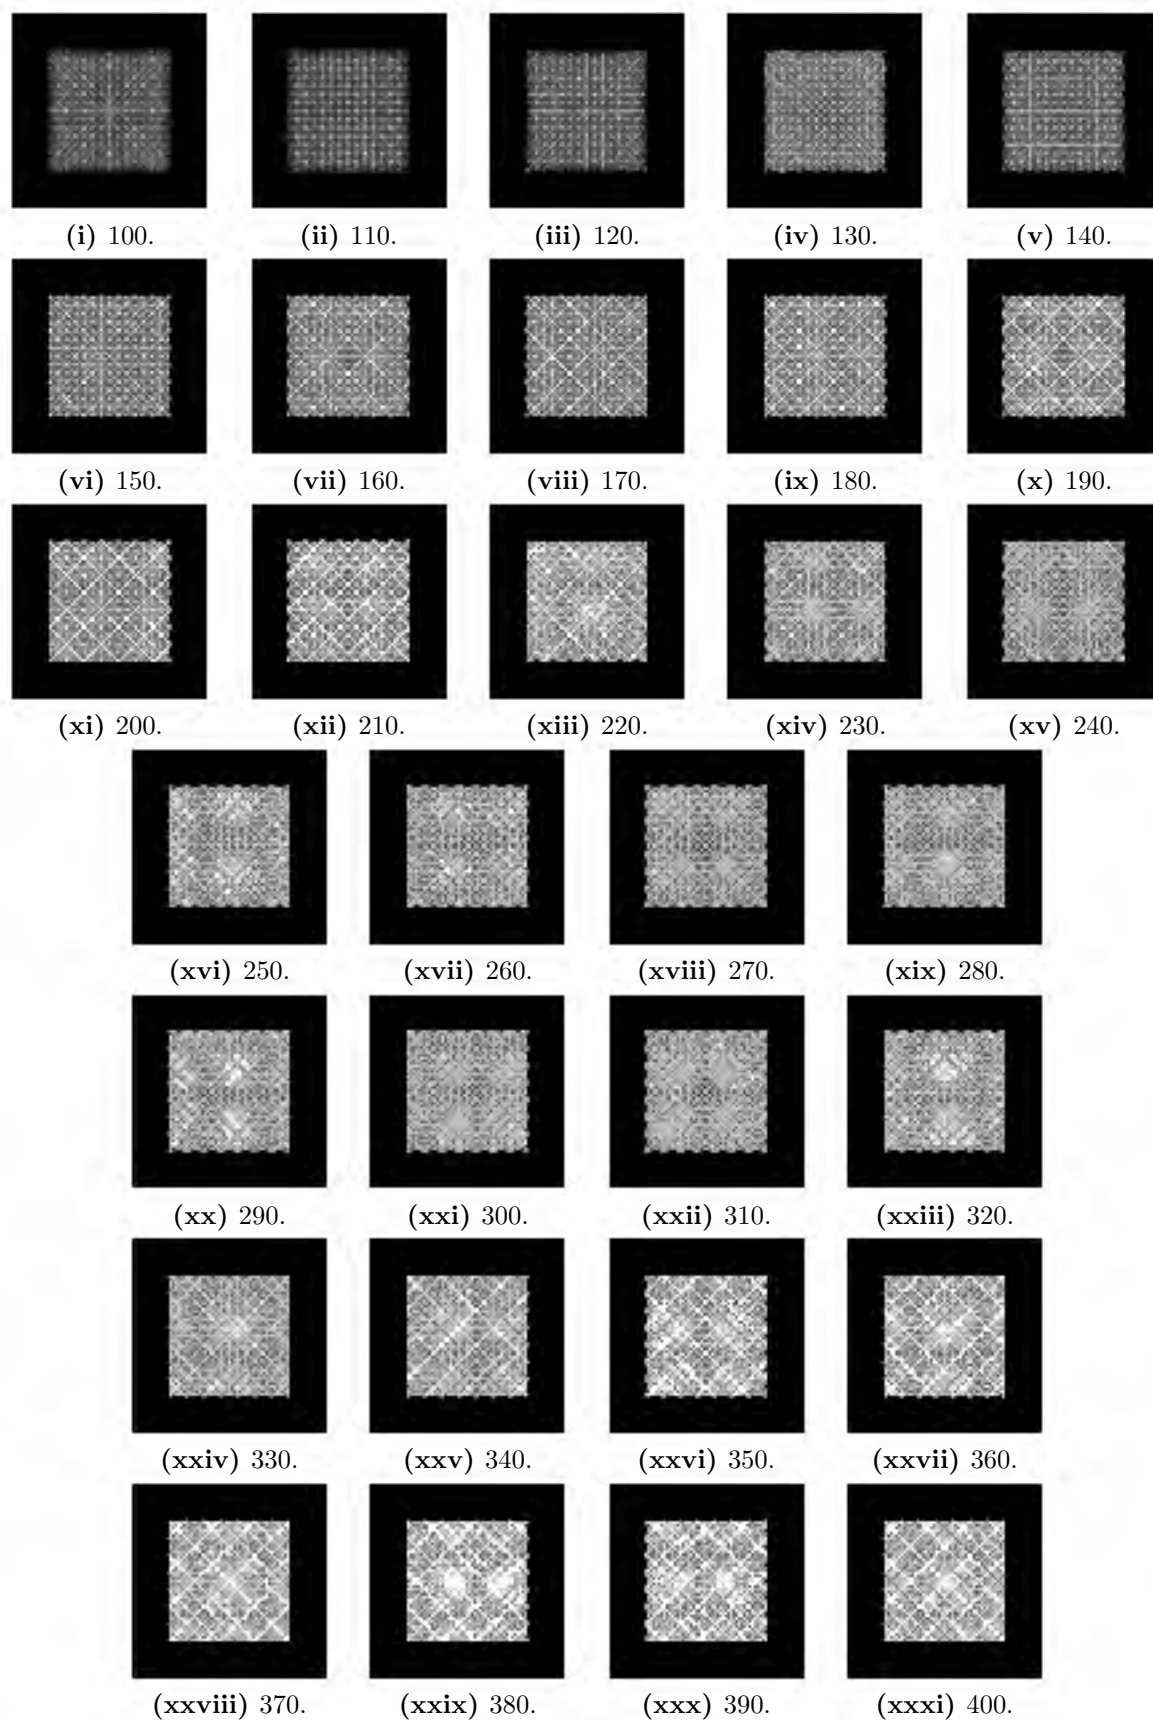


Figure F.20: Rasterisation of OT6-t2 lattice at varying image sizes.

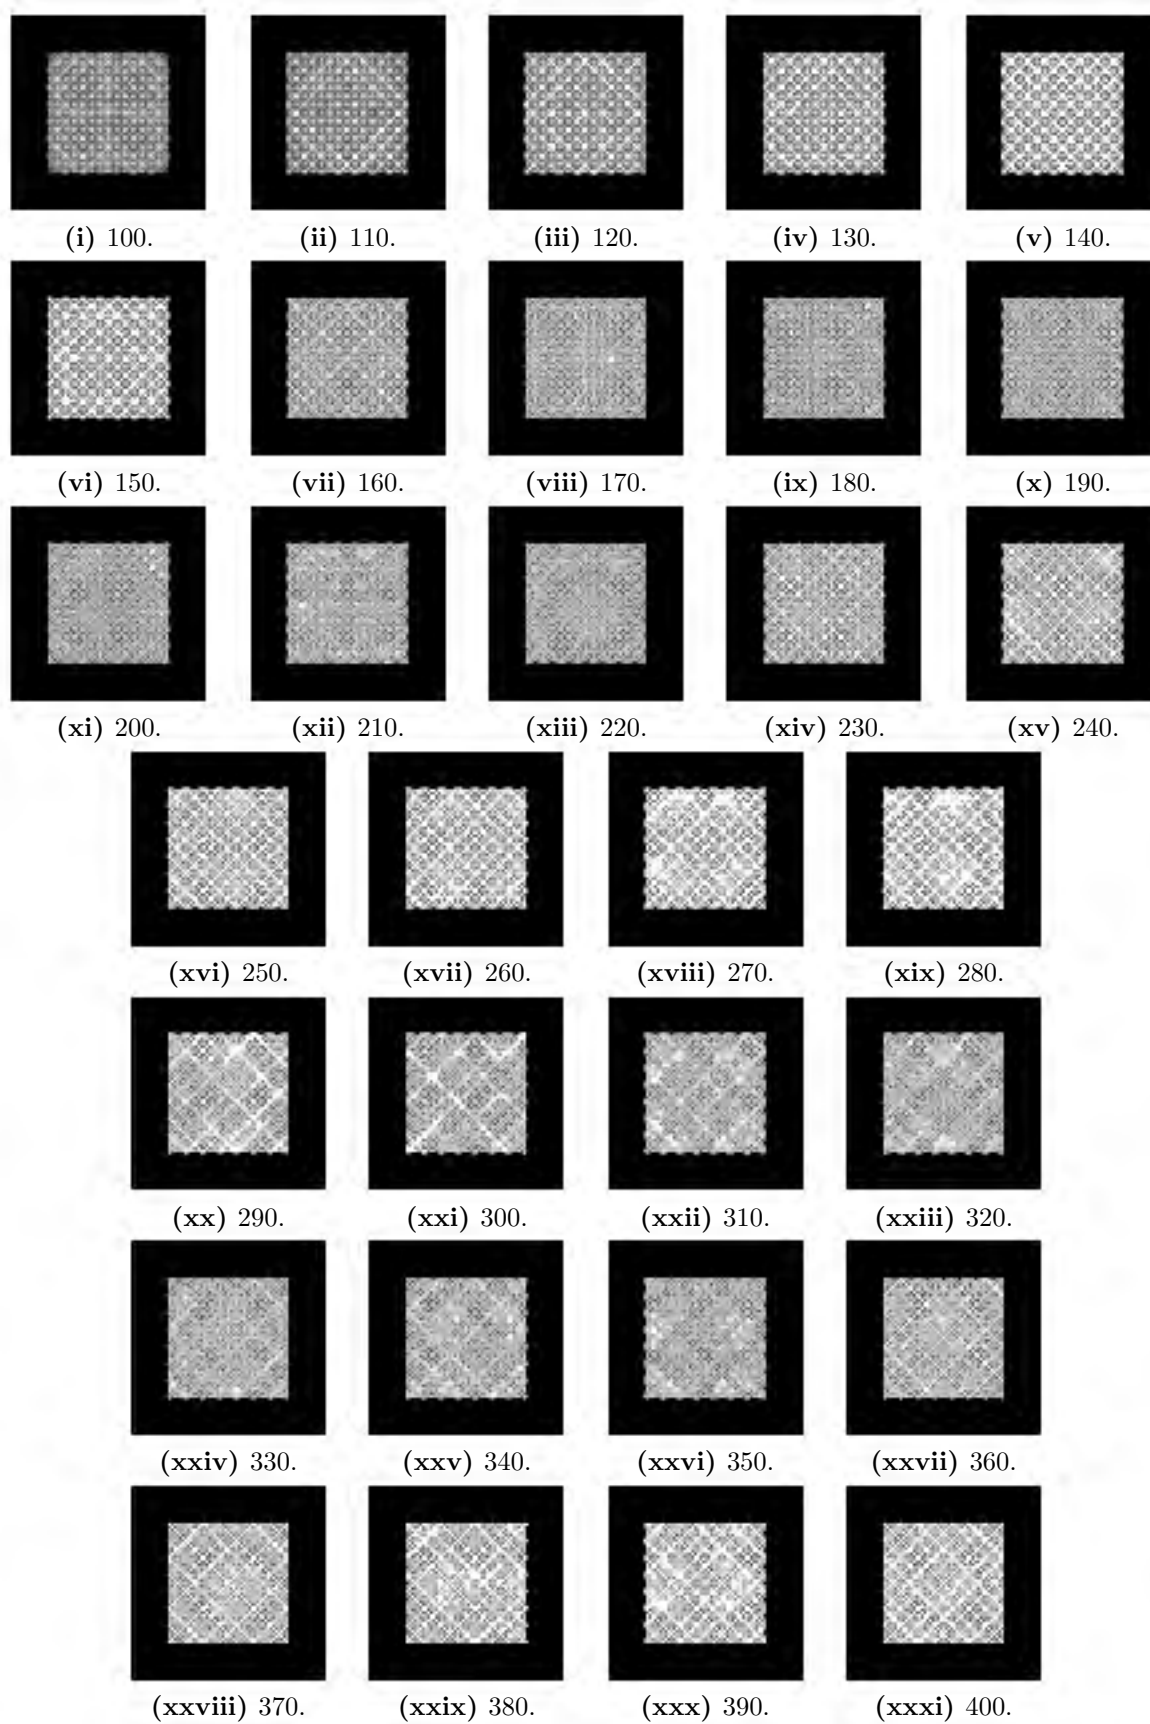


Figure F.21: Rasterisation of OT6-t3 lattice at varying image sizes.

F.3.2 Size 8

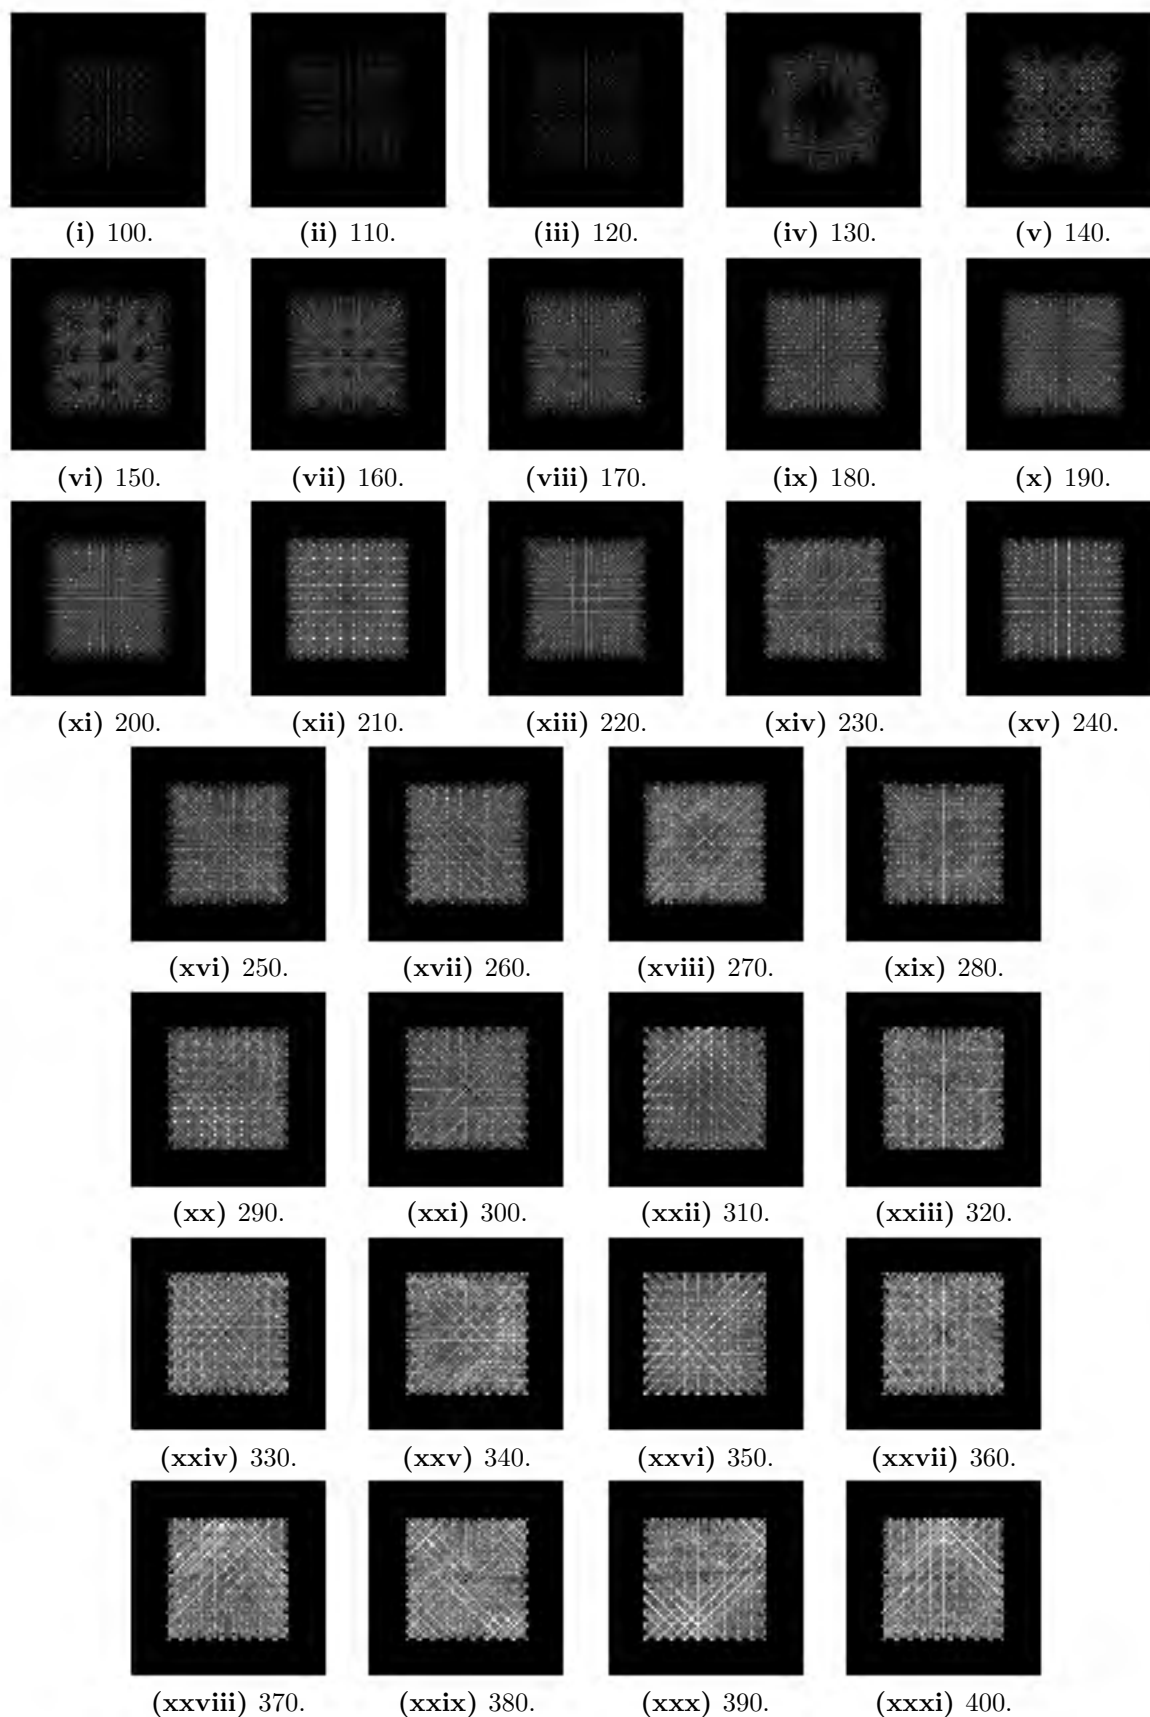


Figure F.22: Rasterisation of OT8-t1 lattice at varying image sizes.

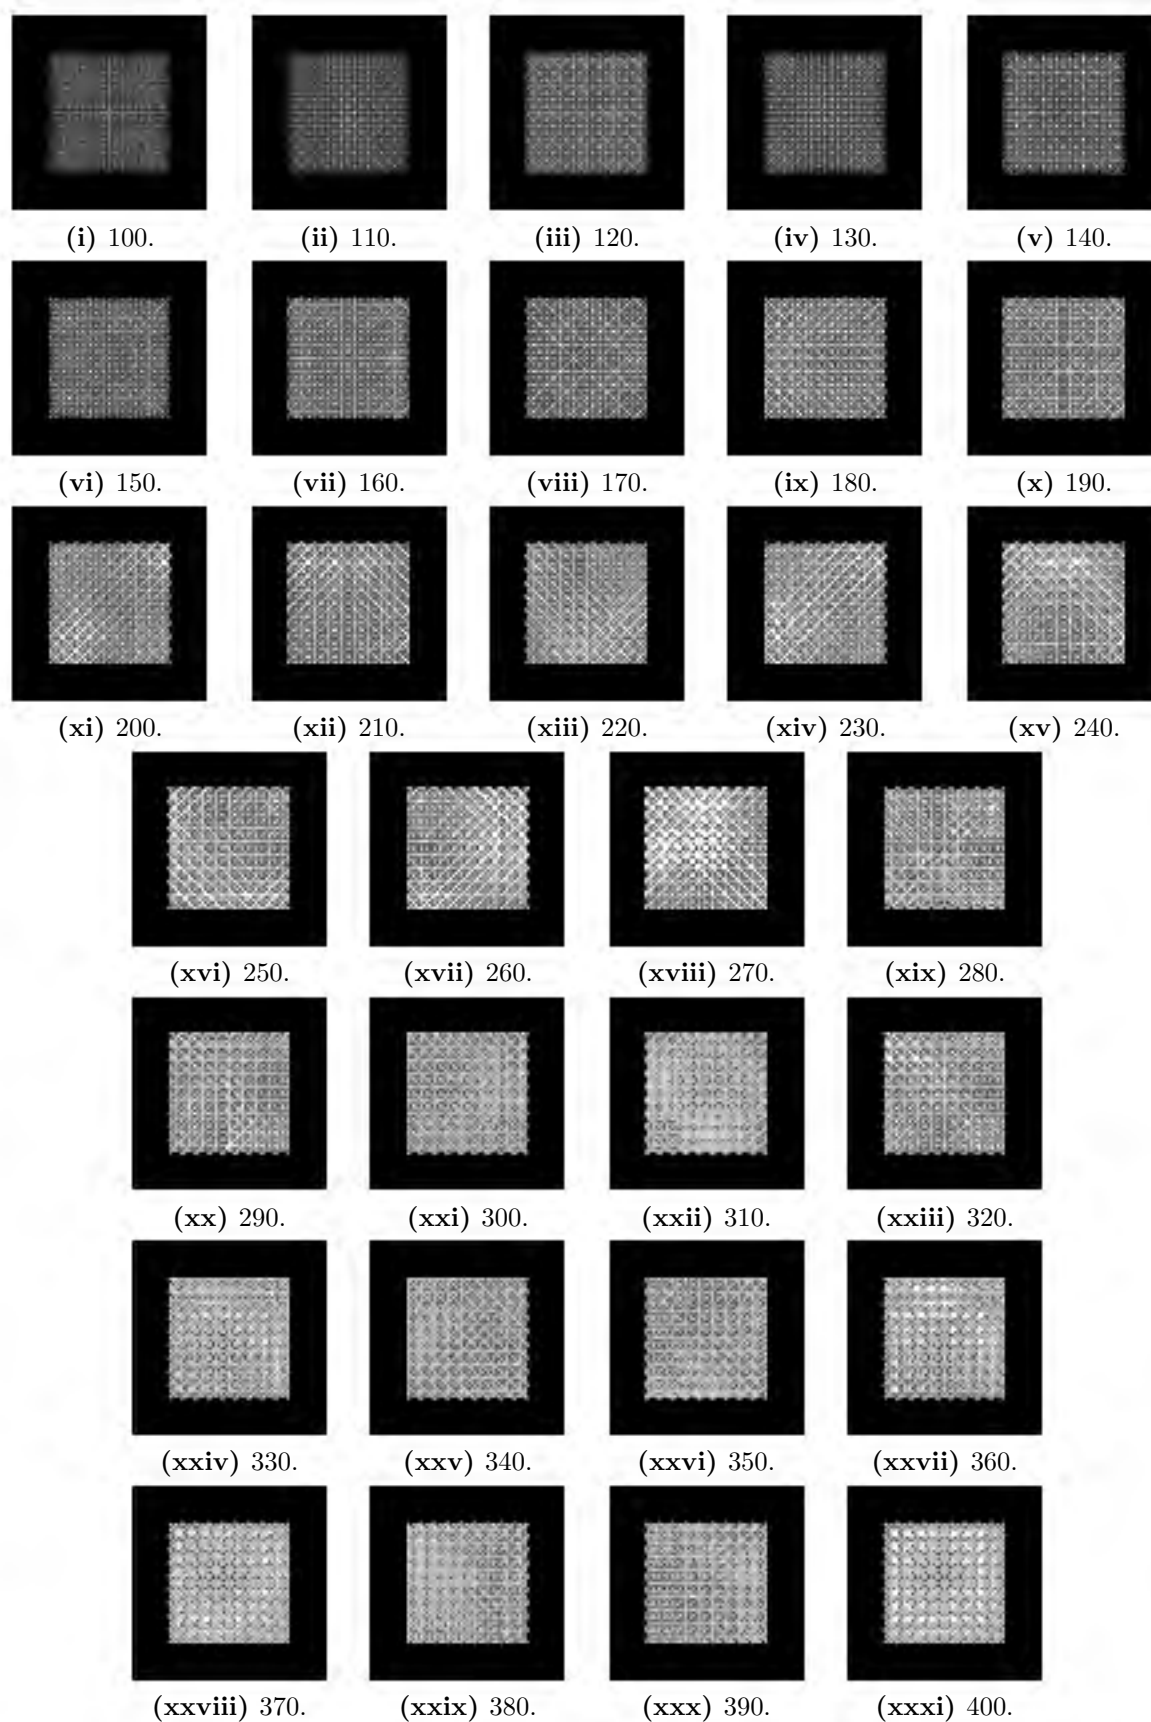


Figure F.23: Rasterisation of OT8-t2 lattice at varying image sizes.

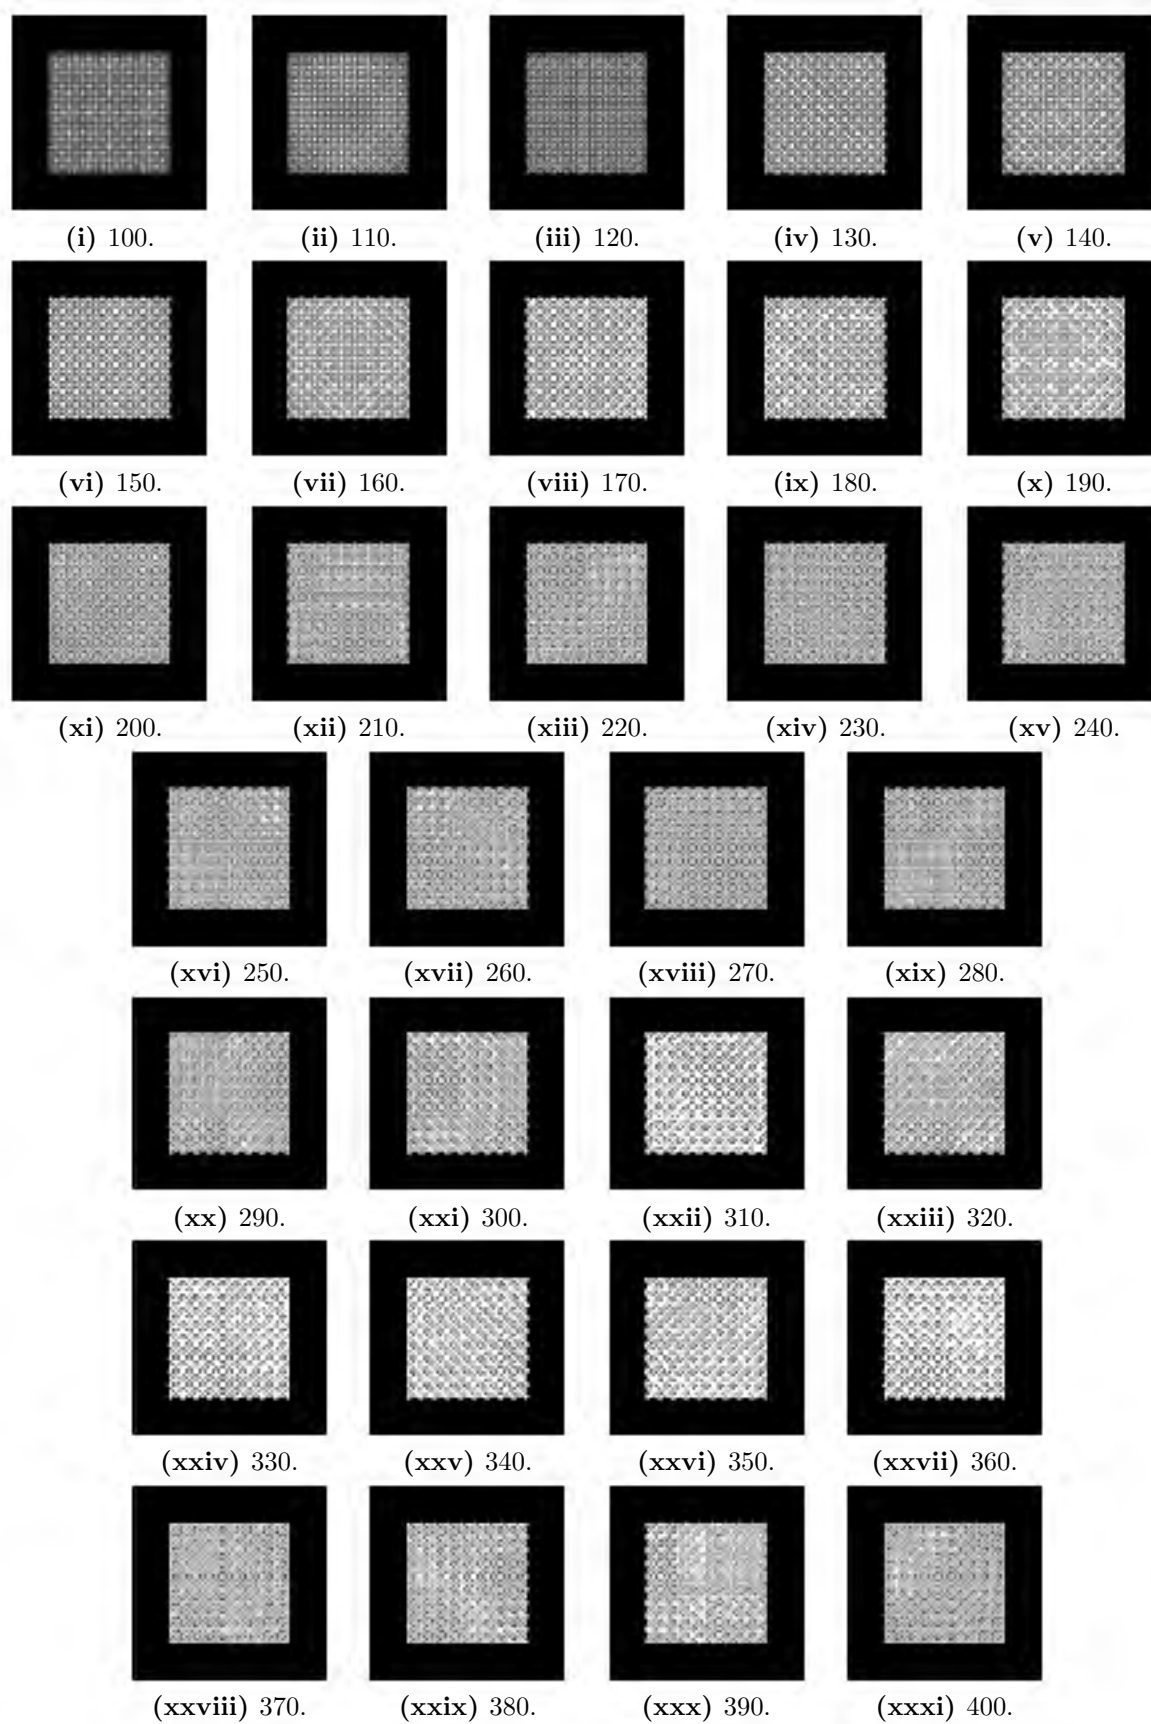


Figure F.24: Rasterisation of OT8-t3 lattice at varying image sizes.

F.3.3 Size 10

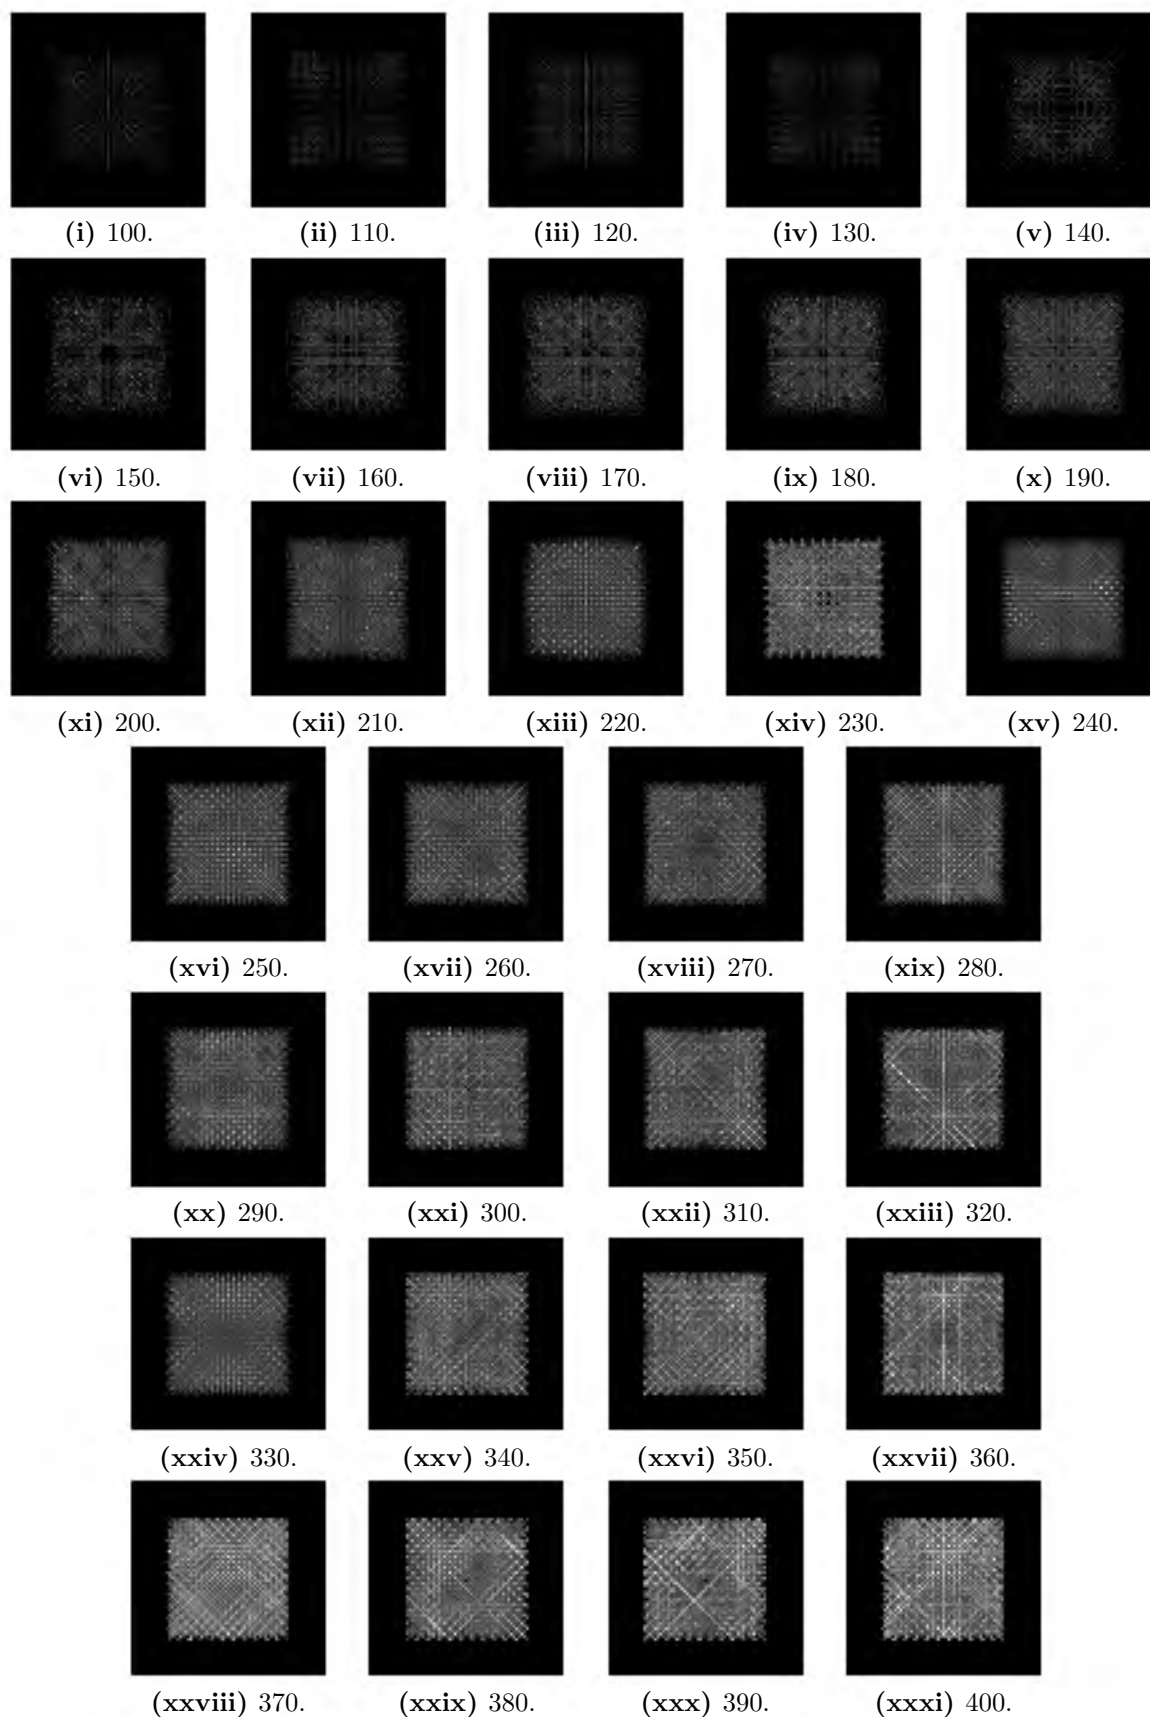


Figure F.25: Rasterisation of OT10-t1 lattice at varying image sizes.

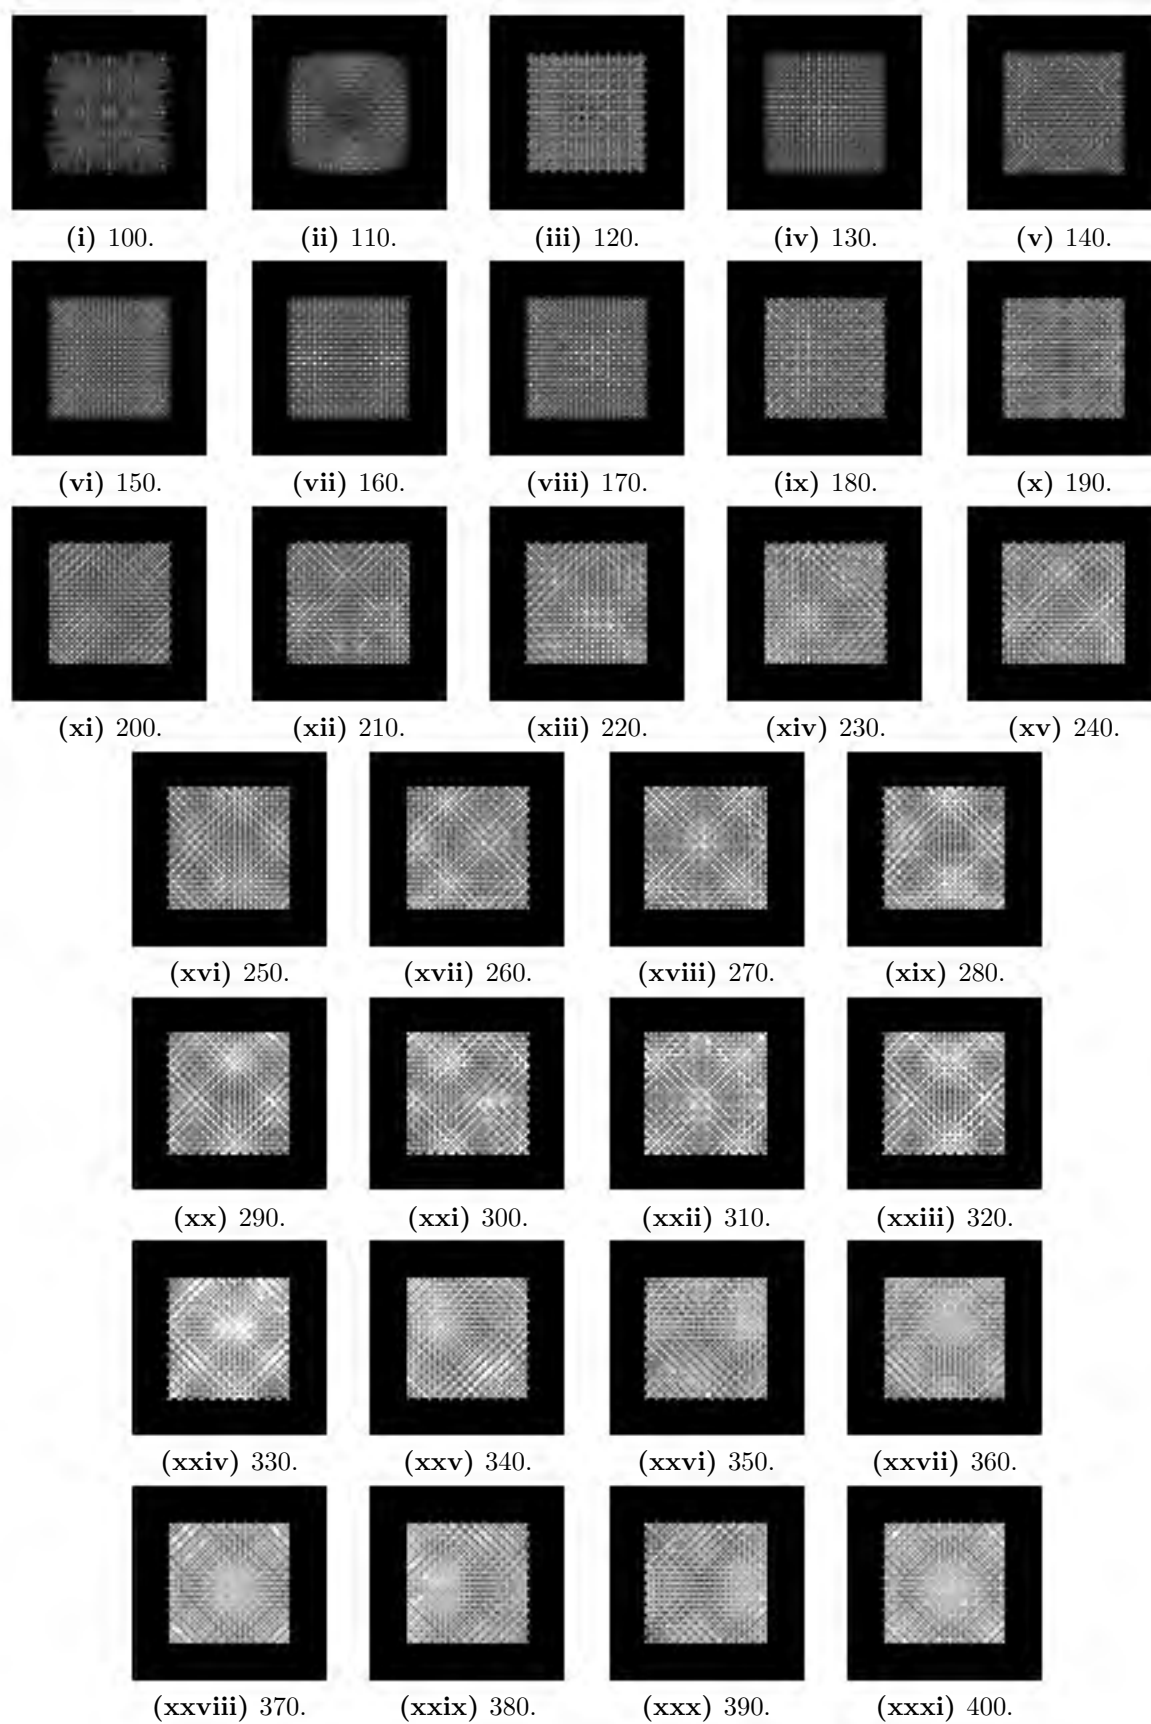


Figure F.26: Rasterisation of OT10-t2 lattice at varying image sizes.




Figure F.27: Rasterisation of OT10-t3 lattice at varying image sizes.

F.4 Angled cubic

F.4.1 Size 8

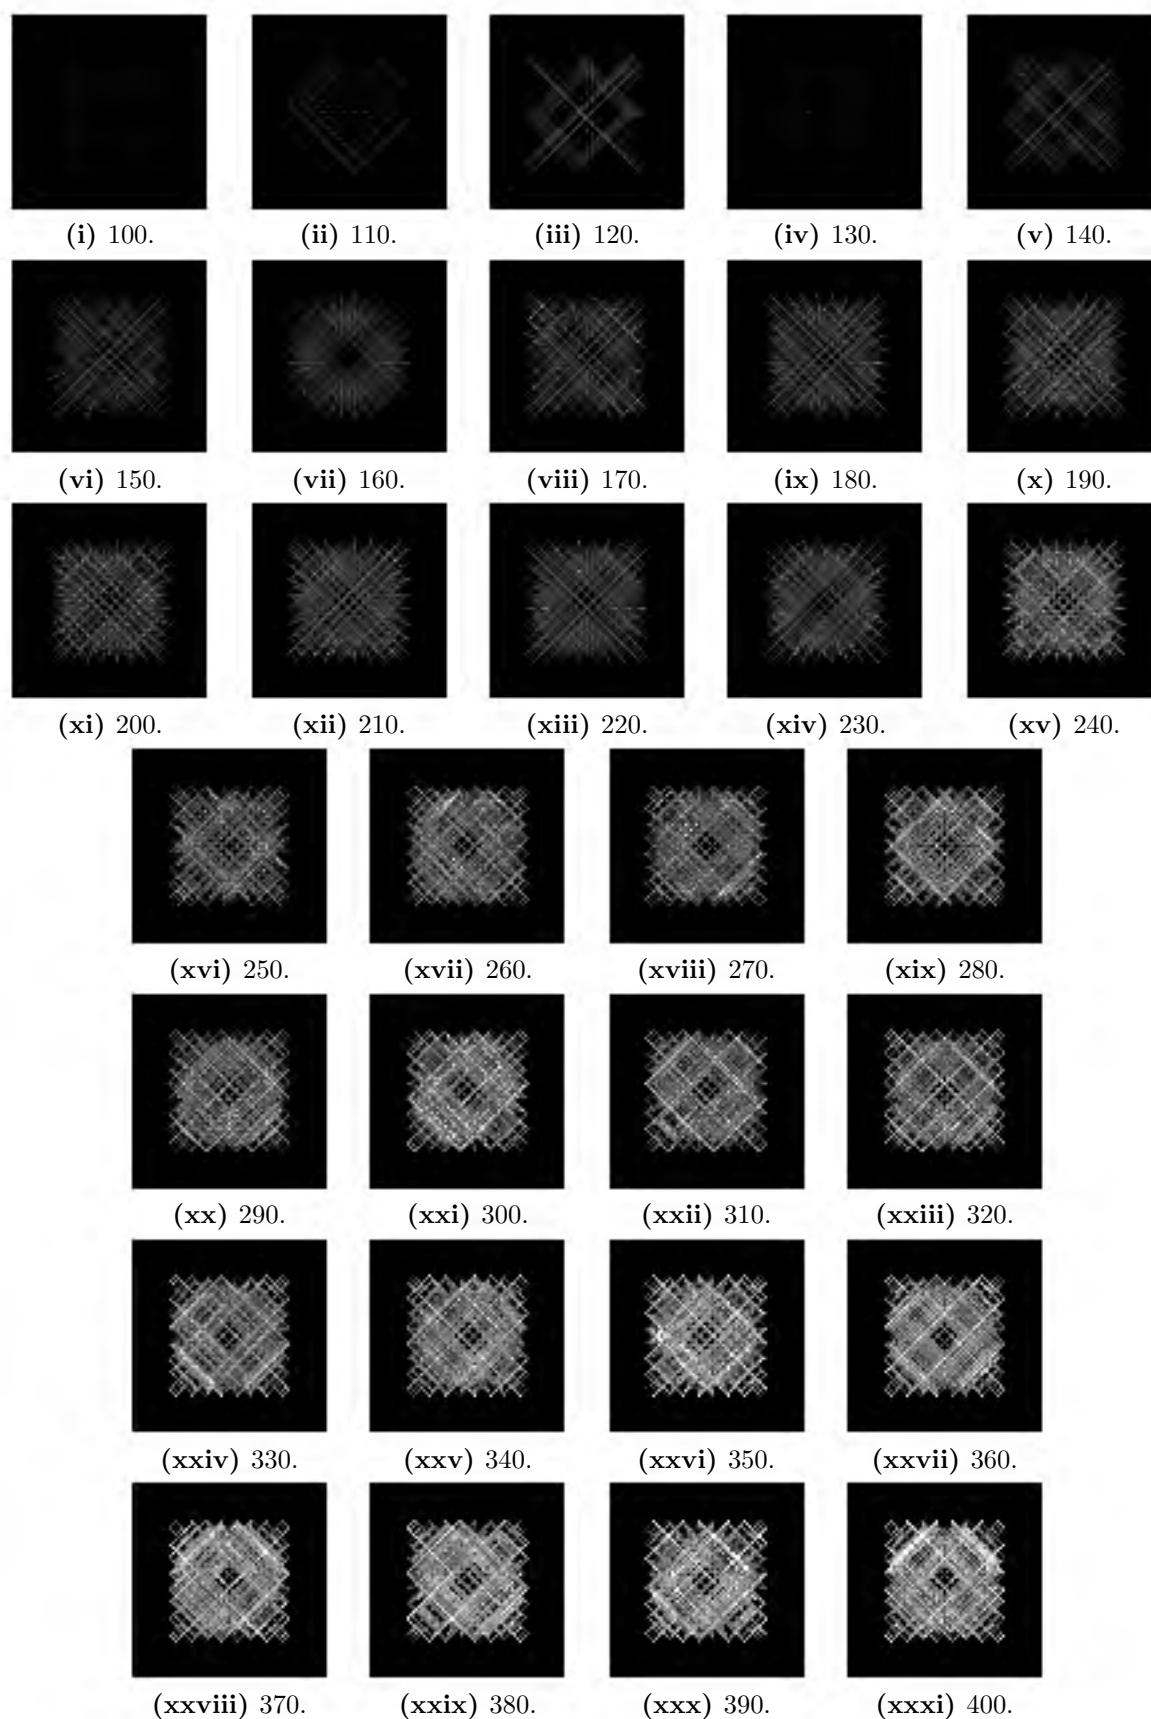


Figure F.28: Rasterisation of angled C8-t1 lattice at varying image sizes.

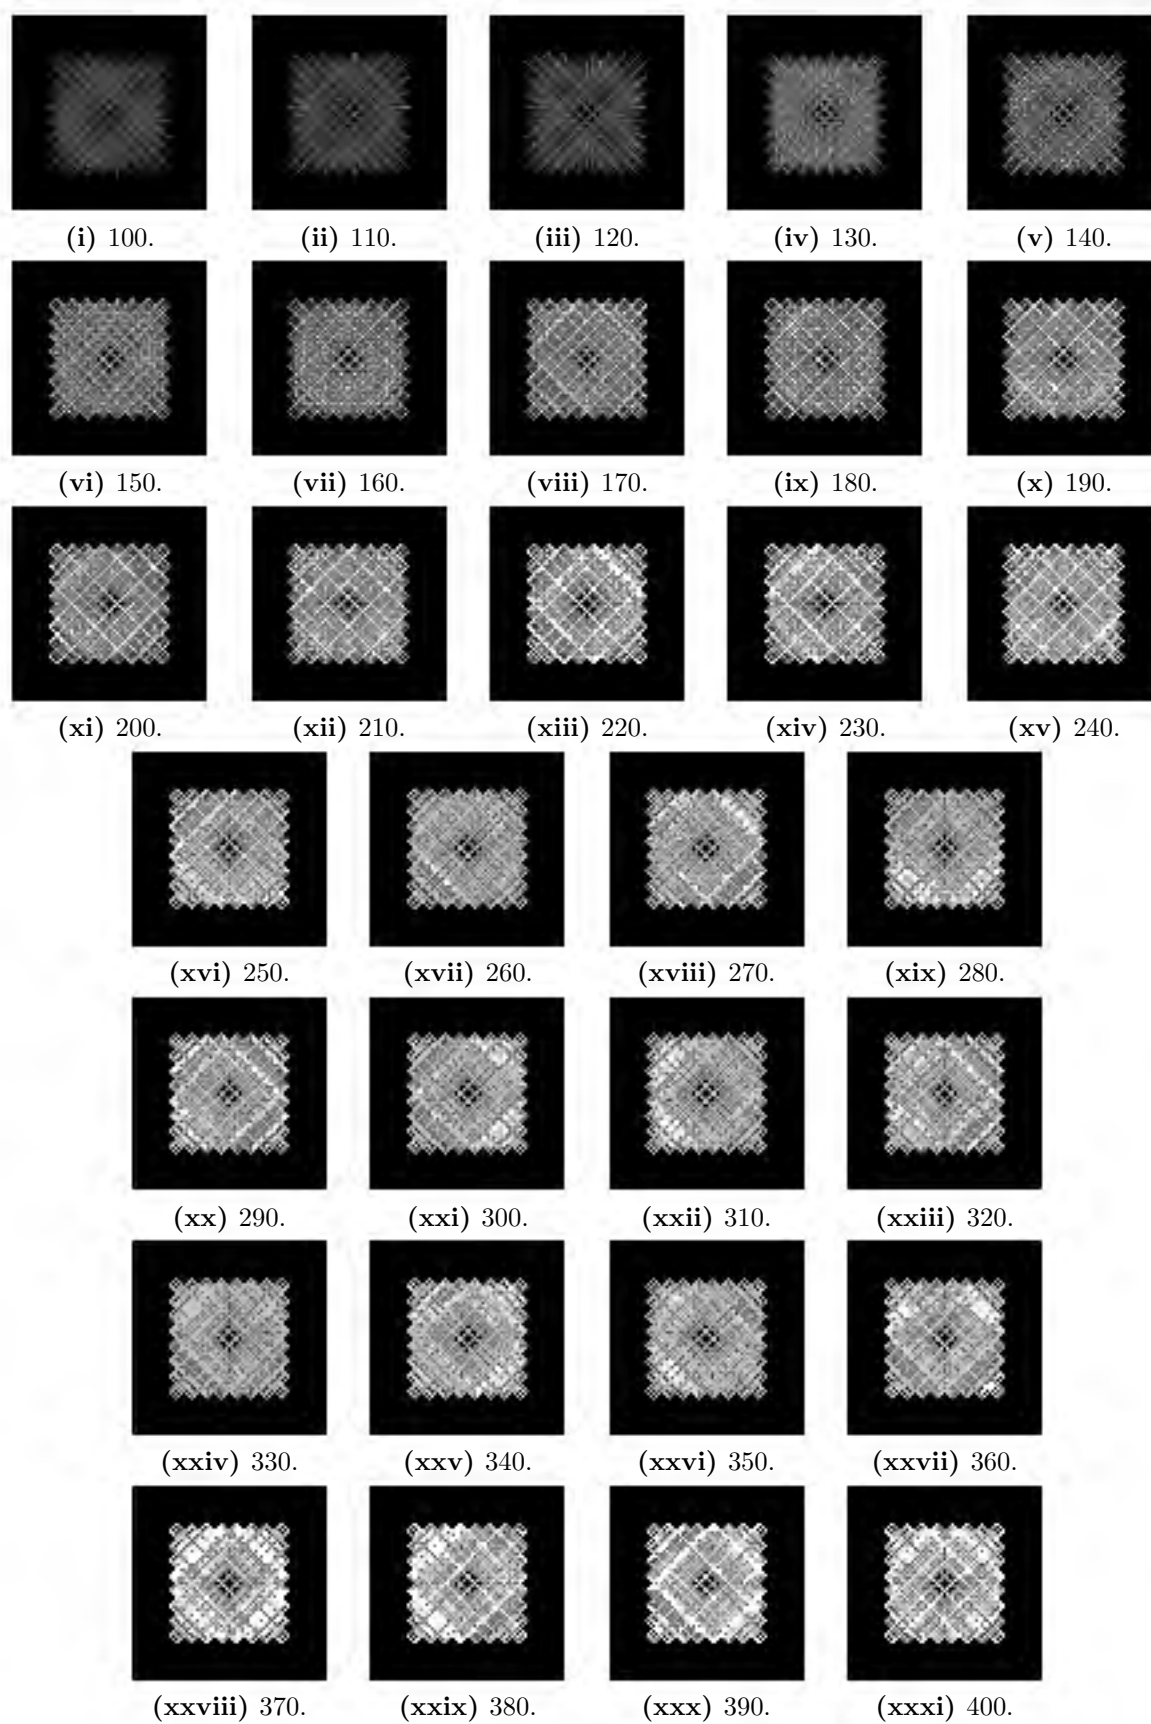


Figure F.29: Rasterisation of angled C8-t2 lattice at varying image sizes.

Appendix G

Inherent VBM boundary condition

In order to discuss the boundary condition inherent in the VBM mesh, it is important to visualise how a trabecula or strut is discretised using the BSM and the VBM. Figure G.1 shows a representation of a strut, with corresponding meshes created using the BSM and VBM.

If a translational constraint in the z -direction is applied to the beam-shell mesh, the strut is able to rotate while maintaining the applied constraint, as shown in Figure G.2. In contrast, if the same translational z -constraint is applied to the VBM, the bottom surface of the strut is prevented from tilting, as shown in Figure G.3, which causes a different deformation mechanism to what was seen in the BSM mesh.

Due to this inherent difference between the BSM and VBM meshes, it can be said that the VBM contains an additional boundary condition, which is inherent to the methodology. This additional constraint must be taken into account when comparing any VBM simulations to baseline results (as in the validation study), or when comparing VBM results to BSM results (as for the bone results).




Figure G.1: Discretisation of a strut using the BSM and VBM.




Figure G.2: Rotation of beam-shell mesh with translational z -constraint.

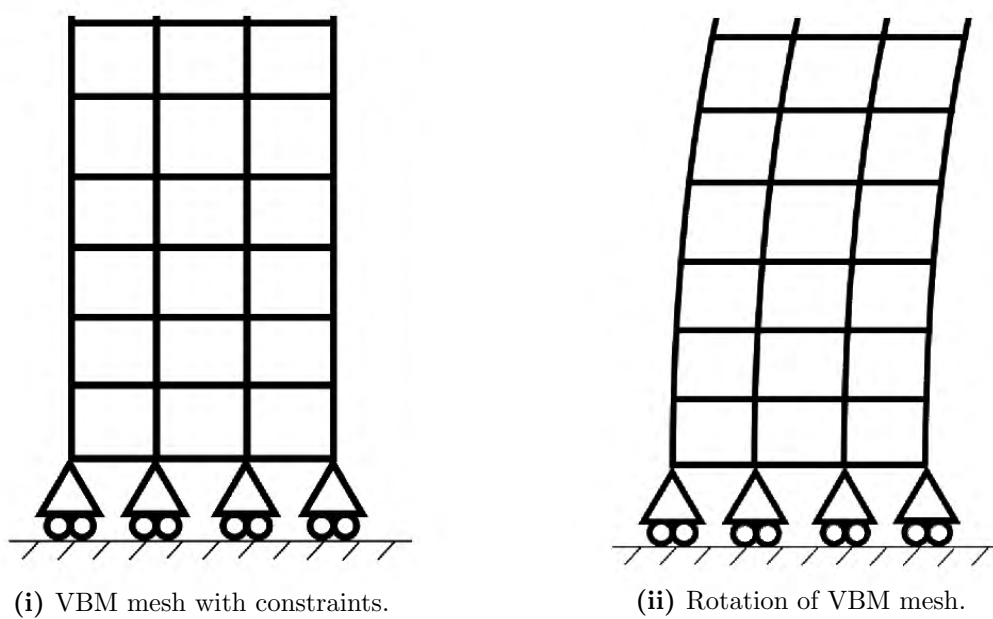


Figure G.3: Rotation of voxel-based mesh with translational z -constraint.

Appendix H

Bone specimens

The bone specimens used for this work were labelled numerically and are referred to as such in the text. The specimen numbers are listed below and grouped into their corresponding experimental test configuration, as discussed in Section 2.6.2.

Table H.1: List of bone specimens. Quantity of specimens is included next to the configuration name in brackets ().

Standard specimens			Large specimens	
Defatted (33)	Marrow (40)	Confined (33)	$\phi 20$ mm (9)	$\phi 28$ mm (8)
13	26	19	498	529
15	105	31	503	530
17	119	33	504	532
25	124	38	506	533
79	136	59	508	534
94	171	77	510	535
114	228	82	512	537
129	230	90	513	538
131	247	97	514	
135	248	137		
138	251	140		
152	270	150		
163	280	161		
165	309	177		
184	319	182		
198	332	205		
204	384	260		
208	402	261		
229	410	264		
255	424	283		
273	427	285		
274	451	324		
277	452	345		
314	453	358		
326	454	361		
366	455	365		
371	456	379		
390	457	385		
400	458	394		
404	461	399		
433	462	434		
438	464	435		
439	465	440		
	466			
	467			
	469			
	493			
	494			
	496			
	502			

Appendix I

Segmented Bone Scans




Figure I.1: Segmentation of bone scans (13–33).




Figure I.2: Segmentation of bone scans (38–124).




Figure I.3: Segmentation of bone scans (129–165).




Figure I.4: Segmentation of bone scans (171–247).




Figure I.5: Segmentation of bone scans (248–283).




Figure I.6: Segmentation of bone scans (285–366).




Figure I.7: Segmentation of bone scans (371–424).




Figure I.8: Segmentation of bone scans (427–455).




Figure I.9: Segmentation of bone scans (456–494).




Figure I.10: Segmentation of bone scans (496–529).




Figure I.11: Segmentation of bone scans (530–538).

Appendix J

Choice of line fit

Initially, a standard linear least squares line fit was used to show the trend between the experimental and simulation results, as displayed in Figure J.1. However, it is clear from the figure that this type of regression analysis does not provide the best fit to the data as the line gradient is too high. This occurs because the linear least squares line fit, $f(x)$, is determined by calculating and minimising the sum of the squares of the residual [110], where a residual is defined as the difference between the data value at a point and the line-fit value at that point, as shown in Equation J.1.

$$R^2 = \Sigma[y_i - f(x_i)]^2, \quad (\text{J.1})$$

This approach using the squared residuals causes the outlying results to be weighted disproportionately to the results which are close to the line [110], which leads to a skewed trend line, as is seen in Figure J.1.

Due to the limitations of the least squares regression, a different analysis technique was required to determine the line of best fit for the data. As previously noted, the gradient of the least squares line is too high. This occurs because the scatter in Figure J.1 is bounded by the x -axis below the trend line, but unbounded above the trend line. If the instantaneous gradients of the data points are calculated, where:

$$G = \frac{E_{sim}}{E_{exp}}, \quad (\text{J.2})$$

the instantaneous gradients will always be positive, as negative apparent moduli can not exist. If a histogram of the instantaneous gradients is plotted as shown in Figure J.2, the mentioned factors combine to produce a dataset which is clearly skewed to the right, similar to a log-normal distribution.

In order to determine if the instantaneous gradients form a log-normal distribution, the natural logarithms of the instantaneous gradients are found after adding a skewness correction constant, c , as shown in Equations J.3.

$$G_{new} = \ln(G + c) = \ln\left(\frac{E_{sim}}{E_{exp}} + c\right) \quad (\text{J.3})$$

If the histogram for the logarithmic data is considered as shown in Figure J.3, it may be seen that the distribution is similar to a normal distribution in shape, however it does not represent a true normal distribution. The mean (μ), median and mode values are within 10% of each




Figure J.1: Least-squares line fit for results from BSM analysis of standard bone specimens.




Figure J.2: Histogram of instantaneous gradients.




Figure J.3: Histogram of natural logarithms of instantaneous gradients.

other¹, which is similar to a normal distribution, in which these three values are equivalent [111]. Additionally, the coefficient of skewness is so small as to be negligible, which is characteristic of the symmetrical normal distribution [111]. However, 41% of the data points lie above the mean, whereas 59% of the data points lie below the mean, which does not correlate with the 50% : 50% split that is expected in a normal distribution.

The logarithmic data displays some attributes which are characteristic of a normal distribution and some attributes which are not, therefore the instantaneous gradients cannot be considered as a true log-normal distribution. However, treating the results as shown in Equation J.3 does correct some of the skewness of the data, which is the main problem associated with the linear least squares method, and hence was deemed useful. If the mean (μ_0), median (\tilde{x}_0) and mode (M_0) of the instantaneous gradients are calculated in the original space [112], as shown in Equations J.4-J.6, each value may be used to create a line of best fit, where the gradient of the line is designated by one of these values, as shown in Figure J.4. Because the mean and median are both affected by the upwards scatter, the mode is the most appropriate value to use in this context, as it provides the line fit which is closest to the majority of the data points, as shown in Figure J.4iii.


$$\mu_0 = e^{\mu+0.5\sigma^2} - c \quad (\text{J.4})$$

$$\tilde{x}_0 = e^{\mu} - c \quad (\text{J.5})$$


$$M_0 = e^{\mu-\sigma^2} - c \quad (\text{J.6})$$

¹The value of the mode changes drastically depending on the number of significant figures used for the calculation. Two significant figures produced the strongest mode, with a total of 10 data points, which was used in this work.


Although the instantaneous data points are not representative of a true log-normal distribution, determining the line of best fit by following this log-normal approach allows the outlying data points to be weighted less heavily than the “good” data points, which provides a more representative line fit than a least-squares approach.



(i) Mean.



(ii) Median.



(iii) Mode.

Figure J.4: Comparison of line fits for results from BSM analysis of standard bone specimens.