



UNIVERSITY OF LEEDS

Development and Evaluation of a Hexahedral Mesh-Morphing Strategy for the Knee Meniscus

Adam Glen Kelly

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

The University of Leeds
Faculty of Engineering
School of Mechanical Engineering

March 2022

Declaration

The candidate confirms that the work submitted is his own and that appropriate credit has been given where reference has been made to the work of others.

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

© 2022 The University of Leeds and Adam Glen Kelly

The right of Adam Glen Kelly to be identified as Author of this work has been asserted by him in accordance with the Copyright, Designs and Patents Act 1988.

Acknowledgements

I would like to express my gratitude to my supervisors, Dr. Marlène Mengoni and Dr. Alison C. Jones, for their guidance, expertise and insight, which have been invaluable. I am indebted to them for their time and support, and truly appreciate the opportunity to work in this field with them.

I would also like to thank Prof. Daniel Ruprecht, Prof. Alejandro F. Frangi, Dr. Claire Brockett and Prof. Ruth Wilcox, for the useful discussions over the course of my PhD. Thank you to Prof. Christophe Geuzaine and Prof. Jean-François Remacle, for providing novel perspectives on programming computational geometry algorithms, and inspiring me to dive into areas I thought were beyond my reach.

I would like to thank Prof. David M. Pierce, Dr. Borja Rodriguez-Vila and Dr. Gavin Day, for supplying and curating crucial data. The validation assessments in this thesis would have been insufficient if it was not for their assistance. Thanks also to my examiners Dr. Zeike Taylor and Prof. Damien Lacroix, and the iMBE team, especially Dr. Robert J. Cooper for his contributions during the early stages of my PhD. I would also like to acknowledge EPSRC, for providing financial support for this research.

Finally, thank you to my family and friends for their encouragement and patience, especially my mother, Patsy. I thank Kate for all her love and support over what became a most arduous journey. This PhD would not have come to fruition without her.

In loving memory of my father,

Paul Glen Kelly

Abstract

The state-of-the-art for generating accurate and high-quality hexahedral meshes of irregular geometries for finite element (FE) simulations is a laborious and time-consuming endeavour. Mesh-morphing is a technique which modifies the vertices of an existing mesh to match the boundary prescribed by another. This technique has been chosen as it has the capacity to automatically create accurate and high-quality hexahedral meshes of irregular geometries from a single template. The meniscus was chosen as it is an important component of the knee and challenging structure to simulate. Hexahedral discretisation is often required due to: multi-body contact, large deformations and complex material properties (nearly-incompressible and anisotropic). In this research, a novel, automatic and general-purpose mesh-morphing strategy was developed. Additionally, several robust and thorough methodologies were developed to assess the sensitivity and validate the performance of the mesh-morphing strategy. This was achieved through performance comparisons against a state-of-the-art procedure – the multi-block method (IA-FEMesh). The performance metrics not only assessed the capacity of the mesh-morphing strategy (i.e. speed, accuracy and mesh-quality), but also the functionality for meaningful FE simulations. The mesh-morphing strategy generated 20 challenging meniscus meshes in under a minute compared to an average of 26 minutes, with comparable surface error and mesh-quality metrics to the state-of-the-art. Also, there was no significant difference in the FE simulation outcomes. The mesh-morphing strategy offers a faster, competitive and automated alternative to the semi-automatic state-of-the-art, and only required one template mesh. The strategy can already be applied to a diverse range of geometries, and with some trivial modifications can approach a larger proportion. The developed tool can be used to improve productivity and automate the development of FE models. This could enable the design of large-scale studies and assist the development of digital twins. A wide-range of industries that require automatic and accurate meshing from 3D scanning technologies could utilise this work.

Contents

Declaration	i
Acknowledgements	ii
Dedication	iii
Abstract	iv
Contents	v
List of Figures	xii
List of Tables	xx
I Background and Development	1
1 Introduction	2
1.1 Background	2
1.1.1 Hexahedral meshing in musculoskeletal finite-element models	2
1.1.2 Mesh-Morphing	3
1.1.3 Meniscus of the knee joint	4
1.2 Aims	5
1.3 Novelty	6
1.4 Thesis Structure	7
2 Literature Review	9
2.1 Meniscus Anatomy and Importance	9
2.1.1 Tissue Structure	9
2.1.2 Tissue Function	11
2.2 Finite Element Modelling of the Tibio-Femoral Joint	13
2.2.1 Purpose	13
2.2.2 Tissue Segmentation and Meshing	15
2.2.3 Parameters of Interest	16

2.2.4	Enforcement of Congruency	17
2.2.5	Clinical Applications and Relevant Outputs	18
2.3	Hexahedral versus Tetrahedral Meshes	21
2.3.1	Element Types	21
2.3.2	Performance of Element Types	23
2.3.3	Comparative Analyses	27
2.3.4	Mesh Types used in the Tibiofemoral Joint	31
2.3.5	Element Quality Metrics	33
2.4	Hexahedral Mesh Generation	36
2.4.1	Overview	37
2.4.2	Advantages and Disadvantages	43
2.4.3	Summary and Conclusions	50
2.5	Mesh-Morphing Algorithms	51
2.5.1	Feature Alignment	52
2.5.2	Surface Parameterisation	54
2.5.3	Volumetric Transformation	66
2.5.4	Conclusions	87
2.6	Applications of Mesh-Morphing	88
2.6.1	Applications	88
2.6.2	Feature Alignment	89
2.6.3	Surface Parameterisation	90
2.6.4	Volumetric Transformation	93
2.6.5	Summary	93
2.7	Conclusions	95
3	Mesh-Morphing Strategy Design	99
3.1	Overview	99
3.2	Centreline Generation	101
3.2.1	Generation of a Signed Distance Function from Radial Basis Functions	101
3.2.2	Minimisation of Signed Distance Function to Centreline	104
3.2.3	Efficiency Optimisation of Centreline Minimisation	106
3.2.4	Clustering	111

3.2.5	B-Spline Parameterisation	113
3.3	Feature Alignment	114
3.4	Surface Parameterisation	117
3.4.1	Centreline-based Projection	117
3.5	Surface Optimisation	120
3.5.1	Laplacian Smoothing	121
3.5.2	Quadrilateral Inversion Detection	121
3.5.3	Quadrilateral Unfolding	122
3.6	Volumetric Transformation	123
3.6.1	Linear Elastic Transformation	124
3.6.2	Laplacian Smoothing	125
3.7	Mesh-Morphing Strategy	125
3.8	Mesh-Morphing Strategy Optimisation	127
4	Performance Metrics and Data	133
4.1	Surface Error Analysis	133
4.1.1	Surface Error	133
4.1.2	Volume Error	135
4.2	Mesh Quality Analysis	136
4.3	Hexahedral Mesh Generation	138
4.3.1	Meniscus	138
4.3.2	Tibial Cartilage	140
4.3.3	Femoral Cartilage	141
4.4	Target Geometries	141
4.5	Finite Element Analysis	144
4.5.1	Tibio-Femoral Joint Model	144
4.5.2	Model Geometry	145
4.5.3	Loading and Boundary Conditions	145
4.5.4	Material Properties	146
4.5.5	Contact Variables	149
4.5.6	Simulation Convergence Rates	150
4.5.7	Mesh Convergence Analysis	150

4.6	Software and Libraries	151
II	Sensitivity Analysis of the Mesh-Morphing Strategy	152
5	Case A: Centreline Quality	153
5.1	Introduction	153
5.2	Methodology	154
5.2.1	Centrelines	154
5.2.2	Template Mesh	156
5.2.3	Target Mesh	156
5.2.4	Outputs of Interest	157
5.3	Results	158
5.3.1	Surface Error Analysis	158
5.3.2	Mesh-Quality Analysis	160
5.4	Discussion	163
5.4.1	Geometry Changes	163
5.4.2	Effect on Mesh-Morphing Strategy	164
5.4.3	Consequences for Finite Element Modelling	166
5.4.4	Improvements and Optimisation of the Centreline Algorithm	166
5.4.5	Limitations	167
5.5	Conclusions	168
6	Case B: Geometric Robustness	169
6.1	Introduction	169
6.2	Methodology	170
6.2.1	Template Mesh	170
6.2.2	Target Meshes	171
6.3	Results	172
6.3.1	Surface Error Analysis	173
6.3.2	Mesh-Quality Analysis	175
6.4	Discussion	178
6.4.1	Template Quality Improvements	178
6.4.2	Effect of Geometric Variations	179

6.4.3	Consequences for Finite Element Modelling	182
6.4.4	Limitations	183
6.5	Conclusions	184
7	Case C: Mesh Sensitivity	185
7.1	Introduction	185
7.2	Methodology	187
7.2.1	Template Meshes	187
7.3	Results	193
7.3.1	Surface Error Analysis	194
7.3.2	Mesh-Quality Analysis	196
7.4	Discussion	200
7.4.1	Sensitivity to Template Mesh Resolution	200
7.4.2	Sensitivity to Template Mesh Quality	202
7.4.3	Sensitivity to Template Mesh Topology	202
7.4.4	Sensitivity to Template Mesh Geometry	204
7.4.5	Consequences for Finite Element Modelling	205
7.4.6	Limitations	205
7.5	Conclusions	208
III	Performance Evaluation	209
8	Case D: Performance Benchmark	210
8.1	Introduction	210
8.2	Methodology	211
8.2.1	Target Meshes	211
8.2.2	Template Mesh	215
8.2.3	Generation of Meniscus Meshes using the Multi-Block Method	215
8.2.4	Generation Times	215
8.3	Results	218
8.3.1	Generation Times	218
8.3.2	Surface Error Analysis	219
8.3.3	Mesh-Quality Analysis	222

8.3.4	Finite Element Analysis	224
8.3.5	Detailed Analysis	225
8.4	Discussion	235
8.4.1	Generation Times	235
8.4.2	Surface Error	235
8.4.3	Mesh-Quality	237
8.4.4	Limitations	239
8.4.5	Multi-Block Method as Gold Standard	240
8.4.6	Causes for Key Biomechanical Differences	241
8.4.7	Consequences of the Mesh-Morphing Strategy	245
9	Case E: Performance Optimisation	248
9.1	Introduction	248
9.2	Methodology	249
9.2.1	Overall Performance	249
9.2.2	Qualitative Analysis of Finite Element Outcomes	250
9.3	Results	250
9.3.1	Surface Error Analysis	250
9.3.2	Generation Times	252
9.3.3	Mesh-Quality Analysis	254
9.3.4	Overall Performance	255
9.3.5	Finite Element Analysis	256
9.3.6	Detailed Analysis	258
9.4	Discussion	265
9.4.1	Surface Error	265
9.4.2	Generation Times	267
9.4.3	Mesh-Quality	268
9.4.4	Consequences of the Optimised Strategy	270
9.5	Conclusions	274
IV	Discussion	275
10	Discussion	276

10.1 Summary	276
10.2 Development and Future Work of Morphing Strategies	277
10.2.1 Centreline Algorithm	277
10.2.2 Mesh-Morphing Strategy Improvements and Future Work	279
10.3 Development and Impact of Methodologies	281
10.3.1 Centreline-based Surface Error Measurement	282
10.3.2 Sensitivity Evaluation of a Mesh-Morphing Strategy	283
10.3.3 Validation Methodologies to Benchmark and Assess Performance	286
10.4 Performance and Impact of Mesh-Morphing Strategies	288
10.4.1 Single Template Mesh-Morphing for Different Anatomies	288
10.4.2 Rapid and Automatic Generation of Complex Simulations	290
10.4.3 Application and Impact of the Mesh-Morphing Strategies	291
10.5 Conclusion	293
Bibliography	294
Appendices	343
A Meniscus Target Geometries	344
B Case D: Summary of Results	355
C Case E: Summary of Results	357

List of Figures

1.1	Stages of FE model development: (a) image acquisition, (b) image segmentation, (c) surface meshing, (d) volumetric meshing, (e) application of boundary conditions, and (f) solution of FE model	2
1.2	Shape and cross-section of a meniscus anatomy (Source: LM_00 [Table A.1]) . .	5
2.1	Meniscus Anatomy: (a) medial and lateral meniscus with associated horn-attachments and ligaments (<i>reproduced from</i> [1]) (b) meniscus cross-section and associated zones (<i>adapted from</i> [2])	9
2.2	Pathologies of the meniscus (arrows show the progressive deterioration)	12
2.3	Common components used to model the tibiofemoral joint (Source: open-knee model [3])	14
2.4	Common linear elements: (a) triangle, (b) quadrilateral, (c) tetrahedral and (d) hexahedral Elements. Integration points are shown for the 3D elements (black) .	22
2.5	Less common elements: (a) a pentahedron and (b) pyramid	23
2.6	Common quadratic elements: (a,b) tetrahedral and (c,d) hexahedral variations .	24
2.7	Frequency of element types used for tissues in tibiofemoral joint models	32
2.8	Ideal or reference hexahedron used to compare against a hexahedral element with respect to the Jacobian, principal axes (P1, P2, P3), dihedral angles (α , β and γ) and ratio of minimum and maximum edge lengths (<i>adapted from</i> [4,5])	34
2.9	A piecewise bijective map (Ψ) of a non-closed ear mesh (\mathcal{M}) and the corresponding planar parameterisation (\mathcal{U}), the red triangle highlights isomorphism (<i>reproduced from</i> [6])	55
2.10	A one-ring neighborhood: a central (\mathbf{v}_i) and adjacent vertex (\mathbf{v}_j) with associated angles (α , β , δ and γ) used to define particular spring weights (<i>adapted from</i> [7])	56
2.11	Interpolation of two meshes using spherical parameterisation (<i>adapted from</i> [8]) .	57
2.12	Planar cross-parameterisation between two non-closed femoral head meshes, which was used to morph a tetrahedral mesh. Linear (Ψ and $\hat{\Psi}$) and inverse maps (Ψ^{-1} and $\hat{\Psi}^{-1}$) could be used to determine direct correspondences between meshes (<i>adapted from</i> [9] ©2008 IEEE)	58

2.13 Spherical cross-parameterisation between two closed genus-0 meshes, used to morph and interpolate triangular meshes. Linear (Ψ and $\hat{\Psi}$) and inverse maps (Ψ^{-1} and $\hat{\Psi}^{-1}$) could be used to determine direct correspondences between meshes (<i>adapted from [8]</i>)	61
2.14 Spherical parameterisation of a closed genus-0 brain mesh used to improve visualisation of white-matter curvatures (<i>reproduced from [10] ©2000 IEEE</i>)	62
2.15 A general transformation operation (Ψ): morphing a template mesh (\mathcal{M}) to define a new mesh (\mathcal{M}') that resembles a target geometry ($\partial\Omega'$)	66
2.16 Linear elastic transformation (Ψ) with Dirichlet (Γ_D) and Neumann boundary conditions (Γ_N)	73
2.17 A general Laplacian smoothing operation (\mathcal{L}): improving element quality by moving a vertex (\mathbf{v}_i) towards the barycentric average	79
2.18 (a) Laplace-Beltrami operator (Δ) for a vertex \mathbf{v}_i on a surface and (b) discrete Laplacian Operator (\mathcal{L}) for the corresponding vertex on a surface mesh (<i>adapted from [7]</i>)	80
3.1 Flow chart of the input data and algorithms that compose the (a) mesh-morphing strategy and (b) optimised mesh-morphing strategy	100
3.2 The RBF-based signed distance function is defined by three layers: the original surface (black, $h(\mathcal{S})$), expanded surface (purple, $h(\mathcal{S}^+)$) and contracted surface (blue, $h(\mathcal{S}^-)$). Points inside the boundary have a negative distance value, and vice versa	102
3.3 The minimisation of the signed distance function collapses a triangular surface mesh towards its centreline (a), with repeated iterations the centreline begins to contract (b)	104
3.4 Edge collapse of the mesh simplification algorithm from \mathbf{v}_i to \mathbf{v}_m . The cost function measures the (a) modified Gaussian curvature (Equation 3.57) and (b) aspect ratio (Equation 3.61) from the associated neighbourhoods before and after a potential edge collapse, respectively	111

3.5	Following the collapse of a triangular surface mesh into a centreline (a) clustering is performed on the vertices which merges the associated edges (b). The shortest edges are selected to determine the adjacent clusters, which are unordered (c). Dijkstra's shortest path algorithm is then used to determine the correct order for B-spline parameterisation (d)	112
3.6	Incremental process to find the closest point on a centreline to a surface vertex .	114
3.7	Ray-triangle intersection was used to project each template surface vertex onto the target	119
3.8	Overview of volumetric transformation: (a) template mesh (blue) and target geometry (red), (b) boundary morphing of template and (c) volumetric transformation to derive a morphed mesh (purple)	124
4.1	Ray-triangle intersection was used to determine correspondences by projecting each target vertex onto the boundary of a hexahedral mesh	134
4.2	The (a) cross-sectional and (b) transversal (superior) view of a coarse template mesh (no subdivisions), with the B-spline curve of the cross-section shown in black	139
4.3	Cross-sectional (a-c) and superior (d-f) view of the (a,d) low-, (b,e) medium- and (c,f) high-resolution meshes	140
4.4	Superior view of the idealised tibial cartilage, showing the superior surface (red) and tibial bone interface (white)	141
4.5	Idealised femoral cartilage hexahedral mesh: (a) sagittal and (b) coronal view, showing the inferior surface (red) and femoral bone interface (white)	142
4.6	The effect of the artefact reduction algorithm: (a) a high-curvature artefact before and (b) after identification and smoothing	143
4.7	A simplified finite element model for a single condyle of a tibio-femoral joint . .	145
4.8	The definition of collagen fibre orientations for the hexahedral meniscus meshes using local element coordinates, with \mathbf{a}_0 the unit vector defining the fibre orientation	148
5.1	The (a,d) low-, (b,e) medium- and (c,f) high-quality centrelines with respect to length (a-c) and noise (d-f)	155
5.2	The superior view of (a) the target geometry and (b) baseline	158
5.3	The smoothing effect causes (a) sharper features of a target mesh to (b) become rounded in a morphed mesh, (c) which causes regions of over- and under-estimation	159

5.4	Distribution of large surface errors for meshes derived from centrelines with varying (a) length and (b) noise qualities	160
5.5	Superior view of surface errors (amalgamated above 5 mm) on meshes derived from (a) low-, (b) medium- and (c) high-quality length centrelines (red is worse and blue is better)	160
5.6	Distribution of low-quality elements for meshes derived from centrelines with varying (a) length and (b) noise qualities	162
5.7	Superior view of the hexahedral scaled-Jacobian on meshes derived from (a) low-, (b) medium- and (c) high-quality length centrelines (red is better and blue is worse)	162
5.8	Cross-sectional view of the hexahedral scaled-Jacobian on meshes from (a) low-, (b) medium- and (c) high-quality length centrelines (red is better and blue is worse)	162
5.9	The meshes derived from shorter centrelines have lower densities (red circle) around the horn-attachments and higher densities (orange circle) away from this region compared (a) to the baseline that had a more even distribution of densities (green circles) (b)	165
5.10	The corners of the template's horn-attachments are sharper (a) than the target's (b), where the large mismatch between angles ($\alpha_M \ll \alpha_T$ and $\beta_M \ll \beta_T$) causes greater correspondence errors and element distortions in the initial projection (c)	165
6.1	The superior view of (a) the target geometry \mathbf{T}_L^+ and (b) the morphed mesh with smoothed features at the horn-attachment corners circled	173
6.2	Inferior view of surface errors on meshes derived from transverse length changes at (a) +20% and (b) -20%, and width changes at (c) +20% and (d) -20% (red is worse and blue is better)	174
6.3	Inferior (a-d) and superior (e-h) view of surface errors on meshes derived from cross-sectional width changes at (a,e) +20% and (b,f) -20%, and height changes at (c,g) +20% and (d,h) -20% (red is worse and blue is better)	175
6.4	Superior view of the hexahedral scaled-Jacobian on meshes derived from transverse length changes at (a) +20% and (b) -20%, and width changes at (c) +20% and (d) -20%. Cross-sectional width changes at (e) +20% and (f) -20%, and height changes at (g) +20% and (h) -20% (red is better and blue is worse) . . .	177

6.5	Cross-sectional view of the hexahedral scaled-Jacobian on meshes derived from transverse length changes at (a) +20% and (b) -20%, and width changes at (c) +20% and (d) -20%. Cross-sectional width changes at (e) +20% and (f) -20%, and height changes at (g) +20% and (h) -20% (red is better and blue is worse) .	177
6.6	Influence of cross-sectional width and height variations on the angle of the wedge ω and upper rim v (superior-periphery): (a) the average geometry of the template mesh, with the (b) increase and (c) decrease of width, and the (d) increase and (e) decrease of height	181
6.7	Cause of correspondence errors for the cross-sectional variations: (a) different centreline shapes between the target (blue) and template (yellow) meshes, (b-c) similarity of correspondence errors between variations with equivalent width-to-height ratios	181
6.8	Correspondence errors associated with different horn-attachment angles (α , β and γ) and centreline shapes: (a) between the average geometry of the template mesh (blue, β) and the variations which (b) decrease (green, α) and (c) increase (red, γ) the transverse length-to-width ratio	182
7.1	Superior (a-c) and cross-sectional (d-f) view of the (a,d) low-, (b,e) medium- and (c,f) high-resolution meshes	188
7.2	Cross-sectional view of the (a) low-, (b) medium- and (c) high-quality meshes . .	189
7.3	Internal (dashed-line) and external (solid-line) vertex singularity differences between the (a) CAD-Swept and (b) CAD-MB meshes	191
7.4	Superior view of the (a) CAD-MB and (b) anatomical-MB meshes	192
7.5	Superior view of surface errors (amalgamated above 3 mm) for the (a) low-, (b) medium- and (c) high-resolution meshes, and the (e) CAD-Swept, (f) CAD-MB and (g) anatomical-MB meshes (red is worse and blue is better)	195
7.6	Inferior view of the hexahedral scaled-Jacobian on the (a) low-, (b) medium- and (c) high-resolution meshes, and the (e) CAD-Swept, (f) CAD-MB and (g) anatomical-MB meshes (red is better and blue is worse)	199
7.7	Cross-sectional view of the hexahedral scaled-Jacobian on the (a) low-, (b) medium- and (c) high-resolution meshes, and the (e) CAD-Swept, (f) CAD-MB and (g) anatomical-MB meshes (red is better and blue is worse)	199

7.8	The proportion of an element occupying a surface irregularity (x, y and z) and the associated distortion (blue) for (a) low-, (b) medium- and high-resolution meshes, where the irregularities for lower resolution meshes represent greater element distortion ($x > y > z$)	201
7.9	Effect of hexahedral Laplacian smoothing on shape change due to singularity locations on the hexahedral mesh derived from the CAD-based (a-b) and multi-block method (c-d)	204
8.1	Flow-chart of what proportion of the 47 geometries were used for training/practice of the multi-block method and used in Case D	212
8.2	The challenging aspects of the ten features used to rank target geometries on their difficulty (the least challenging baseline shown is in the middle)	213
8.3	Procedural comparisons between (a) an automatic mesh-morphing strategy and (b) a semi-automatic multi-block method. Automatic processes shown in green and operator required processes shown in purple.	217
8.4	Mesh generation times for (a,b) individual geometries and the distribution of times for the (c) morphed and (d) multi-block meshes	219
8.5	Comparison of the number of target vertices against the time taken to generate each hexahedral mesh for the (a) mesh-morphing strategy and (b) multi-block method	220
8.6	Surface errors for (a,b) individual geometries and associated distributions for the (c,e) morphed and (d,f) multi-block meshes	221
8.7	(a) Mesh volume comparison between morphed, multi-block and target meshes, (b) target volume difference for each morphed and multi-block mesh	222
8.8	The number of (a) very low and (b) low-quality elements for individual geometries and the low-quality element distributions for the (c) morphed and (d) multi-block meshes	223
8.9	(a) Individual and (b) overall convergence times for simulations derived from the morphed and multi-block meshes	224

8.10	The contact areas for individual simulations on the (a) tibial and (b) femoral cartilage. The absolute contact area differences (c-d) and distributions (e-f) for the morphed mesh simulations on the (c,e) tibial and (b,f) femoral cartilage, respectively	226
8.11	Regions and features of the target geometries which caused (a,c-f) higher and (b) lower surface errors for the morphed meshes (blue=lower surface error; red=higher surface error)	228
8.12	Regions and features of the target geometries which caused high surface errors for the multi-block meshes (blue=lower surface error; red=higher surface error) .	229
8.13	Regions and features of the morphed meshes depicting variations of higher and lower quality elements (blue=lower element-quality; red=higher element-quality)	230
8.14	Regions and features of the multi-block meshes depicting variations of higher and lower quality elements (blue=lower element-quality; red=higher element-quality)	231
8.15	Contact pressure distributions of simulation pairs, comparing morphed meshes with the lowest (a-b) and highest (c-d) RMSE. Contact pressure differences around thin-regions highlighted	233
8.16	Different shaped horn-attachments between some of the (a) morphed and (b) multi-block meshes	233
8.17	Contact area distributions of simulation pairs, comparing the lowest (a-b) and highest (c-d) contact area difference	234
9.1	An example of a mesh-quality vs. surface error comparison with quadrants highlighting differences in performance	250
9.2	RMSE (a) and maximum surface errors (b) for individual geometries	251
9.3	(a) Mesh volume comparison between the unoptimised, optimised, multi-block and target meshes, (b) target volume difference for each mesh	252
9.4	Mesh generation times for (a,b) individual geometries. Comparison of the number of target vertices against the time taken to generate each hexahedral mesh for the (c) optimised and (d) unoptimised mesh-morphing strategies	253
9.5	The number of (a) very low and (b) low-quality elements for individual geometries	254

9.6	Comparison of the root-mean-square error against number of low-quality elements for each method: (a) without enclosing area and (b) with enclosing area, excluding outliers	255
9.7	(a) Individual and (b) overall convergence times for simulations derived from unoptimised, optimised and multi-block meshes	256
9.8	The contact areas for individual simulations on the (a) tibial and (b) femoral cartilage. The absolute contact area differences (c-d) for the unoptimised and optimised simulations	257
9.9	Regions and features of the target geometries which caused surface error variations for the optimised meshes (blue=lower surface error; red=higher surface error)	259
9.10	Regions and features of the optimised meshes depicting variations of higher and lower quality elements (blue=lower element-quality; red=higher element-quality)	261
9.11	Contact pressure descriptor ranges for the (a) smoothness-to-patchiness and (b) continuousness-to-discontinuousness qualitative measures, comparing where the optimised mesh simulations lie compared to those derived from the unoptimised and multi-block meshes	263
9.12	Contact pressure distributions of simulation triplets, comparing optimised meshes with the lowest (a-c) and highest (d-f) RMSE. Contact pressure differences around thin-regions highlighted	264
9.13	Different shaped horn-attachments between some of the (a) optimised and (b) multi-block meshes	264
9.14	Contact area distributions of simulation triplets, comparing optimised mesh simulations with the lowest (a-c) and highest (d-f) contact area difference	265

List of Tables

2.1	Summary of 3D finite elements	23
2.2	Element types used for tissues in tibiofemoral joint models	31
2.3	Method classifications and associated algorithms	43
2.4	Common spring weights used to define surface parameterisations	56
2.5	Speed comparison for common volumetric transformation operators (seconds) [11, 12]	68
2.6	Three meshes with properties that were used to assess the speed of different operators	68
2.7	Common RBF kernels used for shape interpolation and mesh-morphing [13–16] .	76
2.8	Review of mesh-morphing strategies for biomedical applications	92
4.1	Comparison of medial and lateral meniscus volumes for the target meshes	136
4.2	Mesh-quality criteria with respect to scaled-Jacobian values	138
4.3	Transverse and cross-sectional dimensions for the template meshes	139
4.4	Template mesh characteristics with respect to resolution and quality (subdivi- sions are with respect to the coarse mesh, <i>Figure 4.2</i>)	139
4.5	Neo-Hookean material parameters for the menisci [17]	147
4.6	Stiffness of meniscal horn attachments [18]	148
4.7	Convergence analysis on different mesh-resolutions and element sizes	151
5.1	Quality and properties of each centreline with respect to length	155
5.2	Quality and properties of each centreline with respect to noise	156
5.3	Dimensions of the target and template meshes	157
5.4	Summary of results for the performance analyses of sub-cases A.1 and A.2	158
5.5	Comparison of the RMSE and maximum surface error for sub-cases A.1 and A.2	160
5.6	Comparison of the average and minimum scaled-Jacobian, and the percentage of very low-, low- and very high-quality elements for sub-cases A.1 and A.2	161
6.1	Transverse and cross-sectional dimensions for the template mesh	170
6.2	Dimensional changes applied to an average geometry to create the target meshes	171

6.3	Summary of results for the performance analyses of sub-cases B.1 and B.2	172
6.4	Comparison of the RMSE, maximum surface error and percentage of errors above 1 mm for sub-cases B.1 and B.2	174
6.5	Comparison of the average and minimum scaled-Jacobian, and the percentage of lower and very high-quality elements for sub-cases B.1 and B.2	176
7.1	Control parameters and the associated template mesh resolution	188
7.2	Degree of noise added and the associated template mesh quality	189
7.3	Mesh generation attributes and the associated template mesh topology	190
7.4	Mesh generation attributes and the associated template mesh shape	191
7.5	Summary of results for the performance analyses of sub-cases C.1-C.4	193
7.6	Comparison of the RMSE, maximum surface error and percentage of errors above 1 mm for sub-cases C.1-C.4	195
7.7	Comparison of the average and minimum scaled-Jacobian, and the percentage of very low- and low-quality elements for sub-cases C.1-C.4	198
8.1	Features and their descriptors used to classify the difficulty of generating hexa- hedral meshes from target geometries	212
8.2	Comparison of the average, standard deviation, minimum and maximum gener- ation times between the mesh-morphing strategy and multi-block method	218
8.3	Comparison of the average, standard deviation, minimum and maximum surface errors between the mesh-morphing strategy and multi-block method	220
8.4	Comparison of the average, standard deviation, minimum and maximum absolute volume errors between the mesh-morphing strategy and multi-block method . . .	222
8.5	Comparison of the average, standard deviation, minimum and maximum percent- age of very low and low-quality elements between the mesh-morphing strategy and multi-block method	223
8.6	Comparison of the average, standard deviation, minimum and maximum conver- gence times between the mesh-morphing strategy and multi-block method	224
8.7	Comparison of the average, standard deviation, minimum and maximum absolute contact area differences between the mesh-morphing strategy and multi-block method	225

9.1	Comparison of the average, standard deviation, minimum and maximum surface errors between the three methods	251
9.2	Comparison of the average, standard deviation, minimum and maximum absolute volume errors between the three methods	252
9.3	Comparison of the average, standard deviation, minimum and maximum generation times between the three methods	254
9.4	Comparison of the average, standard deviation, minimum and maximum percentage of very low and low-quality elements between the three methods	254
9.5	Comparison of the average, standard deviation, minimum and maximum convergence times between the three methods	256
9.6	Comparison of the average, standard deviation, minimum and maximum absolute contact area differences between the multi-block method and the unoptimised and optimised mesh-morphing strategies	257
A.1	Meniscus target geometries with their feature descriptions and sources	344
B.1	Case D: Summary of results for the performance analyses from the lateral meniscus meshes	355
B.2	Case D: Summary of results for the performance analyses from the medial meniscus meshes	356
C.1	Case E: Summary of results for the performance analyses from the lateral meniscus meshes	357
C.2	Case E: Summary of results for the performance analyses from the medial meniscus meshes	359

Part I

Background and Development

1 Introduction

1.1 Background

1.1.1 Hexahedral meshing in musculoskeletal finite-element models

Finite element (FE) models of the musculoskeletal system are used to study the biomechanics and interactions of a range of joints in the human body [17, 19–25] (*Figure 1.1*). They allow greater understanding of the intricate functionality of a joint and the effects of current and novel clinical therapies and procedures [17, 23–25]. A key challenge for musculoskeletal modelling is to accurately create a 3D representation of biological structures of interest from imaging data, such as magnetic resonance (MR) or computed tomography (CT) images [26, 27]. It is often challenging to determine the boundaries of irregularly shaped tissues when manually segmenting them layer-by-layer from three-dimensional image data.

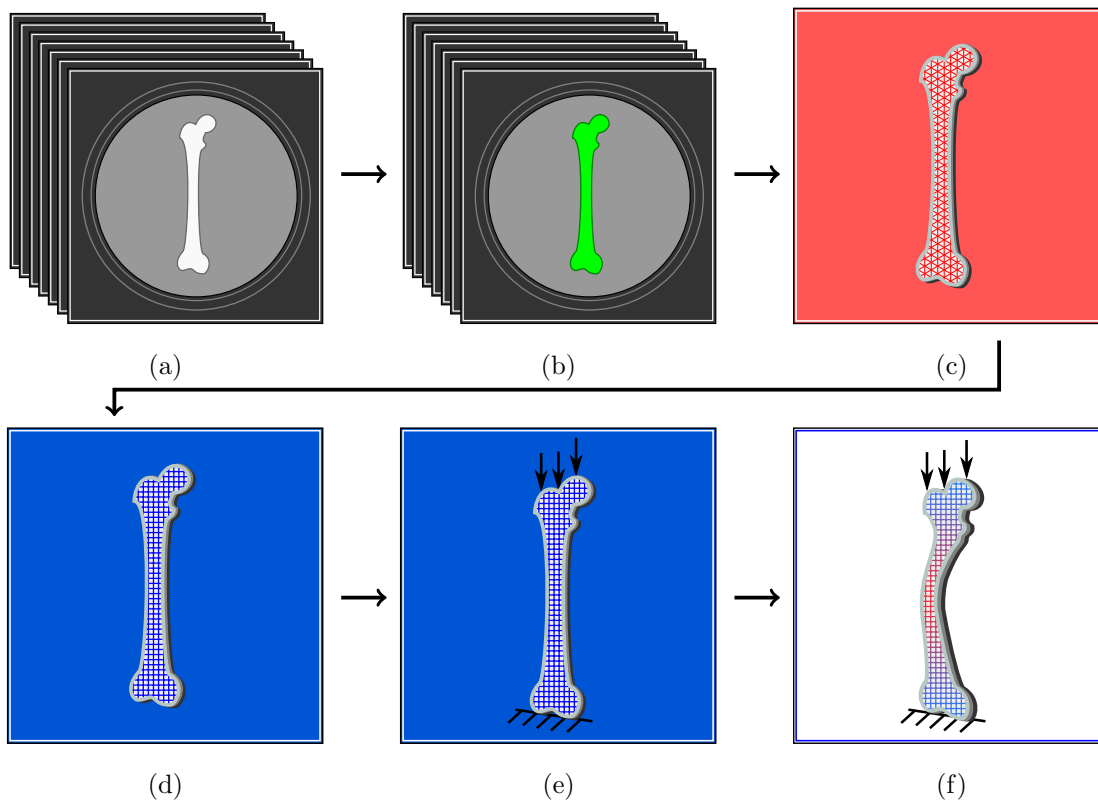


Figure 1.1: Stages of FE model development: (a) image acquisition, (b) image segmentation, (c) surface meshing, (d) volumetric meshing, (e) application of boundary conditions, and (f) solution of FE model

In FE modelling the segmented tissue volume requires discretisation into *elements*, with the most common types being tetrahedral and hexahedral. There are significant advantages for using hexahedral meshes when studying soft biological tissues [28]. Soft tissues display four features which pose challenges for computational simulations: (1) sliding contact between two or more surfaces, (2) large deformation, (3) nearly-incompressible material response, and (4) geometry-specific anisotropic and inhomogeneous material properties (e.g. alignments of collagen fibres) [29–33]. For all four aspects, hexahedral meshes perform better, as tetrahedral meshes can have inherent issues which can reduce accuracies and prevent the FE solver from finding a solution of a model [29, 30, 34, 35]. However, automatic mesh generation algorithms often fail when attempting to generate high-quality hexahedral meshes for more complex geometries, such as those found in the tissues of human joints. The state-of-the-art used to generate high-quality hexahedral meshes of irregular geometries is the multi-block method [3, 19, 21, 22, 36–38]. For optimal outcomes, the multi-block method requires adaption of the input parameters, which demands an adequate level of knowledge and skill in the user [39]. Therefore, the generation of hexahedral meshes for these applications often require manual processing, which can add several hours of laborious and technical work [27, 40–42]. Consequently, the reproducibility of the derived meshes is subject to manual processing errors, such as user bias [43, 44].

The segmentation of anatomies from images and the subsequent generation of an appropriate hexahedral mesh therefore represent two major barriers inhibiting large-scale clinical uptake of FE models. The focus of this research is to address the second challenge: decreasing the technical difficulties and time constraints required to develop high-quality hexahedral meshes.

1.1.2 Mesh-Morphing

Mesh-morphing is a technique used to create a mesh of a geometry through the modification of the vertices from an existing mesh [27, 45]. The number of vertices, elements and associated connectivity (i.e. topology) remain the same following a morphing operation. The existing mesh is commonly referred to as a template mesh. The geometry to be meshed is referred to as the *target geometry* and is often represented as a triangulated surface mesh. Two main goals of a hexahedral mesh-morphing are: (1) to adjust the vertices on the surface of the template mesh to match the surface of the target geometry (surface parameterisation), and (2) to adjust the internal template vertices to optimise the quality of the volumetric elements (volumetric transformation). The core objective of mesh-morphing strategies are to obtain an accurate

representation of the target geometry whilst maintaining a high-proportion of the original mesh-quality. If the template mesh is reasonably close to the target geometry, a high-proportion of the original quality can be maintained [46–48]. Instead of developing a high-quality hexahedral mesh for each geometry of interest, one template is created based on a similar shape.

The main advantage of mesh-morphing is that it can be automated, and provide faster generation times with meshes comparable to semi-automated approaches. As such, manual hours of work can be reduced to a choice of which template to use for a particular target geometry. Another advantage is due to the isotopic (i.e. same topology) nature of the derived meshes. This can be used in two different ways: (1) to automate FE model development, and (2) to serve as points of comparison between models as a form of volumetric parameterisation. In contrast, automated and semi-automated approaches of tetrahedral and hexahedral mesh, respectively, tend to yield meshes with different topologies [29,49]. The main disadvantage of mesh-morphing is that the template mesh must have a reasonably similar shape to a target geometry. Otherwise, element distortion can be induced which can lead to excessive mesh-quality reductions and invalid meshes. In these situations, mesh optimisation steps are often required to improve degraded or invalid elements [30,45].

The majority of existing mesh-morphing strategies focus on the surface parameterisation of triangular meshes [50]. Although, these strategies can accurately calculate bijective maps or one-to-one correspondences, they do not transfer well to hexahedral meshes. Moreover, these strategies can induce high distortions in the surface elements of the hexahedral meshes, due to not being triangular. Additionally, the reviewed volumetric-based mesh-morphing strategies do not provide adequate robustness or validation for morphing highly irregular geometries, such as the soft-tissues found in the musculoskeletal system [9,27,45–48,51–59]. Also, the majority of these strategies are semi-automatic.

1.1.3 Meniscus of the knee joint

The meniscus is a vital part of the biomechanical environment of the knee, which is of great interest due to improving the congruency between the tibial and femoral surfaces, evenly distributing loads and acting as a shock absorber [23]. Additionally, changes to the meniscus due to injury, degeneration or surgery have a profound impact on the biomechanics of the knee and associated symptoms of pain [60]. Typically, meniscus changes can be associated with

accelerated joint degeneration and osteoarthritis [23, 61].

The knee meniscus has been chosen to test the developed mesh-morphing strategy as it displays all four of the challenging features of biological soft tissues [62, 63] and has a shape which provides sufficient algorithmic challenges. The meniscus geometry is characterised by a crescent shape in the transverse plane and a cross-sectional shape similar to a right-angled triangle with rounded features (*Figure 1.2*). Therefore, it has regions of variable global curvatures, with convex and concave sections. As with all subject-specific natural tissue geometries, the meniscus has localised surface irregularities. During either segmentation or the pre-processing of the extracted surface mesh, the geometry is often clipped to avoid explicitly including the attachment with the bone. This leads to localised regions with high curvatures. All of these features provide specific challenges for a mesh-morphing strategy.




Figure 1.2: Shape and cross-section of a meniscus anatomy (Source: LM_00 [*Table A.1*])

1.2 Aims

The aim of this research is to enable biomedical engineers to automatically generate high-quality hexahedral meshes from anatomical sites of interest. Therefore, the proposed solution over a state-of-the-art semi-automatic procedure is an automated mesh-morphing strategy. The research has focussed on creating a robust, automatic and resource efficient mesh-morphing strategy, to enable the rapid and reproducible creation of high-quality hexahedral meshes for biological tissues and anatomies.

The knee meniscus was chosen, to understand the applicability and capacity of the strategy to operate over a complex anatomy which often demands hexahedral discretisation. Meniscus geometries were used in a range of sensitivity analyses and performance assessments, to minimise confounding factors. Furthermore, focussing on one application enabled the isolation of issues

that would not only affect generating meshes for the meniscus but for other tissues with similar features and complexity.

Analyses into the comparative performance and limitations of the strategy were an additional focus to understand the operation of the implemented pipeline of algorithms. The aim of these aspects were to evaluate the strategies strengths and weaknesses, as well as determine routes for optimisation.

Another aim was to create a general-purpose design for two purposes. First, to enable the strategy to be applicable to a range of different biological structures. Second, to allow simple modifications to be implemented or combined with the current algorithm implementations to extend the capabilities to additional anatomical sites of interest.

All meniscus geometries in this thesis have a closed surface (i.e. water-tight) with a genus-0 topology, therefore have no holes (e.g. tunnel loops [64,65]). Although menisci are susceptible to pathologies which can cause holes [66,67], the morphing of geometries with holes are considered beyond the scope of this work.

There are three main aims of this PhD:

1. Develop a novel general-purpose mesh-morphing strategy that creates high-quality hexahedral meshes automatically from existing meshes, which has the potential to be extended to a variety of musculoskeletal tissues
2. Determine the sensitivity, reproducibility and accuracy of the developed strategy
3. Validate the performance of the developed strategy against a state-of-the art method for generating hexahedral meshes

1.3 Novelty

The novelty in this work lies in the development and implementation of a resource efficient strategy that can morph to highly irregular geometries. The sensitivity and performance analyses provide additional novelty, as mesh-morphing strategies reported in literature have not been evaluated to the same degree.

The first aspects of novelty derive from the strategy itself, which combines well-established concepts and algorithms from literature into a coherent and application-adaptable strategy for

highly irregular geometries and anatomies. Furthermore, some ideas and concepts have been improved. The two improvements relate to the generation of centrelines in general, and how those centrelines are used.

In order to validate a mesh-morphing strategy, the performance was compared to meshes derived from a state-of-the-art procedure for generating hexahedral meshes. This provided performance quantifications for a wide-range of metrics against a reasonably large cohort of meniscus anatomies. Additionally, the analyses between the strategy and state-of-the-art procedure highlight areas where the strategy can be improved and adapted to different applications.

For the first time, due to this research, high-quality hexahedral meshes of biological structures can be rapidly created using an automatic, general-purpose and validated procedure.

1.4 Thesis Structure

The subsequent chapters of this thesis are divided into four parts:

Part I: Background, strategy design and performance metrics

Part II: Sensitivity analyses and routes for optimisation

Part III: Performance validation compared to a state-of-the-art method

Part IV: Applicability and usefulness within the biomedical community

In *Chapter 2*, a review of the literature provides information about the development and importance of FE models of the knee and characteristics of meniscus anatomies. Additionally, information is provided about the characteristics, utilisation and generation methods of hexahedral meshes. Finally, the core components and applications of mesh-morphing strategies are detailed, with their associated advantages and disadvantages. In *Chapter 3*, the design and implementation of the unoptimised and optimised strategies are described, and are evaluated in the following chapters and associated cases. In *Chapter 4*, the metrics, implementations and meshes, used to characterise the performance and robustness in subsequent methodologies are defined. These core metrics and criteria are used in several methodologies and serve as the basis of their assessment.

The general sensitivity of the unoptimised strategy is assessed for a variety of input variables in *Chapters 5-7*. The assessments in these chapters aim to determine the most optimised

conditions to compare against a gold-standard procedure. In **Chapter 5** (case A), the sensitivity of the unoptimised strategy is evaluated with respect to two centrelines attributes. In **Chapter 6** (case B), the sensitivity of the unoptimised strategy is evaluated against extreme population variations of meniscus geometries. In **Chapter 7** (case C), the sensitivity of the unoptimised strategy is evaluated against four template mesh attributes.

In **Chapter 8** (case D), the performance of the unoptimised strategy is validated against a gold-standard procedure. This is used to provide a benchmark to assess the optimised mesh-morphing strategy. In **Chapter 9** (case E), the performance of the optimised strategy is validated against the established benchmark.

In **Chapter 10**, the critical observations and analyses from the case studies and their applicability for biomedical engineering and mesh-morphing are discussed. Additionally, the future work and impact of the strategies are described.

2 Literature Review

2.1 Meniscus Anatomy and Importance

2.1.1 Tissue Structure

Geometry

The meniscus is consistently described as having a crescent or semi-lunar shape [1,2,68]. This is the typical shape of a healthy meniscus when viewed in the superior-inferior direction (*Figure 2.1a*). The cross-section of a meniscus has a wedge or right-angled triangle shape (*Figure 2.1b*). This geometry allows the structure to fit between the tibial and femoral articular cartilage surfaces. At the ends of the crescent geometry, horn-attachments exist to secure the tissue to the tibial plateau, which enables a degree of mobility for the tissue [69], during flexion and extension of the tibio-femoral joint [60]. The mobility of the menisci is one of the properties that allows them to maintain congruency between the femur and tibia [69].

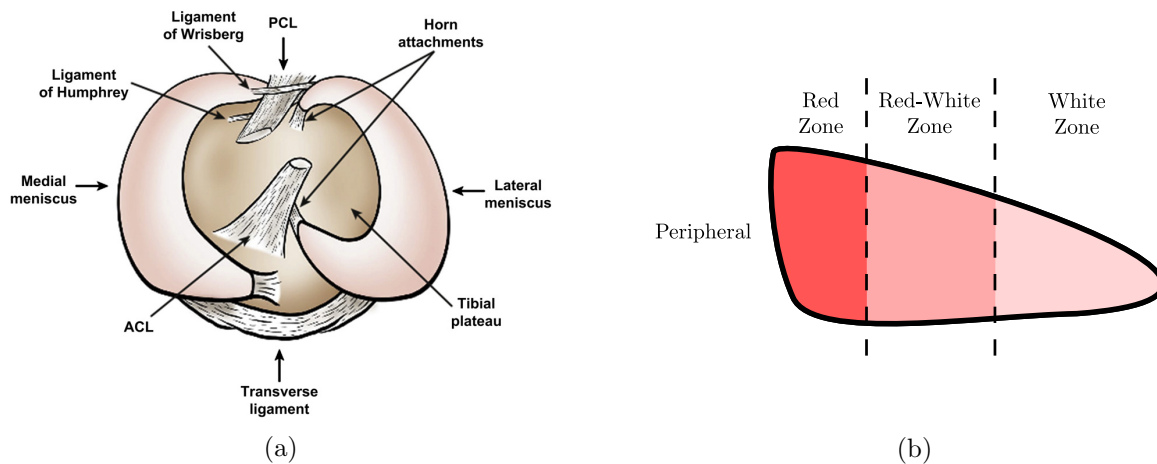


Figure 2.1: Meniscus Anatomy:

(a) medial and lateral meniscus with associated horn-attachments and ligaments (*reproduced from [1]*) (b) meniscus cross-section and associated zones (*adapted from [2]*)

Composition

The meniscus is mainly composed of collagen and is distinguished by three zones that dictate the collagen type and degree of other biological materials. The three zones separate the cross-section of a meniscus from red to white, each defined as: red-red, red-white and white-white [1,2]

(*Figure 2.1b*). The red-red zone is composed of type-I collagen, which is the same type found in tendons, and has vascular and neural components [1,2]. The white-white zone is composed of type-II collagen, this type of collagen is found in hyaline-cartilage, and is avascular and aneural [1,2]. Consequently, a healthy meniscus has hybrid properties resembling the flexibility of tendons at the periphery of the shape and the low-friction properties of cartilage in the opposite direction. The level of vascular and neural tissues dictate what ability particular zones and regions have at healing and the degree of injury permanence [1,2].

Of the reviewed literature for finite element models involving the meniscus, the different compositions and consequently material properties of the aforementioned zones have not been included into the simulations. One study simulating and investigating meniscus tears have declared the three zones with respect to the formation of particular tears but did not specify different properties for each component [70]. The material properties of meniscus tissues and the variance between individuals is beginning to be understood but not down to the level of each zone and transition between zones.

Associated Ligaments

The meniscus has several associated ligaments that provide supplementary stability, in addition to the horn-attachments (*Figure 2.1a*). Interestingly, the presence of some of these ligaments vary between individuals. For example, the Wrisberg and Humphries ligaments are two menisco-femoral ligaments providing attachments to the femur surrounding the posterior cruciate ligament [1, 2, 68, 71] (*Figure 2.1a*). From 26 cadaveric specimens, only 46% were identified as having both; whereas, 100% of the specimens had at least one [1,71]. Additionally, the transverse ligament, which connects the anterior horns of the medial and lateral menisci, has previously been located in 64% of 46 cadaveric specimens [68,72].

Of the reviewed literature, there was only one study that incorporated the menisco-femoral ligaments into finite element models with meniscus anatomies [73]. In contrast, there have been several finite element simulations of the tibio-femoral joint that have included the transverse attachment [19,62,73].

Medial and Lateral Differences

The medial and lateral menisci share the majority of their properties but differ in several aspects. Generally, the lateral meniscus has greater variability and size with respect to shape and thickness [2]. The medial meniscus is longer than the lateral meniscus [2] and has been reported to cover 51-74% of the medial compartment of the tibial plateau [1]. However, the lateral meniscus represents a larger proportion of the lateral tibial plateau, between 73-93% has been reported [1]. Additionally, the lateral meniscus has loose peripheral attachments to the joint capsule only [2] and therefore has greater mobility than the medial counterpart [2,69]. The medial meniscus is relatively immobile compared to the lateral meniscus due to continuous and firm peripheral attachments to the joint capsule and medial collateral ligament, respectively [2]. Despite these differences, the medial and lateral menisci are often modelled with the same number of attachments and degree of mobility [18, 19, 62, 63, 74–76].

2.1.2 Tissue Function

Roles within the knee

The meniscus provides several functions for the healthy articulation and operation of the knee, which include:

1. Load distribution and transmission [23, 68]
2. Shock absorption [23, 68]
3. Improved congruency [23, 68]
4. Joint stability [23, 68]
5. Lubrication (proposed) [23, 68]
6. Nutrient distribution (proposed) [68]
7. Proprioception (proposed) [23, 68]

The menisci provide a range of unique properties that all work together to maintain the functionality and efficiency of the tibio-femoral joint. Improving the congruency and distribution of loads between the femur and tibia are important roles of the meniscus, due to the inherent lack of conformity between articular surfaces of the femoral and tibial cartilage. Although the me-

dial compartments have greater conformity than the lateral, this changes with knee articulation due to the irregular shape of the femoral surface [77]. When the knee articulates, congruency and transmitted loads change and the meniscus is the component that stabilises the joint. The meniscus shape and attachments allow flexible movement to maximise congruency [69]. These features combined with the composition of the types of collagen present in tendons and articular cartilage enables the tissue to be resilient, absorb shocks and evenly distribute loads.

Defects and associated affects on the knee

Meniscus injuries are a very common form of knee injury and have an incidence between 12-14% [78]. The prevalence of meniscus injuries have been reported to be 61 in 100,000, which is around 850,000 surgeries per year in the United States [78]. There are several different types of injuries, degenerative effects and risks that menisci may suffer (*Figures 2.2*):

1. Tears: longitudinal/vertical, oblique/parrot-beak, transversal/radial, bucket-handle and horizontal [2, 23] (*Figures 2.2a-2.2e*)
2. Extrusion [79, 80] (*Figures 2.2f*)
3. Horn-laxity [79, 80]
4. Surgical modification: partial and total meniscectomy [60, 81]




Figure 2.2: Pathologies of the meniscus (arrows show the progressive deterioration)

Several arguments can be put forth that any form of injury or degenerative effect that causes menisci to lose either structure or function can lead to the development osteoarthritis. For example, reduced load transmission and shock absorption capabilities are consequences of menis-

cus tears [82], which in turn is a risk factor for osteoarthritis due to structural changes [83]. However, osteoarthritis can occur through a range of different mechanisms and risk-factors, some of which are not initiated by the meniscus, such as: bone marrow lesions [84], cartilage degeneration [85,86], osteophyte formation [83,85], particular repetitive movements [87], inflammation [85], excessive mechanical forces [87] and malalignment of the tibio-femoral joint [86,87]. The pathways for the development of osteoarthritis are not fully understood [83]. There are complex mechanisms that lead to osteoarthritis and several of which exacerbate each other. For example, the presence of articular cartilage damage has the potential to degrade and extrude the menisci [85]. Meniscus extrusion and tears have been linked to being a possible causal pathway for osteoarthritis [79]. This correlates with results from finite element analyses of meniscus extrusion and tears, indicated by increases in compressive and shear stresses of articular cartilage and meniscus tissues [82]. Additionally, it has been proposed that meniscus extrusion and horn-laxity are intricately associated [80]. Although, one study investigating meniscal extrusion and horn-laxity of 251 knees from 153 patients, could not find consistent evidence to confirm that the presence of these conditions leads to a loss of cartilage [86]. The aforementioned study did find evidence that medial malalignment and lateral meniscus damage does lead to cartilage loss on the femoral and tibial components [86]. This can make it difficult to determine the cause of osteoarthritis and formulate an acceptable treatment or therapy. For example, if damage to the cartilage tissues within the knee have been measured, was it articular cartilage or meniscus degradation that initiated the condition or neither? To further complicate matters, meniscus defects can be caused by both traumatic events and degenerative conditions [66,78]. Based on a cohort of 198 participants with meniscus tears, it is clear is that meniscus degeneration does not heal or recover from a degenerate state and has a 16% chance to continue to become worse over the duration of 8 years [83]. Therefore, intervention is required to manage symptoms, that may include pain, and to reduce the risk of osteoarthritic development within the knee [83].

2.2 Finite Element Modelling of the Tibio-Femoral Joint

2.2.1 Purpose

The purposes of tibio-femoral joint (TFJ) modelling can be broken down into two main categories. The first is the improvement of TFJ models, which can be broken down further into three groups: improvements, optimisations and validations [19–22]. The second is clinical rel-

evance of the studies, where derived conclusions can be used to improve clinical therapies and the outcomes of procedures [17, 23–25]. However, there are sometimes overlap between these categories, particularly between validation and clinically focussed models [25, 36]. Additionally, there is some overlap between optimisation and validation papers [19, 20]. Validation is a common point of overlap as it provides concrete evidence that the FE model was not only developed correctly but produces information which correlates to either in-vivo data or experimental models. Furthermore, it provides a level of confidence of the data and conclusions derived from the FE model.

Commonly used tissues incorporated into TFJ models are the femur, tibia, femoral cartilage, tibial cartilage and menisci [17–25, 36, 62, 75, 88–90] (*Figure 2.3*). Sometimes the menisci are not integrated into TFJ models, where there inclusion is sometimes defined as a tibio-menisco-femoral model [88, 91]. Also, the complexity can be increased through integration of cruciate ligaments as well as other supporting ligaments, tendons and bones in the legs [36, 88, 89].

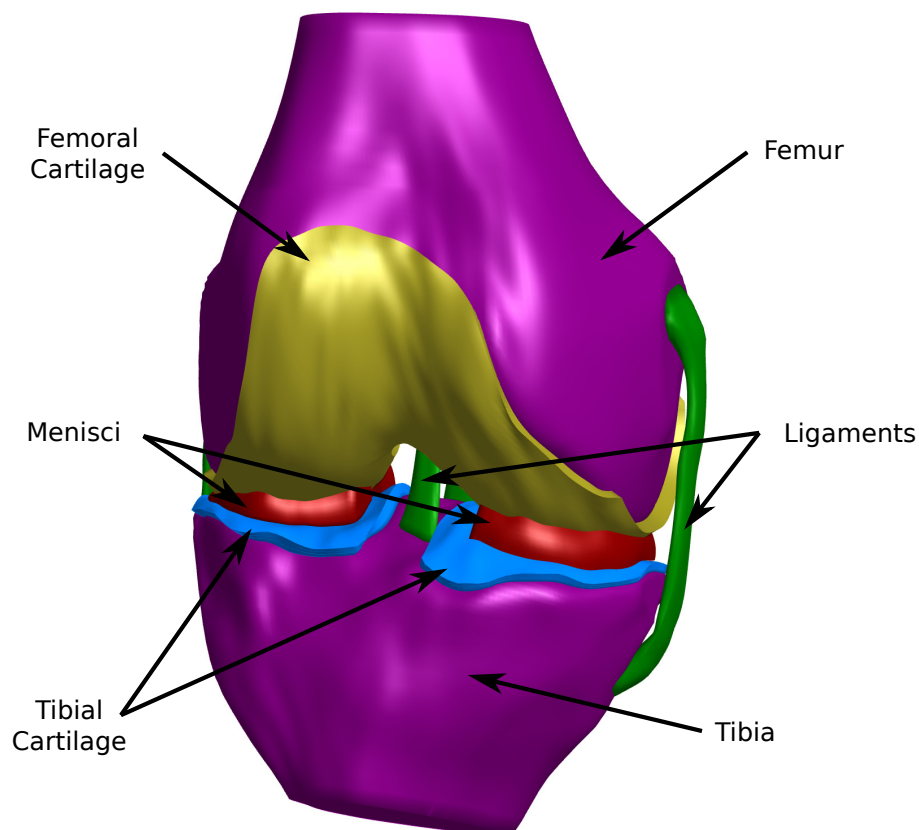


Figure 2.3: Common components used to model the tibiofemoral joint
(Source: open-knee model [3])

Validation Methods

The most common method to validate a TFJ model is to create the model based on a cadaveric joint, which can be built into an experimental model equipped with sensors. Additionally, this requires a sophisticated mechanical joint simulator that can fix the joint and accurately acquire the relevant data, often pressure and reaction forces [19, 22, 25, 36]. However, other validations can be done without cadaveric tissues, through MRI of in-vivo parameters using location-based information [20, 21]. Although, MRI-based validation methods cannot compare in-vivo measures of stresses and forces, which currently cannot be achieved without an invasive surgical procedure [21].

Improvements and Optimisations

Due to the complexity of the TFJ, assumptions are made to simplify models in a number of different ways, which includes material properties, boundary conditions and exclusion of particular tissues. Consequently, there are several improvements that can be made to existing models based on overcoming these earlier assumptions. Sometimes, optimisations are performed to understand and narrow down which parameter ranges of potential kinematic and material property variables are appropriate to improve model accuracies [19, 20]. On the other hand, improvements can be made through understanding and implementing the complex biomaterial properties of particular tissues, such as the biphasic time-dependant nature of cartilage tissues [75, 92].

Clinically-relevant Models

Of the literature reviewed, clinically-relevant models of the TFJ are the most common [17, 18, 23–25, 36, 89]. Studies of this nature focus on understanding the implications of particular TFJ conditions. These conditions can include the effects of meniscal tears, meniscectomies, tissue replacement devices, malalignments, surgical procedures and degenerative pathologies [17, 18, 23–25, 36, 89].

2.2.2 Tissue Segmentation and Meshing

A FE model of the TFJ can be derived using a number of different protocols. The geometries are often acquired in-vivo using MRI; or a combination of MR and CT imaging

[17–21, 23–25, 31, 36, 75, 88, 89, 93, 94]. To derive a geometry from this type of bio-imaging data, segmentation is required to isolate anatomies of interest. Segmentation is currently a manual operation, with the aid of several tools and algorithms to speed up the process and extract a digital representation of the geometries [31]. Additionally, FE models which are being experimentally validated from cadaveric sources acquire the 3D geometries using dissection and laser scanning or direct digitization of the anatomies of interest [22, 62, 90]. Once the geometries have been extracted, soft-tissues such as cartilage and menisci are typically meshed with hexahedral elements and osseous tissue as either hexahedral or tetrahedral [17–23, 25, 36, 62, 75, 76, 88, 90]. However, osseous tissues are often assumed rigid due to stiffness differences and have been modelled with surface elements, such as triangles and quadrilaterals, or simply with fixed constraints on the articular cartilage surfaces [17–21, 23, 36, 62, 75, 76, 81, 88, 90]. Furthermore, the articular cartilage of the osseous tissues can be extrapolated from the distal elements of the meshed tissues and not directly extracted through segmentation and separately incorporated into the model [25]. Additionally, some soft tissues such as ligaments, tendons, attachment sites and other thin connective tissues are often represented by linear and non-linear spring elements [19, 36, 88, 90]. On occasion, cartilage tissues are meshed with both hexahedral and pentahedral elements [20]. The final stage of developing a FE element model includes assembly of the meshed geometries and specification of material properties, boundary conditions, loads and constraints [31].

2.2.3 Parameters of Interest

Differences in the FE model development protocols are usually due to the aims of the investigations. Similarities are found where procedures and data sources are commonly used and/or known to produce accurate results. The purpose and objectives for developing models of a TFJ can vary, which can include the understanding of: biomechanical behaviours; kinematic behaviours; damage modelling; sensitivities of certain model definitions; validity and comparisons of experimental data and/or mechanical joint simulators and the implications of particular prosthetic designs and implant locations. Consequently, there are a wide range of parameters of interest that TFJ models have been used to derive:

1. Maximum contact pressure (MPa)
2. Contact area (% , mm², cm²)
3. Maximum tensile strain (%)
4. Maximum compressive strain (%)
5. Displacement (mm)
6. Reaction force (N)
7. Maximum contact force (N)
8. Gait cycle at max. contact force (%)
9. Fluid pressure (MPa)
10. Fluid support ratio (%)
11. Rate of strain change (%/s)
12. Rotations (°)

The most common parameters of interest are the contact pressure, force, area and compressive stress. Furthermore, the location of where the maximum values for these variables occur are commonly reported; in particular, differences between medial or lateral compartments [17–19, 21, 25, 62, 75, 89, 92, 95].

2.2.4 Enforcement of Congruency

Congruency is not usually enforced and often assumed between the tissues of TFJ models. However, it is known that the meniscus improves congruency between the femoral and tibial cartilage surfaces [18, 23, 24]. Moreover, in the standing position and at higher loads the meniscus changes shape becoming more deformed and increases congruency between articular cartilage surfaces [21, 23]. The differences in meniscus geometry and congruency between weight-bearing and non-weight-bearing states causes a common issue of simulating weight-bearing states after deriving geometries in a non-weight-bearing state [21]. One issue that arises is due to the inaccurate representation of the geometry. Another issue is due to the lack of robustness inherent in the FE contact algorithms simulating non-conforming surfaces [96]. The knee joint is more incongruent than other joints, such as the hip [96]. However, there is still a degree of congruency present due to the location and properties of the meniscus in a weight-bearing state between the femoral and tibial surfaces [18, 23, 24]. Therefore, the lack of accurate representation of these structures could lead to inaccurate simulations and incorrect interpretations of the results.

How congruency between surfaces was achieved is an area that is not well described, where the causes are situated somewhere between geometry acquisition, meshing and FE model creation. Segmenting three separate anatomies and meshing them with hexahedral elements will naturally incur surface differences compared to the original structures. However, the aspects of achieving conformity are not addressed directly and it is up to the assumptions of the reader

to infer why particular FE definitions or smoothing operations were included. For example, one investigation mentions in their discussion that the medial meniscus did not conform and stated minimal load transfer was caused as a consequence. Additionally, this investigation was the only one to include a compliant contact force model between the meniscus and cartilage surfaces and was reportedly used to achieve efficient computation of contact forces. Potentially, this compliant contact model may have improved solver convergence for the interacting surfaces with the low-conformity [19]. One study briefly described a smoothing and re-meshing scheme, with one of the purposes being to reduce surface irregularities [25]. Furthermore, during the segmentation, meshing and assembly processes they employed a *non-manifold algorithm* which aimed to prevent mismatches in contact boundaries and improve solver convergence [25]. The only other mention of conforming surfaces in the reviewed literature was from an investigation where congruency was assumed between the tibial cartilage and menisci [22].

2.2.5 Clinical Applications and Relevant Outputs

Clinical Applications and Outputs

Differences between healthy and degenerative TFJs are common interests for clinically-focused models. These differences can be due to osteoarthritis, injuries, surgical procedures or a combination [17, 18, 23, 25, 36, 89]. Within this area, a study investigating the effects of meniscal tears and meniscectomies has previously been conducted [23]. A range of different meniscus states were investigated which used contact area, maximal compressive stress and contact pressure distributions to determine how each state may affect the TFJ. Compared to their healthy TFJ model, higher compressive stresses were observed in the meniscectomy models. They used preliminary knowledge of the meniscus, such as evenly distributing transmitted forces between articular surfaces; and post-meniscectomy knowledge that cartilage damage and degeneration was observed, to determine that the higher stresses may be the result of the aforementioned damage [23]. This information could be used to determine the appropriate action and therapies required to improve or arrest further cartilage damage and degenerative symptoms. Moreover, clinicians could focus on therapies which would serve to reduce the higher stresses observed within the cartilage and meniscus tissues. Furthermore, they used both maximal compressive stress and contact pressures to confirm their conclusions based on previous literature and validation studies [23]. Their conclusions correlate with previous meniscectomy studies that use similar

parameters and indicate maximal contact stresses and contact areas increase and decrease, respectively, following this procedure [97]. They note that as contact area decreases, an increase in the compressive stresses is observed [23]. Additionally, higher maximal contact stresses have previously been reported to correlate with the progression of degenerative osteoarthritis within the TFJ [89].

Contact Pressure

Contact pressure on the articular cartilage and meniscus tissues are common parameters of interest [17, 21–23, 36, 89]. This variable is influenced by a number of different factors, due to differences in applied loads; proportion and shape of contact between interfaces and current flexion position in the gait cycle [98]. The contact pressure identifies what proportion of forces are being transmitted between cartilage-cartilage and meniscus-cartilage interfaces with respect to the contact area. This variable is of interest for clinical applications for two reasons. First, it is the role of the meniscus to evenly distribute forces over a greater proportion of the articular surfaces [21, 23, 62]. Therefore, it can be used to determine how effective the meniscus is at performing one of its critical roles. Second, increases in contact pressures between these interfaces have been observed in patients with osteoarthritis and similar degenerative conditions [23, 25, 97]. For validation based investigations, the contact pressure is of interest as it can be measured directly from pressure sensors [25, 36, 62]. Therefore, they can be used to provide direct comparisons between experimental and simulated values, within a small degree of uncertainty or error. Additionally, due to the common use of this variable, it is often used to compare to either validation studies or other FE models [19, 20, 23, 63, 75].

Compressive Stress

Another common variable of interest is the compressive stress of the articular cartilage and menisci [18, 21, 23, 75]. This variable is strongly associated with the contact pressure and the applied loads that result in that pressure. It is the material response to the contact pressure and compressive forces transmitted between tissues. Therefore, it provides information on how the measured pressures and forces have particularly affected the tissues of interest. Moreover, it can be used to determine if excessive damage extends into the tissue or what risks may be occurring during certain conditions [23]. These conditions could be a result of injury, surgical procedures or osteoarthritis [18, 23, 75]. Although, correlative with the contact pressure, the

compressive stresses occur throughout each tissue and can be measured beyond the areas of contact between interfacing anatomies [18, 23]. Consequently, it can indicate how an entire anatomy is affected as a result of either a statically or dynamically-loaded scenario, from the standing position to partial (instantaneous moment) and complete portions of the gait cycle. Similarly, the values are often used to compare between studies but not as common for validation studies. Experimentally obtained compressive stress values typically require robust materials testing equipment to facilitate accurate measurements. Although, this could be implemented using strain gauges and measured reaction forces, the compressive stresses of the knee joint are not measured during experimental cadaveric simulations [25, 36, 62].

Contact Area

The contact area between the cartilage-cartilage and cartilage-meniscus interfaces are contact variables with high interest amongst researchers [18, 21–23, 75]. This variable implicitly links with contact pressure and force, and can be used to determine the functionality of the meniscus at distributing loads over a greater proportion of the articular cartilage surfaces. Additionally, it can be used to gain an idea of how congruent the surfaces are between the interfaces and useful to compare between healthy and symptomatic patients [18, 23, 63, 75]. Lower contact areas are associated with osteoarthritis and degenerative conditions, due to higher forces being transmitted through a smaller area leading to higher compressive stresses and contact pressures [23]. Furthermore, accelerated osteoarthritis and tissue damage has been observed through a reduction in contact areas and the associated increases in contact pressures that results from meniscus injuries or surgical procedures, such as meniscectomies or HTO corrections [18, 23, 99].

Validation Studies and Outputs

Studies which are validated use the derived model data differently to clinically-focussed investigations. The parameters of interest derived from these studies are similar and typically focus on contact variables, such as contact area, pressure and force. Validation studies often compare the parameters of interest measured from experimental sources to the predicted values of a FE model using the root-mean-square error (RMSE) [19, 22, 25, 36]. This enables an understanding of how different a FE model is to experimentally measured data. Furthermore, this is useful in determining if an optimisation or improvement has lead to a greater prediction accuracy of a model, identified by a reduction in the measured errors. Additionally, the RMSE or root-mean-

square-normalised error (RMSNE) can be used to combine the errors from several different parameters of interest or contact variables into one metric [22]. This can be a useful approach to quantifying the total accuracy loss or improvement between a range of variables based on the use of a different approach, model or method. Also, this can enable easier comparisons between the accuracy of a different method, comparing the total loss instead of the individual errors contributing to that factor.

2.3 Hexahedral versus Tetrahedral Meshes

2.3.1 Element Types

There are a large variety of elements that can be used to perform finite element analyses. These can first be separated into 2D and 3D elements, and second with respect to their order. For each linear or first-order element there are higher order formulations, most commonly second-order or quadratic. Additionally, there are other element formulations with cubic, quartic and higher interpolations [100]. Some cubic elements are used within the field of biomedical engineering, e.g. the heart [101]. However, they are not used in TFJ models (*Section 2.3.4*), and therefore are beyond the scope of this review.

First-Order Elements

The most common 2D elements are the triangle (TRI3) and quadrilateral (QUAD4) (*Figure 2.4a-2.4b*) [102,103]. For 3D analyses, the linear tetrahedron (TET4) and hexahedron (HEX8) are the most common (*Figure 2.4c-2.4d*) [102,103].

Additionally, there are other less commonly used 3D elements, such as the pentahedral (PENTA6) and pyramid elements (PYRA5) (*Figure 2.5*). Typically, these are combined with and bridge the more common 3D element types (TET4, HEX8) [104]. This is made possible due to the presence of triangular and quadrilateral faces, allowing more diverse connectivity to both tetrahedral and hexahedral elements. Meshes composed of multiple elements types are referred to as hybrid or mixed meshes [104–106] and attempt to combine the positive attributes of particular element types - such as improved surface fit [TET4] with greater computational performance [HEX8]. These elements are able to accommodate joining hexahedra to tetrahedral elements and prevent hanging vertices [38] due to having both quadrilateral and triangular faces. There are several non-standard elements that exist, such as the polyhedral or abstract elements in

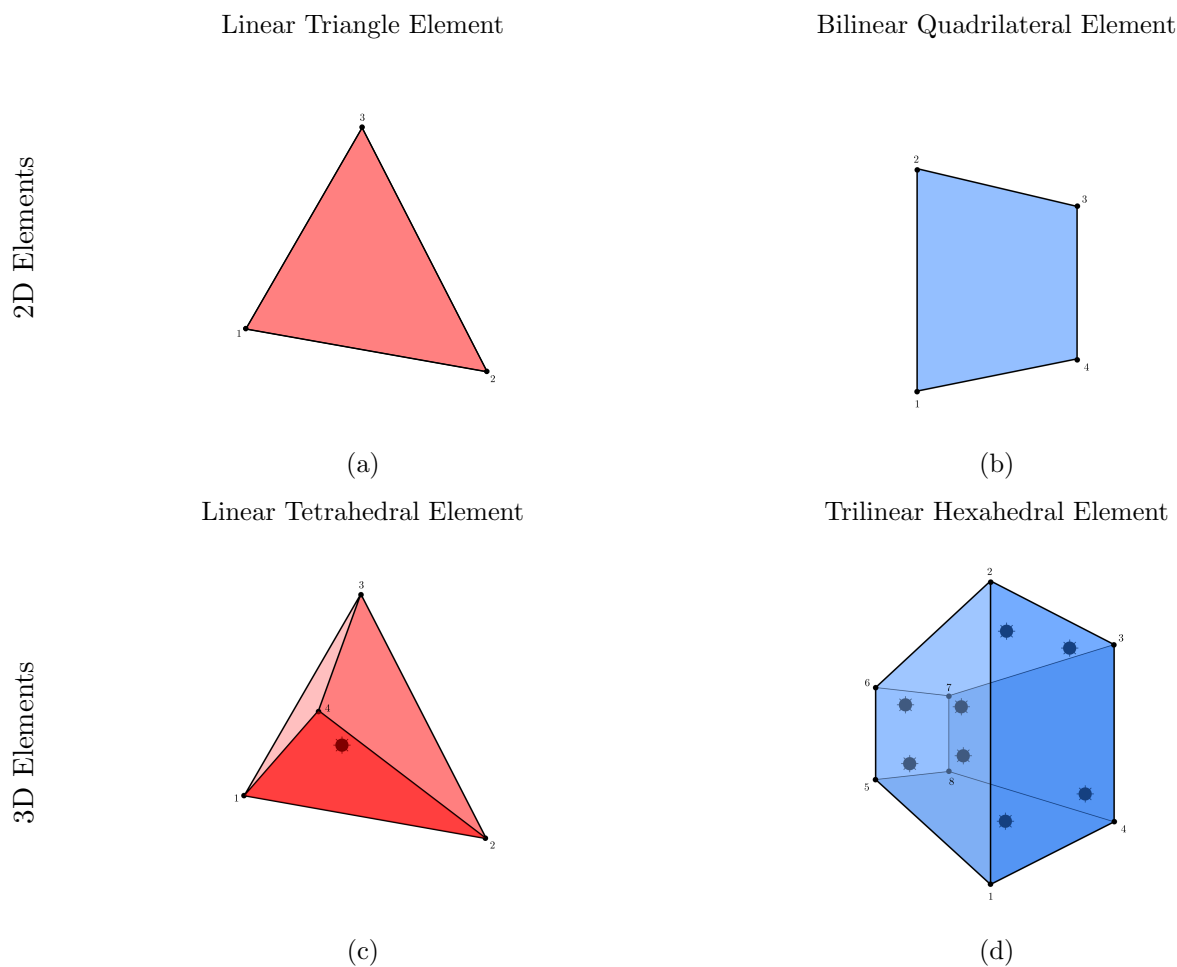



Figure 2.4: Common linear elements: (a) triangle, (b) quadrilateral, (c) tetrahedral and (d) hexahedral Elements. Integration points are shown for the 3D elements (black)

both 2D and 3D, that have been found suitable for particular applications [107]. There is also a knife element [108, 109] which can result from the decomposition of an irregular geometry into a hex-dominant mesh [109] and can be considered a special-case to the wedge element [108] or degenerate case of a hexahedron [109]. However, polyhedral and knife elements are not implemented in commonly used FE packages [110, 111].

Second-Order Elements


The most prevalent quadratic elements are also based on the tetrahedron (TET10; TET15) and hexahedron (HEX20; HEX27) (*Figure 2.6*) [34, 100, 102, 112]. One level of quadratic element (TET10; HEX20) involve the placement of extra vertices on the edges of the linear elements that provide additional degrees-of-freedom (DOF) and the ability to represent those edges as curves [34, 100] (*Figures 2.6a, 2.6c*). A level beyond this adds vertices to the centres of each face and element, further extending the associated DOF (TET15; HEX27) [34, 100] (*Figures 2.6b, 2.6d*).

Linear Pentahedral Element



(a)

Linear Pyramid Element



(b)

Figure 2.5: Less common elements: (a) a pentahedron and (b) pyramid

Quadratic elements require greater computational effort to achieve solutions to their equivalent linear counter-parts [100, 113]. However, fewer second-order elements are required to achieve an adequate surface fit due to more DOF [114, 115], but the specific second-order formulations are not always reported [76, 81]. Although, TET10 is the only quadratic tetrahedral element available for Abaqus [110].

The following table summarizes a range of elements with their associated order, number of nodes, number of faces and possible connectivity in mixed meshes (*Table 2.1*):

Table 2.1: Summary of 3D finite elements

Element	ID	Order	Vertices (#)	Faces (#)	Connectivity
Tetrahedral	TET4	First	4	4	PYRA5, PENTA6
	TET10	Second	10	4	PYRA13, PENTA15
	TET15	Second	15	4	PYRA19, PENTA21
Hexahedral	HEX8	First	8	6	PYRA5, PENTA6
	HEX20	Second	20	6	PYRA13, PENTA15
	HEX27	Second	27	6	PYRA19, PENTA21
Pyramid	PYRA5	First	5	5	TET4, HEX8, PENTA6
	PYRA13	Second	13	5	TET10, HEX20, PENTA15
	PYRA19	Second	19	5	TET15, HEX27, PENTA21
Pentahedral	PENTA6	First	6	5	TET4, HEX8, PYRA5
	PENTA15	Second	15	5	TET10, HEX20, PYRA13
	PENTA21	Second	21	5	TET15, HEX27, PYRA19

2.3.2 Performance of Element Types

There are a number of different factors effecting the convergence of a mesh: element type, element order, integration formulations, material definitions and the size of stiffness matrix

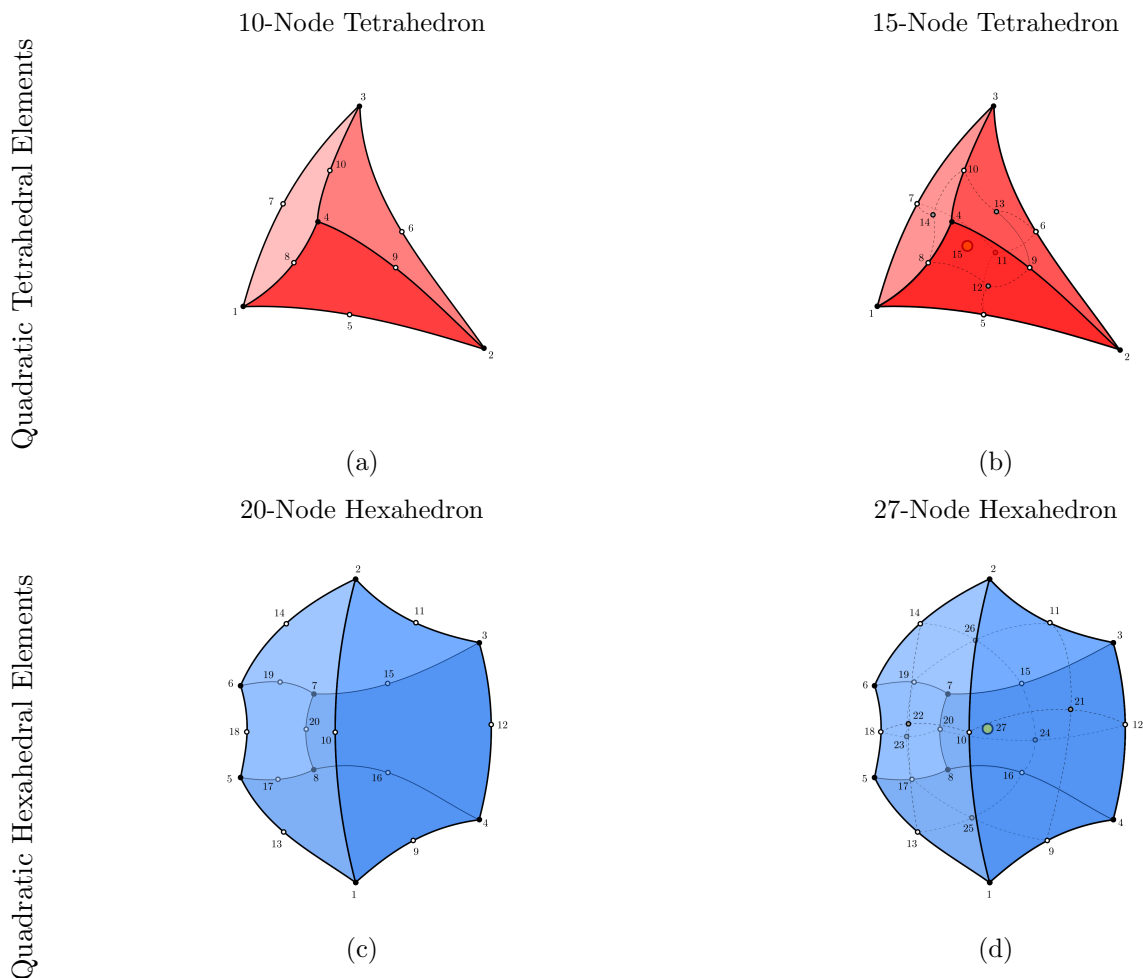


Figure 2.6: Common quadratic elements: (a,b) tetrahedral and (c,d) hexahedral variations

eigenvalues [116].

Stiffness Matrix Eigenvalues

Meshes with smaller stiffness matrix eigenvalues are preferable [116], to reduce the overestimation of material stiffness inherent to the finite element method [112, 116]. It has been observed that TET4 can lead to higher eigenvalues than HEX8, decreasing the accuracy of a solution for a range of applications [4, 116, 117]. The magnitude of the eigenvalue corresponds to the element stiffness or softness, i.e. higher eigenvalues indicate stiffer elements, and vice versa [118]. The stiffness or softness of an element can invoke or overcome certain detrimental phenomena, such as locking [28, 39, 100] and hourglassing [4, 100].

Locking

Locking of elements can occur in two particular situations: (1) material properties approaching nearly-incompressibility [100, 112, 116, 119] and (2) the simulation of thin geometries or bending [100, 103, 112, 120]. The first situation can initiate volumetric and strain locking, whereas, the second can initiate shear locking. Shear locking is also referred to as “tet-locking” due to its prevalence of occurring under large deformations for tetrahedral meshes [39, 121]. Elements with greater element stiffness [28, 39, 100] are most susceptible due to containing fewer DOF [39, 100]. Consequently, the effects are less evident with higher-order elements and less distortion [100, 120]. Particularly, when using quadratic tetrahedral elements over TET4. There have been several attempts to help alleviate the severity of the issues, which include implementing different integration rules [112, 120, 122] and employing constraint relaxing techniques [119, 123].

Hourglassing

In contrast, hourglassing can occur from softer elements and is evident when deformation occurs when there is no strain and can reduce accuracy as well as initiate element inversion [4]. These can occur in HEX8 and quadratic elements but not TET4, which are too stiff [4, 100]. However, higher-order elements can reduce the hourglassing [100]. The prevalence of these effects increases when simulating large deformations [4], dynamic conditions [4] and coarse meshes [122]. Using different integration rules and increasing the element energy and stiffness can overcome these issues [4, 112, 122]. Additionally, some attempts to counter locking effects can exacerbate hourglassing phenomena [122].

Integration Points

Integration points contained within each element are used to evaluate local stresses [34] (*Figure 2.4c - 2.4d*). A reduction in computational effort can be achieved using fewer integration points [34, 112] and can increase accuracy [112] due to reducing overestimated finite-element stiffness effects [112, 116]. However, more integration points can lead to greater accuracy of element representations [34, 124] and reduce hourglassing effects [112]. This is one of the reasons why HEX8 meshes are often more accurate than TET4, due to containing more integration points [29–34] (*Figure 2.4c - 2.4d*). The number of integration points affect the computational effort more than DOF [34]. However, increasing the DOF of an element usually requires an

increase in the number of integration points [34].

Four Core Phenomena

There are four core simulation phenomena that can be difficult to solve:

1. Large Deformations
2. Contact Analyses
3. Nearly-Incompressible Materials
4. Anisotropic and Inhomogeneous Material Properties

The first three phenomena have been commonly reported [29–34], whereas the fourth is less acknowledged. For the first three phenomena, HEX8 meshes have been shown to perform greater than TET4 meshes [29–34]. When imparting large deformations, HEX8 are more reliable, require fewer elements and produce lower errors, whereas, TET4 are more susceptible to shear locking [39]. Similarly, when TET4 are used to define the geometries of nearly-incompressible materials they have a greater tendency to volumetrically lock [119]. When TET4 have been used in contact analyses, they have been shown to be inaccurate which is indicated by noise from their irregular mesh topology [34,113]. Furthermore, they perform poorly during shear and torsion yielding unacceptable errors [116] and have been found to exhibit spurious deformation modes [34], making them unsuitable for contact analyses. Consequently, hexahedral meshes are the state-of-the art for finite element simulations with respect to accuracy and efficiency [5, 34, 116, 125]. They can be more computationally efficient, having a reduced number of elements per vertices compared to tetrahedral meshes [30, 38, 39, 39, 100, 113, 117, 125–128], which allow matrix assembly to be performed faster [125]. Additionally, the intrinsic structure of their basis functions as a tensor-product and efficiency of local pre-conditioners, enable high-performance computations [125].

For structured and semi-structured hexahedral meshes, it is more straight-forward to apply different material properties to particular regions of the geometries, such as the zonal properties of menisci (*Section 2.1.1*) and cartilage layers on bone [19,30]. Another benefit of structured and semi-structured hexahedral meshes for material definitions is that it can be simpler and more accurate to assign fiber directions [111]. For transversely isotropic and orthotropic materials [20, 62, 111, 129], this is done by using the local element structure [20, 81, 111], which

negates the need to define less accurate cylindrical axes [111,129]. In addition to fibers, transversely isotropic definitions can apply to biphasic material properties [130–133]. Ligaments, tendons, articular cartilage and menisci are more accurately and often defined with transversely isotropic and sometimes orthotropic properties [18–20,23,24,26,36,62,73,75,76,81,85,129–139]. Greater accuracy or geometric-specificity can be achieved due to the regular-distribution of the hexahedral elements capturing the irregular geometries and fibers more closely than cylindrical or spherical descriptions. This is not possible with tetrahedral elements which inherently are irregularly-distributed.

2.3.3 Comparative Analyses

There are numerous advantages and disadvantages associated with particular finite elements. Several studies have been reported that highlight and compare key differences in the performance of tetrahedral and hexahedral meshes. The majority of comparison studies compare linear hexahedral meshes against quadratic tetrahedral meshes [34,113,116], as several studies over the past couple of decades have determined that the performance of linear tetrahedral meshes are inferior to linear hexahedral meshes [34,39,113,116]. Overall, it appears there are clear trends in the performance of some element types: TET4 < TET10 < HEX8 < HEX20 [113,116,121]. However, it is still up for debate where TET15 and HEX27 are situated, which require further investigations. Although, based on the current trend, HEX27 would likely achieve the greatest performance.

Element Comparisons

In general, slower running speeds [34,100] and poorer convergence metrics have been observed for TET10/15 elements compared to HEX8, particularly with increasing model complexity [34]. Similar findings indicate HEX8 performs better than TET10 (e.g. stiffer) [100,116], particularly where large deformation and shear stress is dominant [116]. Also, the quadratic equivalents for TET4 and HEX8 perform better [116]. Interestingly, HEX27 has been found to be more efficient than TET15 [100], similar to HEX8 over TET4. Between tetrahedral models, TET15 has been shown to run the fastest with greater convergence behaviour [34], and provides greater robustness and accuracy [100]. At lower resolutions, TET15 has shown poorer convergence behaviour than HEX8, and all tetrahedral variants show irregularities in results due to variations in the mesh topology [34]. At higher resolutions, the quadratic elements produced results similar

to the HEX8, but still showed some mesh-based variations, whereas the TET4 mesh showed significant differences.

Some findings suggest quadratic tetrahedral (QT) meshes are less suitable for applications involving thin regions (shear locking) [100], nearly-incompressibility and frictional contact [113]. These are important findings as biomedical joint simulations can contain thin anatomies (e.g. cartilage), and a large degree of sliding interactions with high shearing forces [140]. Furthermore, particular pathologies of the knee can include severe cartilage degeneration that can lead to increased friction [141] and may initiate bone-on-bone [142] and bone-on-cartilage [141, 143] frictional contact. Additionally, this would include investigations into the consequences of abrasive particles from arthroplasties [142, 144–147] and other sources, such as osteochron-dritic fragments [144, 148]. Research has indicated that during the progression of osteoarthritis there is greater friction of the articular cartilage [149]. This indicates that the accurate modelling of osteoarthritis and arthroplasties would likely require HEX8. However, FE simulations commonly define the material properties of healthy and degenerate menisci as frictionless [18, 20, 23, 62, 73, 76, 85, 150]. Nevertheless, in reality they do impart complex frictional behaviours particularly during dynamic movement, albeit exceptionally low [74, 151].

Limitations of Comparative Studies

There are three general limitations of the comparative studies reviewed: (1) simplicity of models, (2) comparing equivalent number of vertices between different element types, and (3) lack of information regarding mesh generation methods and associated characteristics.

Complex vs. Simple Models

For some applications and conditions, the accuracy of TET10/15 has been shown to be comparable to HEX8 [34, 113]. These applications include simple test cases and biomedical models of a hip [34] and foot [113]. However, patient-specific knee models have greater complexity with respect to three aspects: (1) number of contacting interfaces, (2) irregularity of tissues (shape and surface), and (3) degree of congruency of contacting interfaces (**Section 2.2**). Of the reviewed literature, no comparative study of element types have investigated models with complexities approaching that of a knee model. For example, the hip [34] and foot [113] models were a two-body contact problem with smooth and regular anatomies that have high-congruency. In contrast, knee models can be at least three-body contact problems, depending on anatomy in-

clusion (e.g. meniscus [17] and patella [152]). Additionally, the anatomies are more irregular and have lower congruency. Consequently, it is difficult to interpret how the current findings suitably translate to other biomedical models.

Equivalence of Vertices Between Element Types

Matching the number of vertices and DOF is sometimes used when comparing meshes with different element types [34,117,128,153]. When matching the number of vertices and DOF, this can lead to a reduction in the number of tetrahedral elements, which has led to comparable convergence times for some applications [34]. In practice, there are often more tetrahedral elements than hexahedral for a given geometry [30,38,39,100,113,117,125–128]. Typically, mesh convergence studies are performed to determine optimal mesh resolutions [20,63,85,134,150,154]. Results from comparative studies have shown that fewer HEX8 elements are required to achieve convergent results [34,113]. Additionally, in some cases the highest resolution QT models still have errors greater than 5% compared to HEX8 [34]. Consequently, simply matching the number of vertices and DOF may not be a suitable approach, and could be misleading with respect to comparing convergence times. As such, further research is required to determine the suitability of matching equivalent vertices instead of conducting mesh convergence studies.

Mesh Generation Methods and Characteristics

Only one study has been found that investigated different generation methods (N=12) and quality metrics (N=3) for HEX8 and QT meshes used in biomedical models [117]. This study concluded that HEX8 meshes were the most favourable. In other studies, there was no information regarding the methods used to generate some or all of the meshes [34,113]. Also, there was no information on the surface error or mesh-quality, although this is not often reported [4]. It would be interesting to know whether the tetrahedral meshes were converted from TET4 or were decomposed directly into quadratic formulations to ensure adequate surface fitting. For practical applications, this could have a large difference on surface errors in highly irregular geometries, particularly when matching the DOF. Additionally, the element-quality has a profound impact on the accuracy of simulations (**Section 2.3.5**). There is a wide range of different mesh generation methods available, which have particular advantages and disadvantages (**Section 2.4.2**). As such, it is not clear which is the best approach to generate either QT or HEX8 meshes for certain applications, to obtain the most favourable results and performance.

Limitations of Quadratic Tetrahedral Elements

There are two key limitations related to the use of QT elements: (1) poor availability of TET15 and (2) prerequisite technical knowledge. First, TET10 has been shown capable of producing reasonable results for some applications, but it is TET15 that offers the greatest competition to HEX8 [34, 100]. However, TET15 are not as widely available in other packages outside FEBio [33], e.g. Abaqus [110]. Additionally, only one biomedical simulation has been found to use TET15, which was a comparison study of different element types [34].

Second, there appears to be some requirements of prior technical knowledge, experience and/or trial-and-error needed to develop accurate and convergent models with TET10 formulations [34, 100, 113]. This pertains to the matching of appropriate integration rules and contact implementations. These difficulties and challenges could inhibit the adoption of these elements, and could increase the complexity of model development. On one hand, this provides the flexibility to choose different composite formulations and integration rules, as some have been found to yield improved properties for particular applications [100]. On the other hand, incorrect choices have been shown to lead to highly unstable simulations [34]. Furthermore, three TET10 formulations for a foot model were investigated, where only one formulation led to simulations that converged, C3D10I (Abaqus) [113]. These were compared to the standard C3D8 formulation of HEX8. This suggests that despite the large number of HEX8 elements available [110], the standard implementation yields sufficiently accurate results for the majority for applications. Additionally, there have been concerns that recovering the stresses for certain geometries may require specific integration rules, to prevent inaccuracies near boundaries and around thin regions [34]. This could be problematic for several biomedical models, which can have components that are thin (e.g. femoral cartilage), or have very thin regions (e.g. extruded menisci). Consequently, the complexity of the decision-making process for aspects of FE model development is greater for QT elements compared to HEX8.

Conclusions

In general, there is evidence that suggests some QT formulations are viable alternatives to HEX8 for biomedical simulations [34], but this requires further investigation. It would be interesting to see what differences exist for more complex models and material properties (e.g. biphasic and fibre alignment definitions). Specifically, patient-specific models with challenging geometries

and/or pathologies, with a focus on those which push the simulated environments to the limits of convergence. Until further evidence arises, it could be argued that the convergence behaviours for QT elements continue to decrease or not converge at all with increasing model complexity. A lack of solver convergence has already been reported from a study that used TET10 in a reasonably complex model of the tibiofemoral joint [76].

2.3.4 Mesh Types used in the Tibiofemoral Joint

Depending on the tissue being simulated and application being studied, there is usually a common approach used to define them, with exceptions. For solid tissues, there have been uses of four elements in the reviewed tibiofemoral joint models, linear and quadratic tetrahedral and hexahedral elements (**Table 2.2**). HEX8 were found to be the most common, followed by in descending order of frequency used: TET4, TET10 and HEX20 (**Table 2.2** and **Figure 2.7**). **Table 2.2** summarises the choice of finite elements for particular tissues used within literature for the tibiofemoral joint. TET10 and HEX20 have been assumed where information (quadratic formulation) has not been specified or the FE package used did not support TET15 (i.e. Abaqus [110]).

Table 2.2: Element types used for tissues in tibiofemoral joint models

Tissues	TET4	HEX8	TET10	HEX20
Tibia	[73, 81, 134, 155–157]	[19–21, 36, 62, 63, 75, 92, 95]	[76]	–
Femur	[73, 81, 134, 155–157]	[19–21, 36, 62, 63, 75, 92, 95]	[76]	–
Femoral Cartilage	[156, 157]	[18–21, 23, 36, 62, 63, 74, 75, 85, 92, 95, 134, 150]	[73, 76]	[81]
Tibial Cartilage	[156, 157]	[18–21, 23, 36, 62, 63, 74, 75, 81, 85, 92, 95, 134, 150]	[73, 76]	–
Menisci	[155–157]	[18–21, 23, 36, 62, 63, 75, 81, 85, 92, 95, 134, 150]	[63, 73, 76]	–
Ligaments	[155]	[23, 36, 63, 81, 134, 150]	[23, 36, 63, 81, 134, 150]	–
Tendons	[157]	–	[73]	–
Total	21	70	16	1

Due to the performance limitations inherent with linear tetrahedral meshes, these elements are often limited to structures that have either: relatively stiff or rigid material properties; small deformations or no contact analyses associated. Consequently, hard and osseous tissues are often

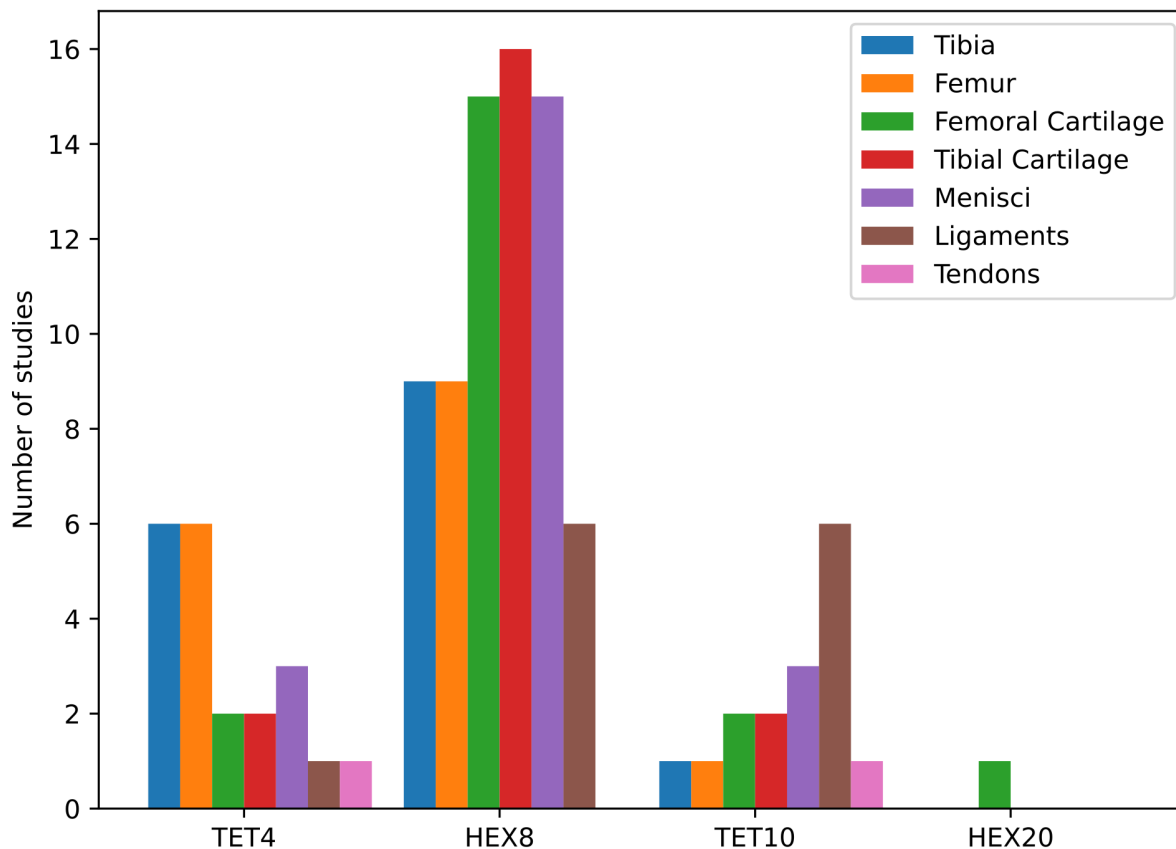


Figure 2.7: Frequency of element types used for tissues in tibiofemoral joint models

defined rigidly by linear tetrahedral elements [22, 62, 158]. However, exceptions exist which have used hexahedral [3] and quadratic tetrahedral elements [76, 134] (*Table 2.2* and *Figure 2.7*). The use of quadratic tetrahedral elements have been found to impart an unnecessary computational burden compared to TET4 and HEX8, when used for modelling the femur [127].

Soft tissues undergo the three phenomena (*Section 2.3.2*) and therefore are often meshed with hexahedral elements [29–34] (*Table 2.2* and *Figure 2.7*). The TFJ encapsulates a variety of soft tissues which include: articular cartilage, meniscus, ligaments, tendons and muscles [36, 73, 157]. Due to the highly-irregular nature of some of these geometries and lack of automated hexahedral mesh generators available, some of these tissues are meshed with quadratic tetrahedral elements [76, 134] (*Table 2.2* and *Figure 2.7*) or represented by spring elements [36]. Since 2015, the adoption of quadratic tetrahedral meshes has begun to be used in some studies, prior to this the use of these elements has not been observed (*Table 2.2*). The use of quadratic hexahedral elements in tibiofemoral joint modelling is more rare, however, there has been one study that have used such elements to mesh femoral cartilage [81] (*Table*

2.2). In that study, the reason for using quadratic hexahedral elements was not given but may be due to a lack of solver convergence or poor surface fitting from the linear hexahedral mesh. Furthermore, quadratic hexahedral elements (HEX20*) were only used for the femoral cartilage, whereas, linear hexahedral elements (HEX8) were used for the tibial cartilage and meniscus, again without justification [81]. Linear hexahedral meshes are most often used for articular cartilage, menisci and ligaments (**Table 2.2** and **Figure 2.7**). Only two examples of tendon meshing have been observed, where linear [157] and quadratic tetrahedral elements were used for the patella tendon [73] (**Table 2.2**).

Of the reviewed literature, there are currently no TFJ models involving HEX27 definitions for any of the tissues (**Table 2.2**). This is most likely due to the challenges of generating hexahedral meshes, increased computational costs and the lack of convergence issues for particular problems.

2.3.5 Element Quality Metrics

The main purpose of element-quality metrics used in FE Models are to determine the validity of a mesh and the likelihood it will lead to a convergent and accurate simulation [4, 30, 159–161]. However, the quality of finite-element meshes are rarely reported for biomedical models in literature [4]. An extensive overview for the reporting of quality assessments for bone models was conducted and showed only 5 in 80 studies used mesh-quality metrics [4]. Furthermore, little information was provided to what metrics and minimum criteria were required, and where low-quality elements reside [4]. This information is important as the general quality of a mesh and position of low-quality elements can cause significant differences to the accuracy of the solution, that cannot be overlooked through mesh-sensitivity analysis [4, 161]. Additionally, low-quality meshes will often reduce the time taken to reach convergence and increase the overall solution times, with the possibility of not converging entirely [161, 162]. Mesh-quality metrics are also used during the solving process, to ensure significant distortion has not been introduced that would prevent solver convergence [111, 159].

There is a natural trade-off between the degree of high-quality elements and the accuracy of an anatomical surface fit [159]. To improve the fitting of anatomical irregularities requires more elements to deviate from element ideals [159] (**Figure 2.8**). This is another important reason why mesh-quality analyses should be reported in biomedical models, due to the simulation of highly-irregular geometries.




Figure 2.8: Ideal or reference hexahedron used to compare against a hexahedral element with respect to the Jacobian, principal axes (P1, P2, P3), dihedral angles (α , β and γ) and ratio of minimum and maximum edge lengths (*adapted from [4, 5]*)

There are a wide array of metrics available and some are not always cross-compatible with other element types [108, 163], which can make comparisons between the performance of different meshes difficult. To complicate matters, different metrics are more relevant than others depending on the particular application and model [161].

The majority of metrics focus on the degree of distortion or deviations from element ideals [4, 30, 117, 160, 163–165]. However, others can indicate purely geometric properties such as volume [38, 160] or topological issues, such as singularities [165–167].

Of the variety of available metrics [108, 163], four common metrics used to analyse the distortion of a hexahedral element are the: Jacobian, aspect-ratio, skewness and dihedral angle [4, 30, 117, 159, 160, 164, 168, 169]. The benefit of these metrics are that they can be used for comparing the suitability between tetrahedral and hexahedral meshes, when undertaking application-specific performance analyses [117, 170].

Determining the validity of tetrahedral meshes can be less trivial than for hexahedral meshes and can be assumed to be valid if they contain a positive volume [5]. Also, accurate analysis of quadratic equivalents for elements are more difficult to assess [5].

Jacobian

The Jacobian or scaled-Jacobian metric is often the standard to determining the validity of a mesh, through the determination that all elements have a positive value [4, 5, 30, 108, 161, 166, 169].

The Jacobian of an element's vertices and centre is calculated from the determinant of the Jacobian matrix (*Equation 2.1*) [4, 108]:

$$J = \det \begin{bmatrix} \frac{\delta x}{\delta \xi} & \frac{\delta y}{\delta \xi} & \frac{\delta z}{\delta \xi} \\ \frac{\delta x}{\delta \nu} & \frac{\delta y}{\delta \nu} & \frac{\delta z}{\delta \nu} \\ \frac{\delta x}{\delta \mu} & \frac{\delta y}{\delta \mu} & \frac{\delta z}{\delta \mu} \end{bmatrix} \quad (2.1)$$

Essentially, this provides a pointwise map of each vertex and centre of a hexahedron within a mesh to the corresponding ideal [108, 117, 163] (*Figure 2.8*). The Jacobian is the minimum value derived from an element [108]. Element information provided by this calculation include distortion, volume, orientation and shape deviations to element ideals [4, 163]. Consequently, several metrics are based on the Jacobian matrix, such as: distortion, shape and shear [108].

A valid element has a positive value (J) at each vertex and centre [108, 166]. However, it was reported that positive values on the corners and edges did not guarantee validity [171]. An element with a non-positive Jacobian is classified as highly-distorted or inverted [4]. The presence of negative Jacobians in a mesh will lead to either nonsensical results or failure achieving convergence [30, 172]. If a non-positive Jacobian is encountered during the solution of the PDEs, FE packages will terminate further computations [111]. The value must be positive because it is fundamental to how finite element solutions are determined using reference mapping and requires the point-wise map to be injective (one-to-one function) [5, 103, 161, 173]. Consequently, if an element is highly-distorted or inverted, injectivity is lost [5] and the solution would be inaccurate or not converge [30, 161], potentially causing non-physical behaviors in the process [161].

The scaled-Jacobian is a normalised equivalent of the Jacobian, achieved through the division of each edge length [108]. Consequently, the Jacobian is influenced by the size of the element, whereas, the scaled-Jacobian will always reside between -1 and 1. The Jacobian is sufficient where concern is only to determine mesh and element validity of a single mesh. However, the scaled-Jacobian is superior when comparing different meshes, regardless of size variabilities. Scaled-Jacobian values between 0.5-1.0 have previously been defined as acceptable [108], with high-quality elements having values above 0.8 [30]. Although, there does not appear to be a consensus to what constitutes a high or low-quality element. Evidence indicates that the average scaled-Jacobian affects the accuracy and convergence of a simulation [166]. However, further

research is required to determine if these metrics alone are sufficient for quantifying hexahedral mesh-quality and validity [5, 166, 171].

Non-Jacobian-based metrics

There are several metrics that are not based on the Jacobian matrix and measures how close an element is to the ideal. These metrics are not calculated with respect to the ideal or reference hexahedron but compared to it based on their associated values.

The skew measures the angle between the three principal axes (P1, P2, P3) of a hexahedron to determine the degree of parallelism between pairs [108] (**Figure 2.8**). The maximum value between the three pairs is used to measure the skewness of a hexahedron to indicate the closeness of the element to the corresponding ideal [108, 117]. For this property, the skew metric has been considered a primary quality measure [117]. The values for this metric exist between 0 and 1, where 0 to 0.5 is the acceptability range and a value close to zero indicates the element of a mesh is close to an ideal hexahedron [108, 117].

The aspect-ratio or edge-ratio of an element is simply the longest edge divided by the shortest edge (**Figure 2.8**) [4, 108]. Several studies have investigated which ranges provide the most accurate results with acceptable ranges reported between 1 and 4, with the quantity of elements above 3 not exceeding 5% [4, 117].

A dihedral angle is one defined from two planes. Each vertex in a hexahedron has three dihedral angles (α , β and γ) [4] (**Figure 2.8**). The dihedral angles for an element would ideally be 90° [4]. However, they have been deemed acceptable if they lie between 45° and 135° [160, 164, 174], or more conservatively between 60° and 120° [4, 175]. Outside these ranges they have been classification as distorted [160, 164]. For highly-irregular geometries, such as those from biological origin, there should be no more than 5% of distorted elements below 20° or above 160° [4]. It has been reported that a simulation may not converge if the dihedral angles approach 0° or 180° [5, 164].

2.4 Hexahedral Mesh Generation

There are a wide variety of algorithms that have been developed to generate hexahedral meshes based on a geometry of interest (**Table 2.3**). However, the problem of generating hexahe-

dral meshes for arbitrary geometries has remained since the 1950s [41]. Since then, solutions have gradually moved from manual to semi-automatic and automatic, which could be classified as either general-purpose or geometry-specific [39, 40, 176, 177]. Unlike tetrahedral mesh generators, none can derive high-quality meshes in a fully-automatic and timely manner for arbitrary geometries [39, 177]. Unfortunately, the techniques that enable the automation of other element types do not translate well to hexahedral meshes and yield comparatively poorer results [178–180]. The generation of high-quality hexahedral meshes represents one of the main barriers to the development of state-of-the-art FE models [176, 181]. The first section provides an overview of the variety and types of generation methods currently available, as well as definitions and criteria that characterise them. This is followed by a description of their advantages and disadvantages, and concluded with a summary of their suitability for biomedical anatomies and applications.

2.4.1 Overview

There are 14 criteria that define an ideal general-purpose mesh generation method. As no hexahedral mesh generator has yet achieved them [177, 182, 183], they aid in distinguishing methods with respect to this idealisation. Additionally, there are several definitions that aid in describing the types of methods and meshes they derive, which are: the regularity, mesh type, meshing approach and method classification.

Evaluation Criteria

Twelve criteria were previously proposed to characterise and compare automatic hexahedral mesh generators [184, 185]. Originally, nine criteria were proposed (C1-9) [184], which were later extended by three (C10-12) [185]. An additional criterion (C13) has now been added to represent a common metric used to determine the accuracy of meshes derived from bioimaging data. Also, one of the original criterion (C10) has now been split into two: generation time (C10) and user-interaction (C14). The 14 criteria with brief descriptions are the following [184, 185]:

- C1. Geometric Domain:** The types of geometries that can be meshed
- C2. Non-geometric Domain:** Additional information, which includes element size requirements and simulation information, such as: loads, constraints and material properties
- C3. Range:** The type of meshes that can be derived, e.g. hexahedral, tetrahedral, etc.

- C4. Locality:** A modification of small regions will lead to localised differences
- C5. Continuity:** Small dimensional modifications have proportional mesh differences
- C6. Orientation-Independence:** Orientation changes will not affect the meshing process
- C7. Representation-Independence:** Geometry representation will not affect the meshing process, i.e. number and type of discretisation, as well as the order of vertices and elements
- C8. Idempotence:** Internal boundary definitions will affect the meshing process
- C9. Local Density Control:** An unrefined mesh is a proper subset following local refinement
- C10. Generation Time:** The amount of time required to derive a mesh
- C11. Mesh-Quality:** Minimum quality criteria required for chosen metrics
- C12. Boundary Sensitivity:** The meshing process derives high-quality boundary elements
- C13. Surface Fit:** The accuracy of a derived mesh's boundary to a target geometry
- C14. User-interaction:** The amount of user-interaction, knowledge and skill required

Essentially, the original nine cover general robustness and consistency, but don't necessarily infer two important characteristics for biomedical simulations: mesh-quality and surface fit. The addition of surface fit (C13) was necessary as it is an important metric for biomedical models [126]. However, for some methods, this characteristic is not a problem for more regular CAD-based geometries, which can have a perfect surface fit. Generation time (C10) was separated into two criteria (C10 and C14) to serve two purposes. First, this would aid in identifying a faster or slower method from automatic and semi-automatic methods. This distinguishes automatic methods that may take as long as semi-automatic methods but with the benefit of not requiring user-interaction. Second, different levels of proficiency and types of skill and knowledge are necessary to derive high-quality meshes from semi-automatic methods.

Despite the thoroughness of the previously reported criteria (C1-12), only two could be found comparing their meshing algorithm to them [176,186]. Others have mentioned a subset of these criteria indirectly and without reference [39,42,180,187,188]. Typically, some criteria are less important, particularly if more critical criteria are not satisfied. For example, if a method generates invalid elements consistently or takes an extraordinary amount of user-iteration then the remaining factors are irrelevant. However, the list is quite complete and shows not only

how methods compare but how far they are from an ideal general-purpose fully-automatic mesh generator.

Arguably, the most important and relevant criteria for biomedical models are: surface fit, mesh-quality and generation time. An adequate surface fit provides greater accuracy for patient-specific models and would be required for state-of-the-art clinical applications [189]. High-quality meshes are also required for accurate solutions and can improve convergence times [190–192]. The rapid generation of high-quality patient-specific meshes (ideally fully-automated) will be required for the widespread adoption of biomedical modelling for clinical applications and large scale studies [181]. Following these three criteria, the next most important factors are: geometric domain and boundary sensitivity. Ideally, a wide-range of anatomies could be meshed using the same or similar process, to approach the large variety of systems operating in the human body. For applications involving contact mechanics, common for musculoskeletal tissues, boundary sensitivity would enable greater solver performance and accuracy [193]. Also, continuity, representation-independence and orientation-independence could be quite useful for some comparative studies. The criteria which relate to the most critical advantages and disadvantages for each method classification will be highlighted.

Regularity

A key distinction between different generation methods is the regularity of the derived hexahedral meshes. The regularity of a mesh can be separated into three categories:

1. **Structured:**

These meshes have a high-degree of regularity and ordering in all directions [41]. Generally, several characteristics can be defined parametrically based on an array or sequence of indices (e.g. i , j and k), which include the elements, vertices and neighbourhoods [41, 126, 194]. Typically, the elements are orthogonal and there are no internal singularities present within these meshes, e.g. a 3D grid or blocks of hexahedral elements [188, 194].

2. **Semi-structured:**

These meshes have less regularity and are only structured in one direction [194]. They can be parameterised similarly to structured meshes (e.g. i , j and k), however the indices can not be used to define all the same characteristics implicitly. Also, there are a larger

quantity of internal singularities in these meshes than those that are structured [195].

3. Unstructured:

There is no regularity in these meshes and have the highest proportion of internal singularities [196]. Also, the characteristics cannot be determined based on parametric definitions and must be defined explicitly [38,194]. However, these meshes can be composed of regions with structured and semi-structured meshes [194].

There are several key distinctions that can be made with respect to mesh regularity. Less regular meshes can often accommodate geometries with greater complexity and reduce excessive distortions, due to the presence of internal singularities [35,197,198]. However, FE solutions can be computed with greater efficiency and accuracy with increasing mesh regularity [39,169,199–201]. Also, unstructured meshes can reduce the accuracy of model definitions that require aligned topologies [35,197,198]. In contrast, meshes with greater regularity can improve the accuracy and simplicity of alignment definitions. These definitions can include boundary alignment, for computational fluid dynamics, and composite materials [177,201,202], particularly collagen fibre orientation for biomedical applications [20,63].

Mesh Types

Another key characteristic of methods are the purity or type of mesh they derive, which are either all-hexahedral (all-hex) or hexahedral-dominant (hex-dominant) [183,185]. All-hex meshes contain only hexahedral elements and can be either structured, semi-structured or unstructured [39,177]. In contrast, hex-dominant defines hybrid meshes that contain a higher proportion of hexahedral elements in addition to others, notably prisms and tetrahedral elements [203]. Hex-dominant meshes can reduce element distortions [177,183] and mesh complex geometries with greater ease [204], but are usually unstructured [183]. However, all-hex meshes often perform with greater solver efficiency and accuracy, due to the lack of less robust elements [177,183,185]. Additionally, all-hex meshes tend to have greater regularity, which can provide several benefits. Consequently, the majority of methods discussed will be based on all-hex meshing.

Meshing Approaches

Hexahedral mesh generation methods can also be classified into three categories based on the approach used to derive a mesh, either: outside-in, inside-out or decomposition [177,185]. Methods related to these approaches can be generalised with respect to several evaluation criteria, particularly: geometric domain (C1), orientation-independence (C8), mesh-quality (C11) and boundary-sensitivity (C12).

Outside-in methods start at the boundaries and progressively work their way towards the core of a geometry, often employing an advancing front algorithm [35,185]. Advancing front approaches can achieve high-quality meshes for triangular, tetrahedral and quadrilateral elements [39,178,179,182]. However, the robustness and mesh-quality benefits do not transfer well to hexahedral elements [178,179,188,205]. Generally, these meshes have higher quality and more regular elements towards the boundaries but the opposite towards the core (boundary-sensitive) [185,206]. Additionally, there is often an *internal void problem* that occurs where the advancing boundaries intersect into a region that either cannot be meshed or yield poor quality elements [35,177,205,206]. This often requires the transition to hex-dominant from a previously all-hex definition to remove an over-constrained void or the excessively distorted elements [177,207,208]. Additionally, the mesh-quality and topology can vary greatly between geometric domains.

Inside-out approaches operate in the opposite direction to outside-in methods, by first defining elements internally and working outwards towards the boundary [35,185]. The characteristics of these methods tend also to be the opposite of the outside-in methods. Notably, they have structured high-quality internal elements with irregular and lower quality elements on the boundary (boundary-insensitive) [35,185].

Decomposition-based methods discretise a geometry into either hexahedral elements, blocks or primitive shapes [177,185]. They represent the class of methods with the largest variety. Generally, they are boundary-sensitive, orientation-independent and can produce high-quality meshes. None of the approaches can provide high-quality meshes for arbitrary geometries.

Method Classifications

There are five categories a mesh generation algorithm can be placed: (1) mapped, (2) direct, (3) mapped-direct, (4) indirect and (5) superposition (**Table 2.3**). The following provides a

brief overview of each method:

1. **Mapped:**

Mapped methods generate hexahedral meshes by mapping a geometry to an existing mesh and are decomposition-based methods [39, 178, 185]. These meshes can be either structured, semi-structured or unstructured, based on how an existing mesh was defined [177, 206]. However, structured and semi-structured meshes are usually derived using these methods due to their benefits. The pre-defined meshes tend to be based on primitive structures, such as a cube or blocks [177, 185]. Although, more complicated shapes can be mapped to similar target geometries [36, 177] or through geometry decomposition [178, 185, 208, 209]. This class represents a sub-set of methods that are currently considered state-of-the-art [39, 177, 188, 197, 210, 211].

2. **Direct:**

Direct methods are algorithms that directly discretise a geometry into hexahedral elements [39]. These approaches either build meshes element by element or employ the structuring (weaving) of hexahedral chords and sheets (dual-methods) [177]. Generally, they are outside-in methods and produce unstructured meshes [185].

3. **Mapped-Direct:**

Mapped-direct methods combine aspects of mapped and direct methods. Specifically, a hexahedral mesh is generated directly with a similar target geometry, which is then mapped onto a target boundary (decomposition-based methods). In general, mapped-direct meshes are more geometry specific with respect to singularities than the other methods. Their greater control over singularities can lead to reduced element distortions [188, 198]. These meshes can be either structured [212, 213] or unstructured [183, 192, 214]. Some of these methods could arguably be considered purely mapped or direct methods [177], but can provide greater control with respect to singularities and automation of an existing mesh.

4. **Indirect:**

Indirect methods involve the conversion of a tetrahedral mesh into a hexahedral mesh [39, 215]. There are two ways this can be achieved: (1) merge tetrahedra into hexahedral

elements (outside-in method) or (2) dice each tetrahedron into four hexahedral elements (decomposition-based method) [215,216]. These meshes are often highly unstructured and tend to have a higher degree of singularities than the methods from the other categories [39,167,183].

5. Superposition:

Superposition methods use grid structures to directly define all or part of a domain [177, 206]. They are typically classified as inside-out methods [177]. These meshes could be considered direct methods, based on that definition, however they are distinctly different with respect to topology and structure. These meshes often have a shell of unstructured hexahedral elements surrounding a structured core, but can be purely structured [177]. Consequently, they deserve their own class as the other methods tend to derive meshes that are only structured, semi-structured or unstructured, whereas these meshes usually represent a particular composition. Subsequently, some consider them direct methods [39], whereas others place them in their own class [178].

Table 2.3: Method classifications and associated algorithms

Mapped	Direct	Mapped-Direct	Indirect	Superposition
Mapping [210, 217]	Whisker weaving [205, 218]	Feature-based [219, 220]	H-Morph [208]	Grid [186, 221]
Submapping [209, 222]	Plastering [182, 223]	Skeleton-based [†] [199, 214]	Dicing [215]	Octree [224, 225]
Sweeping [209, 226]	Receding front [35, 206]	Frame-field* [192, 227]		Voxel [228, 229]
Multi-block [38, 230]		PolyCube [198, 213]		
Mesh-morphing [231, 232]				

[†] Skeleton-based methods include the earlier developed medial-axis algorithms [203]

* Frame-field used as an umbrella term for frame, cross [192] and vector-field algorithms [183]

2.4.2 Advantages and Disadvantages

Mapped methods

The greatest advantage of the mapped methods is that the derived meshes tend to have high-quality elements [11, 38, 209]. Also, they can guarantee equivalent topologies for geometry variations and provide the greatest control over the topology [39, 197]. For biomedical ap-

plications, this can allow the controlled design of aligned topologies, which can improve the accuracy and simplicity of anisotropic material properties [20, 63, 233, 234]. Similar benefits and requirements can be obtained for computational fluid dynamics and the simulation of composite materials [177, 201, 202]. Additionally, this can be useful for design optimisation [13, 235] and comparisons between simulations with similar geometries, to accurately determine differences in output variables [235]. Other advantages for mapped methods are that they are boundary-sensitive, orientation-independent and representation-independent [39, 42, 185]. Additionally, they can provide accurate surface fitting for a sub-set of geometries [39]. For these reasons, they represent some of the gold-standard methods for generating state-of-the-art FE models [178, 197], particularly the multi-block method for biomedical applications [19, 21, 36].

Mapping [210, 217], submapping [209, 222] and multi-block [38, 201] methods are common approaches for producing structured meshes [177], whereas sweeping [209, 226] is generally used to derive semi-structured meshes [206]. For mesh-morphing methods, the degree of regularity is entirely dependent on the template mesh, which could be either structured, semi-structured or unstructured.

Most mapping methods consist of semi-automatic algorithms, where user-interaction can vary quite significantly between methods and their variants. The degree of user-interaction and required technical knowledge have the most profound impact on generation times [211]. The earlier methods (mapping, submapping and sweeping) tend to require a greater amount of user-interaction [209, 210, 236] than the more recent approaches (multi-block and mesh-morphing) [48, 126, 211, 231]. Also, semi-automatic mesh-morphing approaches can require even less user-interaction than a multi-block method [48, 51]. Currently, only the mesh-morphing algorithms have the possibility of providing automatic solutions of specific geometries for this class of methods [231, 232].

In general, the mapped methods represent some of the fastest and efficient methods for generating high-quality hexahedral meshes for certain geometries [46, 208, 209, 235, 237]. However, the semi-automatic methods can take up to several days for more complex models [199]. A key disadvantage of the earlier mapped methods are their limitations with respect to meshing a variety of geometric domains [208, 209]. For mapping, submapping and sweeping, they can only produce high-quality meshes rapidly for simple geometries or those that can be decomposed into primitive shapes [199, 208, 209]. The decomposition of complex objects into simpler shapes has been

used to increase the range of geometric domains for submapping and sweeping methods. However, this comes at the cost of speed, efficiency and technical skill requirements [182, 208, 209]. Mesh-morphing methods can derive high-quality meshes reasonably fast for a wide range of geometric domains, but only for geometries similar to a pre-defined template mesh [237]. In contrast, the multi-block method can derive high-quality meshes for a wide-range of geometric domains. However, they require greater time, user-interaction and technical knowledge for more complex geometries [197, 211, 238].

The multi-block and mesh-morphing methods have been used extensively to derive hexahedral meshes for a wide variety of anatomies. Multi-block methods have successfully meshed the following: vascular structures [201], tibias [3, 21], femurs [3, 21, 239, 240], femoral cartilages [3, 21, 239, 240], tibial cartilages [3, 19, 21], brains with isolated regions [211], kidneys [241], ligaments [3, 238], menisci [3, 19, 21], tendons [242], nerves [242] and functional spinal units [238]. Anatomical mesh-morphing examples include: ventricles [45], vertebrae [48], phalanges [46, 237], patellas [27], tibial cartilages [27], femoral cartilages [27], brains [51], faces [243], fingers [244], kidneys [241] and functional spinal units [245].

Direct Methods

There are few purely direct methods for generating hexahedral meshes: plastering [182, 187, 223], whisker weaving [205, 218, 246] and receding front [206]. The main advantages of these methods and their variants are that they are fully-automatic and general-purpose [187, 206, 246]. Other advantages of these methods are orientation-independence, boundary-sensitivity and the ability to fit surfaces with high-accuracy [39, 177, 206].

The whisker weaving (topological approach) and plastering (geometrical approach) algorithms define outside-in (advancing front) approaches [167, 177], and both suffer from the associated issues [178, 179, 188, 205]. Generally, acceptable mesh-quality can only be derived for simple geometries, which becomes poorer with increasing complexity [35, 183, 205]. Additionally, they require a high-quality quadrilateral surface representation [39], which must have an even number of elements to ensure the derivation of an all-hex mesh [208]. Although high-quality quadrilateral meshing algorithms have now been developed [112, 210, 247], this provides an additional step and dependency for the performance of a given solution. However, regardless of the quadrilateral mesh-quality, the topology will also have a significant impact on the derived hexahedral mesh

and dictate the final hexahedral topology [39]. Also, a degree of post-processing is often required to improve the topology and subsequent mesh-quality [39].

Receding front methods have more recently been established, which operate similar to plastering methods but using an inside-out approach [35, 206]. These methods attempt to overcome the internal void issues and low-quality internal elements of outside-in methods, whilst maintaining boundary sensitivity [35, 206]. These methods seek to determine a compromise between grid-based (inside-out) and advancing front methods (outside-in) [35, 177, 206]. However, as the receding front algorithms are relatively new, further research is required to extend them to more complex geometries [35, 206].

The main disadvantage of these methods are that they cannot achieve high-quality meshes for complex geometries [183, 188, 205, 206]. Another key disadvantage of the established approaches lies in the uncertainty of not knowing a priori if a hex-dominant mesh will be required to overcome the internal void issues [177, 182]. Furthermore, the established methods have poor representation-independence, as they are susceptible to the topology and processing of a quadrilateral mesh [39]. Also, they all derive unstructured meshes which have several shortcomings, as previously mentioned [35, 197, 198].

Mapped-Direct Methods

Currently, this class of algorithms show the most promise and novel ideas to approach the challenge of general-purpose automatic hexahedral meshing. Consequently, they highlight how this challenge has led to approaches with increasing complexity and sophistication. The main advantage of these approaches are the derivation of high-quality hexahedral meshes for a wide range of geometries [169, 177, 188, 198, 212]. Additionally, hexahedral meshes generated by these methods tend to be boundary-sensitive, orientation-independent, representation-independent and provide accurate surface fitting [39, 169, 177, 188, 212]. The frame-field [188, 227] and poly-cube [198, 214] approaches offer the most promising results of this class. Generally, they derive good quality meshes but can still lead to high-distortion mappings [188, 213, 248].

The methods can be split into automatic and semi-automatic methods. The automatic methods have the added benefit of being able to produce meshes reasonably fast, often within several minutes [198, 212, 249, 250]. However, some semi-automatic methods require significant user-interaction and technical knowledge, which can lead to shortcomings similar to the multi-block

method. In particular, user-interaction increases with geometric complexity and can lead to notable mesh-quality differences between users. Two examples include the definition of frame-fields and polycubes, which can be impractical for particular geometries [198,227,251]. However, there has recently been progress on automating these structures [188,198,252]. Although, none of the approaches from this class of methods have perfected the process. This has led to the compromise of hex-dominant meshes, due to regions with high-distortion [183,252,253].

The advantages of these methods stem from the focus on the singularity structure of a particular geometric domain. This enables some to have genus neutral approaches, such as the frame-field [188] and polycube methods [198]. Also, several methods have some have shown impressive results for both CAD-based and natural objects, with irregular geometries and smooth surfaces, e.g. frame-field [188,192,227], polycube [198,213], and general-purpose skeleton-based methods [199,214]. However, few have been tested on highly irregular anatomies derived from bioimaging data. Frame-field methods have achieved great results meshing a skull [227] and hand [192,227]. Polycube methods have also achieved notable results generating meshes for a femur [213], hand [198,213] and skull [251]. Additionally, skeleton-based methods have achieved equally impressive results for vascular structures [169,199,212]. However, the surfaces in these cases were very smooth with minimal to no surface irregularities. It is yet to be seen how these algorithms perform on a wider range of anatomies and those containing surface irregularities. Furthermore, what degree of post-processing would be required for image segmentations and how that would affect simulation performance.

The semi-automatic polycube and frame-field methods could be classified as mapped, as they are similar to the multi-block and submapping methods [213]. However, despite a general solution for these structures not yet complete, there have been several attempts to achieve this [198,249,250]. In contrast, multi-block methods have been classified as mapped methods, due to their wide-spread use as semi-automatic methods [39]. There have been attempts to automate the generation of multi-block structures, with examples using skeleton structures [169] and frame-fields [197]. In this thesis, such methods would be classified as skeleton-based or frame-field approaches, as the multi-blocks for those algorithms were decomposed directly, instead of using pre-existing meshes. Similarly, other methods that derive multi-blocks directly, would be placed in the mapped-direct category. Furthermore, they would likely be categorised based on how they achieved the direct decomposition, similar those methods. Also, the polycube

methods have similarities to superposition methods but do not exhibit their general structured and unstructured regions, as well as other dissimilarities. Typically, polycube methods produce higher quality structured meshes [213].

Although promising, the main disadvantage of these approaches stem from the lack of robustness testing for a diverse range of anatomical geometries. Despite being applicable to a wide range of natural geometries [169, 188, 212], they have been reported to have poor robustness and reliability for geometries with increasing complexity [183, 198, 248]. Furthermore, there is a lack of data regarding the simulation performance. Also, due to the difficulty of deriving high-quality all-hex meshes, approaches deriving the less preferable hex-dominant meshes have been considered [177, 183, 252].

Indirect

There are two diametrically opposed algorithms for indirect meshing: H-morph [208] and tetrahedral dicing [215]. The implementations and associated properties are quite different for these approaches. A key difference between the two approaches resides in their boundary-sensitivity. However, due to relying on a base tetrahedral mesh there are some similar advantages and disadvantages. The main advantages of these methods are that they are automated and can be applied to a wide range of geometric domains [39, 177]. Additionally, they are orientation-independent and provide accurate surface fitting.

Despite the name, H-morph is not a mesh-morphing method but an advancing front algorithm that combines tetrahedral elements to derive an all-hex or hex-dominant mesh [208]. It is similar to the advancing front algorithms described for the direct methods (particularly plastering), but is based on an existing tetrahedral mesh, hence an indirect method. Consequently, it has all the same shortcomings as those methods. Specifically, increasing geometric complexity will lead to poorer quality elements and the associated hex-dominant problems [177, 207, 208].

Dicing algorithms provide great robustness and speed for generating hexahedral meshes [177, 215]. The speed is a result of simple processing operations, tet-to-hex subdivision and re-organisation of the additional vertices [215]. Consequently, the geometric domain range for this algorithm is very large and limited only to the performance of the tetrahedral mesh generator. However, the topology and mesh-quality of the derived meshes are generally very poor for the majority of applications and can lead to highly distorted hexahedral elements [41, 177]. The

element shapes are far from an idealised hexahedron [177] and there are only a few applications where meshes with this type of topology are acceptable [216]. Also, these meshes can have an excessively large number of elements compared to those derived using other methods [254]. Therefore, the speed processing benefits of generating the meshes would come at the detriment of increased convergence times.

Both these types of algorithms have poor representation-independence, as the topology is strongly influenced by the base tetrahedral mesh and not the geometry [183]. Also, similar to the direct methods, H-morph will have a different topology based on the order of processing and combining the tetrahedral elements. Although H-morph can derive higher quality elements with less topological irregularity than dicing, they both tend to derive poor quality elements with irregular topologies compared to other methods [177].

Superposition

The main advantage of these methods is their ability to mesh arbitrary geometric domains automatically [167,255]. Additionally, these methods have strong representation-independence, as only the boundaries of the geometries are considered and not the discretisation. A common approach of these methods is to use either an octree or structured grid of elements to define the majority of a geometry's volume. Then, either project the boundary of those elements [177] or fill the remaining regions between the boundaries with elements [206]. Also, a pure grid of voxels has been considered, which can also be generated directly from 3D imaging data [229].

Despite being straightforward to implement, the octree and grid-based methods have all the problems associated with inside-out methods. Specifically, they have highly irregular and distorted elements on the boundaries [167,177,255]. There are a variety of octree and grid-based variants which all have the same shortcomings [221,224,225]. Consequently, these meshes are generally poor for some applications [256], such as contact analysis, which require high-quality boundaries [193]. Although not historically correct, these approaches could be considered extensions of voxel-based meshing, as they append or project a smooth boundary to an otherwise voxelated mesh.

Voxel-based meshing can generate hexahedral meshes incredibly fast and is a straightforward algorithm to implement [211,229]. However, they contain the most poorly defined geometric boundaries of all the methods. Also, they suffer from poor surface fitting due to a stair-stepping

effect, preventing their use for contact analyses and large deformations [41,211]. However, there are currently new solvers being developed to accommodate their stair-stepped boundaries [257].

Superposition methods also suffer two major disadvantages: (1) poor continuity and (2) orientation-independence [199,213,255]. The position of the reference system (i.e. grid, octree or voxel axes) relative the geometry leads to strong differences in the element cut-off along the boundary. Consequently, subtle translational and rotational changes of a geometry or reference system will lead to mesh differences.

Despite the drawbacks associated with the octree and grid-based methods, they have remained a common choice for industry due to their robustness, simplicity and derivation of all-hex meshes [42,199,213]. Additionally, several anatomies have been meshed using these methods, such as: a brain [221,255], section of cancellous bone [258], femur [224,258], hand [221], head [221,224], hip (ilium) [258], pelvis [224], phalanx [258], liver [224], skull [221] and torso [224].

2.4.3 Summary and Conclusions

Unfortunately, all of the reviewed methods fall into at least one category of inadequacies:

1. Inability to generate high-quality meshes for arbitrary geometries
2. Require extensive time, user-interaction and technical skill
3. Not well-established or thoroughly tested with respect to anatomical range and robustness

A large majority are unable to satisfy the challenging objective of deriving high-quality hexahedral meshes for complex geometries, which includes the following methods: mapping, submapping, sweeping, mesh-morphing, direct, indirect and superposition. The few that have shown the capacity to overcome this challenge have either high requirements for user-interaction or have not satisfactorily proved their robustness for wide-spread adoption. Specifically, the multi-block and semi-automatic variants of the polycube and frame-field methods have yielded impressive results, but at the expense of important resources. In contrast, the automatic mapped-direct methods have provided the most promising solutions for the emergence of a general-purpose fully-automatic hexahedral mesh generator to date. However, due to the lack of adoption (i.e. FE studies) and demonstrated robustness, it is likely too early to declare a victor of this extraordinary challenge. Most of the mapped-direct methods are still works in progress, but potentially have not yet gained popularity due to their recent arrival to the simulation communities. A

compromise of several methods are to derive hex-dominant meshes [177, 183, 252]. However, all-hex meshes can provide improved solver efficiency [39, 169, 199–201] and precise alignment definitions for meshes with greater regularity [177, 201, 202].

A modern comprehensive analysis of the recent methods is required for biomedical applications, similar to that performed by Viceconti et al. over 20 years ago [229]. This should focus on challenging and demanding simulations with the three core phenomena, similar to that performed for a small sub-set of methods more recently [117]. This would be worthwhile, as it is not clear how close possible solutions have approached the problem of automatic hexahedral meshing for particular anatomies and structures. A task of this magnitude may be demanding, as hexahedral meshing and pre-processing applications are rarely designed with biomedical applications in mind [38]. However, contrary to the emergence of novel methods, mesh-morphing would still provide overall development benefits whether used in isolation or in combination with a novel method. A novel method could generate a high-quality decomposition of an anatomical structure, but would not provide model definitions, such as: material properties, boundary conditions, assembly and interacting surfaces. In contrast, a morphing method could transfer model definitions directly into a given decomposition, which would speed up model development times. For these reasons and the overall advantages, the development of an automatic mesh-morphing strategy has been chosen in this thesis, to solve the problem of automatic mesh-generation for anatomies. Two additional methods discussed in this section have been used in this thesis: (1) a parametrised sweeping method and (2) a multi-block method. These are state-of-the-art procedures for generating high-quality hexahedral meshes, and are both used to create template meshes. Additionally, the multi-block method is used in two validation methodologies. The purpose of this is to determine how the strategies' performance compares to a state-of-the-art method for generating hexahedral meshes for biomedical applications.

2.5 Mesh-Morphing Algorithms

The origins of mesh-morphing began in antiquity with cartography, by mapping the earth to 2D objects [259, 260]. The concepts of different projection techniques and their effects have been known and studied for millenia, and have had a significant influence on modern surface parameterisation and mesh-morphing techniques [260]. Today, a mesh-morphing strategy consists of a series of algorithms that modify an existing mesh to match the boundary prescribed by

another. In this thesis, mesh-morphing describes the modification of a volumetric mesh. Typically, a mesh-morphing strategy requires two objects: a target geometry and template mesh. The target geometry is often represented by a triangular surface mesh, although other representations can be used, e.g. bioimaging data [57]. The template is usually based on a high-quality mesh with a similar geometry to a target. In general, there are three core components of a mesh-morphing strategy:

1. Feature alignment: orientation of objects to minimise correspondence differences
2. Surface parameterisation: determination of correspondences between two objects
3. Volumetric transformation: modification of vertices to match the correspondences

Each component builds on the previous, where greater accuracy for one component leads to greater accuracy of the subsequent components. For each component, there are a larger variety of algorithms that can address the problem.

2.5.1 Feature Alignment

Feature alignment, often referred to as rigid registration, is a process used to align or orient two or more objects based on geometry characteristics or descriptors. Feature alignment algorithms generally consist of at least one of three processes: (1) feature identification, (2) feature matching between objects, and (3) feature orientation. There are a wide variety of applications that utilise feature alignment and detections algorithms, which include: geometric morphometric analyses [261], mesh-morphing [46], object detection and tracking [262], path navigation [263], registration and segmentation of bioimaging data [264], statistical shape modelling [265], surface reconstruction [266], and surgery planning and navigation [267]. Different techniques are available depending on the representation of the objects, i.e. polygonal meshes, point clouds, and pixel- and voxel-based images [268, 269]. Numerous techniques are targetted towards point cloud data [268, 270–272], which can be implicitly used on polygonal meshes. A key shortcoming for point clouds is that they can require estimations of curvatures and surface normals [270], which can be inaccurate. In contrast, the calculations for similar metrics are more straightforward and accurate for polygonal meshes. As such, results reported for point clouds can be less accurate than if they had been applied to polygonal meshes.

Depending on how features are detected, alignment techniques can also be separated into two

categories: automatic and manual detection. There are several different automatic techniques that can be employed to detect unique features or general similarities, such as: curvatures and surface normals [272]. The most established developments are based on algorithmic approaches, however there have recently been advances using machine learning techniques [273]. Over the coming years machine learning algorithms are expected to become more common, as this emerging field continues to produce high-quality results in other areas. Automatic approaches can be fast and robust for some applications, but can lead to erroneous alignments, particularly if suboptimal control parameters are selected [268, 272, 274, 275]. Manual feature detection typically involves a skilled-operator interacting with an object using a graphical user interface (GUI) to select and identify particular features of interest. An advantage of manual techniques is the greater accuracy, but results can be time-consuming, laborious, error-prone and inconsistent (user-biases and -drift) for complex geometries and noisy datasets [232, 261].

Rigid registration refers to the alignment of objects without changing the relative differences between vertices defining each object. Typically, these processes produce a linear transformation matrix, which is used to translate, rotate and/or scale geometries into alignment [268, 276]. They can be broken down into two categories: local or global optimisation, where the majority of methods operate using local optimisations [268]. A key disadvantage of local optimisers are that they may not find the optimal alignment globally [272]. In contrast, global optimisers have greater computational complexity [268].

Commonly used rigid registration algorithms are: iterative closest point (ICP) methods [268], random sample consensus (RANSAC) methods [274] and Procrustes methods [276, 277]. These methods differ by operating over known (Procrustes) or unknown (ICP and RANSAC) correspondences. For general applications, ICP is the most popular rigid registration algorithm [268]. However, ICP is a local optimiser and has several shortcomings, such as: requiring a close initial alignment [47, 269], requiring strong shape similarities [47], sensitivity to outliers [271], large computational complexity [268, 270], and local convergence instead of global [270]. There have been several attempts to improve ICP [268]. In contrast, RANSAC tends to be more robust and tolerant of outliers [269], but also has several shortcomings, some of which are associated with global optimisers [268, 275, 278].

2.5.2 Surface Parameterisation

The parameterisation of geometries is a core component and problem in computer graphics and mesh-processing [6, 279]. The use of parameterisation techniques have been employed to solve several common applications, which include: texture mapping [6, 260, 280], surface fitting [6, 281], mesh-morphing [282], remeshing [6, 260, 283], mesh repair [260, 281], mesh compression [260, 281], scattered data fitting [260], detail mapping [6, 280], detail transfer [281] and biomedical visualisation [281, 284]. A key motivation in the graphics community for the earlier developments of parameterisation techniques was texture mapping, which is a laborious and time-consuming manual operation [6, 260]. Further developments in this field were strongly influenced by requirements to remesh, compress and approximate complex surfaces, particularly those acquired by 3D scanning technologies [260].

Although the intended applications are different, the fundamental objectives and general problem is similar amongst several fields. Essentially, the aim of these techniques are to determine one-to-one correspondences between similar but different geometries, which can range from four-legged animals to patient-specific anatomies [8, 9].

Typically, parameterisation involves mapping a surface to a homeomorphic domain [260, 279]. A mapping provides one-to-one correspondences to a suitable domain, which could be unit disc for non-closed mesh patches and a unit sphere for closed, genus-0 meshes [6, 260, 279]. This can make the process of finding one-to-one correspondences to another object mapped to the same domain simpler, particular for morphing, remeshing and texturing. In general, a surface parameterisation consists of two triangular meshes, one is a geometric representation of an object (\mathcal{M}) and the other resides in a parametric domain (\mathcal{U}). Each triangular mesh $\mathcal{T} = (\mathcal{V}, \mathcal{K})$ is defined by a set of vertices \mathcal{V} and an abstract simplicial complex \mathcal{K} , where $\mathcal{V} = \{V_1, \dots, V_N\}^\top$ and \mathcal{K} contains two integer subsets of $\{1, \dots, N\}$ which specifies topological data. The two subsets define the edges $\mathbf{e} = \{i, j\} \in \mathcal{K}$ and triangle elements $\mathcal{E}_{\mathcal{T}} = \{i, j, k\} \in \mathcal{K}$. A set of the vertex indices defining a one-ring neighborhood around vertex V_i is denoted by $\mathcal{N}(i) = \{j \mid \{i, j\} \in \mathcal{K}\}$. Additionally, \mathcal{M} and \mathcal{U} are isomorphic, i.e. have the connectivity between vertices. Each vertex (V_i) of \mathcal{M} and \mathcal{U} are denoted by \mathbf{x}_i and \mathbf{u}_i , respectively. The two meshes form a piecewise linear bijective map (Ψ), where any location in one mesh can be determined in the corresponding mesh (**Figure 2.9**) [6]:




Figure 2.9: A piecewise bijective map (Ψ) of a non-closed ear mesh (\mathcal{M}) and the corresponding planar parameterisation (\mathcal{U}), the red triangle highlights isomorphism (*reproduced from [6]*)

The key goal of surface parameterisation is to derive a piecewise linear map with minimal distortion. The distortion energy (\mathbf{E}) between \mathcal{M} and \mathcal{U} at particular locations can be measured by comparing differences in the one-ring neighborhoods [6,260]. Different distortion energies can be defined with respect to the one-ring neighborhoods. Two common forms of parameterisations related to these energies are conformal (Dirichlet) and authalic. Conformal mapping refers to the minimisation of angular distortion energies (i.e. angle-preserving), whereas, authalic mapping minimises the area distortions (i.e. area-preserving) [6,260]. A parameterisation which preserves both angle and area is called an isometric map, which rarely occurs and are impossible for the majority of geometries [6, 260]. Additionally, an isometric map would be able to preserve distances between meshes [6]. The discretisation of distortion energies for triangular meshes takes the following general structure [279]:

$$\mathbf{E}(\Psi, \mathcal{M}) = \frac{1}{2} \sum_{i=1}^N \sum_{j \in \mathcal{N}(i)} \kappa_{ij} \|\Psi(\mathbf{x}_i) - \Psi(\mathbf{x}_j)\|^2 \quad (2.2)$$

or

$$\mathbf{E} = \frac{1}{2} \sum_{i=1}^N \sum_{j \in \mathcal{N}(i)} \kappa_{ij} \|\mathbf{u}_i - \mathbf{u}_j\|^2 \quad (2.3)$$

The minimisation process seeks to find a solution to the following equation:

$$\nabla \mathbf{E} = \frac{\partial \mathbf{E}}{\partial \mathbf{u}_i} = \sum_{j \in \mathcal{N}(i)} \kappa_{ij} (\mathbf{u}_i - \mathbf{u}_j) = 0 \quad (2.4)$$

Where, κ_{ij} refers to the distortion weights. The weights or spring constants define the type of energy and associated mapping to be minimised, common options are reported in **Table 2.4**. They are defined based on the angles within a one-ring neighborhood (**Figure 2.10**):

Table 2.4: Common spring weights used to define surface parameterisations

Type	Weights (κ_{ij})
Uniform [279]	1
Conformal [6]	$\cot \alpha_{ij} + \cot \beta_{ij}$
Authalic [6]	$\frac{\cot \delta_{ij} + \cot \gamma_{ij}}{\ \mathbf{x}_i - \mathbf{x}_j\ ^2}$
Mean Value Coordinates [260]	$\frac{\tan(\delta_{ij}/2) + \tan(\gamma_{ij}/2)}{\ \mathbf{x}_i - \mathbf{x}_j\ }$
Edge Length [285]	$\frac{1}{\ \mathbf{x}_i - \mathbf{x}_j\ }$




Figure 2.10: A one-ring neighborhood: a central (\mathbf{v}_i) and adjacent vertex (\mathbf{v}_j) with associated angles (α , β , δ and γ) used to define particular spring weights (*adapted from [7]*)

An exhaustive list of weights with their associated properties and conditions has previously been reported [281]. Additionally, these weights and associated energies can be extended to non-triangular meshes, such as quadrilaterals [286]. Parameterisations between two surfaces can easily be used for morphing or blending operations with simple linear interpolation (**Figure**

2.11) [8, 50]:

$$\mathbf{V}_i^I = (1 - t)\mathbf{V}_i^0 + t\mathbf{V}_i^1 \quad \forall i \in \{1, \dots, N\} \quad (2.5)$$

Where, \mathbf{V}_i^I defines an interpolated vertex and t specifies the degree of transition. \mathbf{V}_i^0 and \mathbf{V}_i^1 are corresponding vertices from two isomorphic meshes.




Figure 2.11: Interpolation of two meshes using spherical parameterisation (*adapted from [8]*)

Three common approaches to apply these principles for mesh parameterisation are planar, mesh partitioning and spherical. However, mesh parameterisation is not limited to these three and there are other more exotics forms, particularly to higher-genus domains (e.g. torus and double-torus). An extensive background on the wide scope of surface parameterisation techniques have been previously been report [281].

Planar Parameterisations

For planar parameterisation, the vertices of a non-closed mesh $\mathcal{M} \subset \mathbb{R}^3$, denoted by $\mathbf{x}_i = (x_i, y_i, z_i)^\top \in \mathcal{M}$, are mapped onto the vertices of a planar mesh $\mathcal{U} \subset \mathbb{R}^2$, denoted by $\mathbf{u}_i = (u_i, v_i)^\top \in \mathcal{U}$. A planar mesh is often constrained to either a disc or square, which allows simplified coordinates associated with a particular domain to be employed. Additionally, this can simplify a remeshing process and is an efficient technique used by high-quality meshing packages, such as GMSH [283, 287].

There have been several biomedical applications of planar parameterisation techniques. A common objective is the high-quality remeshing of anatomies for simulations [283, 288]. The triangular meshes can be used as shell elements or be volumetrically discretised into tetrahedral elements. A remeshing process provides higher quality surface elements for tetrahedra to be derived from. Consequently, the tetrahedra situated at the mesh boundaries will have a higher quality due to the initial remeshing. This has been shown to provide greatly improved simulation accuracy and convergence [288]. Examples of planar anatomies remeshed using these techniques

are a femoral head and femoral artery [283, 288]. For mesh-morphing, it has been used as an initial surface parameterisation step to determine correspondences between two femoral heads [9] (**Figure 2.12**). For enhanced visualisation, it has been used to flatten a heart to display MRI and electrical data, as well as the position of a surgical tool [284].

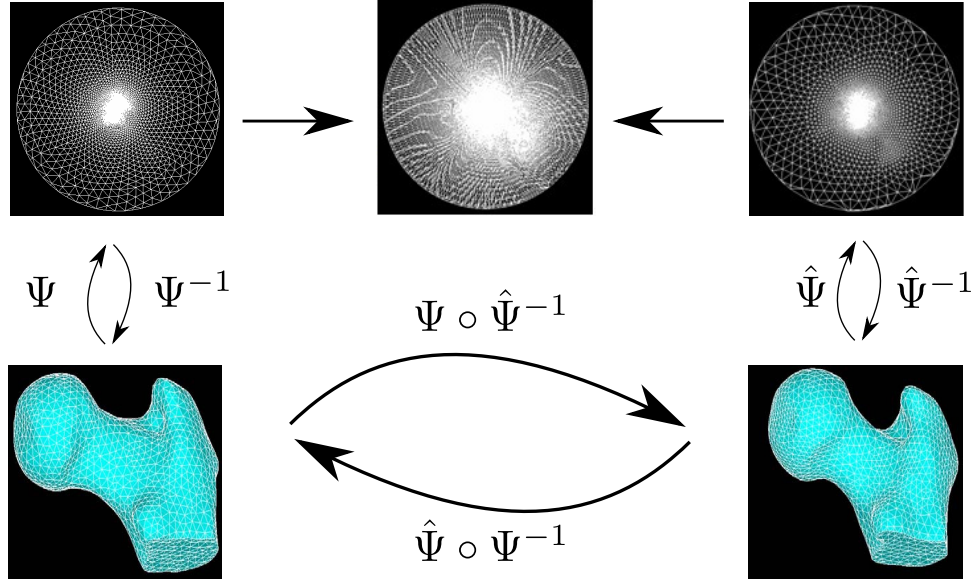


Figure 2.12: Planar cross-parameterisation between two non-closed femoral head meshes, which was used to morph a tetrahedral mesh. Linear (Ψ and $\hat{\Psi}$) and inverse maps (Ψ^{-1} and $\hat{\Psi}^{-1}$) could be used to determine direct correspondences between meshes (*adapted from* [9] ©2008 IEEE)

Planar parameterisations can be used with or without boundary constraints and does not have to be mapped to a base domain. [6, 282]. Natural boundaries, or natural conformal maps, refers to a planar parameterisation that does not constrain the boundaries. This allows the boundary of a mesh to take a more natural and less distorted form. Also, the overall distortion is lower, which can achieve better results for texturing [6]. Although, it would hinder mesh-morphing applications, due to having boundaries that lack one-to-one correspondences and therefore no bijectivity overall. For cross-parameterisation and mesh-morphing applications mapping two geometries to an appropriate base domain would allow one-to-one correspondences to be calculated between them. It is possible to map a non-closed mesh with multiple boundaries using the same techniques, however, it does increase the complexity of the process. In this case, one boundary is constrained to a base domain with additional boundaries either left natural or constrained to a particular shape as sub-domains.

These parameterisations can be calculated using **Equation 2.3** and either of the weights in

Table 2.4. For constrained parameterisations, the linear system $\mathbf{M}\mathbf{U} = \mathbf{C}$ is defined by the following [6]:

$$\begin{bmatrix} & \tilde{\mathbf{M}} \\ 0 & I \end{bmatrix} \begin{bmatrix} \mathbf{U}^{internal} \\ \mathbf{U}^{boundary} \end{bmatrix} = \begin{bmatrix} 0 \\ \mathbf{C}^{boundary} \end{bmatrix} \quad (2.6)$$

where

$$\tilde{\mathbf{M}}_{ij} = \begin{cases} \kappa_{ij} & \text{if } j \in \mathcal{N}(i) \\ -\sum_{k \in \mathcal{N}(i)} \kappa_{ik} & \text{if } i = j \\ 0 & \text{if } j \notin \mathcal{N}(i). \end{cases}$$

Preconditioned conjugate gradients can be used to solve this sparse linear system. The symmetric successive over-relaxation (SSOR) method has previously been used for preconditioning this system [6, 284]. Any number of boundaries can be constrained with this technique. For natural boundaries, these start as constrained parameterisations and are then optimised on a plane without constraints. This relieves distortions at the boundary and overall achieves a natural structure relative to the original geometry [6].

Mesh Partitioning

Mesh partitioning refers to the process of cutting a closed mesh to enable the use planar parameterisation techniques [281]. The partitioning process is performed for two reasons. First, it is not possible to map two objects that are not homeomorphic, e.g. a sphere to a disc. Whereas, it is possible to cut a sphere into two hemispheres and parameterise each to a disc. Also, it could be performed on planar meshes with multiple boundaries to achieve single boundary meshes. The second reason is to reduce the surface complexity of an object. Regardless of the parameterisation technique, distortion tends to increase proportional to surface complexity. Therefore, cutting a mesh can reduce the complexity of individual components allowing them to have reduced distortion and higher quality parameterisations [281]. In general, mesh partitioning is used to map the vertices of a closed mesh $\mathcal{M} \subset \mathbb{R}^3$, denoted by $\mathbf{x}_i = (x_i, y_i, z_i)^\top$, onto a non-closed mesh (or series of meshes) $\mathcal{U} \subset \mathbb{R}^2$ with vertices $\mathbf{u}_i = (u_i, v_i)^\top$ in a planar base domain.

There are two types of cutting techniques: segmentation and seam generation [281]. Segmentation techniques partition a closed mesh into a series of planar meshes, referred to as charts [281, 282]. Each chart is then assembled into a texture atlas that can be used for textur-

ing, remeshing, compression or morphing [280–282,289]. Segmentation techniques use cutting as a pre-processing stage to achieve multiple planar parameterisations. A post-processing stage is performed to stitch, or pack, the charts together into either a planar or closed domain [281,282]. Seam generation techniques use cuts to create a single non-closed mesh that can be mapped directly to a planar domain [281]. Therefore, seam generation techniques could be considered a pre-processing stage that directly enables the use of planar parameterisation for closed meshes.

The main advantage of mesh partitioning is that they allow the parameterisation of closed meshes with any genus. Additionally, they can derive higher quality parameterisations of planar meshes [288]. Similar to planar parameterisation techniques, they have been used for remeshing [288,290]. Examples of anatomies remeshed using these techniques include: a foot (planar, genus-0), a femoral artery (planar, genus-0), an aorta (planar, genus-0), an upper jaw (closed, genus-0), a hemipelvis (close, genus-1), a skull (closed, genus-2) and a pelvis (closed, genus-9) [288,290]. However, a key challenge is to minimise the number and length of required cuts [281]. A common disadvantage of this approach is that discontinuities can form between partition boundaries [282]. Another challenge is automating where to partition a mesh to minimise distortion and maximising one-to-one correspondence accuracy for cross-parameterisations. Consequently, this technique is not commonly used for morphing applications and is usually used for texturing and remeshing [280–282]. There are two additional problems for mesh-morphing: (1) partitioning the meshes in similar locations and (2) having the same number of partitions. One solution to these problems were to define the partitions manually [289]. To automate a solution to the two problems would be a difficult task in itself. Combined with minimising the distortion of those partitions exacerbates the challenge further.

Spherical Parameterisations

For spherical parameterisation, the vertices of a closed genus-0 mesh \mathcal{M} , denoted by $\mathbf{x}_i = (x_i, y_i, z_i)^\top \in \mathbb{R}^3$, are mapped onto the vertices of a closed mesh \mathcal{U} , denoted by $\mathbf{u}_i = (u_i, v_i, w_i)^\top \in \mathbb{R}^3$, which lie on a unit sphere $S^2 \subset \mathbb{R}^3$. Once mapped to a sphere, the meshes can be cross-parameterised similarly to planar parameterisations but using the 3D Cartesian coordinates on a spherical surface (**Figure 2.13**). In some cases, polar coordinates have sometimes been used to define the spherical coordinates [291]. The use of spherical parameterisation has become more common over the past decade, with applications ranging from graphics processing [281] to biomedical visualisation [10].

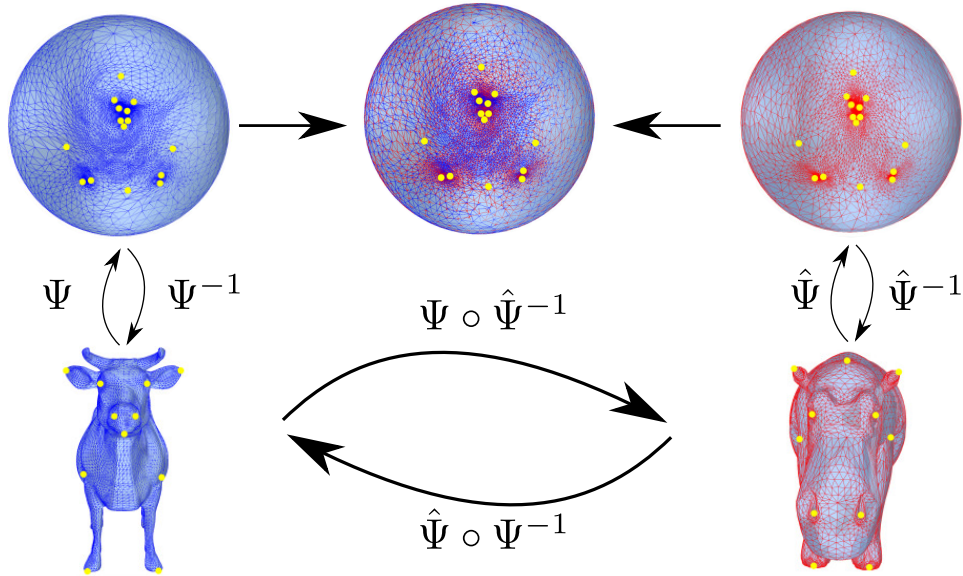


Figure 2.13: Spherical cross-parameterisation between two closed genus-0 meshes, used to morph and interpolate triangular meshes. Linear (Ψ and $\hat{\Psi}$) and inverse maps (Ψ^{-1} and $\hat{\Psi}^{-1}$) could be used to determine direct correspondences between meshes (*adapted from [8]*)

A major benefit of spherical parameterisations is that they can negate the need to pre-process and post-process the parameterisation of genus-0 geometries that would be required for mesh partitioning. Of course, this only applies to techniques that do not use partitioning during a spherical parameterisation process. These additional operations tend to have a negative effect on the quality of a parameterisation and are difficult to automate, to minimise the additional effects. A key disadvantage is that despite significant efforts a piecewise bijective parameterisation may still contain excessive distortion and not allow adequate sampling of a surface. Distortions can sometimes induce sliver elements that represent very thin and elongated shapes. These can heavily skew cross-parameterisations used for morphing, as they do not exist in either of the original meshes [292]. In general, sliver elements and other distortions lead to low-quality parameterisations and morphed objects.

Spherical parameterisation is restricted to closed genus-0 meshes, as they can be considered deformed spheres. Therefore, a process involves removing the deformations whilst minimising the degree of mesh distortion incurred, to sufficiently sample the entire surface. Despite this restriction, the technique is quite favourable as an extensive range of geometries are represented by closed genus-0 meshes. In particular, a large proportion of anatomical structures fit this criterion, especially those free from injury and certain pathologies. However, conditions arise which can cause anatomical geometries that would naturally be genus-0 to become genus-1 and

higher, such as cartilage lesions or certain types of meniscal tears. However, meshes with a higher genus could be parameterised using this approach with additional modifications, such as partitioning or hole-filling. This would make an algorithm designed for those susceptible anatomies more challenging to generalise for both healthy patients and those with underlying pathologies. In general, closed meshes are often derived from the post-processing of 3D bioimaging data and are required for volumetric discretisation prior to FE analyses. Therefore, anatomies that could be parameterised using this method are widely available. There have been several applications mapping the brain to a sphere due to it already closely resembling that geometry [10, 281, 293]. Therefore, it can be mapped with relatively minimal distortions compared to more irregular geometries. Spherical parameterisation techniques have allowed greater visualisation of the intricate folds and neural activity in particular regions of the brain (*Figure 2.14*) [10].




Figure 2.14: Spherical parameterisation of a closed genus-0 brain mesh used to improve visualisation of white-matter curvatures (*reproduced from* [10] ©2000 IEEE)

In general, biological structures are highly irregular objects and do not easily lend themselves to a unit sphere, unlike like a brain. The graphics community have attempted irregular geometries, like animals and humanoid figures, but the key challenge of minimising distortion is difficult to overcome [282]. Also, there is no theorem that can guarantee inversion free parameterisations and achieving this for some geometries can be challenging. Consequently, there are a wide range of techniques that attempt to solve this problem, as well as controlling the degree of distortions that would lead to favorable cross-parameterisation.

Spherical parameterisations start with an initial projection step, which places the vertices of a mesh onto the surface of a unit sphere. The centroid of a geometry (\mathbf{O}) is often used to define

the projection vectors and will then serve as the origin of a parameterised sphere. The simplest way to achieve this is to position the centroid of a mesh at the Cartesian origin and normalise each vertex using the following equations [294]:

$$\mathbf{O} = \frac{1}{N} \sum_{i=1}^N \mathbf{V}_i \quad (2.7)$$

$$\mathbf{V}_i = \mathbf{V}_i - \mathbf{O} \quad \forall \mathbf{V}_i \in \mathcal{V} \quad (2.8)$$

$$\mathbf{u}_i = \frac{\mathbf{V}_i}{\|\mathbf{V}_i\|} \quad \forall i \in \{1, \dots, N\} \quad (2.9)$$

For the majority of geometries, this creates a highly folded mesh that would not define a piecewise linear affine map. Therefore, the aim of spherical parameterisation techniques are to remove these folded elements from this initial projection to achieve a bijective map in the shape of a sphere. One approach is to minimise a distortion energy, similar to planar parameterisations, using **Equation 2.3**. However, spherical constraints must be applied to prevent the parameterisations collapsing to a single point. There are several different ways this has been achieved with varying success. A simple approach has been to minimise a planar distortion energy (**Equation 2.3**) iteratively and re-project all the vertices (**Equation 2.9**) after a complete minimisation step [50]. Similarly, each vertex could be re-projected immediately after an individual iteration, defining a spherical distortion energy ($\tilde{\mathbf{E}}$) [294]:

$$\frac{\partial \tilde{\mathbf{E}}}{\partial \mathbf{u}_i} = \frac{\sum_{j \in \mathcal{N}(i)} \kappa_{ij} (\mathbf{u}_i - \mathbf{u}_j)}{\left\| \sum_{j \in \mathcal{N}(i)} \kappa_{ij} (\mathbf{u}_i - \mathbf{u}_j) \right\|} \quad (2.10)$$

\therefore

$$\tilde{\mathbf{E}} = \frac{1}{2} \sum_{i=1}^N \frac{\sum_{j \in \mathcal{N}(i)} \kappa_{ij} \|\mathbf{u}_i - \mathbf{u}_j\|^2}{\left\| \sum_{j \in \mathcal{N}(i)} \kappa_{ij} \|\mathbf{u}_i - \mathbf{u}_j\|^2 \right\|} \quad (2.11)$$

In some cases, these approaches achieved smooth parameterisations, however, they are still susceptible to the parameterisation collapsing. This is a known issue with applying spherical constraints using re-projection techniques because minimisation algorithms can sometimes find solutions by enlarging a particular triangle. This triangle eventually wraps around the spherical domain and initiates the collapse of a mesh into a point [295]. One approach to prevent

the collapse of the parameterisation is to constrain particular points, which may introduce further distortions [295]. This introduces the problem of how many and which particular vertices to constrain. Therefore, another approach has successfully solved this problem using soft constraints by application of the Lagrange-Newton-Krylov-Schur (LNKS) method [296]. This technique required defining a Lagrangian penalty function (P) (**Equation 2.14**) and imposing the soft constraints using Lagrange multipliers ($\tilde{\lambda}$) [279]. The distortion energy was defined using **Equation 2.3** and the soft constraints (\mathbf{c}) were defined similar to **Equation 2.9**:

$$\mathbf{c}_i = \|\mathbf{u}_i\|^2 - 1 = 0 \quad \forall i \in \{1, \dots, N\} \quad (2.12)$$

This technique is more complicated and requires assembling a Hessian matrix (W) of a Lagrangian and a Jacobian matrix of the constraints (A); in addition to the gradient of a distortion energy ($\nabla \mathbf{E}$) [296]. Therefore, the second derivatives of a Lagrange function (L) (**Equation 2.13**) must be determined, which could be straightforward or complex depending on the choice of weights.

$$L(\mathbf{u}, \tilde{\lambda}) = \mathbf{E} - \tilde{\lambda}^\top \mathbf{c} \quad (2.13)$$

$$P(\mathbf{u}, \tilde{\lambda}) = \|\nabla \mathbf{E} - A\tilde{\lambda}\|^2 + \|\mathbf{c}\|^2 \quad (2.14)$$

The implementation of the LNKS method is more sophisticated and challenging than the other techniques. The Hessian and Jacobian matrices are assembled to form the following linear system:

$$\begin{bmatrix} W & -A \\ -A^\top & 0 \end{bmatrix} \begin{Bmatrix} \delta_{\mathbf{u}} \\ \delta_{\lambda} \end{Bmatrix} = - \begin{Bmatrix} \nabla \mathbf{E} - A\tilde{\lambda} \\ -\mathbf{c} \end{Bmatrix} \quad (2.15)$$

This system can be solved using Cholesky decomposition, to determine the search directions $\delta_{\mathbf{u}}$ and δ_{λ} for the spherical vertices and Lagrange multipliers, respectively. A line search is performed to determine if the search directions require correction with a lower step size. Once suitable search directions have been acquired the vertices and Lagrange multipliers are updated (**Equations 2.16-2.17**). This process continues iteratively until reaching convergence,

achieving a valid spherical mapping [279].

$$\mathbf{u} = \mathbf{u} + \alpha\delta_{\mathbf{u}} \quad (2.16)$$

$$\tilde{\lambda} = \tilde{\lambda} + \alpha\delta_{\lambda} \quad (2.17)$$

Due to the challenges and complexity involved in prescribing spherical constraints, an unconstrained method has been proposed using the PETSc framework [297]. This technique also defines the distortion energies using *Equation 2.3* [295]. The unconstrained minimisation creates upper and lower bounds for a distortion energy to operate within. The upper and lower energies would define the energies outside and inside the spherical boundary, respectively. The lower bound energy uses the original planar distortion energy, as vertices inside a flat triangle would exist within the boundary defined by a spherical triangle. To define the upper bound energy, the planar energy is appended with a distance term (d_{min}), defined as the minimum distance between each triangle and the spherical origin:

$$\mathbf{E}(\Psi, \mathcal{M}) \leq \tilde{\mathbf{E}}(\Psi, \mathcal{M}) \leq \frac{\mathbf{E}(\Psi, \mathcal{M})}{d_{min}^2} \quad (2.18)$$

This term projects a planar triangle so that it is tangent to a sphere. Hence, the upper bound is achieved as each triangle lies just outside a sphere. The term also prevents the parameterisation collapsing, as there is a higher energy associated with large triangles enveloping the hemisphere that would initiate a collapsed configuration [295].

There are more approaches to achieving spherical parameterisations, which includes: mesh partitioning [298], minimisation of stretch metrics [282] and spherical harmonics (SPHARM) [291, 299]. The mesh partitioning approach is an example that uses all three surface parameterisation techniques described in this section, in an attempt to produce high-quality maps [298]. First, a mesh was split into two even components with each mapped to a disc. Next, each disc was mapped to a hemisphere and stitched together to form a spherical parameterisation. The final spherical parameterisation still requires further processing to remove excess distortions. Clearly, the challenge of achieving minimally distorted parameterisations of complex geometries is an endeavour likely to continue for some time. In general, these key problems highlight the

difficulties achieving high-quality bijective parameterisations of irregular geometries, that are defined by minimal distortions, no inversions and can provide uniform sampling of a geometry.

2.5.3 Volumetric Transformation

For some applications, volumetric transformation could be considered the definition of mesh-morphing in a general sense. The aim of volumetric transformation is to preserve a large proportion of a template's mesh-quality and prevent element inversions [11]. The need to apply geometric modifications to an existing mesh occurs in several areas: animation [300], automatic mesh generation [52], crash simulations [300], dynamic biomedical simulations [56, 300], design/shape optimisation [13, 301], high energy deposition physics [301], metals forging [301], projectile penetration studies [301], propellant burn [301] and surgery optimisation [54].

A volumetric transformation Ψ deforms a domain Ω defined by a boundary $\partial\Omega$, that will define a new domain Ω' with a particular boundary of interest $\partial\Omega'$. For an arbitrary volumetric mesh \mathcal{M} with geometry $\partial\Omega$, mesh-morphing seeks to find the vertex adjustments that will define an isomorphic mesh \mathcal{M}' prescribed by a different geometry $\partial\Omega'$ [11]. A template mesh would be defined by \mathcal{M} and a target mesh would be described by $\partial\Omega'$ with a morphed mesh denoted by \mathcal{M}' (**Figure 2.15**):

$$\begin{aligned}\Psi : \Omega &\rightarrow \Omega' \\ \partial\Omega &\rightarrow \partial\Omega' \\ \mathcal{M} &\rightarrow \mathcal{M}'\end{aligned}\tag{2.19}$$




Figure 2.15: A general transformation operation (Ψ): morphing a template mesh (\mathcal{M}) to define a new mesh (\mathcal{M}') that resembles a target geometry ($\partial\Omega'$)

There are four common volumetric morphing techniques: FE-based, radial basis function (RBF) interpolation, smoothing and barycentric interpolation. All the techniques can be used for both surface and volumetric morphing. Although the focus here is on volumetric meshes, examples of

surface morphing will be provided to explain some processes and simplify equivalent descriptions. Unfortunately, no volumetric morphing operation can guarantee inversion free transformations, defined by the induction of negative Jacobians. However, some techniques are more stable and robust than others, to significant geometric deformations. Regardless of technique, the propensity for element inversions increases with larger deformations. Therefore, it is possible to derive a mesh (\mathcal{M}') that is unsuitable for FE simulations.

Each technique can be separated into two categories, mesh-based or meshless. For mesh-based transformations, the solution depends on the topology and elements of a mesh. Whereas, meshless transformations operate purely on vertex positions and do not incorporate aspects vertex connectivity. Smoothing and FE-based operations are examples of mesh-based transformations. RBF and barycentric interpolations are considered meshless transformations. FE-based transformations and some smoothing operations are associated with specific element definitions. It can be considered an advantage of meshless operators that the same implementation can immediately be used for any arbitrary mesh, including exotic polygonal meshes [13]. This could be advantageous for morphing FE models composed of multiple elements types. However, a strong motivation for some mesh-morphing strategies in biomedical applications is to enable the rapid use of all-hexahedral meshes, as apposed to hexahedral-dominant meshes which may contain prisms and tetrahedra. All of the reviewed investigations focussed on one particular element type (**Table 2.8**). An FE-based morphing operation would require the element components to be defined for a mesh prior to morphing. However, it would be a fairly non-trivial task to modify or automate mesh-based algorithms to include other element types. An issue could occur when morphing a non-standard volumetric elements. Although any arbitrary element can be defined for the FEM, only common elements are typically implemented. Elements beyond standard formulations would technically require implementation for an operator. However, this is not an issue for the majority of common applications, which operate over standard finite elements.

There are several key differences between the operation techniques with respect to speed, robustness and precision. Techniques dependent on the solution of large matrices typically have lower speeds. RBF interpolation requires the solution of a dense matrix and consequently yields the slowest results. Similarly, smoothing and FE-based techniques require a solution of a sparse matrix. This yields faster results than RBF interpolation but slower than techniques that do not require solutions of large matrices. Barycentric interpolation does not require a large matrix to

be solved and yields the fastest speeds. Speed comparisons between methods for three meshes (*Table 2.6*) have been reported in the following table [11, 12]:

Table 2.5: Speed comparison for common volumetric transformation operators (seconds) [11, 12]

Model	FEMWARP	RBF	Smoothing		Barycentric (Simplex)	
			Optimisation	LBWARP	S. Linear	S. Natural
Bore	0.29	2.3-8.0	0.75	1.36	0.40	0.38
Pipe	0.32	1.3-4.5	1.19	1.50	0.28	0.29
Courier	4.89	22-104	12.27	18.71	1.50	7.45

Table 2.6: Three meshes with properties that were used to assess the speed of different operators

Model	Vertices (#)	Elements (#)	Complexity
Bore	11,904	15,190	Simple
Pipe	11,520	8,532	Simple
Courier	101,817	83,934	Complex

The range provided for RBF interpolation is the result of comparing five different solvers. The fastest solver was a GPU-based solver (MAGMA) and the second fastest was a CPU-based solver (MKL), both have proprietary licenses. The third fastest was based on a custom incremental QR solver. The slowest two were derived from open-source solvers from LAPACK [13]. In comparison, FEMWARP and LBWARP were solved using the open-source package Trilinos (Amesos KLU) [12, 302]. Consequently, the times are not fully comparable due to solver differences but highlight that the speed for matrix-based methods are solver dependant. Also, potentially faster speeds could be derived for FEMWARP and LBWARP using proprietary or GPU-based solvers.

A volumetric morphing operation can use two approaches: absolute or relative. Absolute refers to morphing a mesh in one step, whereas, relative defines an operation broken down into a series of steps. Relative morphing can achieve improved mesh-qualities at the expense of speed [13]. For large deformations, some operations must be relative to prevent element inversions [13]. However, it is not clear how to correctly implement relative morphing, as no procedures have reported an optimal approach with respect to both speed and quality [13]. The reported speeds in *Table 2.5* are based on absolute morphing.

Finite Element Method

The FEM is a technique to solve partial differential equations (PDEs) pertaining to a particular problem of interest, which include: structural [303], thermal [303], mechanical [304],

physics [102], electromagnetic [304], chemical [305] and biological phenomena [306]. FE-based transformations are a common volumetric morphing operation [12, 45, 54, 237, 307–309, 309, 310]. To achieve this, the FEM is used to prescribe boundary conditions for a domain, where the solution calculates the displacement field for the internal vertices. Boundary conditions can be defined after correspondences between two geometries have been determined. There are several approaches to FE-based transformations, which includes: Laplacian [311], biharmonic [312], diffusion [313], and elasticity [45, 314]. The difference between these methods lie in the response of a mesh to a deformed boundary. The FEM falls into the category of variational methods [11]. Additionally, other variational methods similar to the FEM have been developed for volumetric transformations [189]. One example incorporated components of optical flow, from MRI images, and the Laplacian operator to derive a displacement field. This displacement field was used to morph a template mesh using a third-order variational method [189].

Laplacian-based approaches, often referred to as FEMWARP, are defined by a second-order harmonic PDE [11, 311, 312]:

$$\Delta \mathbf{v} = \nabla^2 \mathbf{v} = 0 \quad (2.20)$$

Where, \mathbf{v} is the vertices of a mesh. This approach often represents a mesh as a spring model, where connections between vertices (or mesh edges) are defined as springs [311, 312]. There are several ways to solve the Laplacian of a mesh with varying degrees of numerical robustness (*Section 2.5.3: Smoothing Operations*). Solving the Laplacian of a mesh using the FEM has been reported to be superior to smoothing approaches with respect to robustness [11]. Biharmonic-based methods are similar to Laplacian-based methods [11, 312]. The difference lies in using a fourth-order biharmonic PDE to define a system. Consequently, a biharmonic approach is performed by solving a higher order Laplacian (bi-Laplacian) [312]:

$$\Delta^2 \mathbf{v} = \nabla^4 \mathbf{v} = 0 \quad (2.21)$$

Diffusion-based methods employ Navier-Stokes equations to define a fluidic response of vertices [313].

For elastic transformations, the FEM has been used as a tool to morph a mesh by defining it as a deformable elastic material. Many different material properties can be defined for FE

analyses and similar transformations. A hyperelastic (nonlinear elastic) Mooney-Rivlin material definition has previously been used for a volumetric transformation operation [315]. The most simple and common material definition used by the FEM is a linear elastic material [102, 120, 303, 316]. Linear elasticity refers to the linear correlation between stress and strain responses of a material under deformation. This relationship has been well studied and is exemplified by Hooke's law for a linear spring (**Equation 2.22**). Where, a spring's force F is linearly proportional to its stiffness k and displacement \mathbf{u} (**Equation 2.22**). This principle has a direct analogue to the elastic deformation of a body. In continuum mechanics, this relationship is described with respect to stress (σ), strain (ϵ) and stiffness (\mathbf{s}) tensors [45]:

$$F = k\mathbf{u} \quad (2.22)$$

$$\sigma = \mathbf{c}\epsilon \quad (2.23)$$

or more accurately [316]:

$$\sigma = \langle \mathbf{c}, \epsilon \rangle \quad (2.24)$$

Where, $\langle \cdot, \cdot \rangle$ denotes the inner product. Linear elastic material properties can be defined using Young's modulus and Poisson's ratio, which describe a material's stiffness and behaviour to tension and compression, respectively. The material properties are encoded into fourth-order stiffness tensors using Young's modulus (E) and Poisson's ratio (ν). However, Lamè's parameters (λ and μ) are often used to define their relationship, as they simplify the descriptions [120]:

$$\lambda = \frac{E\nu}{(1+\nu)(1-2\nu)} \quad (2.25)$$

$$\mu = \frac{E}{2(1+\nu)} \quad (2.26)$$

Therefore, Hooke's law can be expressed with respect to the material properties of a given body [316]:

$$\begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{33} \\ \sigma_{23} \\ \sigma_{13} \\ \sigma_{12} \end{bmatrix} = \begin{bmatrix} 2\mu + \lambda & \lambda & \lambda & 0 & 0 & 0 \\ \lambda & 2\mu + \lambda & \lambda & 0 & 0 & 0 \\ \lambda & \lambda & 2\mu + \lambda & 0 & 0 & 0 \\ 0 & 0 & 0 & \mu & 0 & 0 \\ 0 & 0 & 0 & 0 & \mu & 0 \\ 0 & 0 & 0 & 0 & 0 & \mu \end{bmatrix} \begin{bmatrix} \epsilon_{11} \\ \epsilon_{22} \\ \epsilon_{33} \\ 2\epsilon_{23} \\ 2\epsilon_{13} \\ 2\epsilon_{12} \end{bmatrix} \quad (2.27)$$

which can be simplified into the following notations:

$$\begin{bmatrix} \sigma_{11} & \sigma_{12} & \sigma_{13} \\ \sigma_{12} & \sigma_{22} & \sigma_{23} \\ \sigma_{13} & \sigma_{23} & \sigma_{33} \end{bmatrix} = 2\mu \begin{bmatrix} \epsilon_{11} & \epsilon_{12} & \epsilon_{13} \\ \epsilon_{12} & \epsilon_{22} & \epsilon_{23} \\ \epsilon_{13} & \epsilon_{23} & \epsilon_{33} \end{bmatrix} + \lambda \mathbf{I}(\epsilon_{11} + \epsilon_{22} + \epsilon_{33}) \quad (2.28)$$

or [316–318]

$$\sigma = 2\mu\epsilon + \lambda \operatorname{tr}(\epsilon)\mathbf{I} \quad (2.29)$$

Where, $\operatorname{tr}(\cdot)$ is the trace operator and \mathbf{I} is the identity matrix. The strain tensor ($\epsilon = \epsilon(\mathbf{u})$) is described with respect to displacements (\mathbf{u}) [318]:

$$\epsilon(\mathbf{u}) = \frac{1}{2}(\nabla\mathbf{u} + (\nabla\mathbf{u})^\top) \quad (2.30)$$

where [317]

$$\nabla\mathbf{u} = \left[\frac{\partial\mathbf{u}}{\partial x}, \frac{\partial\mathbf{u}}{\partial y}, \frac{\partial\mathbf{u}}{\partial z} \right]^\top \quad (2.31)$$

Therefore, the stress of a domain (*Equation 2.29*) can be expressed purely in terms of displacements and material properties:

$$\sigma = \lambda(\nabla \cdot \mathbf{u})\mathbf{I} + \mu(\nabla\mathbf{u} + (\nabla\mathbf{u})^\top) \quad (2.32)$$

where

$$\operatorname{tr}(\epsilon) = \nabla \cdot \mathbf{u} \quad (2.33)$$

Where, $\nabla \cdot$ is the divergence operator. Boundary conditions must be defined to solve *Equation 2.32* and derive meaningful results [102]. Therefore, the problem seeks to determine the displacements of a domain Ω with linear elastic material properties subject to deformations on the boundary $\partial\Omega$ [102, 318]:

$$-\nabla \cdot \sigma = \mathbf{f} \quad \text{in } \Omega \quad (2.34)$$

\therefore

$$-\nabla \cdot (\lambda(\nabla \cdot \mathbf{u})\mathbf{I} + \mu(\nabla\mathbf{u} + (\nabla\mathbf{u})^\top)) = \mathbf{f} \quad \text{in } \Omega \quad (2.35)$$

Where, \mathbf{f} defines the optional body forces of a domain, e.g. gravity [102]. The variational

approach to solving the problem consists of defining test functions \mathbf{w} that vanish towards the boundaries [120]. These are incorporated by inserting **Equation 2.29** into **Equation 2.34**, taking the inner product with respect to the test functions and integrating with respect to the domain [319]:

$$\int_{\Omega} \langle -\nabla \cdot [\lambda \operatorname{tr}(\epsilon(\mathbf{u}))\mathbf{I} + 2\mu\epsilon(\mathbf{u})], \mathbf{w} \rangle d\Omega = \int_{\Omega} \langle \mathbf{f}, \mathbf{w} \rangle d\Omega \quad (2.36)$$

Using Green's theorem the following equation can be derived [102, 319]:

$$\int_{\Omega} \lambda \operatorname{tr}(\epsilon(\mathbf{u})) \operatorname{tr}(\epsilon(\mathbf{w})) + 2\mu \langle \epsilon(\mathbf{u}), \epsilon(\mathbf{w}) \rangle d\Omega = \int_{\Omega} \langle \mathbf{f}, \mathbf{w} \rangle d\Omega + \int_{\Gamma_N} \langle \mathbf{g}, \mathbf{w} \rangle d\Gamma_N \quad (2.37)$$

Where, \mathbf{g} defines the prescribed traction forces operating on the boundary Γ_N [317, 318]. To derive a FE formulation (Galerkin method) of **Equation 2.37**, basis functions ϕ are used to approximate \mathbf{u} and \mathbf{w} [319]:

$$\mathbf{u} \approx \sum_i u_i \phi_i \quad (2.38)$$

$$\mathbf{w} \approx \sum_j \phi_j \quad (2.39)$$

Where, u_i defines unknown coefficients [319]. Substituting these approximations into **Equation 2.37** results in the following derivation which can be solved linearly using LU decomposition [102, 319]:

$$\sum_i u_i \int_{\Omega} \lambda \operatorname{tr}(\epsilon(\phi_i)) \operatorname{tr}(\epsilon(\phi_j)) + 2\mu \langle \epsilon(\phi_i), \epsilon(\phi_j) \rangle d\Omega = \int_{\Omega} \langle \mathbf{f}, \phi_j \rangle d\Omega + \int_{\Gamma_N} \langle \mathbf{g}, \phi_j \rangle d\Gamma_N \quad (2.40)$$

or

$$\mathbf{K}\mathbf{u} = \mathbf{f} + \mathbf{g} \quad (2.41)$$

Where, \mathbf{K} is the infamous stiffness matrix of the FEM [102]. This final equation of linear elasticity draws strong parallels to the definition of a linear spring (**Equation 2.22**).

Boundary conditions are used to define the limits of a solution and specify what values the solution must achieve (**Figure 2.16**). There are several types of boundary conditions: Dirichlet (essential), Neumann (natural), Robin, Cauchy and mixed [318, 320, 321]. Dirichlet and

Neumann boundary conditions are commonly used in FE analyses [322–324] and associated transformations [45, 237]. The other three are different combinations of the first two. The Dirichlet boundary condition (Γ_D) defines what values a domain boundary ($\partial\Omega$) must be in a solution. For solid mechanics, they specify the boundary displacements of a domain and are sometimes referred to as fixed boundary conditions as a result. Neumann boundary conditions (Γ_N) define what derivative values a domain boundary must be in a solution. They are described as natural due to their spontaneous occurrence in the weak formulation (**Equation 2.37**). For solid mechanics, this refers to the application of traction forces on a domain boundary. The strong form of Dirichlet and Neumann boundary conditions are specified by the following, respectively [317, 318]:

$$\mathbf{u} = \mathbf{u}_D \quad \text{on } \Gamma_D \quad (2.42)$$

$$\boldsymbol{\sigma} \cdot \mathbf{n} = \mathbf{g} \quad \text{on } \Gamma_N \quad (2.43)$$

Where, \mathbf{u}_D defines the prescribed displacements of the boundary Γ_D and \mathbf{n} is the unit outward normal to the boundary Γ_N [317, 318].




Figure 2.16: Linear elastic transformation (Ψ) with Dirichlet (Γ_D) and Neumann boundary conditions (Γ_N)

There are several advantages to using FE-based transformations. First, original element qualities can be well-preserved [11, 12]. Also, in a review of volumetric morphing operators, it achieved the best results; which was evident from sustaining larger deformations prior to inducing negative Jacobians [12]. This could be attributed to the use of a transformation method which incorporates the same degrees of freedom (DOF) as the simulation method. Therefore, meshes morphed by variational methods can intrinsically maintain the quality and regularity of a mesh due to operating on the same principles [189, 313]. Second, there are a variety of FE packages and applications that can be used to simplify the development process. Therefore, it

is not required to implement complicated formulations, which can be error-prone due to the sophisticated nature of the mathematics. These formulations typically include a variety of element basis functions and material definitions. Additionally, FE-based transformations are a relatively simple operation compared to other FE applications, therefore the implementation can be a straightforward process. Third, they can provide greater numerical robustness over other methods [11]. In a benchmarking investigation of several volumetric morphing operators, the FE-based approach provided the best results for hexahedral and tetrahedral meshes [12]. Fourth, the morphing process can be a reasonably fast operation. Although, the solution is computationally expensive, there are a variety of well-developed FE frameworks and applications. These can be highly optimised, GPU-capable and operate efficiently on high-resolution meshes [325]. Fifth, due to the nature of the FE process, a variety boundary conditions can be defined. This provides additional flexibility for a strategy over simple Dirichlet boundary conditions allowing more exotic strategies to be formulated. The remaining volumetric techniques described in this section are only capable of displacement-based transformations, whereas FE-based transformations can use various combinations of displacement and force-based transformations [237].

For elasticity-based transformations, one aspect could be considered either an advantage or disadvantage. The morph quality is highly dependent on the material properties, i.e. compliancy [45]. This could be an advantageous as it could be used to provide greater control over a transformation. Application-specific parameter tuning could be used to tailor a transformation. Alternatively, the material properties could be defined differently for each element based on particular metrics, such as volume [54]. This would involve defining an anisotropic material and has previously been used to reduce the degree of deformation on highly distorted elements and vice versa [54]. However, this can be seen as a disadvantage as it could introduce an optimisation problem.

RBF Interpolation

Radial basis functions (RBFs) are used in a wide variety of applications and fields. They have diverse functionality and can be used for: geometry intersection [326], geometry reconstruction [327], geometry minimisation [328], mesh-morphing [13, 46], ray-casting [329], signed-distance field interpolation [327, 330] and scattered data interpolation [13, 327]. Analogous to Dirichlet boundary conditions, the displacements of a boundary domain are prescribed as constraints.

Each surface vertex $\mathbf{s}_i \in \mathcal{V}$ of a volumetric mesh represents a constraint. RBF interpolation creates a deformation field based on these constraints that implicitly morphs an entire mesh. The deformation field is used to smoothly interpolate the internal vertices into a modified domain. Additionally, the boundary vertices are implicitly interpolated by the deformation field to positions identical to the prescribed displacements: $\mathbf{s}'_i - \mathbf{s}_i$.

A deformation function $\mathbf{d}(\mathbf{v}_i) : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is defined by kernel functions $\phi(\mathbf{r}) : \mathbb{R} \rightarrow \mathbb{R}$ that are combined linearly. Additionally, the kernels are radially symmetric ($r = \|\mathbf{v}_i - \mathbf{v}_j\|$), centered around each vertex $\mathbf{v}_i \in \mathcal{V}$ [13]. Historically, a linear polynomial function ($\pi(\mathbf{v}_i)$) has been included, which ensures the linear precision of a deformation function with respect to the affine transformation [13, 14]. Additionally, a solution of the linear system (**Equation 3.5**) is guaranteed to be unique, with an added polynomial [14]. The solution determines the weights $\mathbf{w}_j \in \mathbb{R}^3$ and $\mathbf{q}_k \in \mathbb{R}^3$. The kernels are weighted by \mathbf{w}_j and the polynomials are weighted by \mathbf{q}_k . The deformation function (**Equation 2.44**) can then be used to morph any mesh vertex: $\mathbf{v}_i = \mathbf{v}_i + \mathbf{d}(\mathbf{v}_i)$.

Lemma 2.1. *Solve a linear system of equations to determine the weights (\mathcal{W} and \mathcal{Q}) for a deformation function derived from a matrix of RBFs (Φ):*

N = number of surface vertices

$i, j = 1, \dots, N$

$$\mathbf{d}(\mathbf{v}_i) = \sum_{j=1}^N \mathbf{w}_j \phi_j(\mathbf{v}_i) + \sum_{k=1}^4 \mathbf{q}_k \pi_k(\mathbf{v}_i) \quad (2.44)$$

where

$$\phi_j(\mathbf{v}_i) = \phi(\|\mathbf{v}_i - \mathbf{v}_j\|) \quad (2.45)$$

$$\pi_k(\mathbf{v}_i) = (\mathbf{v}_i^x, \mathbf{v}_i^y, \mathbf{v}_i^z, 1) \quad (2.46)$$

The linear system for the RBF-based deformation function is defined as follows:

$$\begin{bmatrix} \phi_1(\mathbf{s}_1) & \cdots & \phi_N(\mathbf{s}_1) & \pi_1(\mathbf{s}_1) & \cdots & \pi_4(\mathbf{s}_1) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \phi_1(\mathbf{s}_N) & \cdots & \phi_N(\mathbf{s}_N) & \pi_1(\mathbf{s}_N) & \cdots & \pi_4(\mathbf{s}_N) \\ \pi_1(\mathbf{s}_1) & \cdots & \pi_1(\mathbf{s}_N) & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \pi_4(\mathbf{s}_1) & \cdots & \pi_4(\mathbf{s}_N) & 0 & \cdots & 0 \end{bmatrix} \begin{bmatrix} \mathbf{w}_1 \\ \vdots \\ \mathbf{w}_N \\ \mathbf{q}_1 \\ \vdots \\ \mathbf{q}_4 \end{bmatrix} = \begin{bmatrix} \mathbf{d}(\mathbf{s}_1) \\ \vdots \\ \mathbf{d}(\mathbf{s}_N) \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad (2.47)$$

or

$$\boxed{\begin{bmatrix} \Phi & \Pi \\ \Pi^\top & 0 \end{bmatrix} \begin{bmatrix} \mathcal{W} \\ \mathcal{Q} \end{bmatrix} = \begin{bmatrix} \mathcal{D} \\ 0 \end{bmatrix}} \quad (2.48)$$

where

$$\mathbf{d}(\mathbf{s}_i) = \mathbf{s}'_i - \mathbf{s}_i \quad (2.49)$$

LU and LDL[⊤] decomposition can be used to solve the non-positive symmetric linear system [13].

There are a variety of kernels commonly employed for RBF interpolation (**Table 2.7**). The shape parameter (ϵ) can be used to control the degree of local and global deformation [53].

There are three common types of kernels: infinitely smooth (1-4), compact (5) and piecewise smooth (6) [15]. An extensive overview of weights has previously been reported [14–16].

Table 2.7: Common RBF kernels used for shape interpolation and mesh-morphing [13–16]

	Type	Radial Function $\phi(r)$
1	Multiquadratic	$\sqrt{1 + (\epsilon r)^2}$
2	Inverse multiquadratic	$\frac{1}{\sqrt{1 + (\epsilon r)^2}}$
3	Inverse quadratic	$\frac{1}{1 + (\epsilon r)^2}$
4	Gaussian	$e^{-(\epsilon r)^2}$
5	Wendland	$(1 - r)^6 + (35r^2 + 18r + 3)$
6	Polyharmonic spline in \mathbb{R}^d	$\begin{cases} r^{2k} - d, & d \text{ odd,} \\ r^{2k} - d \log(r), & d \text{ even.} \end{cases}$

RBF interpolation presents several advantages and disadvantages to the problem of volumetric mesh-morphing. Most importantly, it has been shown to outperform FE-based transformations in several cases, for numerous parameter changes, i.e. deformations [13]. Although, in a comparison study for mesh-morphing phalanges of the hand, a FE-based strategy performed better than a RBF-based strategy [46]. Another advantage of RBF interpolation is that it is relatively simple to implement and straightforward to set up and operate. Furthermore, if customisations and optimisations are not required, some packages already provide implementations for some kernels, e.g. VTK (thin-plate spline only) [331] and PyGem (several kernels) [332]. A final advantage is that once the weights have been determined, the morphing operation is *embarrassingly parallel*, as each kernel can be computed independently [333]. Therefore, this part of the algorithm could be computed very rapidly, as it is simply a linear combination of weights and

kernels.

There are several disadvantages associated with RBF interpolation. The most critical disadvantage is that the overall operation is very computationally expensive. Generally, the use of RBFs has been deemed inappropriate for systems requiring more than 2,000 interpolants, regardless of application [327]. This is due to the large system of linear equations that are required to be solved to determine the interpolant weights. This threshold could easily be exceeded for a large proportion of meshes representing complex geometries, particularly anatomical structures derived from bioimaging data. Therefore, it may be an inefficient approach for morphing the volume of large meshes. Furthermore, in comparison to smoothing and FE-based approaches, it is reportedly more expensive [13].

The use of RBF interpolation presents an optimisation problem, to determine which kernel provides the best quality morphing operation. Additionally, there are conflicting reports about which kernel provides optimal results. One study found the triharmonic spline to provide the best mesh-quality over biharmonic, Gaussian, thin-plate and multi-quadratic kernels [334]. This was corroborated by another study which showed the triharmonic spline provided greater mesh-quality results for two out of three geometries, compared to the biharmonic and quadharmonic splines [13]. In contrast, a study that investigated a range of kernels reported the triharmonic spline to yield unacceptable results [51]. Moreover, the biharmonic spline provided the best mesh-quality with a scaled multi-quadratic variant yielding the next best results [51]. Additionally, another study which employed a biharmonic spline [335] stated it guarantees minimal quality degradation based on work that used biharmonic equations for the FEM [312]. This could be considered poor justification, as the FE study compared the fourth-order operator to a second-order operator, the Laplacian. Additionally, a conclusion from that research stated a higher order function was better, which could point to the triharmonic spline and beyond [312]. Other studies using RBF-based mesh-morphing strategies have employed the Gaussian [9, 53], inverse multi-quadratic (surface morphing) [53] and thin-plate spline [46, 52] kernels without justification.

Despite the relative simplicity of the implementation, optimising the efficiency of the method could pose some complications. Due to the high computational costs, this method should be optimised, where possible. Attention should initially be focussed on the solution of the linear system, as is it the most expensive aspect of the method. However, there are a variety of solving

procedures available that can achieve efficiency improvements to varying degrees. Although this is favourable, the accuracy and quality of the morphing operation is susceptible to different solver implementations [13]. This presents an additional optimisation problem, including which solver parameters yield optimal results.

One aspect could be regarded as an advantage or disadvantage depending on the perspective. There are few parameters that control a RBF interpolation, although this can depend on the choice of kernel. From an optimisation perspective, this could indicate there are fewer parameters to tune. However, this could present an issue and limit possible avenues of adjustment, if a morphing operation does not yield the desired results. From this perspective, the choice of interpolant could also be viewed as an equivocal aspect, providing additional routes to optimise the process.

Smoothing Operations

Mesh smoothing operations consist of various methods which are used to optimise the positions of mesh vertices with the objective of improving the quality of the associated elements [13, 336]. These methods are often performed as a post-processing operation following an initial mesh generation procedure [30, 336]. In some cases, smoothing operators are used to define the Dirichlet energy of a mesh for surface parameterisations [6]. For volumetric transformations, they can operate similar to the Dirichlet boundary conditions described for the FEM [311, 337]. Where displaced boundary vertices are constrained and the internal vertices are repositioned using a smoothing operator [338]. A key difference of smoothing methods is that they attempt to obtain a valid mesh, prescribed by a new boundary, by solving for new positions directly. Whereas, FE-based and RBF-based methods solve for vertex displacements. There have been a variety of different smoothing operators reported for polygonal meshes: Laplacian [13, 336, 339, 340], Poisson [336, 337, 341, 342] and metric optimisation [13, 340, 343, 344].

Laplacian smoothing is the most commonly employed mesh improvement operation [342, 345]. The definition and implementation bears remarkable similarities to surface parameterisation methods (*Section 2.5.2*) and barycentric coordinates (*Section 2.5.3: Barycentric Interpolation*). Additionally, it highlights the typical behaviour of a smoothing operator, where the use of vertex or element adjacency is used to optimise the quality of a mesh. Which is achieved by placing vertices towards the centre of their barycentric neighborhoods (*Figure 2.17*).




Figure 2.17: A general Lalacian smoothing operation (\mathcal{L}): improving element quality by moving a vertex (\mathbf{v}_i) towards the barycentric average

In general, the Laplacian (Δ) defines the sum of the second partial derivatives of a surface at \mathbf{v}_i , which is the local curvature. More accurately, this represents the divergence ($\nabla \cdot$) of the first partial derivatives of a vertex [286, 337]:

$$\Delta \mathbf{v}_i = \nabla^2 \mathbf{v}_i = \nabla \cdot \nabla \mathbf{v}_i = \frac{\partial^2 \mathbf{v}_i}{\partial x^2} + \frac{\partial^2 \mathbf{v}_i}{\partial y^2} + \frac{\partial^2 \mathbf{v}_i}{\partial z^2} \quad (2.50)$$

where $\nabla \mathbf{v}_i$ describes the first derivative (or gradient) of a surface, which defines the tangent to a surface at \mathbf{v}_i :

$$\nabla \mathbf{v}_i = \left(\frac{\partial \mathbf{v}_i}{\partial x}, \frac{\partial \mathbf{v}_i}{\partial y}, \frac{\partial \mathbf{v}_i}{\partial z} \right) \quad (2.51)$$

The discrete Laplacian operator (\mathcal{L}), sometimes referred to as the Laplace-Beltrami operator (**Equation 2.52**), approximates coordinate differentials for a piecewise linear function \mathcal{M} (**Figure 2.18**) [7, 286]. It consists of solving the Laplacian for a given domain (Ω) and can be discretised for a mesh describing that domain using the following [7, 286, 346–348]:

$$\Delta \mathbf{v}_i = - \lim_{diam(A) \rightarrow 0} \frac{\nabla A(\mathbf{v}_i)}{A(\mathbf{v}_i)} = -H(\mathbf{v}_i) \mathbf{n}_i \in \mathbb{R}^3 \quad (2.52)$$

$$\Delta \mathbf{v}_i = \lim_{|\gamma| \rightarrow 0} \frac{1}{|\gamma|} \int_{\mathbf{v} \in \gamma} (\mathbf{v}_i - \mathbf{v}) dl(\mathbf{v}) \quad (2.53)$$

Where, A is an area surrounding a point \mathbf{v}_i , $diam(A)$ is the diameter of that area, $H(\mathbf{v}_i)$ is the mean curvature and \mathbf{n}_i is the surface normal, together defining the mean curvature normal at \mathbf{v}_i [7, 346–348]. γ is a closed surface curve surrounding \mathbf{v}_i and $|\gamma|$ is the length of that curve [7]. A generalised version of the discrete Laplacian operator is defined by the following:

$$\mathcal{L}(\mathbf{v}_i) \approx \Delta \mathbf{v}_i \quad (2.54)$$

$$\mathcal{L}(\mathbf{v}_i) = \alpha \mathbf{v}_i - \beta \sum_{j \in \mathcal{N}(i)} w_{ij} \mathbf{v}_j \quad (2.55)$$

or

$$\mathcal{L}(\mathbf{v}_i) = \sum_{j \in \mathcal{N}(i)} w_{ij} (\alpha \mathbf{v}_i - \beta \mathbf{v}_j) \quad (2.56)$$




Figure 2.18: (a) Laplace-Beltrami operator (Δ) for a vertex \mathbf{v}_i on a surface and (b) discrete Laplacian Operator (\mathcal{L}) for the corresponding vertex on a surface mesh (*adapted from [7]*)

Equation 2.55 is often solved iteratively using a point Jacobi scheme [30, 168, 338]. Additionally, a solution of **Equation 2.56** in matrix form can be derived using same the construction for planar parameterisations (**Equation 2.6**) [7]. This is due to the systems being based on the same fundamental principles. Although these principles are based on surface smoothing for triangular meshes, **Equations 2.55-2.56** can smooth any surface or volumetric mesh.

There are several non-mutually exclusive variants of the Laplacian operator, which could be defined by six categories: absolute [342, 349, 350], relaxed [168], interpolated [351], smart [345], Taubin [348] and weighted [12, 338, 352]. Absolute smoothing simply places a vertex at the average of the one-ring neighbors ($\alpha = 0$ and $\beta = 1$) [351]. Relaxed smoothing applies a

relaxation factor to reduce the influence of the one-ring neighbors ($\alpha = 1$ and $0 \leq \beta \leq 1$) [168]. Interpolated smoothing interpolates the position between a vertex and the barycentric average ($0 \leq \alpha \leq 1$ and $\beta = 1 - \alpha$). Interpolated smoothing has been shown to reduce shrinking [351]. Smart Laplacian smoothing operators check the quality of a mesh before each vertex adjustment. A vertex relocation is only applied if it leads to an improvement with respect to a chosen mesh quality metric [345]. Taubin smoothing is used to shrink ($\alpha = 1$ and $\beta > 0$) and inflate ($\alpha = 1$ and $\beta < 0$) a surface in two steps [348]. Weighted smoothing refers to the use different weights (w_{ij}) to control and improve the operation. There are a variety of different weights that can be employed for **Equations 2.55-2.56**. For triangular meshes, these same weights are also used for surface parameterisations (**Table 2.4**). For volumetric meshes, weights can be determined through optimisation methods, such as the log barrier [352]. The log-barrier weights have been developed specifically for mesh-morphing applications and are more complex to derive [12, 352]. Meshes morphed using this operator (LBWARP) have achieved a better mesh-quality than other smoothing methods [12]. A commonly used weight for arbitrary meshes is based on a barycentric average (**Equation 2.57**) of the uniform spring weight ($w_{ij}^u = 1$) [351, 353]:

$$\sum_{j \in \mathcal{N}(i)} w_{ij} = 1 \quad (2.57)$$

\therefore

$$w_{ij} = \frac{w_{ij}^u}{\sum_{k \in \mathcal{N}(i)} w_{ik}^u} = \frac{1}{|\mathcal{N}(i)|} \quad (2.58)$$

This partition of unity approach has been applied to other weights for morphing applications, by replacing w_{ij}^u for other weights in **Table 2.4**, e.g. mean value coordinates [353–355].

Poisson operators can be used for smoothing, parameterisation and surface reconstruction [7, 280, 337]. Poisson smoothing operations are based on the solution to the Poisson equation [7, 337, 356]. They bear some similarities to Laplacian operators and can also be discretised using local neighborhoods [7, 336, 337, 342, 356]. Additionally, weighted Poisson-variants exist with the most well-known being the Winslow operator [340, 341, 356]. The Winslow operator is commonly used for smoothing structured meshes [340]. Poisson-variants exhibit a trade-off between preserving vertex spacing and element inversion resistance [341, 356, 357]. Winslow smoothing is able to provide greater resistance to element inversions than Laplacian operators

but with poor space preservation [341,357]. Other Poisson-variants have been found to preserve the spacing between vertices but are less robust and have greater susceptibility to inversions [341,357].

Optimisation-based methods minimise or maximum an objective function that assesses a mesh quality metric [301,340,345]. They are more effective than standard Laplacian smoothing at the cost of computational efficiency [345]. An example of a minimisation function for either vertex or element mesh-quality metrics is given in the following equation [340]:

$$\mathcal{F}(\mathbf{x}_i) = \sum_{i \in E} (q_i(\mathbf{x}))^2 \in \mathbb{R} \quad (2.59)$$

Where, \mathcal{F} is a mesh-quality objective function, $q_i(\mathbf{x})$ is a vertex or element quality for a given metric q_i and E is a set of elements [340]. A method called the target-matrix paradigm developed to optimise meshes has a similar objective function [301]. However, the objective function is not quadratic and employs the power mean of a quality metric. Smoothing operators based on these types of functions are more sophisticated and challenging to implement than the majority of Laplacian and Poisson-variants. However, both objective functions and several others are freely available in MESQUITE [340].

The smoothing methods present various advantages and disadvantages. Due to the diversity of implementations within this class of morphing algorithms, some advantages and disadvantages do not correspond to all. The most prominent advantage is held by the optimisation-based operators which represent comparatively fast and efficient morphing operators [12]. Additionally, this smoothing variant scales reasonably well with an increasing number of vertices, second only to the barycentric approaches. For one model investigated in a review of methods, an optimisation-based smoothing technique achieved the fastest time [12]. Another advantage is the availability for some of these variants. There are several sophisticated Laplacian, Poisson and optimisation-based operators freely available in packages like MESQUITE [340]. Also, some Laplacian-based variants are simple to implement but this strongly depends on the choice weights. An advantage for the Laplacian-based variants is that they offer flexible approaches to performing the morphing operations, either by solving a linear system or iteratively. This flexibility could prove advantageous for customised applications and testing aspects of a method in particular scenarios.

Despite these advantages there are several disadvantages that discourage the use of this class of methods. The most critical disadvantage is the decision trade-off when choosing a particular smoothing operation for volumetric transformation. In a review of volumetric morphing operators, the optimisation-based approaches are fast but produce some of the worst results. Whereas, the LBWARP operator can yield good results but in the slowest speeds. Also, the LBWARP operator scaled the worst with increasing vertices and was the most unpredictable with respect to speed. It was shown to produce drastically different speeds for meshes with a similar number of vertices. Consequently, it would be difficult to reliably predict how long a morphing process would take to complete for an application based on this type of operator. Additionally, the computations could take an excessive amount of time for larger meshes, despite yielding good results. Similar to FE-based methods, the speed penalties occur from the requirement to solve a sparse matrix of linear equations. However, it has been reported that FE-based methods can be improved using parallel computing techniques [12]. Also, due to the extensive use of the FE method it has received several optimisations over the years [325,358,359]. For the LBWARP operator, it may not be possible for this bespoke algorithm to directly gain from these benefits and may require algorithm specific optimisations.

Another trade-off lies in the implementation differences between operators. The more simplistic or readily available operators produce poorer results than other methods. In contrast, the more sophisticated operators with challenging implementations are required to achieve more reasonable results. The LBWARP operator is an example of a more sophisticated weighted-Laplacian variant as the implementation for the log-barrier weights are complicated [12]. Additionally, it is not available in any software applications, frameworks or packages. Another drawback for the LBWARP operator is that it is subject to convergence and tolerance issues [12]. Similarly, some weighted-Laplacian variants, such as LBWARP, cannot operate on meshes with inverted elements, whereas, all the other morphing operators mentioned in this review can. The ability to morph inverted elements could be useful in particular applications or intermediate stages of volumetric morphing. Therefore, the inability to use such operators in these conditions implies a limitation. For the optimisation-based approach, it has been shown to have poor robustness, producing large differences in mesh quality that are not always proportional to the degree of parameter changes [12]. An equivocal aspect of the Laplacian-based variants is related to the wide choice of associated weights. This could either present an optimisation problem or be used as a parameter to refine a process that yields undesirable results. However, there are other

volumetric morphing operations that can yield more favourable results with respect to speed, quality and ease of access.

Barycentric Interpolation

There are a variety of morphing techniques based on the principles of barycentric coordinates [12, 360]. Barycentric coordinates in their simplest form can be used to interpolate vertex data (e.g. location and colour) within a triangular domain [360,361]. Generalised barycentric coordinates consist of an affine or convex combination of vertices on a given boundary. This is used to interpolate points and associated data within a prescribed domain. The basic principles can be extended to interpolate vertex data in polygons and polyhedra [355,360]. The use of barycentric coordinates in computer graphics are widespread and include: mesh parameterisation [360,362], mesh-morphing [362], shape deformation [360], smoothing [336], vertex shading [360] and ray-tracing [363].

Given a domain $\Omega \subset \mathbb{R}^3$ defined by a closed triangular mesh \mathcal{T} on the boundary $\partial\Omega$, any point \mathbf{p} within the domain can be expressed as a combination of those vertices and a series of non-negative functions $\xi_i(\mathbf{p})$ or coefficients [360,364]:

$$\sum_{i=1}^N \xi_i(\mathbf{p}) \mathbf{v}_i = \mathbf{p} \quad (2.60)$$

where

$$\sum_{i=1}^N \xi_i(\mathbf{p}) = 1 \quad (2.61)$$

Also, this expression can be used for interpolating points on a planar mesh [354]. It is a natural extension of barycentric interpolation for points within a one-ring neighborhood [50]:

$$\sum_{i \in \mathcal{N}} \xi_i(\mathbf{p}) \mathbf{v}_i = \mathbf{p} \quad (2.62)$$

A common approach to normalising the coefficients, to ensure they sum to 1 (*Equation 2.61*),

is by dividing each coefficient by the sum of the coefficients [360]:

$$\xi_i(\mathbf{p}) = \frac{w_i}{\sum_{j=1}^N w_j} \quad (2.63)$$

This depicts a similarity to Laplacian smoothing operations. There are many ways to calculate the coefficients, which include: mean-value coordinates, harmonic coordinates and Wachspress coordinates [13]. Again, some of these coefficients are identical to the surface parameterisation weights presented in **Table 2.4**. The fact that the same weights can be used for both mesh parameterisation, Laplacian smoothing and barycentric interpolation highlight that these expressions are fundamental in geometry processing. An extensive overview of these coefficients have previously been reported [361].

The use of mean value coordinates are a common approach to defining the coefficients for barycentric coordinates and have been employed for a variety of applications [354, 364]. One common application, previously mentioned in **Section 2.5.2**, is mesh parameterisation, where the same weights can be used to derive piecewise affine maps for parametric domains [260, 360]. Additionally, mean value coordinates can be generalised to non-convex shapes, unlike some other coordinate definitions, e.g. Wachspress [360].

Once all the coefficients have been calculated for each internal point, a deformation or morphing operation can be performed. Similar to Dirichlet boundary conditions, the vertices of the boundary are displaced to achieve a desired target geometry [360]. The morphed internal vertices can now be calculated by applying the associated coefficients to the displaced vertices:

$$\sum_{i=1}^N \xi_i(\mathbf{p}) \mathbf{v}'_i = \mathbf{p}' \quad (2.64)$$

For a volumetric mesh, \mathbf{p} could be represented by any internal vertex ($\mathbf{v}_i^{\text{internal}} \subset \mathcal{V}$). Depending on the chosen weights, extrapolation of vertices outside the geometry is possible, e.g. mean-value coordinates [360].

A common way to use utilise barycentric coordinates is to define a control mesh, or cage [360]. This control mesh encompasses a mesh of interest to be morphed. The surface vertices of a mesh could be used to define a control mesh. Alternatively, to increase the efficiency of the process a cage with fewer control vertices is specified. Therefore, fewer control points are required to

be adjusted and used for interpolation. To achieve this, the control mesh is often constructed manually [365]. However, if there are too little control vertices the morphing operation will not be smooth and conformal [365]. A control mesh could be defined automatically using a mesh decimation operation on a slightly expanded surface. At the expense of regularity, this could speed up the process whilst being able to specify a number of evenly distributed control vertices. Similar automated approaches have previously been reported [366]. There is a variation of this morphing operation that uses a triangular mesh of a boundary differently. For hexahedral meshes, the vertices of a mesh boundary are tessellated to acquire a triangular representation. The vertices of that triangular mesh are then used to define tetrahedral elements that fill the domain. For all internal vertices, the enclosing tetrahedra are determined. Interpolation of the morphed location can be performed using the four vertices of a tetrahedron. This method is referred to as simplex-linear and is more efficient than the cage methods [12]. Simplex-natural neighbor expands this criterion, using barycentric coordinates defined by vertices associated with the neighboring tetrahedra, not just the one enclosing tetrahedron [12]. This is an attempt to increase accuracy and achieve better quality meshes at the expense of increased computational cost. This could be thought of as a sub-cage compared to using the entire cage in those associated methods.

There are several advantages and disadvantages associated with volumetric morphing operators based on barycentric coordinates. In a review of volumetric morphing operators, two techniques based on barycentric coordinates (simplex-linear and simplex-natural neighbor) were the fastest algorithms. Simplex-linear yielded faster results than simplex-natural neighbor [12]. Furthermore, the two operators were incredibly efficient with respect to scaling. For simplex-linear, there was almost no difference in the speed of the algorithm when morphing meshes between 10,000 - 100,000 vertices. All meshes which had the number of vertices within this range were morphed in under a second [12]. Not only were they the fastest operators but also produced respectable results. Although, the results were not as good as the FE-based operator, which was recommended by the authors of that review. The authors concluded that the simplex-linear operator should also be considered for applications involving volumetric morphing, due to the impressive efficiency and good quality characteristics [12]. Additionally, these two operators were the most predictable with respect to the scaling [12]. This would provide greater reliability when predicting how long a process will take to complete for a given mesh.

One disadvantage for these methods is that the geometries require a triangulated boundary. However, this pre-processing step is fairly non-trivial due to the ubiquitous availability of high-quality Delaunay tessellators. Another disadvantage is due to the complexity and poor availability of morphing operators based on barycentric coordinates. Although the concepts of barycentric coordinates are relatively simple, the simplex-linear and simplex-natural neighbor techniques are not particularly straightforward to implement. Furthermore, these algorithms are not available in any commonly used packages and frameworks. Similar to weighted-Laplacian variants, the complexity of the implementation for some operators can depend on the choice of weights. However, the implementation of morphing operators based on barycentric coordinates could be eased by using a computational geometry frameworks, such as libigl [367] and CGAL [368]. Additionally, there are a wide choice of weights that can affect the quality of the results and therefore present an additional optimisation problem. Despite achieving good results, there have been concerns addressed that this class of morphing operators may not preserve element qualities due to a lack of intrinsic smoothness of the algorithms [12, 13].

2.5.4 Conclusions

In conclusion, there are several candidates within each core component that can enable the development of an automated mesh-morphing strategy, which can produce accurate and high-quality meshes. Regarding the feature alignment algorithms, the strongest candidate for accurate alignment are the Procrustes methods, if they can be combined with an automated and robust method of determining known correspondences. A compelling option for this would be the end-points of skeleton or centreline decompositions. For the surface parametrisation algorithms, there are a variety of well-established techniques, concepts and principles that have been developed for triangular meshes. A key shortcoming of these approaches for mesh-morphing is the mapping to a base domain (indirect parameterisation), and minimisation of the associated distortions unrelated to a volumetric mesh. However, the general concepts and principles can be extended in numerous ways to develop novel algorithms more suitable for quadrilateral and hexahedral meshes, through direct parameterisation. For the volumetric transformation operators, FE-based approaches offer the best compromise between mesh-quality, speed and flexibility; specifically the elasticity-based approaches. Consequently, these aspects have been incorporated into the mesh-morphing strategies developed in this thesis.

2.6 Applications of Mesh-Morphing

The term mesh-morphing can also be referred to as mesh-matching [55], mesh-metamorphosis [369], mesh-deformation [370], mesh-warping [371], mesh-moving [372], non-rigid registration [373] and deformable registration [47]. There are a variety of different ways to develop mesh-morphing strategies using the three core components: feature alignment, surface parameterisation and volumetric transformation (*Table 2.8*). A mesh-morphing strategy combines the core components in a particular manner to achieve a desired result. *Table 2.8* provides a summary of mesh-morphing strategies used for biomedical applications and the three key components from which they are composed.

There is greater consistency over the methods used for feature alignment and volumetric transformation. In comparison, there is little consistency for surface parameterisations. The feature alignment and volumetric transformation components can be applied reasonably straightforward, following the derivation of landmarks or surface correspondences, respectively. Consequently, the surface parameterisation component is arguably the most important, having the strongest impact on the performance and type of strategies. Additionally, this component strongly influences the degree of user-interactivity required to operate a strategy. Although all work together synergistically, larger errors between correspondences will lead to greater distortion and surface errors. In contrast, minor alignment issues can affect a surface parameterisation, but most methods are able to overcome them as they are generally expected. Currently, most volumetric transformations are based solely on the surface parameterisation step, as it defines the boundary conditions.

2.6.1 Applications

There are several reasons mesh-morphing strategies have been developed:

1. **Mesh Generation:**

The most common is the rapid and simplified generation of anatomical meshes from target geometries, particularly surface meshes. In addition to speed, an added benefit over traditional mesh generation techniques are that the attachment of tissues and prosthetic components can be transferred to a new model [52,374].

2. Image Segmentation:

This objective attempts to address the precursor to the target geometries [27,57,375,376]. These morphing operations aim to derive anatomical meshes directly from bioimaging data, incorporating segmentation and meshing into one combined process.

3. Surgical Optimisation:

The optimisation of surgical procedures and clinical therapies is an important problem that needs to be addressed [3,377]. One example has used mesh-morphing to rapidly optimise the positioning of cementless total hip replacements [54]. These types of applications could guide surgeons to more successful patient outcomes and fewer implant failures for a variety of procedures.

4. Dynamic or Moving Meshes:

Mesh-morphing has also provided solutions for moving meshes, which is common for dynamic simulations in cardiology [59,300,378]. Moving meshes must be updated at each time step, in order to accurately simulate a dynamic boundary [300]. The key problem is maintaining the original topology and mesh-quality between boundary changes, which is also a core objective of mesh-morphing in general.

2.6.2 Feature Alignment

Target and template meshes are often aligned using rigid registration techniques. A large proportion of strategies have used landmark-based rigid registration, such as a Procrustes algorithm [47] or Rodrigues' formula [45]. Most take advantage of the fact manually defined landmarks have already been acquired and therefore is a non-trivial step using existing algorithms and data. Interestingly, only two have used an automatic approach for anatomical rigid registration, which was ICP [46,58]. However, one of those strategies had used Procrustes alignment first, due to ICP yielding insufficient results in isolation [58]. Another key limitation of ICP is that the geometries must be reasonably similar [47]. In contrast, registration based on manual landmarks can achieve greater accuracy than automatic algorithms, at the cost of speed and required user-interaction. Surprisingly, no strategies have used RANSAC, which is likely due to similar shortcomings to ICP, despite overall improvements.

In some cases, alignment is not required for one of three reasons:

1. The alignment operation is incorporated into the volumetric transformation [51]
2. The target is a morphed surface from a template mesh [54, 59]
3. The target is interpolated from a template mesh [56]

Usually, alignment is performed to aid the surface parameterisation step. However, where correspondences are already known, displacement-based volumetric morphing operations are able to naturally incorporate such alignment transformations. Additionally, in relation to reasons 2 and 3 above, the target meshes were already aligned based on how they were derived. These were based on either optimisation [54] or moving mesh [56, 59] applications, which were able to take advantage of these freedoms. In contrast, the majority of other applications cannot, especially the automatic approaches.

2.6.3 Surface Parameterisation

Amongst the strategies, there are varying degrees of user-interactivity requirements defining the semi-automatic procedures. The majority of strategies (8/15) reviewed in **Table 2.8** used manual landmarks to determine one-to-one correspondences. The choice of using manual landmarks can offer greater accuracy at the expense of labor. Additionally, an increased number of landmarks are required for higher degrees of accuracy, particularly for complex geometries, which could exacerbate time, user-interaction and knowledge requirements. Although, there is research attempting to automate the selection of landmarks [232, 334]. Operations that require a large number of manual landmarks will naturally start to incur the negative effects of rater-bias and drift, which some have attempted to measure [53, 55]. Manual landmarks are often used to both align and parameterise a target geometry and template mesh. An excellent use of these landmarks in some strategies has been to semi-automate the segmentation process, as well as the generation of volumetric meshes [27, 375].

Apart from manually defining landmarks, there is no common approach to determining the surface correspondences. Seven of the reviewed strategies used an automated approach to define correspondences [45–47, 52, 57–59]. Of these, three used surface projection techniques [45–47]. The vertices were either projected based on surface normals [46, 47] or from a centreline [45]. For surface normal projection, geometries must be reasonably similar for these techniques to accurately achieve surface correspondences and minimise element distortion. Strategies based on this approach could be susceptible to local irregularities which could be problematic for

regions that are thin or have high curvature. Although, this could be overcome with the use of an iterative procedure [46,47]. The use of a centreline would likely improve the accuracy of the projections as global information about the geometry would be encoded into the projection vectors.

The remaining four non-projection based methods employed either: statistical shape modelling [58], an image-based deformation field [57,59] or a surface wrapping technique [52]. A SSM has been developed that contains the template correspondences for each shape [58]. The associated shape, shape parameters and correspondences were then identified for the target, by minimising a function based on the SSM and biharmonic parameterisation. A key issue with this method is that the target geometry must be present within the SSMs population, in order to derive accurate results. The first image-based method employed registration-based propagation techniques to determine the optical flow between image frames of a cardiac cycle [59]. A triangular mesh of the geometry in the initial frame was registered to the other frames sequentially based on the optical flow, providing correspondences for each frame. In contrast to optical flow, a related method used image-based forces to morph the surface of a template mesh, which was based on voxel coordinates and a Gaussian smoothing kernel [57]. A limitation of this method is that it requires images with a high contrast to yield accurate results. The last non-projection based approach employed an auxiliary surface to wrap the template and target meshes using energy minimisation techniques [52]. This could be considered similar to spherical parameterisation approaches, as the auxiliary surfaces are analogous to parameterised spheres. Interestingly, the use of traditional parameterisation techniques have shown minimal adoption in the biomedical modelling community, unlike their wide-spread use in the graphics community. Consequently, only two strategies used a traditional parameterisation technique: planar [9] and spherical parameterisation [379]. However, both required the definition of manually defined landmarks to improve the accuracy of the parameterisations. In the field of computer graphics, the combined use of spherical parameterisation and centrelines has provided an automatic approach for triangular meshes [369]. This type of strategy uses the centreline end-points to automatically define landmarks for the alignment of spherical correspondences prior to morphing. Although a promising automatic solution, it is yet to be seen how this approach operates over volumetric meshes. However, this approach could cause issues, due to a reduced set of skeleton-based landmarks compared to those manually or anatomically defined.

Table 2.8: Review of mesh-morphing strategies for biomedical applications

Reference	Component of Mesh-Morphing Strategy			Mesh Type	Anatomy	Simulated?
	Feature Alignment	Surface Parameterisation	Volumetric Transformation			
Behdadfar et al. [45]	Rigid registration (Rodrigues' formula)	Centreline projection	LET (Dirichlet BC)	HEX8	Ventricle (heart)	✓
Baldwin et al. [27]	Manual transformation	ML	Linear interpolation	HEX8	Patella, tibial + femoral cartilage	✓
Li et al. [51]	NR - applied in RBF	ML	RBF (biharmonic)	HEX8	Brain	✗
Grosland et al. [47]	Rigid registration (ICP)	Normal Projection	LET (mixed BC)	HEX8	Phalanx (hand)	✗
Magnotta et al. [46]	Rigid registration (ICP)	Normal Projection	RBF (biharmonic)	HEX8	Phalanx	✗
Hadagali et al. [48]	Rigid registration	ML + dual kriging + manual adjustment	Multi-block decomposition	HEX8	Vertebra	✗
Sigal et al. [52]	Not specified	ML or energy minimisation	RBF (biharmonic)	TET4	Vertebra	✓
Grassi et al. [53]	Rigid registration	ML + RBF (IMQ) + Laplacian smoothing	RBF (Gaussian)	TET4	Femur	✓
Bah et al. [54]	NR - already aligned	NR - same mesh topology	LET (Dirichlet BC)	TET4	Femur	✓
Salo et al. [55]	Rigid registration	ML	Manual Mapping	TET4	Pelvis	✓
Xu et al. [56]	NR - interpolated mesh	ML	RBF (triharmonic)	TET4	Ventricle	✓
Baghdadi et al. [57]	Manual transformation	Image-based deformation + Laplacian smoothing	Laplacian smoothing	TET4	Artery (carotid)	✗
Hraiech et al. [9]	Rigid registration	Planar parameterisation + ML + RBF (Gaussian)	RBF (Gaussian)	TET4	Femur	✗
Lauzeral et al. [58]	Rigid registration (Procrustes then ICP)	Statistical shape model	LET (Dirichlet BC)	TET4	Femur + liver	✗
Upendra et al. [59]	NR - same mesh	Image-based deformation	LBWARP	TET4	Ventricle	✗

ML - Manual landmarks

NR - Not required

BC - Boundary condition

LET - Linear elastic transformation

IMQ - Inverse multiquadratic

2.6.4 Volumetric Transformation

RBF interpolation was the most popular choice amongst the reviewed strategies for the volumetric transformation step, with little consistency over kernel choices [9, 46, 51–53, 56] (as highlighted in *Section 2.5.3: RBF Interpolation*). The biharmonic or thin-plate spline was the most commonly used, followed by Gaussian and then triharmonic. Two interesting concepts were employed by some studies to define a RBF transformation: sparse data [51, 56] and affinity control [9, 53]. The use of sparse data reduces the computational requirements to solve a transformation, but is also necessary when using manual landmarks. When using sparse data, there were examples of controlling the affinity of mesh vertices relative to their distance from landmarks, to reduce distortion and improve parameterisation. Essentially, vertices closer to landmarks would move a greater proportion than those more distant. The second most common volumetric transformation method was the linear elastic transformation [45, 47, 54, 58]. The boundary conditions used for those transformations were either Dirichlet or mixed. The method that employed mixed boundary conditions was Neumann-dominant and described as a force-based approach [47]. Force-based boundary conditions led to reduced distortions and sensitivity for larger geometric differences when compared to a displacement-based approach (RBF interpolation). This makes sense as it would allow some flexibility to alleviate correspondence errors derived from a surface parameterisation step. Furthermore, this highlights how FE-based transformations can offer a greater variety of boundary conditions compared to solely Dirichlet. Also, hierarchical refinement has been implemented to reduce computational requirements [47], similar to the use of sparse data for RBF interpolation. Additionally, similar affinity control can be achieved for linear elastic transformations by specifying inhomogeneous material properties (element-wise), which can also serve the benefit of minimising the distortion of lower quality elements [54, 58]. Currently, no non-linear elastic transformations have been found for biomedical mesh-morphing applications.

2.6.5 Summary

Most strategies employ manual landmarks, which can provide greater accuracy at the expense of time and user-interaction. Automated surface parameterisation methods can be split into four categories: projection-based (inside-out), image-based (inside-out), SSM-based (direct) and wrapping-based (outside-in). Although interesting and novel, methods based on bioim-

ages and SSMS have several shortcomings. Bioimaging data for non-moving meshes require high-contrasting boundaries and are subject to artefacts and noise. In contrast, SSMS must represent a large proportion of the population and acquiring such data could be expensive. The methods that offer the greatest flexibility are based on wrapping (outside-in) and projection algorithms (inside-out). Furthermore, projection based on centrelines would likely perform better than surface normal vectors. Global geometric information is encoded in centrelines, which could prevent an inaccurate parameterisation of irregular features between meshes. An algorithm that combines aspects of wrapping with centreline-based projection could provide greater robustness by utilising the aspects of both outside-in and inside-out algorithms. Consequently, it is proposed that novel surface parameterisation algorithms based on these concepts, and extended principles of traditional parameterisation techniques, would provide greater accuracy and robustness for morphing hexahedral meshes. Additionally, the use of centrelines would combine favourably with Procrustes methods, to automatically determine landmarks for accurate and robust geometry alignment. As such, these designs have been incorporated into the strategies developed in this thesis.

RBF interpolation is commonly used for volumetric transformations and has even been used to improve the parameterisation of manual landmarks [9, 53]. However, following closely behind is the linear elastic transformation, which could be more complex to implement but offers greater flexibility and control. Specifically, there are three fundamental benefits of linear elastic transformations:

1. **Piecewise element stiffness control:**

The control over the stiffness of individual elements has been used to prevent further distortion and negative Jacobians [54, 58]. RBF interpolation is a meshless approach and is therefore unaware of elements.

2. **Variety of boundary conditions:**

There is greater flexibility to the types and combinations of boundary conditions that can be defined, which can improve results [47]. In comparison, RBF interpolation can only perform displacement-based boundary conditions.

3. **Intrinsic robustness:**

Linear elastic transformations incorporate information regarding element quality, topol-

ogy and associated DOF, which can provide greater robustness [189]. In contrast, RBF interpolation has no awareness of these aspects, most importantly element-quality [58].

For these reasons and the greater overall advantages, linear elastic transformation has been selected for the strategies developed in this thesis.

2.7 Conclusions

FE modelling of the TFJ provides an important tool for studying and learning about knee biomechanics [3]. Furthermore, they enable the optimisation and assessment of several clinical therapies. For accurate patient-specific modelling, a greater complexity of model definitions are required, which include a variety of supporting tissues, as well as their associated material properties and kinematics. However, simplified models can be useful for determining isolated aspects of the knee biomechanics and reducing assumptions [18]. Additionally, simplified models could more beneficial than complex models for rapidly comparing different development protocols. These analyses would focus on the differences between common output variables, such as contact area and contact pressure.

Although useful and important, the development of high-quality FE models are laborious and time-consuming, due to several stages of processing bioimaging data. In general, these stages are difficult to perform accurately and efficiently, and this is exacerbated by the requirements of hexahedral meshes. Finite element applications typically do not discriminate between tetrahedral and hexahedral meshes. Ideally, they solve a PDE problem with the most appropriate methods available. However, it is evident that they do solve hexahedral elements with greater accuracy and performance. In some cases these differences are more subtle; nevertheless, there are still applications where hexahedral meshes are required to achieve accurate and convergent simulations [113]. These are five situations where hexahedral meshes are generally required to yield greater accuracy:

1. Large deformations [116]
2. Contact analyses [29]
3. Nearly-incompressible materials [29]
4. Definition of fibre-alignments for anisotropic material properties [35]

5. Definition of zonal properties for inhomogeneous material properties [40]

Consequently, the state-of the art and most commonly used meshes for the soft tissues of the knee are hexahedral meshes, which includes the meniscus (*Figure 2.7*).

The meniscus has an irregular geometry and is an important component of the knee, as it strongly influences the biomechanics. Additionally, the meniscus is a perfect example of an anatomy that can be simulated with greater accuracy with respect to each of the five situations above. Despite the impressive results of TET15 meshes for the hip [34], further research is required to determine their suitability and accuracy for knee models. Additionally, their use is more novel and they are less available than their TET10 counterparts, which have been shown to perform poorer than HEX8 meshes [29,116].

However, the generation of hexahedral meshes remains a challenging endeavour, particularly for irregular geometries and anatomies. Currently, the most viable and well-established options for biomedical applications are the multi-block and mesh-morphing methods. As discussed, multi-block methods have proven their capacity for deriving high-quality meshes for a wide range of anatomies and have deserved their state-of-the-art status [178,197]. However, with increasing popularity and adoption of biomedical simulations for clinical applications, their time and user constraints impose severe barriers to current and future biomedical innovations. In contrast, mesh-morphing methods are able to generate high-quality meshes for an equally vast array of anatomies, but require a pre-defined high-quality template mesh. However, the problem for biomedical applications is slightly different to other fields, we know approximately what the shape will be and the type of mesh that would best suit particular anatomies. If a large proportion of similar geometries require hexahedral meshing, the initial time and user constraints would be insignificant compared to the rapid and full-automation of high-quality meshes. This is essentially an analogue to the mass-manufacturing process of the injection moulding industry.

Consequently, several mesh-morphing strategies have been developed (*Table 2.8*), to improve the ease and speed of generating high-quality hexahedral meshes of anatomies. However, crucial information pertaining the three important criteria (previously outlined in *Section 2.4.1*) and simulation validation are often lacking critical analyses. Of the reviewed strategies, only one [53] evaluated all four of the following criteria:

1. Mesh-quality:

Hexahedral mesh-quality metrics are often reported, however they are inconsistent between studies, for example using: the scaled-Jacobian [48, 51, 241, 245], aspect-ratio [45, 48, 241, 245], skewness [48, 51], warpage [48, 51], distortion [46, 47] and volume [46, 47] metrics. The importance of mesh-quality is evident, however, these inconsistencies could be due to the difficulty of defining and characterising this attribute [161]. Locational information and the associated importance for particular applications are rarely provided.

2. Surface fit:

Similar to mesh-quality, surface error metrics are often reported with neither adequate information for particular regions (i.e. as overall measures of loss), nor how this would affect a given application. The chosen metrics are less inconsistent than those for mesh-quality, which includes: the root-mean-square error [27, 48, 53, 55, 57], maximum error [47, 48, 53, 55, 57], mean absolute difference [47, 55, 59] and median error [52].

3. Generation time:

Generation times and comparisons between previous methods are often not provided, even when there are claims to a decrease compared to other development procedures [46, 52]. Additionally, any scenarios that could affect time differences are generally not described.

4. Simulation validation:

This is another important criterion to determine the success of a morphing strategy, which could be added to the 14 evaluation criteria (*Section 2.4.1*). Simulations should serve two purposes: (1) proof of convergence and (2) proof of accuracy (i.e. no significant changes compared to gold-standard procedures). However, less than half of the reviewed studies performed simulations. In some cases the simulations were not compared, perhaps to only evaluate convergability [45, 56]. Furthermore, the clinical implications from simulation differences or their cause were rarely discussed.

Greater results have arisen from strategies that blend together the aspects of surface parameterisation and volumetric transformation [46, 47, 57]. Essentially, both processes become iterative and interwoven, by gradually parametrising and morphing a volumetric mesh and associated boundary in unison. In isolation, iterative (relative) volumetric transformations have been

shown to improve mesh-quality compared to non-iterative (absolute) equivalents [12]. However, by breaking down the distinct boundary between surface parameterisation and volumetric transformation, the diametric problem associated with surface error and element degradation can be approached simultaneously. Surface parameterisation techniques have worked well for surface meshes, but greater flexibility is required to accommodate volumetric meshes. This is important because perfect correspondences (i.e. low surface errors) could still lead to invalid or poor-quality elements [46, 47].

Skeleton-based methods have yielded promising solutions for several natural geometries and anatomies. However, skeleton structures have also been incorporated and combined with other methods, such as polycube [214] and multi-block methods [169], to approach a wider range of geometric domains and complexity. This shows the flexibility and value that skeleton structures can provide to different approaches, as observed for the more established medial-axis methods. Skeleton-based methods can be separated from other methods that use skeleton information by their direct decomposition from them apposed to using them to provide geometric correspondences (e.g. mesh-morphing [45]) and singularity information [214]. The multi-informative nature of skeleton structures, providing geometric feature descriptors, correspondence information and singularity information [214], results in the flexible use of some or all of this information for improving mesh generation algorithms.

There is a drive towards full-automation for mesh-morphing applications. Additionally, there appears to be a drive for a more general-purpose mesh-morphing strategy, that could be applied to a variety of anatomies through the use of different templates. Currently, there are no mesh-morphing strategies developed or tested on the meniscus. There was one for other soft tissues in the knee, but it was highly user-interactive and based on bioimages [27]. In general, the long-term objectives are that rapid and cost-effective solutions would have a profound impact for clinical applications, benefiting a wide range of patients. However, only three of the reviewed mesh-morphing strategies could claim to be fully automatic and general-purpose [46, 47, 58].

3 Mesh-Morphing Strategy Design

3.1 Overview

Two mesh-morphing strategies have been developed. The first strategy is composed of four phases: (1) centreline generation, (2) surface parameterisation, (3) surface optimisation, (4) volumetric transformation (*Figure 3.1a*). The second strategy was designed as an optimisation of the first, and consists of three phases: (1) initialisation step, (2) volumetric optimisation and (3) Laplacian smoothing (if inverted elements exist) (*Figure 3.1b*). The focus of the overall designs for the strategies was to select automatic, robust and accurate algorithms. Consequently, the three core components of the first strategy were the following:

1. **Feature alignment:** Centreline-based rigid registration
2. **Surface parameterisation:** Centreline projection and mesh-wrapping
3. **Volumetric transformation:** LET (Dirichlet BC) and Laplacian smoothing

The method used to align the features is novel, and was selected due to being simple to implement, accurate and fast (automated). Justification for the choices of using centreline projection, mesh-wrapping and linear elastic transformation over alternatives was provided in *Section 2.6.5*. Briefly, the novel design of the first strategy extended and combined several basic principles associated with centreline-based mesh-morphing [45] and surface parameterisation techniques, such as mesh-wrapping [52] and spherical parameterisation [295]. The three core components of the optimised strategy are difficult to distinguish. First, no alignment is required, as the morphed mesh from the first strategy is used (i.e. already aligned). Second, the surface parameterisation and volumetric transformation were interwoven into one iterative process (volumetric optimisation). The justification for combining these components was provided in *Section 2.7*. The aim of this strategy was to incorporate and extend concepts of traditional surface parameterisation techniques for volumetric meshes.

The strategies were implemented in C++ [380] with the CMake build system (GCC compiler) [381], and made extensive use of the Armadillo linear algebra library [382] and the visualisation toolkit (VTK) [331]. Additionally, the open computer language (OpenCL) [383] was used

to provide performance improvements through parallel graphics processing unit (GPU) computation. The choice of implementing the strategies in this manner were due to the superior speed and flexibility offered, compared to other approaches. This is emphasised by the iteration-heavy operations using linear algebra for mesh-processing algorithms, where C++ (with Armadillo) has a clear advantage [315,384]. Also, this approach provides greater flexibility for target operating systems, the availability of libraries, and integration with and porting to other languages (e.g. python). The importance of speed cannot be understated, as the key objectives for the strategies were to be fast and automated. This allows their use for real-world applications and large resolution meshes, and to compete with existing procedures. Additionally, for the majority of alternative languages, the computational complexity can compound into much slower execution times. The strategies were implemented on a Dell XPS 9560 (2.8GHz Intel i7 CPU, 32GB RAM and an Nvidia GTX 1050 GPU with 4GB RAM) running Debian (Linux).

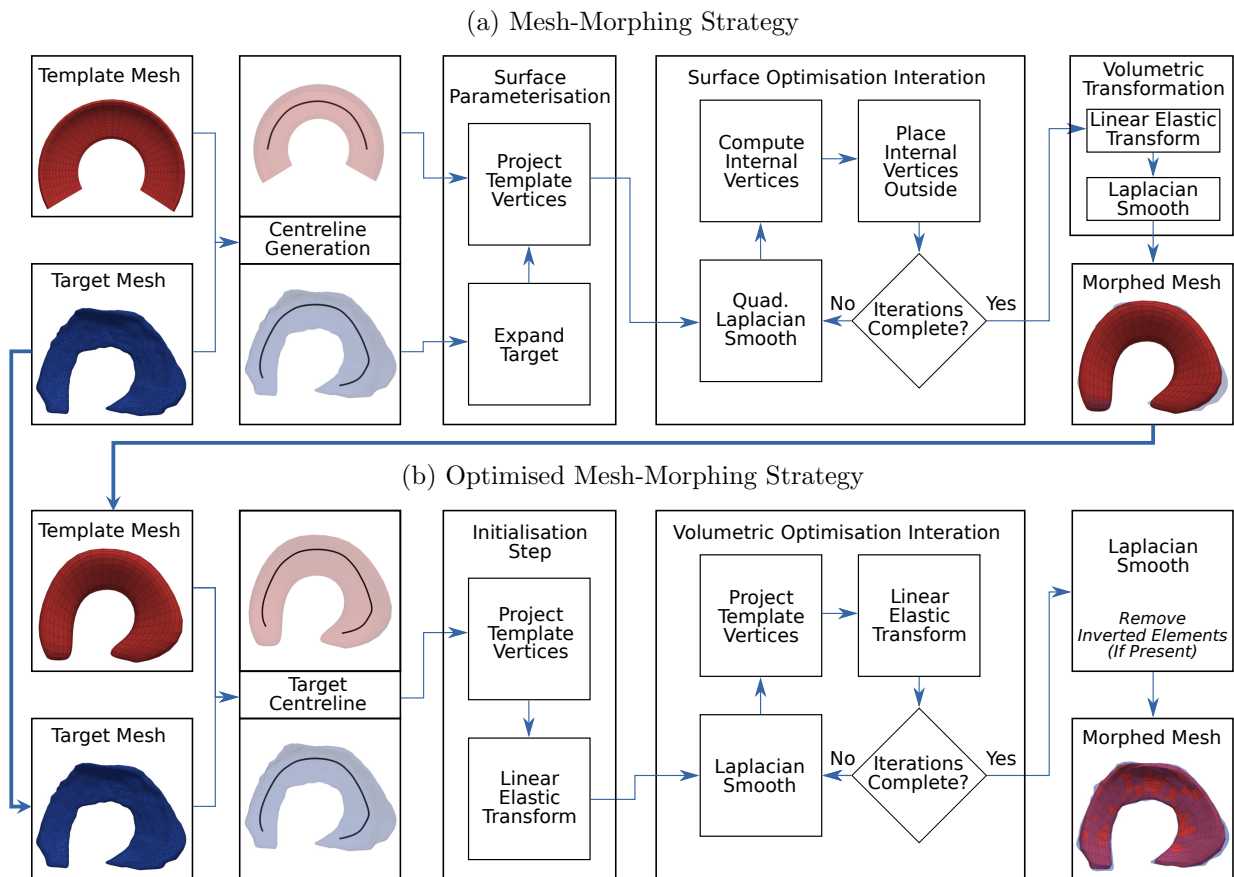


Figure 3.1: Flow chart of the input data and algorithms that compose the (a) mesh-morphing strategy and (b) optimised mesh-morphing strategy

Four algorithms contained components of existing code: (1) the shortest path algorithm (*vtkDijkstra* [331, 385], **Section 3.2.4**), (2) B-spline computation (Dierckx [386], **Section 3.2.5**) (3)

the ray-triangle intersection algorithm (Möller-Trumbore algorithm [387] re-written in OpenCL, *Section 3.4.1*), and (4) solution to the linear elastic equations (MFEM [388], *Section 3.6.1*). All other algorithms were implemented from first principles and cited where relevant.

3.2 Centreline Generation

The centrelines are generated automatically using three algorithms: (1) the minimisation of a signed distance function, (2) clustering of the collapsed surface into a line mesh and (3) B-spline parameterisation of the line mesh. The signed distance function (SDF) was defined using radial basis functions (RBFs), similar to the process outlined in *Section 2.5.3: RBF Interpolation*. Following the minimisation of this function, the surface vertices collapse into the shape of a centreline. To enable cross-parameterisation between centrelines and determine correspondences, vertices were clustered within a given radius and then connected into a series of lines. To improve the accuracy and precision of the curve representation and cross-parameterisation, the clustered vertices of the centrelines were mapped into a B-spline curve. A large proportion of the centreline generation algorithm was developed based a previous study [328]. In particular, the generation and minimisation of an RBF-based SDF to collapse a surface mesh into a centreline. However, these processes were computationally expensive with respect to the number of surface vertices and were optimised to improve the speed of computation. Additionally, improvements to the RBF-based SDF implementation were based on a slightly different definition found in a previous study that were using SDFs to reconstruction and represent 3D surfaces [327].

3.2.1 Generation of a Signed Distance Function from Radial Basis Functions

To generate a RBF-based SDF the surface is first modelled implicitly. An implicit surface is defined by a function where all the points that lie on the surface are equal to 0 (*Equation 3.1*).

$$F(x, y, z) = 0 \tag{3.1}$$

A signed distance function is used to calculate the distance of a point (P) in space from a boundary, where a positive value is outside the boundary and a negative value is within the boundary. Here, the signed distance function is defined from the vertices of three layers generated from a given surface (*Figure 3.2*). The three layers consist of the actual surface, as well as the normal expansion and contraction of that surface. The distance of the actual sur-

face vertices is assigned to 0, whereas the expanded and contracted vertices are assigned +1 and -1, respectively. Essentially, this defines the arbitrary distances of the expanded and contracted boundaries from the original for the SDF. These are the constraints (c_i) of the system, which are analogous to the vertices (\mathbf{v}_i) used in deformation operator described in **Section 2.5.3: RBF Interpolation (Equation 2.44)**. Consequently, the number of constraints (N) are three times the number of surface vertices (n) for the original surface. Formulating a SDF from the three layers of vertices and their associated distance values involves solving a scattered data problem [327]. RBFs are used to solve this problem and interpolate the vertices to their respective distance values. Once the system of RBFs have been determined, they can be minimised to find the points in space which are furthest away from their boundary, i.e. skeleton or centreline, depending on the given boundary. The following derivation includes the use of RBF regularisation using λ_I . Regularisation of the RBF system has been shown to encourage the convergence towards typical level set behaviour and improve the directions of the gradient field vectors towards local minima [328]. Essentially, the regularisation parameters modify the diagonal of the RBF matrix, adding different magnitudes to the surface (e.g. 0.01) and external vertices (e.g. 1.0).




Figure 3.2: The RBF-based signed distance function is defined by three layers: the original surface (black, $h(\mathcal{S})$), expanded surface (purple, $h(\mathcal{S}^+)$) and contracted surface (blue, $h(\mathcal{S}^-)$).

Points inside the boundary have a negative distance value, and vice versa

Lemma 3.1. *Solve a linear system of equations to determine the weights (\mathcal{W} and \mathcal{Q}) for a SDF ($h(c_I)$) derived from a matrix of RBFs (Φ):*

N - number of constraints

$I, J = 1, \dots, N$

$$h(c_I) = \underbrace{\sum_{J=1}^N w_J \phi_J(c_I)}_{\text{RBF}} + \underbrace{\sum_{k=1}^4 q_k \pi_k(c_I)}_{\text{Polynomial}} \quad (3.2)$$

where

$$\phi_J(c_I) = (\|c_I - c_J\|^2 - \epsilon^2)^{-0.5} \quad (3.3)$$

ϵ - control radius

$$\pi_k(c_I) = (c_I^x, c_I^y, c_I^z, 1) \quad (3.4)$$

The linear system for the RBF-based SDF ($h(c_I)$) is defined as follows:

$$\boxed{\begin{bmatrix} \hat{\Phi} & \Pi \\ \Pi^\top & 0 \end{bmatrix} \begin{bmatrix} \mathcal{W} \\ \mathcal{Q} \end{bmatrix} = \begin{bmatrix} \mathcal{H} \\ 0 \end{bmatrix}} \quad (3.5)$$

where

$$\hat{\Phi} = \Phi + \Lambda \quad (3.6)$$

$$\Phi_{IJ} = \phi_J(c_I) \quad (3.7)$$

$$\Pi = (\pi_1(c_I), \pi_2(c_I), \pi_3(c_I), \pi_4(c_I)) \quad (3.8)$$

$$\mathcal{W} = (w_1, \dots, w_N)^\top \quad (3.9)$$

$$\mathcal{Q} = (q_1, q_2, q_3, q_4)^\top \quad (3.10)$$

$$\mathcal{H} = (h_1, \dots, h_N)^\top \quad (3.11)$$

$$\Lambda = \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_N \end{bmatrix} \quad (3.12)$$

To determine the weights (\mathcal{W} and \mathcal{Q}) for the RBF-based SDF ($h(c_I)$), the linear system $\mathcal{A}\mathcal{X} = \mathcal{B}$ can be solved using singular value decomposition or LU factorisation [328].

$$\mathcal{A} = \begin{bmatrix} \hat{\Phi} & \Pi \\ \Pi^\top & 0 \end{bmatrix} \quad (3.13)$$

$$\mathcal{X} = \begin{bmatrix} \mathcal{W} \\ \mathcal{Q} \end{bmatrix} \quad (3.14)$$

$$\mathcal{B} = \begin{bmatrix} \mathcal{H} \\ 0 \end{bmatrix} \quad (3.15)$$

The RBF constraints (c_I) are defined by three levels with respect to the surface nodes (\mathcal{S}_i) and their normal vector directions (\vec{n}_i) and magnitude (β); consequently, $N = 3n$:

n - number of surface nodes

$$i = 1, \dots, n$$

$$c_I = (\mathcal{S}_1, \dots, \mathcal{S}_n, \mathcal{S}_1^+, \dots, \mathcal{S}_n^+, \mathcal{S}_1^-, \dots, \mathcal{S}_n^-) \quad (3.16)$$

where

$$\mathcal{S}_i^+ = \mathcal{S}_i + \beta \vec{n}_i \quad \forall i \in \{1, \dots, n\} \quad (3.17)$$

$$\mathcal{S}_i^- = \mathcal{S}_i - \beta \vec{n}_i \quad \forall i \in \{1, \dots, n\} \quad (3.18)$$

Example 3.1. Empirically determined distance (h_I) and regularisation (λ_I) values used to initialise the RBF-based SDF system:

$$h(\mathcal{S}_i) = 0 \quad \forall i \in \{1, \dots, n\} \quad (3.19)$$

$$h(\mathcal{S}_i^+) = -1 \quad \forall i \in \{1, \dots, n\} \quad (3.20)$$

$$h(\mathcal{S}_i^-) = +1 \quad \forall i \in \{1, \dots, n\} \quad (3.21)$$

Regularisation values:

$$\lambda_i^{\mathcal{S}} = 0.01 \quad \forall i \in \{1, \dots, n\} \quad (3.22)$$

$$\lambda_i^{\mathcal{S}^+} = 1 \quad \forall i \in \{n+1, \dots, 2n\} \quad (3.23)$$

$$\lambda_i^{\mathcal{S}^-} = 1 \quad \forall i \in \{2n+1, \dots, 3n\} \quad (3.24)$$

3.2.2 Minimisation of Signed Distance Function to Centreline

To compute the centreline approximations, the gradient of the SDF was minimised over an empirical number of iterations ($L = 10$) and a step parameter (α_s) of 0.1. During the minimisation process, the surface vertices are used as the initial guess and collapse into the structure of a centreline (**Figure 3.3**).

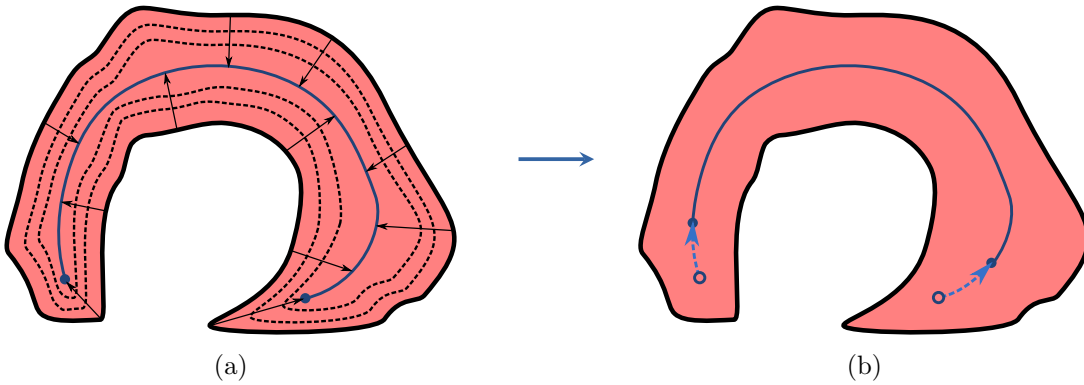


Figure 3.3: The minimisation of the signed distance function collapses a triangular surface mesh towards its centreline (a), with repeated iterations the centreline begins to contract (b)

Lemma 3.2. Calculate the gradient of an RBF-based SDF ($h(\mathcal{S}_i)$) with respect to the surface nodes ($\frac{\partial h(\mathcal{S}_i)}{\partial \mathcal{S}_i}$). Perform gradient descent to find the local maxima (\mathcal{V}_i^L).

n - number of surface nodes

N - number of constraints

$i = 1, \dots, n$

$J = 1, \dots, N$

$$\frac{\partial h(\mathcal{S}_i)}{\partial \mathcal{S}_i} = \sum_{J=1}^N w_J \frac{\partial \phi_J(\mathcal{S}_i)}{\partial \mathcal{S}_i} + \sum_{k=1}^4 q_k \frac{\partial \pi_k(\mathcal{S}_i)}{\partial \mathcal{S}_i} \quad (3.25)$$

where

$$\frac{\partial \phi_J(\mathcal{S}_i)}{\partial \mathcal{S}_i} = -(\mathcal{S}_i - c_J) \cdot (\|\mathcal{S}_i - c_J\|^2 + \epsilon^2)^{-1.5} \quad (3.26)$$

$$\frac{\partial \pi_k(\mathcal{S}_i)}{\partial \mathcal{S}_i} = \left(\frac{\partial \mathcal{S}_i^x}{\partial \mathcal{S}_i}, \frac{\partial \mathcal{S}_i^y}{\partial \mathcal{S}_i}, \frac{\partial \mathcal{S}_i^z}{\partial \mathcal{S}_i}, \frac{\partial (1)}{\partial \mathcal{S}_i} \right) \quad (3.27)$$

\therefore

$$\frac{\partial \pi_k(\mathcal{S}_i)}{\partial \mathcal{S}_i} = (1, 1, 1, 0) \quad (3.28)$$

and

$$\boxed{\frac{\partial h(\mathcal{S}_i)}{\partial \mathcal{S}_i} = \sum_{J=1}^N w_J \frac{\partial \phi_J(\mathcal{S}_i)}{\partial \mathcal{S}_i} + \sum_{k=1}^3 q_k} \quad (3.29)$$

The gradient descent is described by the following:

L - maximum number of iterations

$l = 0, \dots, L$

$$\mathcal{V}_i^{l+1} = \mathcal{V}_i^l + \alpha_s \frac{\partial h(\mathcal{V}_i^l)}{\partial \mathcal{V}_i^l} \quad (3.30)$$

α_s - step parameter

$$\mathcal{V}_i^0 = \mathcal{S}_i \quad (3.31)$$

\therefore

$$\boxed{\mathcal{V}_i^{l+1} = \mathcal{V}_i^l - \alpha_s \left(\sum_{J=1}^N \frac{w_J \cdot (\mathcal{V}_i^l - c_J)}{\sqrt{(\|\mathcal{V}_i^l - c_J\|^2 + \epsilon^2)^3}} - \sum_{k=1}^3 q_k \right)} \quad (3.32)$$

3.2.3 Efficiency Optimisation of Centreline Minimisation

Radial Basis Function Simplifications

The computations required to calculate a centreline approximation of a geometry can be quite expensive, depending on the number of vertices used to define a boundary. To solve the system of linear equations, the number of required floating point operations has been reported to be $N^3/6 + \mathcal{O}(N^2)$ [327]. Where, N in this case refers to the number of constraints, which is three times the number of vertices. Furthermore, to evaluate the system or in this case minimise a single vertex for one iteration has $\mathcal{O}(N)$ complexity [327]. Also, M iterations of a single vertex yields $\mathcal{O}(NM)$ complexity and for the complete centreline solution has a complexity of $\mathcal{O}(N^2M)$. Therefore, they are an expensive class of algorithms and despite offering great accuracy are often limited to cases not involving an excessive number of points [327].

Due to the high computational complexity of these algorithms, three optimisations were implemented to improve the efficiency of the algorithm and reduce the computation time as well as the required resources. The first optimisation was used to eliminate the requirement to solve the large system of linear equations. The second optimisation was to compute the solution in parallel on a GPU. The third optimisation was to simplify the meshes to reduce the number of vertices. It has previously been reported that despite the high accuracy offered by RBFs there use for solving problems with over 2,000 points is inappropriate [327]. As the meshes for meniscus geometries can exceed this value, the optimisations have been implemented to enable the use of RBFs to be more computationally appropriate.

Eliminating the need to solve the linear equations removes the $N^3/6 + \mathcal{O}(N^2)$ cost. Simplifying the geometries does not improve the $\mathcal{O}(N^2M)$ computational complexity but reduces the number of associated with N by reducing the total constraints. Evaluating the gradient of the RBF system in parallel and on a GPU does not improve computational complexity but does lead to more efficient and rapid processing of the solutions.

The solution of the linear equations is referred to as *the fitting process* and is used to determine the weights of the RBF functions. To remove the fitting process, two assumptions were made to estimate the weights of the function that would enable the derivation of centreline approximations. The first assumption was to disregard constraints outside a given radius. This is similar to efficiency improvements proposed using Fast Multipole Methods, that evaluate

near and far field points differently [327, 389]. This assumption exploits the knowledge that the infinite precision of RBFs are not always required [327]. However, for this algorithm the radius was chosen to be small enough to only include the constraint and no others. The second assumption was that the polynomial term was not required to approximate the centrelines and therefore was removed. A system RBFs can be augmented with a polynomial term, which is done to ensure orthogonality and that the system of linear equations derives a unique solution, if they are conditionally positive definite [14]. Additionally, the polynomial term can improve the accuracy of the RBF interpolation at boundaries [14]. Although this parameter can be used to improve convergence, they can sometimes yield inferior accuracies [15]. As the accuracy loss of interpolation at the boundaries is not a major concern, the polynomial term was removed. This could be considered a form of simplified RBF from the lack of a polynomial term and use of points within a given radius [390]. Additionally, these changes result in a RBF that no longer defines a SDF or implicit function but this was not important for the purposes of generating a centreline. For example, the distance of a point inside the boundary may not be negative (and vice versa), nor have a value relative to the distance from the boundary.

These two assumptions simplify the solutions of the weights to the either 0, $+\epsilon$ or $-\epsilon$, depending on whether the constraint was a surface vertex, positive expansion or negative expansion, respectively. Also, the application of these assumptions could perhaps be considered an unusual form of regularisation ($\tilde{\Lambda}$). Furthermore, another aspect that could be considered unusual is the fact the first assumption was only made for the solution of the weights and was not applied to the computation of the gradient vectors. Therefore, all constraints were used to minimise each individual vertex into a centreline approximation (*Equation 3.54*).

$$\tilde{\Lambda} = \begin{bmatrix} \Lambda & 0 \\ 0 & 0 \end{bmatrix} \quad (3.33)$$

The two assumptions can be applied through the element-wise multiplication (\odot) of the new regularisation matrix against the RBF matrix:

$$\left(\tilde{\Lambda} \odot \begin{bmatrix} \hat{\Phi} & \Pi \\ \Pi^\top & 0 \end{bmatrix} \right) \begin{bmatrix} \mathcal{W} \\ \mathcal{Q} \end{bmatrix} = \begin{bmatrix} \mathcal{H} \\ 0 \end{bmatrix} \quad (3.34)$$

\therefore

$$(\Lambda \hat{\Phi}) \mathcal{W} = \mathcal{H} \quad (3.35)$$

and

$$\begin{bmatrix} \Phi_{11}\lambda_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \Phi_{NN}\lambda_N \end{bmatrix} \begin{bmatrix} w_1 \\ \vdots \\ w_N \end{bmatrix} = \begin{bmatrix} h_1 \\ \vdots \\ h_N \end{bmatrix} \quad (3.36)$$

\therefore

$$\Phi_{II}\lambda_I w_I = h_I \quad (3.37)$$

As $I = J$,

$$\Phi_{II} = \phi_I(c_I) = (\|c_I - c_I\|^2 - \epsilon^2)^{-0.5} \quad (3.38)$$

$$\Phi_{II} = (\|0^2\| - \epsilon^2)^{-0.5} \quad (3.39)$$

$$\Phi_{II} = (0 - \epsilon^2)^{-0.5} \quad (3.40)$$

$$\Phi_{II} = (-\epsilon^2)^{-0.5} \quad (3.41)$$

$$\Phi_{II} = -\frac{1}{\epsilon} \quad (3.42)$$

\therefore

$$-\frac{\lambda_I w_I}{\epsilon} = h_I \quad (3.43)$$

$$w_I = -\frac{\epsilon h_I}{\lambda_I} \quad (3.44)$$

As $h_I = 0$ where $I \in [1, \frac{N}{3}]$. Therefore,

$$w_I = -\frac{0\epsilon}{\lambda_I} \quad (3.45)$$

$$w_I = 0 \quad (3.46)$$

As such, only values of I above $\frac{N}{3}$ require calculation due to the weights of the surface vertex constraints equalling zero (where $I \in [1, \frac{N}{3}]$). Therefore, where $I > \frac{N}{3}$:

$$w_I = -\frac{h_I \epsilon}{1} \quad (3.47)$$

$$w_I = -h_I \epsilon \quad (3.48)$$

Therefore, for constraints outside the boundary, where $I \in [\frac{N}{3} + 1, \frac{2N}{3}]$:

$$w_I = -(+1)\epsilon \quad (3.49)$$

$$w_I = -\epsilon \quad (3.50)$$

and for constraints inside the boundary, where $I \in [\frac{2N}{3} + 1, N]$:

$$w_I = -(-1)\epsilon \quad (3.51)$$

$$w_I = +\epsilon \quad (3.52)$$

$$(3.53)$$

With these modifications to the original algorithm, the skeleton or centreline of a geometry can be calculated by gradient descent using the following equation:

$$\mathcal{V}_i^{l+1} = \mathcal{V}_i^l - \alpha_s \left(\sum_{J=1}^N \frac{w_J \cdot (\mathcal{V}_i^l - c_J)}{\sqrt{(\|\mathcal{V}_i^l - c_J\|^2 + \epsilon^2)^3}} \right) \quad (3.54)$$

Mesh Simplification

A technique to reduce the computational cost of the RBF computation is called centre reduction [327]. Typically, RBF interpolation is conducted on a full set of data, however, with centre reduction, interpolation is performed on a subset. A common approach to centre reduction is to average points within a given radius. As this could alter and lose some key defining features of the geometry, mesh decimation was used to reduce the mesh vertices instead. Mesh decimation or simplification refers to the process by which the number of vertices of a mesh are reduced but are still an adequate representation of the geometry. This was performed to address the issue where evaluating a RBF or its derivatives are reported as inappropriate above 2,000 points. This is another aspect that exploits the fact that the infinite precision of RBFs is not always required, especially as the centreline will always be an approximation. The implemented mesh decimation algorithm is based on the maximisation of a cost function, which favors particular mesh characteristics. Each iteration until the desired number of vertices is acquired, the vertex with the lowest associated cost is removed. Vertices were removed using the edge collapse technique to ensure all the remaining vertices were based on original positions and not averaged, which could lead to significant geometry alterations. Edge collapse merges two vertices and associates the neighbors of the removed vertex with the one remaining (**Figure 3.4**). To minimise geometry losses, the cost function was designed to preserve regions with high Gaussian curvature and triangles with good aspect ratios. Therefore, the maximisation of the cost function selectively removed vertices that were on flat regions, where the aspect ratios of the resulting triangles would be improved. This results in a mesh with a smaller number of vertices that are more evenly distributed around the entire geometry, with the overall structure retained and the recession of geometric features minimised. The more evenly distributed vertices have been found to be more favorable for RBFs in general [391].

The cost function (\mathcal{C}) used in the implemented mesh decimation operator is the modified Gaus-

sian curvature (κ_G) multiplied by the aspect ratio (\mathcal{R}_a):

$$\mathcal{C} = \kappa_G \mathcal{R}_a \quad (3.55)$$

The modified Gaussian curvature was calculated based on the current angular defect of a particular vertex, also referred to as the discrete Gaussian curvature [8, 392] (**Figure 3.4a**). In contrast, the aspect ratio was calculated based on the new triangles that would be present following the vertex removal (**Figure 3.4b**). This is almost the inverse of a previously reported cost function which aimed to remove curvature and maintain high-aspect ratios [292]. However, here the objective was to preserve curvature and maintain high-aspect ratios.

The discrete Gaussian curvature (γ) can be defined using the Gauss-Bonnet scheme with respect to the internal angles (α_{ijk}) of the triangles in a one-ring neighbourhood surrounding a particular vertex [280, 392] (**Figure 3.4a**). The modified discrete Gaussian curvature (κ_G) corrects for densely concentrated regions and prevents small triangles reducing the influence of the angular defect [8, 292]. These variables were calculated with the following equations:

$$\gamma = 2\pi - \sum_{j \in \mathcal{N}(i)} \alpha_{ijk} \quad (3.56)$$

$$\kappa_G = \frac{3\gamma}{\sum_{j \in \mathcal{N}(i)} A_{ijk}} \quad \text{where } k = j + 1 \quad (3.57)$$

Where, $\mathcal{N}(i)$ is a set of vertex indices defining a one-ring neighbourhood around V_i , α_{ijk} is the adjacent angle opposite an edge (\mathbf{e}_{jk}) defined by V_j and V_{j+1} , and A_{ijk} is the area of a triangle defined by V_i , V_j and V_{j+1} .

The aspect ratio (\mathcal{R}_a) of a triangle is defined as the circumradius (r_c) divided by twice the inradius (r_i), which simplifies as the following:

$$\mathcal{R}_a = \frac{r_c}{2r_i} \quad (3.58)$$

where

$$r_c = \frac{abc}{\sqrt{(a+b+c)(b+c-a)(c+a-b)(a+b-c)}} \quad (3.59)$$

$$r_i = \frac{1}{2} \sqrt{\frac{(b+c-a)(c+a-b)(a+b-c)}{a+b+c}} \quad (3.60)$$

\therefore

$$\mathcal{R}_a = \frac{abc}{(b+c-a)(c+a-b)(a+b-c)} \quad (3.61)$$

Where, a , b and c are the three sides of a triangle.




Figure 3.4: Edge collapse of the mesh simplification algorithm from v_i to v_m . The cost function measures the (a) modified Gaussian curvature (**Equation 3.57**) and (b) aspect ratio (**Equation 3.61**) from the associated neighbourhoods before and after a potential edge collapse, respectively

3.2.4 Clustering

The clustering algorithm provides another key difference to how the centrelines were generated compared to the method reported in literature [328]. It was observed that during the minimisation of the RBF-SDF, the vertices collapse onto a centreline then move along the curve towards the local minima (**Figure 3.3b**). Therefore, instead of spending additional computational resources and time calculating the local minima that may lead to sub-optimal representations of centrelines, the algorithm would cease after a given number of iterations. This causes the vertices to not naturally cluster into few local minima but into many along the centreline curve. This allowed the author of the original centreline generation algorithm to re-use the RBF-SDF to find the centreline positions between local minima [328]. That was achieved by constraining points on a plane perpendicular to a linear segment that was defined between local minima (i.e. locate minima within that plane).

It was found during preliminary testing of the modified algorithm that the collapsed centreline

vertices were already adequate representations of a centrelines and required no further minimisation. As such, to derive a centreline curve defined as a sequence of linear segments, clustering and ordering with respect to adjacent neighbors was required (**Figure 3.5**). First, the clustering algorithm combines vertices of a collapsed centreline (**Figure 3.5a**) within a given radius (**Figure 3.5b**). During this step, the associated edges were combined leading to a curve defined by a series of lines where the vertices had multiple connections to the adjacent vertices (**Figure 3.5c**). There is the possibility that the edges defined between clustered vertices had connections outside the adjacent neighbors. Therefore, to find the true adjacent neighbors, the two edges with the shortest distances were associated as the adjacent vertices. In order to establish the end vertices of a centreline, each vertex was checked to determine if two adjacent neighbors were closer to each other than the vertex. This was used to indicate that one neighbor was on the other side of an adjacent neighbor and therefore not a true adjacent neighbor.

Following these steps, a series of unordered edge pairs were calculated between the vertices and associated adjacent neighbors; two for each non-end vertex and one for the end vertices. These edges were used to construct a mutable directed graph in VTK (using the *vtkMutableDirectedGraph* class). To derive the ordered path of vertices along the centreline, Dijkstra's shortest path algorithm was employed on the graph (using the *vtkDijkstraGraphGeodesicPath* class) [385]. The shortest path algorithm can be used to find the shortest path between two nodes of a graph. The shortest path between the end vertices would derive the ordered vertices of a centreline (**Figure 3.5d**).

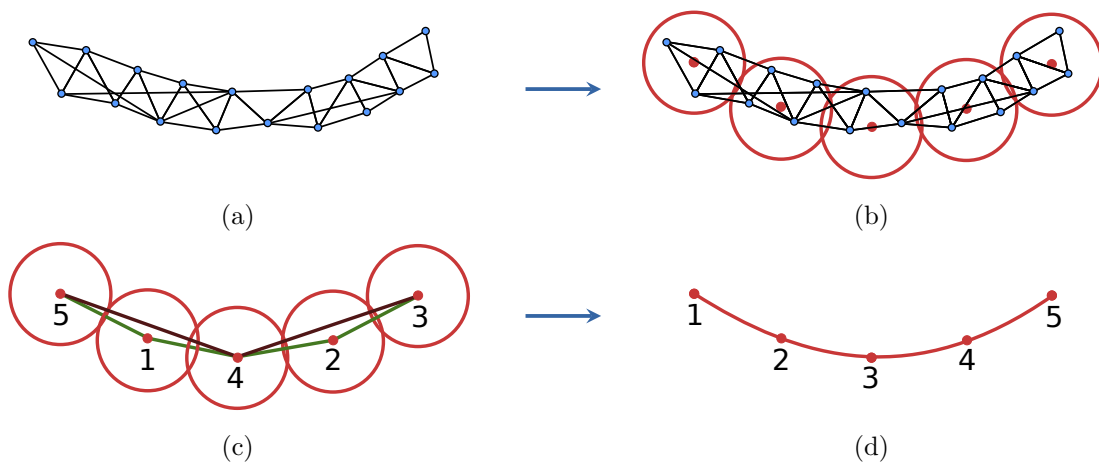


Figure 3.5: Following the collapse of a triangular surface mesh into a centreline (a) clustering is performed on the vertices which merges the associated edges (b). The shortest edges are selected to determine the adjacent clusters, which are unordered (c). Dijkstra's shortest path algorithm is then used to determine the correct order for B-spline parameterisation (d)

3.2.5 B-Spline Parameterisation

Following the clustering algorithm, a series of lines defining a curve is then parameterised using B-splines. The clustering process served two main functions: (1) reducing the collapsed triangular topology into a line mesh or curve and (2) providing an ordered sequence of vertices along that curve. The ordered structure of a centreline was important, as it was required to calculate a B-spline parameterisation (*Figure 3.5d*). The B-spline representation of the centrelines provided several benefits for the surface parameterisation process and other mesh-morphing steps. In particular, it allowed a direct one-to-one parameterisation from 0 at the start of the centrelines to 1 at the end. This allowed corresponding locations to be interpolated along a curve with greater accuracy. Also, the B-spline interpolates the curvature between the linear segments, allowing corresponding locations to be determined on a smoother and continuous representation of a centreline. Another benefit was that the parameterisation simplified the algorithm that was used to find the closest point on a curve to a particular vertex located on a template or target boundary. There were several advantages associated with the storage and computation of these structures (knots and coefficients), as thousands of centreline vertices could be generated rapidly and were not required to be stored nor require the re-computation of vertex interpolations between linear segment.

A B-spline can be re-parameterised to any required number of vertices, by specifying the discretisation values (U) between 0 and 1 ($0 \leq U \leq 1$). Additionally, a B-spline can be sampled or re-parameterised between any values within 0 and 1; e.g. 0.475 or between 0.45 and 0.55, respectively. To reduce the effects of centreline resolution on the projection algorithms, the closest position can be derived within a given tolerance. To achieve this, the closest point on the curve is first calculated from the initial discretisation ($N = 1000$). Then, the adjacent vertices of the curve are then chosen as the bounds of interest with this curve segment being parameterised further. This process then repeats, finding the closest point within the newly generated vertices until the difference between the current (C_p^i) and next (C_p^{i+1}) point are within a given tolerance (*Figure 3.6*). Combined with the continuous curve representation of the centreline, this allows the derived surface projection vectors to be situated in smoother, continuous and more unique locations with respect to each other.

The Dierckx fortran package was employed to calculate the B-spline representations of the centrelines [386]. This package is well-established and provides automatic knot selection from




Figure 3.6: Incremental process to find the closest point on a centreline to a surface vertex

the input of ordered vertices. The smoothness of the B-spline can be controlled with this package, as well as the ability to ensure all input vertices lie on the derived curve. The fortran package was wrapped in C++ for use within the mesh-morphing framework.

3.3 Feature Alignment

A novel centreline-based rigid registration algorithm was used to align the features of a target and template mesh. This following description of the algorithm is specific to crescent-shaped geometries, such as a meniscus. However, the basic principles could be applied to any geometry and associated centrelines.

Following the alignment, the target and template geometries were both aligned with respect to a reference position. This ensured both meshes had their centres located at the Cartesian origin. Additionally, the transverse plane of the geometries were aligned with respect to the x-y plane of the Cartesian coordinate system and the internal to external direction of each meniscus was aligned with the +Y direction (arbitrary choice). The lateral and medial meniscus geometries were both aligned into the same reference system. This imposed several consequences to the orientation of the geometries with respect to the anatomical and Cartesian axes. As such, the horn-attachments always had positive y-coordinates. Additionally, the anterior and posterior horns of the lateral meniscus geometries always had positive and negative x-coordinates, respectively. Due to the same alignment system, the medial meniscus geometries had the opposite, where the x-coordinates were positive for the posterior horn and negative for the anterior horn. Consequently, the superior surface would face towards the +Z direction and inferior surface

would face towards the $-Z$ direction.

The geometries were aligned with respect to three locations, (1) the centres of their geometry and (2-3) horn-attachments. An unusual and challenging aspect of the meniscus geometries is due to the centroid lying outside the boundary. Typically, the centroid is calculated by averaging all the vertices defining the boundary (*Equation 2.7*). For most geometries, this point is located within the boundary and provides a reasonable approximation to the centre of the geometry. Instead of using the centroid calculation to define centre of the geometries, the mid-point of the centreline ($U = 0.5$) was chosen. This was a more suitable representative as it defined the central location along the central curve of the geometries and was guaranteed to be inside the meniscus boundaries. Consequently, the end-points of the centrelines were used to define the centres of the horn-attachments ($U = 0$ and $U = 1$).

The three central locations formed an alignment-triangle which was used to orientate the geometry onto a reference triangle. The reference triangle was an isosceles triangle constructed from the Cartesian origin and two points in $+x$ and $-x$ quadrants: $(0, 0, 0)$, $(+1, +1, 0)$ and $(-1, +1, 0)$. A variant of the Kabsch algorithm was used to compute the affine transformation of the geometries to the reference position [393]. The Kabsch algorithm can be used for rigid registration between two point sets and derives the optimal rotation and translation to align the point sets using a least squares approach. The implementation here uses single value decomposition (SVD) to calculate the least squares solution [394]. The Kabsch algorithm is closely related to the solution of the orthogonal Procrustes and Wahba's problem [395, 396]. It has been reported that a quaternion algorithm could be 50% faster for 3 points [394]. However, the process is already very rapid and allows the capacity to be scaled up to additional registration points, if were required.

This algorithm was used to derive the rotation matrix that would not only orientate the alignment-triangle onto the reference triangle but the entire geometry into this reference system. The rigid registration of the template and target geometries would be completed following the transformation into the reference system. The transformation was applied to the geometry by simply computing the matrix multiplication of the mesh vertices against the rotation matrix.

The first step was to centre all of the geometries with respect the origin using the mid-point of their centerlines (*Equation 2.8*). This would require only the optimal rotation (\mathcal{R}) to be derived to align the geometries, following this initial translation. Using the reference triangle

(\hat{T}) and derived alignment-triangles (T) as point sets, the optimal rotations would be calculated using **Equations 3.62-3.66**. First, the outer products of the corresponding triangle points to be aligned were summed, which derived a 3×3 matrix (E) (**Equation 3.62**). Second, SVD was used to extract the U , Σ and V^\top matrices (**Equation 3.63**). Third, to calculate the optimal rotation matrix for most cases, V and U^\top are multiplied. However, there is the possibility in some rare cases that this may result in a reflection instead of an optimal rotation. This can happen if either of the point sets are coplanar, which is the case given that three points define a plane. A reflection can be identified by calculating the determinant of the rotation matrix, where a result of +1 indicates the optimal rotation and -1 indicates a reflection [394]. There is a straightforward solution to derive the optimal rotation if the determinant of the rotation matrix is -1. The V matrix must be adjusted by flipping the sign of the values in the third column [394]. Once adjusted, the original formula can be used to calculate the optimal rotation. As the determinant of the rotation matrix will be either +1 or -1, this can be incorporated into the final equation to flip the signs of the third column if encountered (**Equations 3.64-3.66**).

$$E = \sum_{i=1}^3 \hat{T}_i \otimes T_i \quad (3.62)$$

$$E = U \Sigma V^\top \quad (3.63)$$

$$D = \det V U^\top \quad (3.64)$$

$$S = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & D \end{bmatrix} \quad (3.65)$$

$$\mathcal{R} = V S U^\top \quad (3.66)$$

Where, \otimes defines the outer product. The point sets of the reference triangle (\hat{T}) and an alignment-triangle (T) are defined as:

$$\hat{T} = \left\{ \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} +1 \\ +1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ +1 \\ 0 \end{bmatrix} \right\} \quad (3.67)$$

$$T = \left\{ \begin{bmatrix} x_1 \\ y_1 \\ z_1 \end{bmatrix}, \begin{bmatrix} x_2 \\ y_2 \\ z_2 \end{bmatrix}, \begin{bmatrix} x_3 \\ y_3 \\ z_3 \end{bmatrix} \right\} \quad (3.68)$$

Once the rotation matrix is derived they are applied to the centred vertices (V_c) to derive the aligned vertices (V_a) using following equation:

$$V_a = \mathcal{R}V_c \quad (3.69)$$

where

$$V_c = \begin{bmatrix} x_1 & \dots & x_n \\ y_1 & \dots & y_n \\ z_1 & \dots & z_n \end{bmatrix} \quad (3.70)$$

3.4 Surface Parameterisation

The surface parameterisation phase is composed of two steps: target surface expansion and centreline-based surface projection. The steps occur following the generation of centrelines and geometry alignment step, which improves the degree of one-to-one correspondences between the target and template geometries. The centrelines of the target and template meshes are used to generate projection vectors from the template onto the target. Following the centreline-based projection algorithm, a template's surface vertices are located on the boundary of a target geometry. The centreline-based projection ensures all the surface vertices of the template lie on the expanded target boundary.

3.4.1 Centreline-based Projection

The purpose of the centreline-based projection steps are to project the surface vertices of the template mesh to corresponding locations on the target boundary. First, the template projection vectors are calculated from the vectors defined between each surface vertex and the closest position on a centreline to that vertex (**Figure 3.6**). To calculate the target projection vectors, the corresponding locations on the target's B-spline are determined for each template projection vector and translated to these locations. The target projection vectors are then projected onto the surface of the target geometry (**Figure 3.7**).

Surface Detection

A ray-projection algorithm was used to calculate the projected locations of the template's surface vertices on the target geometry. Specifically, the surface of the target geometries are detected using the Möller-Trumbore ray-triangle intersection algorithm [387]. There are a va-

riety of intersection algorithms that can be used to detect the intersection of several geometric primitives. However, this type of geometry intersection algorithm was the most suitable for two reasons. First and foremost, the main benefit was due to the target geometries being composed of genus-0 (water-tight) triangular meshes. The second reason is due to the known robustness and efficiency of this type of well-tested and established algorithm. Due to the wide-spread utilisation of ray-casting used in the field of graphics rendering, there are a wide availability of different algorithms that could be later evaluated for potential performance and efficiency improvements [363, 397]. The Möller-Trumbore intersection algorithm was chosen due to being well-established and considered a standard, particularly for comparison purposes of new algorithms [363, 397]. By specifying appropriate tolerance values, this algorithm can guarantee all ray vectors intersect with the triangular mesh. Inappropriate selection can induce rare anomalies, where ray intersections are failed to be detected along particular ray-edge (triangle edges) combinations, what the graphics industry have coined ‘fireflies’ [398]. There are examples of algorithms that exist which can guarantee the prevention of ‘fireflies’ [397]. However, they are more computationally expensive and currently have not been found to be necessary from the preliminary results.

The effect of the tolerance parameter is to slightly enlarge the triangles, to ensure the *watertightness* of the ray-projection algorithm. A tolerance of zero maintains the original size, therefore, it is common practice to select a very low number [397]. The value chosen in this implementation was 10^{-7} . Due to this slight enlargement and the potential of being detected multiple times if passing through highly irregular regions, the closest intersection to the centreline origin was selected.

To perform the ray-projection of a single vector onto a triangular mesh, all triangles (N_t) must be evaluated with respect to that ray, to determine if an intersection occurs between them - requiring N_t calculations (**Figure 3.7**). As such, to project the surface vertices (N_v) of a template geometry onto a triangular mesh requires $N_v \times N_t$ calculations. Therefore, the cost of computing the projection of the template vertices onto a target boundary imposes a complexity of $\mathcal{O}(N_t N_v)$. The cost to project a single surface vertex of the template onto a target is $\mathcal{O}(N_t)$. Consequently, the ray-projection algorithm is very computationally expensive. However, as determining the potential intersection of a given ray-triangle combination can be computed in isolation, without the need to process other triangles, this type of algorithm is a

perfect candidate for parallel GPU computations. Modern day GPU's have hundreds of cores with more recent high-end products having orders of magnitude higher (tens of thousands). Following the GPU processing, data aggregation at the end allows the determination of the intersection locations for each template vertex.

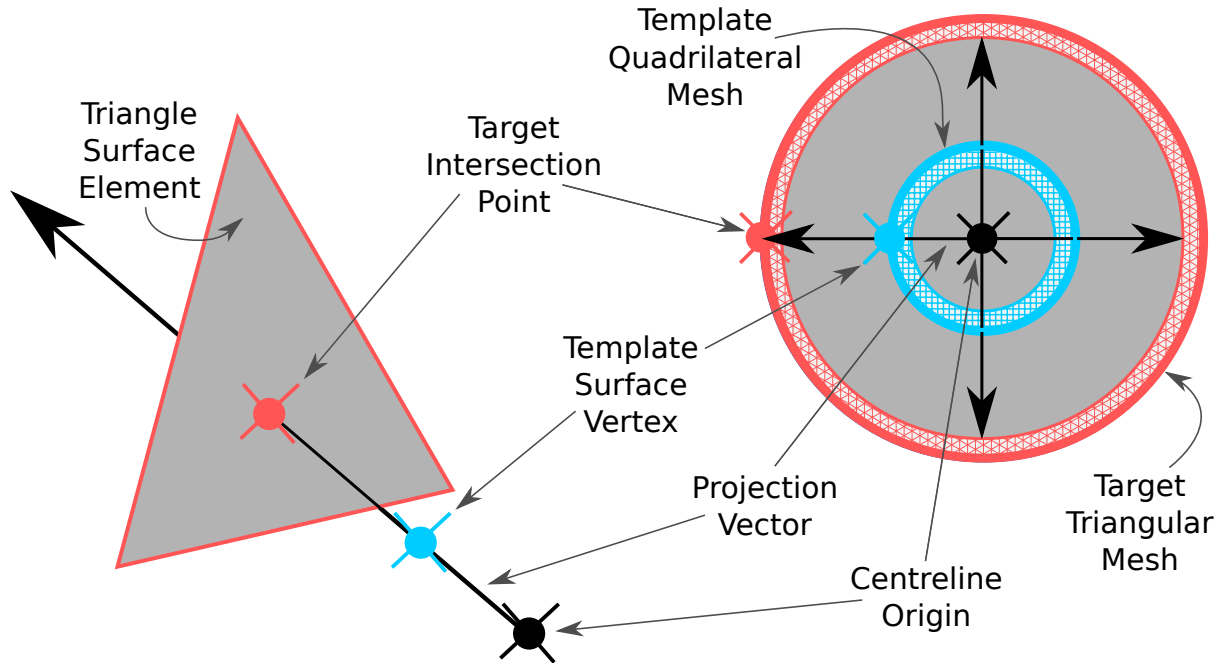


Figure 3.7: Ray-triangle intersection was used to project each template surface vertex onto the target

Ray Casting

Additionally, the ray-triangle intersection algorithm is a fundamental, multi-purpose operation. Consequently, the same ray-triangle intersection algorithm was used to implement a ray-casting algorithm, which was used during the surface optimisation step. The ray-casting algorithm can be used to determine if a point is inside or outside a geometry using the even-odd rule [399,400]. If a point inside a geometry is projected to a location outside the boundary, the number of times it will cross that boundary will be odd. In contrast, if the point is on the outside of the geometry, the number of boundary intersections will be even. To implement the ray-casting algorithm, a point is projected to infinity, outside the geometry. Each surface intersection is calculated using the algorithm outlined above. Instead of locating the closest intersection for a point of interest, the difference lies simply in quantifying the number of intersections. Potential duplicates from overlapping triangles are resolved by specifying a tolerance factor, that assumes multiple points that are very close refer to the same intersection.

3.5 Surface Optimisation

The purpose of the surface optimisation step is to unfold the quadrilateral mesh from the surface parameterisation step and improve the shape of the quadrilateral elements. This aimed to improve the quality and shape of the hexahedral elements during the volumetric transformation step, which is based on the positions of the quadrilateral mesh. The quality of the quadrilateral elements have a direct effect on the hexahedral elements. Essentially, this step aimed to improve the morphed hexahedral elements indirectly by providing a well-shaped quadrilateral mesh that would be used to define the boundary conditions in the volumetric transformation step.

The surface parameterisation step was performed onto a slightly expanded surface which served to provide an upper bound to the optimisation step. The optimisation loop exploits the shrinking effect of Laplacian smoothing to collapse the slightly expanded surface onto the target geometry, the lower bound. Additionally, the Laplacian smoothing of the quadrilateral would also improve the shape of the quadrilateral elements. The idea was to provide an upper and lower boundary between which the optimised quadrilateral mesh could unfold and improve whilst retaining the shape of the target geometry. The presence of any folded or inverted quadrilateral elements would lead to inverted hexahedral elements in the volumetric transformation step. Therefore, it was necessary to remove them to ensure the morphed hexahedral mesh retained an ideal proportion of the original template mesh quality. Additionally, distortion could be imparted from the surface parameterisation phase. Therefore, improvements to the quadrilateral element shapes would also aid in retaining some of the original hexahedral element qualities.

The use of upper and lower bounds has previously been used to control an unconstrained energy minimisation process for spherical parameterisation [295]. This technique allows for the vertex repositioning and movement within the upper and lower bounds without the need to constrain vertices, whilst preventing their collapse from Laplacian shrinking effects. This type of approach has previously been shown to unfold highly inverted triangular meshes of irregular geometries projected onto a sphere. The algorithms here have adapted the process to operate over highly irregular geometries directly instead of a spherical domain and unfold a quadrilateral mesh instead of a triangular mesh.

3.5.1 Laplacian Smoothing

The Laplacian smoothing operator is used to improve the shape of elements in a given mesh. This is achieved by moving the vertices defining those elements into more appropriate and optimal locations. Generally, this is achieved by averaging the positions of the neighboring vertices **Figure 2.17**. A relaxation factor can be incorporated to control the rate of smoothing with respect to the original location. The Laplacian operator can be used for any mesh type and in the developed mesh-morphing strategies the same function has been used for both quadrilateral and hexahedral meshes using **Equation 3.71**. In addition to the mesh vertices (V_i) and a set of the associated neighbors ($\mathcal{N}(i)$), this function requires two parameters: number of iterations (I_{max}) and relaxation factor (α_r).

$$V_i^{I+1} = V_i^I + \frac{\alpha_r}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} V_j^I \quad (3.71)$$

Where, I is the current iteration, i is the index of the vertex to be smoothed ($i \in \{1, \dots, N_v\}$ with N_v specifying the number of vertices) and j defining the index of an adjacent vertex. Also, $|\cdot|$ refers to the cardinality of a set, here representing the total number of neighbors associated with a vertex. This is a variant of the Laplacian smoothing operator aimed to minimise global shrinking of a mesh's geometry [30, 351].

3.5.2 Quadrilateral Inversion Detection

It was important to detect the presence of inverted quadrilateral elements to determine the correct parameters that would lead to the gradual removal of inversions each iteration. If the parameters were chosen incorrectly, too little smoothing would be performed and element inversion would not be removed. In contrast, if the parameters caused excessive smoothing then the quadrilateral mesh would collapse immediately onto the target boundary, preventing further smoothing and the removal of inverted elements. The quadrilateral inversions were detected using two algorithms. (1) The positive (external) normals for each element were calculated, then (2) the normals were ray-casted (**Section 3.4.1**) to determine if they were inside or outside the target boundary. If they were inside the boundary, this would indicate the presence of an inverted element.

The Jacobian metric could not be used to determine if a quadrilateral element was inverted as the

element could be valid but face the wrong direction. Therefore, the following algorithm was designed to be able to detect where folding and inversions occurred. First, the element normals were calculated by converting the quadrilateral mesh to a triangular mesh by subdividing each quadrilateral into two triangles, whilst maintaining the original winding from the quadrilateral elements. This process results in a triangular mesh with the same number of vertices and twice the number of elements than the quadrilateral mesh. The vertex normals were first calculated by exploiting the winding of the triangles using **Equation 3.75**. All quadrilateral and derived triangle elements had a clock-wise winding. Therefore, the calculation of their normals would lead to a vector pointing away from the surface. To calculate the normals of a quadrilateral element, the normals from each vertex defining that element was averaged. The magnitude of the normals were then adjusted to a small length specified by parameter ζ (**Equation 3.76**). The reason for this was to prevent large normal vectors penetrating other regions of the boundary and causing inaccurate results. The vector directions were then inverted and ray-casted to determine if they were inside or outside the boundary, which would identify if an element was pointing inwards or outwards. As all the normals should point outwards, a normal pointing inwards would declare the element as inverted.

$$\vec{v}_0 = V_0^T - V_1^T \quad (3.72)$$

$$\vec{v}_1 = V_0^T - V_2^T \quad (3.73)$$

$$\vec{n} = \vec{v}_0 \times \vec{v}_1 \quad (3.74)$$

$$\hat{n} = \frac{\vec{v}_0}{\|\vec{v}_0\|} \quad (3.75)$$

$$\vec{n}_\zeta = \zeta \hat{n} \quad (3.76)$$

Where, V_0^T , V_1^T and V_2^T are the vertices of a triangle, \vec{v}_0 and \vec{v}_1 are the vectors defined by these vertices, \vec{n} is the normal of the triangle, \hat{n} is the unit normal of the triangle, and \vec{n}_ζ is the normal of the triangle with a magnitude ζ and $\|\cdot\|$ is the L^1 -norm operator.

3.5.3 Quadrilateral Unfolding

The quadrilateral unfolding algorithm consists of an iteration loop with three key steps. First, the quadrilateral mesh from the expanded surface parameterisation would be smoothed using the Laplacian operator (**Equation 3.71**). Second, a ray-casting algorithm was used to determine

the presence of vertices inside the boundary of the original target geometry. Third, the vertices within the target boundary would be repositioned back to their previous location, outside the target boundary. This would result in a new quadrilateral mesh with some of the element inversions removed and their shapes being improved. This process would then repeat using the new quadrilateral mesh as the starting point and continue to remove element inversions and improve the overall element shapes, whilst preventing the quadrilateral mesh shrinking into the target geometry. At the end of each iteration, the number of inversions, geometry intersections and norm between the previous and current vertex positions were calculated. This was initially used to determine if the unfolding process was working and what number of iterations were typically required to reach convergence. In future development cycles, these variable could be used for terminating or tuning the process. Instead of stopping the unfolding process when all inversions were removed, the number of iterations were enforced, as this would lead to further improvements of element shapes. Additionally, this enforcement would prevent excessive iterations where convergence was slow, as inversions could be alleviated with an unconstrained smoothing step following this process. The key goal of the iterative loop was to remove a large proportion of quadrilateral inversions. After completing a specified number of iterations, the quadrilateral mesh would then undergo another smoothing operation to ensure all elements inversions were removed and the presence of element distortions were minimised. The final smoothing operation was performed with a slightly higher number of iterations and larger relaxation constant than what was used during the unfolding process.

3.6 Volumetric Transformation

Once the quadrilateral surface of the template mesh had been optimised, the internal vertices and elements could be morphed using volumetric transformation. The volumetric transformation step is used to evenly distribute and uniformly interpolate the hexahedral elements into the boundary prescribed by the morphed quadrilateral surface (*Figure 3.8*). This step was broken down into two algorithms: linear elastic transformation and Laplacian smoothing. The aim of these algorithms were to position the non-surface vertices within the morphed boundary that will yield the highest hexahedral element quality. Following this step, the hexahedral template will have been morphed into a mesh resembling the target geometry.

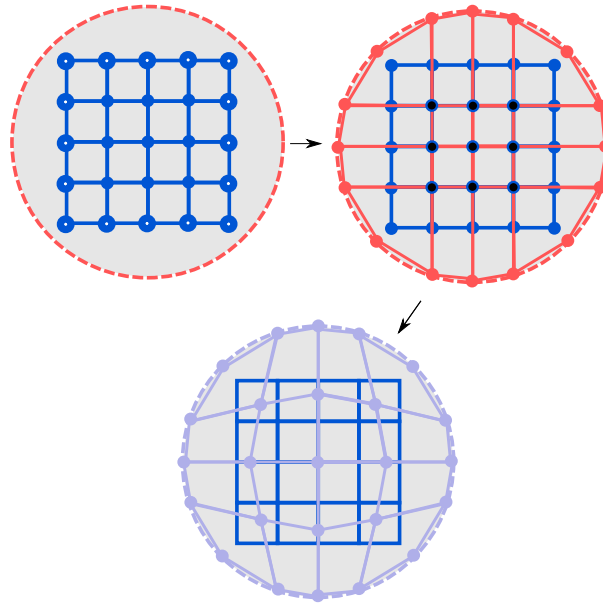


Figure 3.8: Overview of volumetric transformation: (a) template mesh (blue) and target geometry (red), (b) boundary morphing of template and (c) volumetric transformation to derive a morphed mesh (purple)

3.6.1 Linear Elastic Transformation

Linear elastic transformation can be initiated using two approaches, using either Dirichlet or Neumann boundary conditions (*Section 2.5.3*). Dirichlet boundary conditions use fixed displacements, whereas Neumann boundary conditions consists of applying force vectors. The Dirichlet boundary conditions were selected as it would maintain the original positions of the optimised boundary and can be completed in one step. Neumann boundary conditions provide less control over the final boundary, however, gives the solution greater degrees of freedom with respect to complete vertex movement.

To calculate the fixed displacements for the Dirichlet boundary condition, the vertices of the optimised quadrilateral mesh were subtracted from their corresponding positions in the original template. These displacements were then applied to the original hexahedral template in one step. Lamé parameters were used to define the material properties, which were empirically determined ($\lambda = \mu = 50$). The solution to this finite element problem was a morphed hexahedral mesh, where the boundary matches the optimised quadrilateral surface.

The Modular Finite Element library (MFEM) was used to implement this algorithm [388].

3.6.2 Laplacian Smoothing

Despite the optimised quadrilateral mesh containing no inverted elements, this does not guarantee that the morphed hexahedral mesh will not contain any element inversions. Additionally, some of the fixed displacements from the quadrilateral mesh may impart some distortion to some of the hexahedral elements, particularly at singularity locations. To remove the presence of any element inversions and optimise the element shapes globally, Laplacian smoothing was performed on the morphed mesh.

3.7 Mesh-Morphing Strategy

The first novel design was associated with the use of two centrelines, for a target and template mesh. Specifically, the novel contribution was the cross-projection of centreline-based projection vectors. This process transferred a projection vector from a template's centreline to the corresponding location on a target's centreline. The transformed vector was then projected onto the target, which was initially defined between a template's surface vertex and its closest centreline position (*Figure 3.7*). This was implemented to overcome the large geometric differences between target and template meshes, where one centreline would have created excessive distortions and correspondence errors.

Previously, one centreline had been used for both target and template meshes [45], similar to what was implemented for the optimised strategy. However, use of one centreline is only acceptable where target and template meshes are reasonably similar with a strong degree of overlap. This was not the case for the unoptimised strategy, as meniscus anatomies were highly irregular and a standard CAD-based template with one centreline would not accommodate their large geometry variations. An alternative approach could have been to use a statistical average [401] or anatomically derived template [241]. However, this would not have encompassed all geometries, but statistically could have been more similar to a larger proportion of the general population [401]. Additionally, this would have incurred several challenges, such as:

1. Developing two high-quality SSMs of medial and lateral menisci, which represent a significant proportion of the population
2. Generating a high-quality hexahedral template mesh of an irregular statistical average
3. The potential to require more than one template for target geometries too dissimilar from

a statistical average

In contrast, the effort of extracting a second centreline (for a template mesh) is trivial using the developed algorithm. Furthermore, an approximation of a template's centreline could have been achieved through CAD-based approaches by taking the average point from each swept slice. One clear advantage of using a SSM is that the correspondences for target geometries, represented by the parameter's subspace, could be calculated rapidly and accurately [58, 401]. However, this would not guarantee a valid mesh would be derived following a volumetric transformation. Furthermore, a target geometry which could not be defined by a parameter's subspace could not be morphed with that technique. One promising study tested their method against femur (N=29) and aorta anatomies (N=4), which reported at least three valid meshes for the femora [401].

The second novel aspect was the use of B-splines to parameterise and represent the centrelines, which served three purposes:

1. Represent the centrelines as smooth continuous structures
2. Ease the identification of corresponding locations between two centrelines
3. Ease the determination of more unique centreline (closest-point) positions relative to a template's vertices

The first point negated the need to repeatedly discretise and calculate distances along coarse centreline segments. This would have made the algorithms for points two and three more tedious and time-consuming. The second point allowed the rapid identification of corresponding positions, based on the value of the U parameterisation variable. The third point was important to minimise stepping-effects, where a coarsely discretised centreline would lead to large correspondence errors in the initial projection. Also, the stepping-effect was observed when a coarse discretisation was used for surface error measurements. Both of these issues were resolved with increased refinement.

The third novel aspect was the unfolding algorithm (surface optimisation) which uniquely combined aspects of mesh-wrapping and spherical parameterisation. This combination of techniques eliminated the correspondence errors from the initial projection and improved element shapes prior to volumetric transformation. The majority of biomedical mesh-morphing strategies reviewed (**Table 2.8**) consisted of three key components, particularly surface parameterisation

and volumetric transformation. Less than half of these strategies attempted to improve the shape of elements defining the surface parameterisation prior to volumetric transformation. The other strategies relied solely on the volumetric transformation and post-processing operations to resolve these deficiencies. In part, this could be due to the higher relative similarities between their target and templates, which were all anatomically derived. In general, strategies which did incorporate some form of surface optimisation were morphing to anatomies with large variations and complexity [47, 52]. The mesh-morphing strategy is summarised by the **Algorithm 3.1**, where the number of iterations (N_I) performed by the surface optimisation was 15.

Algorithm 3.1: Mesh-Morphing Strategy

Input: Template Mesh (\mathcal{H}_0) and Centreline ($\mathcal{C}_\mathcal{H}$), Target Mesh (\mathcal{T}) and Centreline ($\mathcal{C}_\mathcal{T}$)

Output: Morphed Mesh (\mathcal{H}_1)

```

1  $\mathcal{T} \leftarrow \text{ExpandTarget}(\mathcal{T});$ 
2  $\mathcal{B}_0 \leftarrow \text{ProjectTemplateVertices}(\mathcal{H}_0, \mathcal{T}, \mathcal{C}_\mathcal{H}, \mathcal{C}_\mathcal{T});$ 
3 for  $i \leftarrow 1$  to  $N_I$  do
4    $\mathcal{B}_0 \leftarrow \text{LaplacianSmooth}(\mathcal{B}_0);$ 
5    $\mathcal{V}_I \leftarrow \text{ComputeInternalVertices}(\mathcal{B}_0, \mathcal{T});$ 
6    $\mathcal{B}_0 \leftarrow \text{PlaceInternalVerticesOutside}(\mathcal{B}_0, \mathcal{V}_I);$ 
7  $\mathcal{H}_1 \leftarrow \text{LinearElasticTransform}(\mathcal{B}_0, \mathcal{H}_0);$ 
8  $\mathcal{H}_1 \leftarrow \text{LaplacianSmooth}(\mathcal{H}_1);$ 

```

3.8 Mesh-Morphing Strategy Optimisation

Justification for Optimisation

In general, the overarching goal of mesh-morphing strategies are to strike an acceptable balance between minimising two quadratic energies: surface (\mathbf{E}_s) and volume (\mathbf{E}_v):

$$\mathbf{E} = \underbrace{\left[\sum_{i=1}^{N_v^\Omega} \epsilon_i^2 \right]}_{\mathbf{E}_s} + \underbrace{\left[\sum_{i=1}^{N_e} (1 - \tilde{J}_i)^2 \right]}_{\mathbf{E}_v} \quad (3.77)$$

Where, \mathbf{E} is the morphing energy, ϵ_i is the surface error per vertex, N_v^Ω is the number of boundary vertices, \tilde{J}_i is the scaled Jacobian per element and N_e is the number of elements.

Minimising \mathbf{E} is the objective function of all mesh-morphing strategies. This could be solved either directly (using an energy minimiser, e.g. [279, 295, 402]) or indirectly, which is the case

for all the reviewed strategies (**Table 2.8**). Additionally, this could be achieved through either the dynamic assignment of boundary constraints (i.e. will be changed during the operation):

$$\mathcal{V}^* = \arg \min_{\mathcal{V}} \mathbf{E}(\mathcal{V}) \quad (3.78)$$

or, the fixed assignment of boundary constraints (i.e. won't be changed during the operation):

$$\begin{aligned} \mathcal{V}^* &= \arg \min_{\mathcal{V}} \mathbf{E}_{\mathbf{v}}(\mathcal{V}) \\ \text{s.t. } V_i^\Omega &= V_i^{*\Omega} \quad \forall i \in 1, \dots, N_v^\Omega \end{aligned} \quad (3.79)$$

Where, \mathcal{V}^* is a set of optimal vertices representing a morphed mesh, \mathcal{V} is a set of vertices representing a template mesh, V_i^Ω is a vertex on a morphed mesh's boundary and $V_i^{*\Omega}$ is the optimal position (or projection) of a template's vertex on a target's boundary.

Reducing the morphing energy to zero would yield a perfect surface fit with no element distortion, which for most geometries is impossible. Also, any acceptable distortion metric, criteria or combination could be used to define a volume energy. However, the scaled-Jacobian has been found to be the most suitable (**Section 2.3.5**).

Following the analyses from **Chapter 8**, it became apparent that the unoptimised strategy discriminated the performance in favour of mesh-quality over surface error. The reasons for this are evident following the analysis of the unoptimised strategy's morphing energy (\mathbf{E}^u):

$$\mathbf{E}^u \approx \underbrace{\sum_{i=1}^{N_v^\Omega} \sum_{j \in \mathcal{N}_i^\Omega} \alpha_{ij}^\Omega \|V_i^\Omega - V_j^\Omega\|^2}_{\mathbf{E}_s^u \text{ (contractive)}} + \underbrace{\sum_{i=1}^{N_v} \sum_{j \in \mathcal{N}_i} \alpha_{ij} \|V_i - V_j\|^2}_{\mathbf{E}_v^u \text{ (contractive)}} \quad (3.80)$$

$$\underbrace{\text{s.t. } \mathcal{M} \cap \mathcal{T} = \emptyset}_{\text{(expansive)}} \quad (3.81)$$

where

$$\alpha_{ij}^\Omega = \frac{1}{2|\mathcal{N}_i^\Omega|} \quad (3.82)$$

and

$$\alpha_{ij} = \frac{1}{2|\mathcal{N}_i|} \quad (3.83)$$

Where, α_{ij}^Ω is a set of parameterisation coefficients for the boundary elements, \mathcal{N}_i^Ω is set of the vertex indices defining a one-ring neighborhood (quadrilateral) around a boundary vertex, V_j^Ω is a one-ring neighbour of a boundary vertex, \mathcal{N}_i is set of the vertex indices defining a one-ring neighborhood (hexahedral) around a vertex, α_{ij} is a set of parameterisation coefficients for the volumetric elements, V_i is a morphed mesh vertex, V_j is the one-ring neighbour of a vertex, \mathcal{M} represents a morphed mesh, \mathcal{T} represents a target mesh, \cap represents the set theoretic intersection operator, \emptyset represents an empty set and $|\cdot|$ represents the cardinal operator.

The surface (\mathbf{E}_s^u) and volume (\mathbf{E}_v^u) energies consisted of boundary and volumetric Laplacian smoothing components, respectively. These applied contractive forces to the surface and volume of the morphed meshes. Also, this explains the high smoothness and mesh-quality associated with these meshes. The constraint, or energy barrier, preventing boundary intersection (*Equation 3.81*) provided the only means to oppose the contractive energies and reduce surface error. Consequently, \mathbf{E}^u failed to accurately address the surface error component (\mathbf{E}_s) of the general energy function \mathbf{E} (*Equation 3.77*), which explains the large surface errors. Also, it should be noted that these energies were solved in two stages separately and indirectly (not using traditional minimisers).

This led to the development of a second novel mesh-morphing strategy, which aimed to optimise the meshes derived from the first mesh-morphing strategy with respect to surface error. The surface and volume of the meshes were morphed together iteratively. This was different to both relative and absolute morphing, where a template is deformed to match a target boundary in several steps or just one, respectively [12]. Essentially, the minimisation approach aimed to determine the optimal surface correspondences ($V_i^{*\Omega}$) whilst minimising distortion. Based on this definition, some of the reviewed strategies could also be defined as minimisation-based approaches [46,47]. The optimised morphing energy (\mathbf{E}^o) reflected this difference and correctly addressed the surface energy component (\mathbf{E}_s^o), whereas the volume energy remained the same ($\mathbf{E}_v^o = \mathbf{E}_v^u$):

$$\mathbf{E}^o \approx \underbrace{\sum_{i=1}^{N_v^\Omega} \|V_i^\Omega - V_i^{*\Omega}\|^2}_{\mathbf{E}_s^o \text{ (expansive)}} + \underbrace{\sum_{i=1}^{N_v} \sum_{j \in R_i} \alpha_{ij} \|V_i - V_j\|^2}_{\mathbf{E}_v^o \text{ (contractive)}} \quad (3.84)$$

These minimisation approaches to mesh-morphing, employed by the optimised and unoptimised strategies, took principles from traditional surface parameterisation techniques. To the author's knowledge, these were the first time they have been applied to 3D elements instead of 2D elements in this manner. However, the spherical parameterisation techniques employed for the optimised strategy differed from those used in the unoptimised strategy. A variation of the more common origin projection techniques was utilised [8,282,291,293,298] instead of employing energy bounds [295]. In general, origin projection techniques involve iteratively projecting and smoothing mesh vertices to achieve unfolded parameterisations.

The surface energy (\mathbf{E}_s^o) now provided an expansive force and could be measured by the magnitude of the surface error. Additionally, this force naturally counteracted the contractive force from the volume energy (\mathbf{E}_v^o). The expansive force was implemented using centreline-based projection, which was applied absolutely each iteration. However, it could have been applied relatively with a control parameter. Essentially, the surface error was corrected iteratively, at the expense of increased element distortion. This matches widespread reports that indicate a trade-off between surface error and mesh-quality [45, 53, 58, 307, 403]. The contractive volume energy was again measured by the degree of element distortion and applied using the Laplacian smoothing operator. This shrunk the entire mesh, improving element quality whilst increasing surface error. The linear elastic transformation step between the surface projection and smoothing operations could be viewed as an interfacial energy or tension that transferred the expansive forces to the bulk.

The balance between the two energies and forces was dictated by control parameters. The surface energy was not controlled as a reduction would cause surface losses. Therefore, only the contractive component could be adjusted to reach equilibrium. The volume energy was controlled by the degree of Laplacian smoothing using the associated parameters: the number of iterations and spring stiffness. Empirical tuning of these parameters was used to find a general equilibrium between the opposing forces. Too much or too little shrinking would cause higher degrees of distortion or stagnation (no change), respectively.

The final results operated similar to traditional parameterisation techniques, but for the bijective parameterisation of volumetric elements. Essentially, there were enough perturbations to minimise a mesh into a target's shape, whilst maintaining a large proportion of the original element quality. This led to overall performance improvements highlighted by *Figure 9.6*. The

performance plot compared the RMS surface error and the number of low-quality elements, which was analogous to their morphing energy \mathbf{E} . The values of \mathbf{E} would provide a similar trend between methods, where lower values would correspond to the lower-right quadrant and larger values to the upper left. However, that plot would not be able to distinguish between high surface errors and element distortions, unlike the performance plot.

Summary of Optimisation

The unfolding process from the first strategy generates a high-quality hexahedral mesh at the expense of large surface errors. However, as the morphed meshes were reasonably close and high-quality representations of the target geometry, they could be used as the template mesh for an optimisation process (*Figure 3.1*). Essentially, the morphed mesh provides an initial guess to the target geometry allowing the optimisation process to adjust the vertices and elements to more closely resemble the target geometry. The optimisation process aims to reduce the surface errors, whilst maintaining a large proportion of the element qualities. Additionally, this operation focussed on optimising both the surface fit and mesh-quality simultaneously. This was achieved by reusing the previously developed algorithms and combining them into a different operation.

The key differences to this strategy lie in the repeated smoothing, re-projection and morphing of all the hexahedral elements in an iterative process. Instead of optimising the surface alone then proceeding with a single volumetric transformation step, these steps are combined to iteratively unfold the hexahedral mesh into the target geometry. This is similar to research into spherical parameterisation [279] with two key differences: (1) volumetric elements are used instead of surface elements, and (2) projection onto an irregular geometry instead of a sphere.

The reuse of components in the optimised strategy were separated into three phases (*Figure 3.1b*). First, an initialisation step projects the surface vertices of an unoptimised morphed mesh onto a target geometry, and transforms the volume elements. This would induce element inversions and distortions, but fewer than what would have occurred in the initial projection step in the unoptimised strategy. Second, an iterative volumetric optimisation process was performed ($N_I=25$), which consists of three components: (1) Laplacian smoothing of a hexahedral mesh, (2) re-projection of the surface vertices onto a target, and (3) morphing of the hexahedral elements. This would begin to unfold any inverted hexahedral elements and reduce the degree

of distortions present. The aim was not to remove the inversions in a single iteration but incrementally to find the global minimum of the vertices, where a target geometry can be accommodated with the best possible surface fit (*Equation 3.84*). The re-projection and morphing would again add a degree of distortion and potentially some elements inversions. However, the total effect of a single iteration would lead to a net improvement and approach the global minimum. The magnitude of smoothing controlled by the parameters were selected to minimise large scale vertex movement, that could cause the morphing process to cancel out any improvements. However, the process needed to be balanced because if too little smoothing was performed each iteration the optimisation would either not reach convergence or take too long to achieve it. Third, if any inverted elements were present following the optimisation loop, the hexahedral elements were smoothed until there were no inversions. This step can shrink the geometry of the mesh, if there are numerous inversions and significant distortions present. However, as the smoothing parameters were selected to impart a small smoothing effect, if there were few inversions the geometry was minimally shrunk and the mesh became valid. Additionally, the Laplacian smoothing algorithm was a variant that minimises shrinking (*Section 3.5.1*). The optimised mesh-morphing strategy is summarised by the *Algorithm 3.2*.

Algorithm 3.2: Optimised Mesh-Morphing Strategy

Input: Template Mesh (\mathcal{H}_0), Target Mesh (\mathcal{T}) and Target Centreline ($\mathcal{C}_{\mathcal{T}}$)

Output: Morphed Mesh (\mathcal{H}_1)

```

1  $\mathcal{B}_1 \leftarrow \text{ProjectTemplateVertices}(\mathcal{H}_0, \mathcal{T}, \mathcal{C}_{\mathcal{T}});$ 
2  $\mathcal{H}_1 \leftarrow \text{LinearElasticTransformation}(\mathcal{B}_1, \mathcal{H}_0);$ 
3 for  $i \leftarrow 1$  to  $N_I$  do
4    $\mathcal{H}_1 \leftarrow \text{LaplacianSmooth}(\mathcal{H}_1);$ 
5    $\mathcal{B}_1 \leftarrow \text{ProjectTemplateVertices}(\mathcal{H}_1, \mathcal{T}, \mathcal{C}_{\mathcal{T}});$ 
6    $\mathcal{H}_1 \leftarrow \text{LinearElasticTransformation}(\mathcal{B}_1, \mathcal{H}_1);$ 
7  $\mathcal{Q} \leftarrow \text{CalculateScaledJacobian}(\mathcal{H}_1);$ 
8 while  $\min(\mathcal{Q}) \leq 0$  do
9    $\mathcal{H}_1 \leftarrow \text{LaplacianSmooth}(\mathcal{H}_1);$ 
10   $\mathcal{Q} \leftarrow \text{CalculateScaledJacobian}(\mathcal{H}_1);$ 

```

4 Performance Metrics and Data

A wide range of metrics were used to assess the sensitivity and performance of the mesh-morphing strategies. Additional metrics were incorporated into some of the cases to provide a complete portrayal of comparative performance, which include In total, there were six metrics assessed: surface error, volume error, mesh-quality, FE simulation precision, FE simulation convergence rates and generation times. Consequently, the four core assessment criteria outlined in *Section 2.7* were addressed for the developed mesh-morphing strategies. The following sections provide descriptions on the implementation and use of the metrics, design and generation of the CAD-based hexahedral meshes, and the development of a FE model of the tibio-femoral joint. Throughout this thesis, averages (Avg.) and standard deviations (Std.) are presented in the format Avg \pm Std.

4.1 Surface Error Analysis

The surface error analysis consists of two components: surface and volume error. The volume error infers similar information to the surface error regarding the surface fit (i.e. magnitude of change), but has other implications. Notably, the volume error addresses the aspect of overall size changes, i.e. whether the surface errors lead to larger or smaller meshes. These errors are measured with respect to the target surface mesh, which may deviate from the true values that are dependent on image resolution and accuracy of surface extraction from images.

4.1.1 Surface Error

The surface error was defined by the absolute difference between the boundaries of a target surface and corresponding hexahedral mesh. The correspondences between the two meshes were calculated using the centreline projection algorithm (*Section 3.4.1*). However, the process was inverted as the projection vectors were defined between each target vertex (as opposed to the template vertices) and their closest centreline position (*Figure 4.1*). Therefore, a correspondence was defined by the location where a projected vector intersected the boundary of a hexahedral mesh.

For this to be achieved, the quadrilateral surface of the hexahedral meshes were subdivided

into a triangular mesh. All centrelines were checked to ensure they had an appropriate length to sufficiently capture each geometry. In some instances, a centreplane-based approach provided greater accuracy for measuring the surface errors of geometries with very thin-regions, but caused negligible differences otherwise. Consequently, a centreplane was used only if the centreline-based approach yielded unsatisfactory results. The correspondences and associated projection vectors for the centreplane-based approach were determined using the same specification as the centreline-based approach. The centreline and centreplane structures were generated using the RBF minimisation algorithm on the target geometries (*Section 3.2.2*). The step parameter from *Equation 3.54* was increased in order to derive a centreplane instead of a centreline. Following RBF minimisation, the centrelines were processed with the same methods outlined in *Sections 3.2.4-3.2.5*. The centreplanes were smoothed and subdivided to provide a more continuous representation that would be analogous to the B-spline curves of the centrelines. The purpose of these post-processes were to prevent stepping-effects that would be present in coarser discretisations. The stepping-effects would cause patches of correspondence errors in the same pattern as the centreplane's topology.

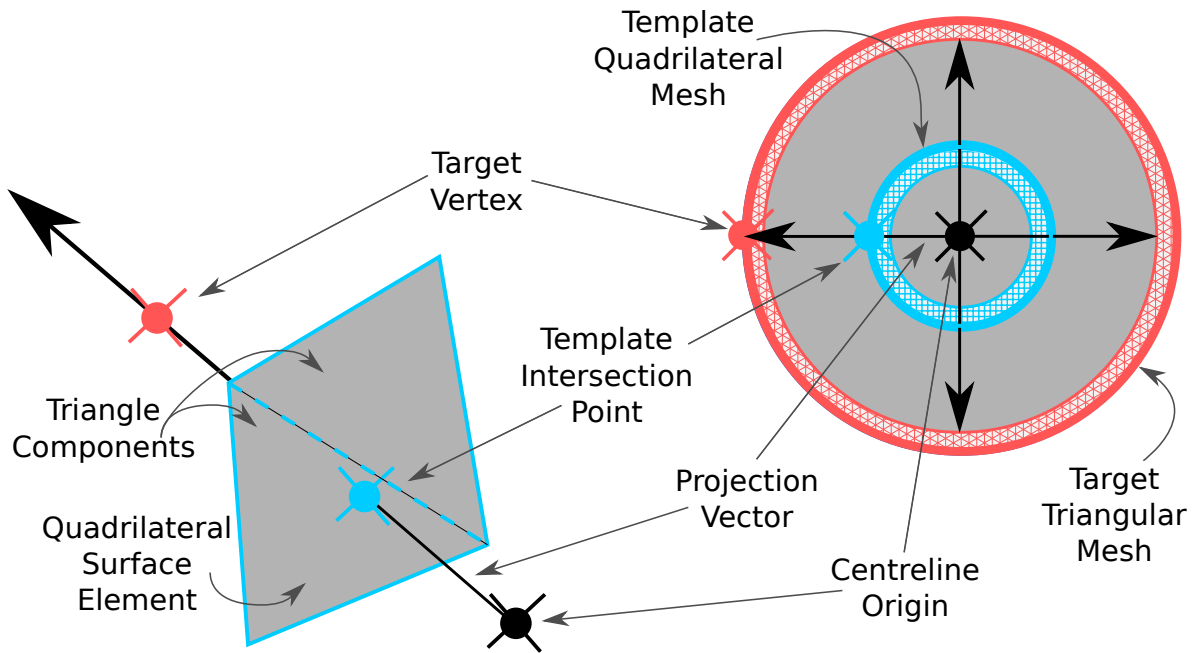


Figure 4.1: Ray-triangle intersection was used to determine correspondences by projecting each target vertex onto the boundary of a hexahedral mesh

Once the correspondences (\hat{V}_i^*) had been calculated using either method the surface error calculations were identical. The pointwise surface error (ϵ_i) was determined by calculating the Euclidean distance (L^1 -norm) between each target vertex (V_i) and their corresponding location

(V_i^*) (**Equation 4.5**). These errors were then used to calculate the maximum surface error (ϵ_{max}), root-mean-square error (RMSE) and standard deviation (σ) for a hexahedral mesh. These metrics were used in all the case studies to assess the accuracy of the surface fit. The individual surface errors were used to generate histograms to compare the error distributions between meshes in some cases. Also, they were mapped on the target geometries, to determine the distribution of errors and identify which features caused greater losses. The maximum surface error (**Equation 4.2**), RMSE (**Equation 4.3**) and standard deviation (**Equation 4.4**) of a surface fit was calculated using the following equations:

$$\epsilon_{max} = \max_{i \in \{1, \dots, N_v\}} \epsilon_i \quad (4.1)$$

$$\text{RMSE} = \sqrt{\frac{\sum_{i=1}^{N_v} \epsilon_i^2}{N_v}} \quad (4.2)$$

$$\sigma = \sqrt{\frac{\sum_{i=1}^{N_v} (\epsilon_i - \mu)^2}{N_v}} \quad (4.3)$$

$$(4.4)$$

Where, N_v represented the number of vertices, e_i was the individual vertex error and μ was the average surface error:

$$\epsilon_i = \left\| V_i - V_i^* \right\|_1 \quad \forall i \in \{1, \dots, N_v\} \quad (4.5)$$

$$\mu = \frac{\sum_{i=1}^{N_v} \epsilon_i}{N_v} \quad (4.6)$$

4.1.2 Volume Error

The volume error defined the volumetric difference between a target and corresponding hexahedral mesh. The volume of each mesh was calculated using meshplex which used the modified Gram-Schmidt process [404]. As the algorithm was designed for triangular surface meshes, the boundaries of the hexahedral meshes were converted to triangular meshes, the same as the surface error calculation. Two metrics were calculated from the volume data: the signed (δ) and absolute volume difference ($|\delta|$):

$$\delta = \overset{\circ}{V}_h - \overset{\circ}{V}_t \quad (4.7)$$

Where, $\overset{\circ}{V}_t$ and $\overset{\circ}{V}_h$ are the volumes of a target and associated hexahedral mesh, respectively.

Consequently, a negative value would indicate a size reduction and vice versa. The absolute difference provided a measure of the overall change, which was more useful when comparing the average change ($|\bar{\delta}|$) between sets of hexahedral meshes (i.e. derived from different methods):

$$|\bar{\delta}| = \frac{\sum_{i=1}^{N_m} |\delta_i|}{N_m} \quad (4.8)$$

Where, N_m represents the number of meshes in each set.

Volume measurements are commonly performed in clinical practice to determine size differences for meniscus tissues [405,406]. This measure has been found useful as the irregular geometry of a meniscus makes it difficult to determine size changes purely by their dimensions (i.e. bounding box). To evaluate the accuracy of the method, the average volumes of the medial and lateral meniscus meshes were found to be within range of values from the literature [405,406] (**Table 4.1**).

Table 4.1: Comparison of medial and lateral meniscus volumes for the target meshes

Source	Region	Average Volume (mm ³)	Medial-to-Lateral Ratio	Samples (#)	Reference
Xu et al.	Medial	1343.2±320.5	1.18	723	[406]
	Lateral	1130.0±263.2		721	
van Elst et al.	Medial	1928.9±560.7	1.16	98	[405]
	Lateral	1681.7±504.8			
Targets	Medial	1637.8±398.4	1.29	10	Table A.1
	Lateral	1264.9±364.2			

4.2 Mesh Quality Analysis

The hexahedral scaled-Jacobian (\tilde{J}) metric was used to determine the quality of each element of the template and morphed meshes. This metric is defined by the minimum determinant of the Jacobian matrix for a hexahedral element (**Equation 2.1**). Essentially, it measures the distortion of an element compared to an ideal hexahedron (**Figure 2.8**). The justification for the choice of this metric was provided in **Section 2.3.5**. The Verdict library was employed to

measure the scaled-Jacobian for the hexahedral meshes [108]. Verdict is a high-quality library that contains a variety of subroutines to measure different element-quality metrics for a range of mesh types.

The values for an element's scaled-Jacobian lie between -1 and $+1$. A value of $+1$ indicates an element with perfect quality and no distortion, whereas a value of -1 indicates an element with the worst possible quality that is completely inverted. Essentially, a higher element-quality is identified by a larger value and vice versa. Additionally, an element is only classified as valid for a value above zero. Scaled-Jacobian values closer to zero can cause mathematical problems for the solution of the finite element method, whereas values above 0.5 are generally more acceptable. Consequently, a total of six criteria have been defined to describe the element-quality (**Table 4.2**). However, the most important criteria are those which could have a negative effect on the solution of a FE model: invalid, very low-quality and low-quality. Any elements within these three criteria were the main points of concern. As such, the mesh-quality analyses focussed on comparing their proportions (%). Additionally, the minimum (\tilde{J}_{\min}) and average (\tilde{J}_{avg}) scaled-Jacobian were calculated and served as points of comparison between meshes (**Equations 4.9-4.10**). In summary, five metrics were calculated to quantify the quality of a hexahedral mesh:

1. The percentage of invalid elements (%)
2. The percentage of very low-quality elements (%)
3. The percentage of low-quality elements (%)
4. The minimum scaled-Jacobian (\tilde{J}_{\min})
5. The average scaled-Jacobian (\tilde{J}_{avg})

$$\tilde{J}_{\min} = \min_{i \in \{1, \dots, N_e\}} \tilde{J}_i \quad (4.9)$$

$$\tilde{J}_{\text{avg}} = \frac{\sum_{i=1}^{N_e} \tilde{J}_i}{N_e} \quad (4.10)$$

Where, N_e is the number of elements.

The scaled-Jacobian metric was used in all of the case studies to assess mesh-quality. Similar

Table 4.2: Mesh-quality criteria with respect to scaled-Jacobian values

Quality Criterion	Scaled-Jacobian Value (\tilde{J})
Invalid	$\tilde{J} \leq 0.0$
Very Low-Quality	$0.0 < \tilde{J} < 0.1$
Low-Quality	$0.0 < \tilde{J} < 0.5$
High-Quality	$0.5 \leq \tilde{J} < 1.0$
Very High-Quality	$0.9 \leq \tilde{J} < 1.0$
Perfect	$\tilde{J} = 1.0$

to the surface errors, the distribution of the scaled-Jacobians were determined by mapping the values on the elements of the hexahedral meshes. These were also used to understand the hexahedral quality differences between the boundary and internal elements.

4.3 Hexahedral Mesh Generation

Several hexahedral meshes were generated in this thesis using either a CAD-based or multi-block method. A CAD-based method was primarily used to generate several templates meshes. Also, similar methods were used to generate meshes of idealised tibial and femoral cartilage for the comparative simulations. Some of the hexahedral meshes used in this thesis were created using the semi-automatic multi-block method built into IA-FEMesh (MIMIX, The University of Iowa, IA, USA) [38]. This is a gold-standard procedure for generating hexahedral meshes and was used to create meniscus meshes from target geometries.

The following sections outline the CAD-based methods used to generate the meniscus, tibial and femoral cartilage meshes. Additionally, they define their mesh characteristics and the boundary conditions that were incorporated into their design to automate aspects of FE model development. The CAD-based methods were all developed using VTK.

4.3.1 Meniscus

The majority of meniscus template meshes were created with a parametric CAD-based method using morphological parameters extracted from a published SSM [94] (**Table 4.3**). The cross-sectional parameters were fit to a B-spline curve and then discretised with quadrilateral elements. This step created a flat quadrilateral mesh which was subdivided and smoothed to define different cross-sectional resolutions. New boundary vertices created by subdivision were projected onto the B-spline curve. The quadrilateral topology was designed manually and incor-

porated principles from literature used to create hexahedral meshes of meniscus anatomies [40] (*Figure 4.2a*). The quadrilateral mesh was then swept around a 240° arc in the transverse plane to create a hexahedral mesh reminiscent of a meniscus (*Figure 4.2b*).

Table 4.3: Transverse and cross-sectional dimensions for the template meshes

Reference	Cross-sectional		Transverse	
	Width (mm)	Height (mm)	Length (mm)	Width (mm)
Template Mesh	8.30	5.15	29.97	23.35

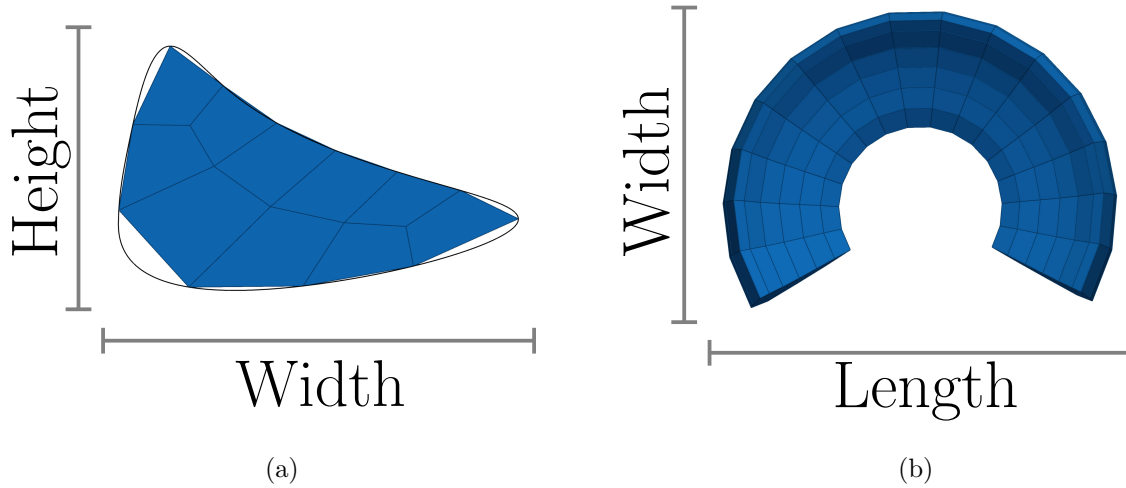


Figure 4.2: The (a) cross-sectional and (b) transversal (superior) view of a coarse template mesh (no subdivisions), with the B-spline curve of the cross-section shown in black

The number of sweeps defined the longitudinal mesh resolution. This parameter in combination with the number of cross-sectional subdivisions was used to define template meshes with higher resolutions (*Figure 4.3*). The hexahedral meshes were annotated with FE model information, which were transferred directly to the morphed meshes. This included the local element directions for the definition of anisotropic material properties and the end vertices used to define spring elements for the horn-attachments. Three template meshes with different resolutions (low, medium and high) were used throughout this thesis. Their resolutions and mesh-quality characteristics are presented in *Table 4.4*, with their depictions shown in *Figure 4.3*:

Table 4.4: Template mesh characteristics with respect to resolution and quality (subdivisions are with respect to the coarse mesh, *Figure 4.2*)

Reference	Mesh-Resolution				Mesh-Quality	
	Sweeps (#)	Subdivisions (#)	N_e	N_v	\bar{J}_{\min}	\bar{J}_{avg}
Coarse [4.2a, 4.2b]	15	0	150	272	0.68	0.83 ± 0.08
Low [4.3a, 4.3d]	30	1	1,200	1,643	0.57	0.85 ± 0.08
Med [4.3b, 4.3e]	60	1	2,400	3,233	0.57	0.86 ± 0.08
High [4.3c, 4.3f]	60	2	9,600	11,285	0.55	0.89 ± 0.08

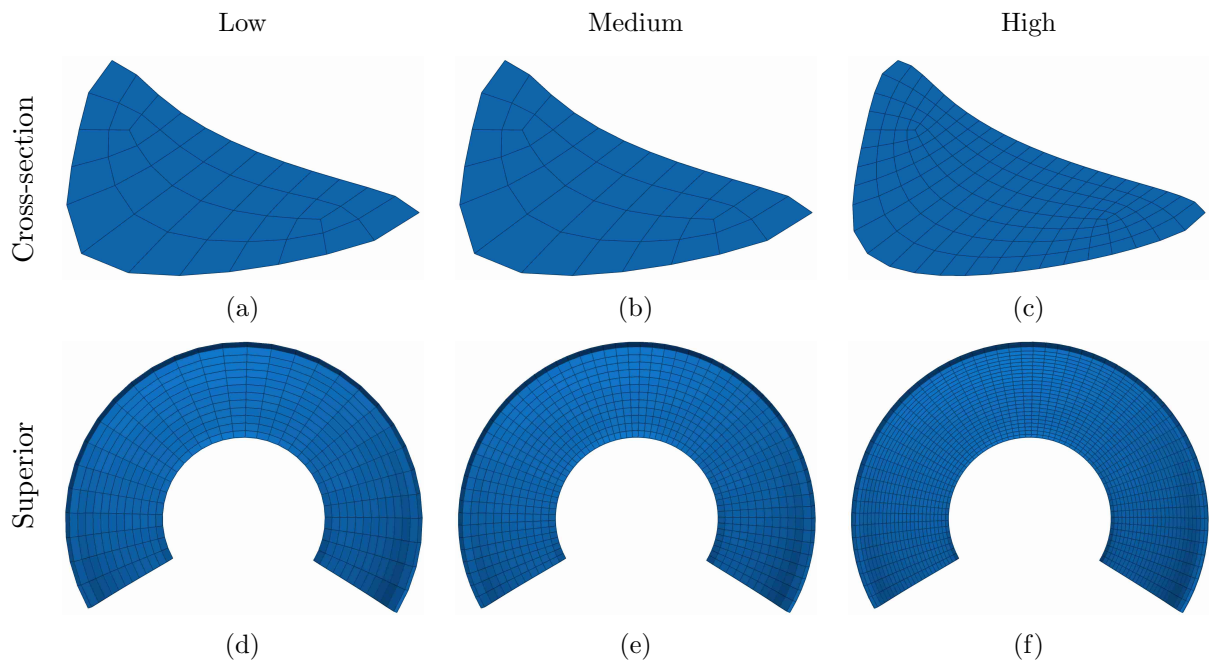


Figure 4.3: Cross-sectional (a-c) and superior (d-f) view of the (a,d) low-, (b,e) medium- and (c,f) high-resolution meshes

4.3.2 Tibial Cartilage

The medial and lateral compartments of the tibial plateau have two different shapes. The medial compartment has been reported to have a concave surface [77, 407–409]. In contrast, the lateral component has been reported to have a convex surface [77, 407, 409]. However, a study investigating the articular geometries of 55 subjects, reported that despite the lateral compartment being convex, it is flat within the main articular region for most of the subjects [407]. Additionally, both compartments have reduced curvature in the coronal plane than the sagittal [409].

To simplify the simulation environment and minimise effects of patient-specificity, one tibial cartilage mesh was used, which was flat for both medial and lateral compartments (*Figure 4.4*).

The hexahedral mesh of an idealised tibial cartilage geometry was created by extruding a square quadrilateral mesh into six hexahedral layers (*Figure 4.4*). The initial square mesh was evenly discretised to derive a mesh with 44×44 quadrilateral elements. Consequently, all 11616 elements had a perfect scaled-Jacobian value of 1. This mesh was annotated with information to automate the development of the finite-element models. Specifically, two vertices were identified which connected to the anterior and posterior springs defining the meniscus horn-attachments.




Figure 4.4: Superior view of the idealised tibial cartilage, showing the superior surface (red) and tibial bone interface (white)

Also, the boundary conditions defining the contacting interface for the superior surface and the fixed constraints for the inferior surface were annotated.

4.3.3 Femoral Cartilage

The idealised geometry of the femoral cartilage mesh was designed to reflect the convex nature of the anatomy in both medial and lateral compartments [409]. This mesh was created by sweeping a wide U-shaped quadrilateral mesh through a 240° arc in the sagittal plane. The U-shaped mesh was isomorphic to 24×6 quadrilateral mesh, i.e. 26 elements along the coronal plane and 6 elements thick. This mesh was swept 60 times which resulted in a mesh with a total of 8640 hexahedral elements (*Figure 4.5*). The idealised femoral cartilage mesh had a very high-quality with an average scaled-Jacobian of 0.91 ± 0.0047 , and a minimum value of 0.79. This mesh was also annotated with information regarding the FE boundary conditions.

4.4 Target Geometries

Several target geometries were used during this thesis. These were acquired from three different sources, one internal and two external. Eight (four pairs) meniscus geometries were acquired from within the Institute of Medical and Biological Engineering (iMBE, University of Leeds, Leeds, UK) [410]. The first external source was the Open Knee project [411] which provided five menisci (two pairs and one lateral). Four of those (two pairs) were provided from publications that detailed the image acquisition processes (one cadaveric [19] and one in-vivo [80]). The second external source was obtained directly from Prof. D.M. Pierce (Department of Biomedical Engineering, University of Connecticut, USA), who provided 32 (16 pairs) meniscus meshes




Figure 4.5: Idealised femoral cartilage hexahedral mesh: (a) sagittal and (b) coronal view, showing the inferior surface (red) and femoral bone interface (white)

directly. Half of those geometries (eight pairs) were used in a study [40], with four (two pairs) of them made available publicly [412]. The publicly available meshes had undergone smoothing and remeshing operations. In contrast, the 28 meshes that were not publicly available had not undergone post-processing operations and had more irregular surfaces. These geometries were segmented from medical images available from the Osteoarthritis Initiative database (OAI) [413, 414], whose imaging protocols have previously been documented [415]. From the three sources, a total of 47 (23 pairs + one lateral) meniscus geometries were acquired, as genus-0 triangular meshes in the stereolithography format (STL). However, as four (two pairs) were duplicates with different post-processing, the unique total was 43. Following any additional post-processing, all geometries were transformed into a standard frame of reference, inline with the process described in *Section 3.3*. Geometries that did not have flat attachment sites were clipped using VTK.

Several geometries had to undergo some additional post-processing (remeshing and artefact reduction) prior to mesh-morphing. Six geometries were remeshed to reduce the number of surface vertices. This was done to determine what effect the number of vertices and elements had on the morphing strategy. Following remeshing, 45% of geometries had a lower resolution (<6000 vertices) and 55% had a higher resolution (>8000). These ranges were selected based on the general distribution of resolutions provided by the original meshes.

Some geometries contained segmentation artefacts that caused severe surface irregularities (*Figure 4.6*). Consequently, a local high-curvature removal algorithm was developed to re-

duce the presence of highly irregular sharp features resulting from image-processing. This algorithm smoothed very sharp and concentrated spikes that would likely be removed during standard post-processing operations (*Algorithm 4.1*). This was done to ensure the geometries were inline with what would be reasonably expected of a meshed segmentation. However, they still maintained a larger degree of surface irregularity to test what effect geometries with less post-processing have on the strategies. The first step of the algorithm was to calculate the mean curvature for the target meshes. A common method used to determine the discrete mean curvature (\mathbf{K}) is based on the summation of internal angles of the one-ring surrounding a vertex [346] (*Equation 4.11*). Where a mean curvature above 10 mm^{-1} was found (value empirically determined to be sufficient for identifying severe irregularities), surface meshes were iteratively ($N=20$) smoothed using localised Laplacian smoothing (*Algorithm 4.1:L8*). The number of iterations were also determined empirically.

$$\mathbf{K}(V_i) = \frac{1}{2\mathcal{A}} \sum_{j \in \mathcal{N}(i)} (\cot \alpha_{ij} + \cot \beta_{ij})(V_i - V_j) \quad (4.11)$$

Where, V_i is a vertex of a triangular mesh, V_j is a neighbouring vertex surrounding V_i , α_{ij} and β_{ij} are the angles opposite the edge defined by V_i and V_j , $\mathcal{N}(i)$ is a set of vertex indices defining a one-ring neighbourhood around V_i and \mathcal{A} is the area of the one-ring region (*Figure 2.10*).




Figure 4.6: The effect of the artefact reduction algorithm: (a) a high-curvature artefact before and (b) after identification and smoothing

Algorithm 4.1: Curvature Smoothing

Input: A Triangular Mesh**Output:** A Smoothed Triangular Mesh

```

1  $V \leftarrow \text{GetVertices}(Mesh);$ 
2  $R \leftarrow \text{GetVerticesNeighbors}(Mesh);$ 
3  $\mathbf{K} \leftarrow \text{CalculateMeanCurvature}(V,R);$ 
4  $N_v \leftarrow \text{Length}(V);$ 
5 for  $i \leftarrow 1$  to 20 do
6   for  $j \leftarrow 1$  to  $N_v$  do
7     if  $|\mathbf{K}_j| > 10$  then
8        $V_j = \text{Average}(V[R_j]);$ 

```

4.5 Finite Element Analysis

The main objective for creating mesh-morphing strategies are to enable a reduction in the time required to develop high-quality clinical FE models. Consequently, the purpose of the FE models developed in this thesis were to determine if the morphed meshes behaved differently to gold-standard meshes in clinically-relevant simulations.

4.5.1 Tibio-Femoral Joint Model

A simplified knee model was developed for each meniscus mesh (*Figure 4.7*), which were solved using FEBio 2.9 (Musculoskeletal Research Laboratories, University of Utah, Salt Lake City, UT, USA) [33]. The models were based on a single condyle with no supporting ligaments for either the medial or lateral menisci. Previously, a similar simplified model of an isolated compartment yielded sufficient accuracy, despite anatomical exclusions [20]. In this thesis, the scenario the models simulated was a fully extended knee in the standing position, without flexion or rotations taking place. The standing position is more neutral to the stability and effects of the ligaments [21]. Furthermore, the same idealised femoral and tibial cartilage meshes were used in each simulation, enabling a reduction to the influencing variables. Attachments rooting the posterior and anterior horns of the menisci to the bone were represented by spring elements, which are a commonly used implementation and assumption [17, 19, 21, 22]. The following sections describe the design of the developed FE models.

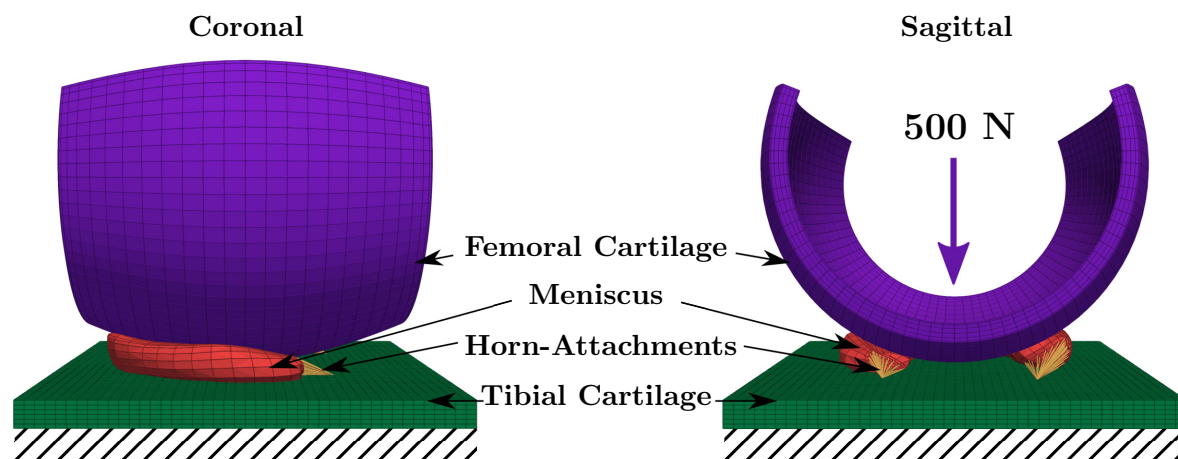


Figure 4.7: A simplified finite element model for a single condyle of a tibio-femoral joint

4.5.2 Model Geometry

The models were composed of three meshes which represented the meniscus, tibial and femoral cartilage. Idealised tibial and femoral cartilage meshes were used for two reasons. First, this enabled easier construction and minimised the number of confounding factors. Second, some of the derived meshes did not have the tibial and femoral cartilage geometries or hexahedral meshes available. The meniscus geometries were resized to allow the cartilage layers to appropriately represent the typical dimensions and relative ratios. To ensure consistency, the meniscus meshes of each simulation triplet (unoptimised mesh-morphing, optimised mesh-morphing and multi-block) had the same transformation. The simulations with the multi-block meshes were created first with the associated transformation matrices recorded. Each transformation matrix was calculated from the manual alignment with the tibial cartilage, ensuring three points of contact between the anterior horn, posterior horn and central body. These transformations were then automatically applied to the morphed meshes. However, some of the unoptimised meshes were slightly larger than the target. Those meshes had to be translated to prevent intersection with the tibial cartilage. The optimised meshes required no further adjustments and their model configurations were exactly the same as the multi-block simulations.

4.5.3 Loading and Boundary Conditions

Due to geometry idealisations there was a degree of congruency missing. To overcome this lack of congruency and to ensure the nature of the FE models maintained a level of realism, an initial displacement was applied to the femoral cartilage. This initial displacement served two

functions: (1) provided closer replication of the standing congruencies between the meniscus and articular cartilage, and (2) provided some pre-tension to the horn-attachments. The displacement was the same for each simulation triplet but different between triplets. The femoral cartilage was first placed at a distance l_0 above the tibial cartilage. This distance was the height of the meniscus (h_m) multiplied by 1.5, and the initial displacement \mathbf{u}_0 was half that distance:

$$l_0 = \frac{3}{2}h_m \quad (4.12)$$

$$\mathbf{u}_0 = \frac{1}{2}l_0 = \frac{3}{4}h_m \quad (4.13)$$

Following the initial displacement, a load was applied to the femoral bone interface to reach a total of 500 N. For simplicity and due to the fact the compartments were isolated, the percentage of the total load transmitted through each condyle was assumed to be 50% [24]. Previously, this has been observed in clinically-relevant models [24] and experimental investigations of the knee [98]. The load was 500 N because the total load applied to a complete knee would have been 1000 N [75].

The inferior surface of the tibial cartilage was fixed with respect to all degrees of freedom [75]. To maintain the fully extended knee position, the flexion-extension angle of the femur was fixed at 0° [23, 75]. Additionally, the superior surface of the femur was fixed with respect to anterior-posterior and medial-lateral translations [75, 76]. This added some of the stability that the supporting ligaments and tendons would have provided. Similarly, abduction-adduction remained fixed for the joint [23]. However, internal-external rotations remained unconstrained [75]. With respect to the contact analysis, the tibial-meniscus, femoral-meniscus and femoral-tibial interfaces were defined as *sliding-elastic* [3] and frictionless [18, 20, 21, 23, 36, 62, 63, 85, 89]. These boundary conditions and applied loads have previously been employed to simulate the knee in full extension [23, 75].

4.5.4 Material Properties

Due to the static simulation conditions the articular cartilage layers were modelled as an isotropic linear elastic material [17–21, 23, 25, 36, 75]. The Young’s modulus and Poisson’s ratio were defined as 15 MPa and 0.47, respectively [17, 19, 21, 36, 89]. Due to the main point of

comparison between models stemming from the different methods used to derive the meniscus meshes, an accurate material model for the meniscus was desired. Consequently, the menisci were defined as nearly incompressible and hyperelastic, as this was more representative than linear elastic definitions [3, 17, 95]. This ensured the FE models were simulating challenging conditions. Additionally, the meniscus was defined as transversely isotropic [17–21, 25, 36, 89], using a neo-Hookean material model with the following parameters [17]:

Table 4.5: Neo-Hookean material parameters for the menisci [17]

C_1 (MPa)	C_3 (MPa)	C_4 (-)	C_5 (MPa)	κ (MPa)	λ^*
4.6115	0.12	150	400	227.5	1.02

Where, C_1 is the bulk constant with respect to the shear modulus, C_3 is the exponential stress coefficient, C_4 is related to the collagen fibre uncrimping rate, C_5 is the elastic modulus of the collagen fibres when straightened, κ is the bulk modulus and λ^* defines the stretch of the straightened collagen fibres [17, 111].

The transversely isotropic Neo-Hookean material model incorporates properties contributed from both the ground substance ($\boldsymbol{\sigma}_{\text{bulk}}$) and collagen fibres ($\boldsymbol{\sigma}_{\text{fibres}}$) [416]. The following equations were used to define the Cauchy stress ($\boldsymbol{\sigma}$) of this material model [17, 111, 316, 416]:

$$\boldsymbol{\sigma} = \boldsymbol{\sigma}_{\text{bulk}} + \boldsymbol{\sigma}_{\text{fibres}} \quad (4.14)$$

$$\boldsymbol{\sigma} = 2C_1 J^{-\frac{5}{3}} \left(\mathbf{B} - \frac{1}{3} I_{1(B)} \mathbf{I} \right) + \kappa (J - 1) \mathbf{I} + \lambda \frac{\partial F_2}{\partial \lambda} \mathbf{a} \otimes \mathbf{a} \quad (4.15)$$

where

$$\lambda \frac{\partial F_2}{\partial \lambda} = \begin{cases} 0 & \text{if } \lambda \leq 1 \\ C_3 (e^{C_4(\lambda-1)} - 1) & \text{if } 1 < \lambda < \lambda^* \\ C_5 \lambda + C_6 & \text{if } \lambda \geq \lambda^* \end{cases} \quad (4.16)$$

$$C_6 = \frac{1}{\lambda^*} \left[C_3 (e^{C_4(\lambda^*-1)} - 1) - C_5 \right] \quad (4.17)$$

$$\lambda \mathbf{a} = \mathbf{F} \cdot \mathbf{a}_0 \quad (4.18)$$

Where, \mathbf{F} is the deformation gradient, J represents the Jacobian of that deformation gradient, \mathbf{B} is the left Cauchy-Green deformation tensor, $I_{1(B)}$ is the first invariant of that tensor, \mathbf{I} is the identity matrix, λ defines the fibre stretch, F_2 represents the collagen fibre contribution,

$\lambda \frac{\partial F_2}{\partial \lambda}$ defines the stress of the fibres, \mathbf{a} is the direction of the deformed fibres, C_6 defines the continuous stress at λ^* , \mathbf{a}_0 is the initial direction of the fibres and \otimes represents the tensor outer product.

The fibre stress has three conditions which depend on the fibre stretch (*Equation 4.16*) [17]. The first condition shows the fibres have no contribution to stress if they are compressed ($\lambda \leq 1$). The second condition describes the exponential increase in fibre stiffness if stretched below their straightened state ($1 < \lambda < \lambda^*$). The last condition defines the linear increase in fibre stiffness if stretched beyond their straightened state ($\lambda \geq \lambda^*$). The distribution of fibres in a mesh are defined using the unit vector \mathbf{a}_0 [416] (*Figure 4.8*). The use of structured and semi-structured hexahedral meshes enables the most accurate definition of these vectors, through the aligned geometry of local elements. Consequently, the models in this thesis used local element definitions, as the hexahedra of the meniscus meshes were aligned along the appropriate fibre orientations.

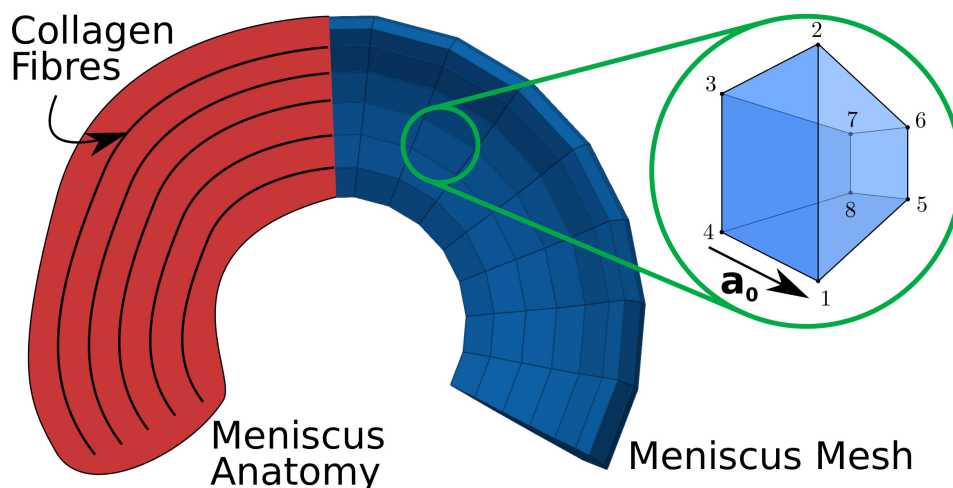


Figure 4.8: The definition of collagen fibre orientations for the hexahedral meniscus meshes using local element coordinates, with \mathbf{a}_0 the unit vector defining the fibre orientation

The horn-attachments were defined using linear springs [17–21, 36, 62, 90]. Previously, the stiffness of each attachment site (S) has been reported [18]:

Table 4.6: Stiffness of meniscal horn attachments [18]

Meniscal Horn	Attachment Stiffness (N/mm)
Lateral Anterior	324.0
Lateral Posterior	197.8
Medial Anterior	255.1
Medial Posterior	337.0

The stiffness per spring (k_i) was calculated from the attachment stiffness (S_i) using the following equation [18]:

$$k_i = \frac{S_i}{N_i} \quad \forall i \in \{1, \dots, |S|\} \quad (4.19)$$

Where, N_i is the number of springs used to represent an attachment and $|\cdot|$ defines the cardinality of a set. Each attachment was defined by 40 linear springs using the stiffness values reported in **Table 4.6**.

4.5.5 Contact Variables

Two commonly reported contact variables were derived from each simulation: contact pressure and contact area [21, 23, 62, 63]. Specifically, five aspects have been selected for comparison, which are the following:

1. Contact Area (mm²)
 - (a) Total Area
 - (b) Distribution
2. Contact Pressure (MPa):
 - (a) Maximum
 - (b) Location of Maximum
 - (c) Distribution

Contact areas were measured between the femoral-meniscus and tibial-meniscus interfaces. The maximum contact pressure between the same interfaces were calculated. The mean contact pressure was not used as it was found to be an insensitive variable. The location of the maximum contact pressure was assessed qualitatively by determining if the maximum value occurred in the anterior, middle or posterior regions of the meniscus, and which interface it occurred between. The distribution of contact pressures and areas were displayed on the tibial and femoral cartilage meshes for qualitative comparisons, where appropriate. Additionally, the contact area differences between the morphed and multi-block simulations were reported for each interface. The detailed analysis for the contact pressure and area distributions consisted of reporting the main patterns and comparative differences between the morphed and multi-block

mesh simulations.

4.5.6 Simulation Convergence Rates

In addition to the contact variables, the convergence times between the morphed and multi-block simulations were compared. The average, minimum and maximum convergence times were reported for each method. This was to determine if there were any distinctions between the speed of computing the associated simulations.

4.5.7 Mesh Convergence Analysis

Convergence analysis was performed for the developed model, to determine at which mesh-resolutions the simulation accuracy was acceptable (*Table 4.7*). Six test models were evaluated with increasing mesh-resolutions, where the simulation with the highest resolution was used as the reference. The mesh-density and -resolution ranges explored were consistent with previous convergence studies for FE models of the TFJ [134,150,154]. The accuracy of models A-E were determined by comparing their percentage change in contact area between two consecutive, where absolute changes under 5% were classified as accurate [134,150]. Absolute changes above 5% classified the associated meshes as non-optimal. In an ideal scenario, the mesh-resolution from model C would have been chosen as optimal, as it provided the highest computational efficiency (lowest resolution) amongst the accurate simulations [134,150].

However, the non-optimal mesh-resolution from model D was chosen when comparing these models with meshes derived from the multi-block method. This was because the resolution of the meniscus meshes was limited by the ability to create higher resolution multi-block meshes. Although the morphed meshes could have satisfied the accuracy requirements (higher resolutions), they all had the same number of elements to reduce confounding factors for the comparative simulations. Consequently, the resolution of the tibial and femoral cartilage meshes was specified to be consistent with the meniscus meshes, and have edge lengths within a similar range. This was done to reduce the loss of accuracy that could occur from mesh-density mismatches [417,418]. Nevertheless, the average edge length of the model was found to be consistent with previously reported TFJ models, which were all 1 mm [20,63,85,134,150,154]. All of these studies performed convergence analyses to arrive at that element size, which in their models was found to yield sufficiently accurate results. Although the model presented in this thesis has fundamentally different geometries to the aforementioned studies, it is inline with the

density and resolutions reported for other convergent models.

Table 4.7: Convergence analysis on different mesh-resolutions and element sizes

Model	Number of Elements (N_e) and Element Size (S_e , mm)						Change in Contact Area (%)
	Tibial Cartilage		Meniscus		Femoral Cartilage		
	N_e	S_e	N_e	S_e	N_e	S_e	
Reference	30,492	0.5	9,600	0.5	30,720	0.5	—
Model A	30,492	0.5	9,600	0.5	8,640	1	1.25
Model B	30,492	0.5	2,400	0.75	8,640	1	4.67
Model C	11,616	1	2,400	0.75	8,640	1	3.18
Model D	11,616	1	1,200	1	8,640	1	6.28
Model E	5,445	1.5	1,200	1	5,760	1.5	5.19

4.6 Software and Libraries

All 2D graphical representations were generated with Matplotlib [419] using the Python programming language (Python Software Foundation, Wilmington, Delaware, USA) [420]. The Visualization Toolkit (VTK) [331] and Paraview (Kitware Inc., Clifton Park, New York, USA) [421] were used to provide 3D graphical representations of meshes and their associated data.

Part II

Sensitivity Analysis of the Mesh-Morphing Strategy

5 Case A: Centreline Quality

5.1 Introduction

The mesh-morphing strategy outlined in *Section 3.8* requires two centrelines; one for a target geometry and one for a template mesh. This information is used to approximate correspondences between their surfaces (*Section 3.4.1*). Once the correspondences are determined, the volumetric transformation can be computed (*Section 3.6*). The approach to extracting centrelines starts with the surface mesh from a template and target (*Section 3.2*). First, the centrelines are generated using an automated RBF minimisation algorithm [328]. These centrelines are approximations of a theoretical ideal and have the same topology as the initial surface meshes. Next, a clustering algorithm is used to reduce the topology of the centrelines into a series of vertices that can define linear segments along the minimised mesh, i.e. a curve. To define a parameterised B-spline curve, the final step requires a divide-and-conquer algorithm between clustered vertices to determine the sequential order from the start to the end of the centreline. The main component of the automated centreline generation algorithm is the RBF minimisation. Depending on the size of the step parameter this sub-algorithm can lead to two different phenomena:

1. Contracting or shorter centrelines due to a large step size
2. Deviations in vertex positions along the centreline due to a small step size

Consequently, centrelines generated from the same geometry can have different shapes depending on the selected parameters. Additionally, some geometries tend towards shorter or deviated centrelines, which can be prevented by changing the parameters. The incorrect choice of parameters for a geometry can lead to centrelines which are not consistent with their theoretical ideal, i.e. a smooth curve which encompasses a significant proportion of a geometry. Therefore, the quality of an extracted centreline can be defined using two attributes: length and noise.

In this chapter, these attributes of a target's centreline are investigated to determine their effect on the mesh-morphing strategy. This is important to understand, as the centrelines are the first data structures computed and any inaccuracies they cause will propagate down the morphing pipeline. Consequently, this study is indirectly investigating the effect of the RBF

algorithm's control parameters on the performance of the mesh-morphing strategy. The aim for this investigation and the hypothesis to be tested are the following:

Aim: To determine the sensitivity of the mesh-morphing strategy on the quality of different centreline attributes

Hypothesis: Lower quality centreline attributes have a negative effect on the performance of the mesh-morphing strategy

5.2 Methodology

Two sub-cases were developed to determine the effects of length and noise. Three target centrelines were designed for each sub-case to understand how a progressive change in quality affects the performance of the mesh-morphing strategy. A target's centreline was modified as it would be expected that a high-quality template mesh would have a high-quality centreline. In each sub-case, a template mesh was morphed to match the geometry prescribed by a target, using the designed centrelines. The first sub-case investigated how centreline length affects the morphing strategy, whereas the second examined the effect of noise. The analysis compared the performance of the six morphed meshes (three from each sub-case) using surface error and mesh-quality metrics.

5.2.1 Centrelines

Three centrelines with different qualities were derived for each sub-case (A.1 and A.2), defined as: low, medium and high (*Figure 5.1*). In each sub-case, the high-quality centreline was the same, which was derived from the centreline algorithm and determined to have an appropriate length and no deviations. The low- and medium-quality centrelines were modifications of the high-quality centreline, which had increasing amounts of length removed (A.1) or noise added (A.2). The magnitude of length removed and noise added was determined based on the observations from a preliminary investigation. The preliminary investigation tested the centreline algorithm on a range of geometries with different control parameters. Some centrelines were found to be shorter or deviated depending on the parameters. The low-quality centrelines from both sub-cases were worse than what the centreline algorithm would typically produce. This was done to estimate the limits of the strategy and the worst-case scenario of incorrect parameter selections.

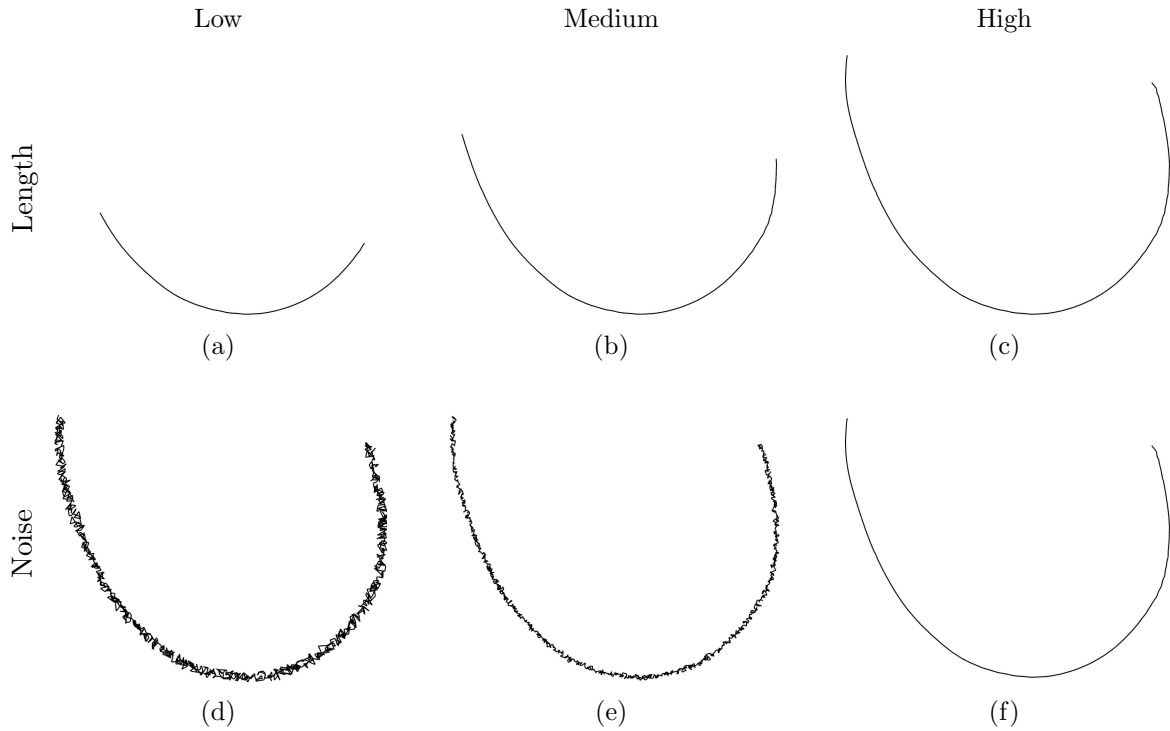


Figure 5.1: The (a,d) low-, (b,e) medium- and (c,f) high-quality centrelines with respect to length (a-c) and noise (d-f)

A.1: Centreline Length

When the shortened centrelines were examined, the limit of contraction did not exceed 50% of the original length. Therefore, this limit was designed into the centrelines, which were reduced manually and not through the control parameters. The low-quality centreline was set to this limit and the medium-quality centreline was designed to be half of that, which provided a smooth transition between the length of each centreline. The final centreline lengths and percentage reductions are presented in *Table 5.1*.

Table 5.1: Quality and properties of each centreline with respect to length

Reference	Quality	Length (mm)	Length Reduction (%)
Small-Length Mesh	Low	30.52	50
Medium-Length Mesh	Medium	45.78	25
Baseline	High	61.04	0

A.2: Centreline Noise

In some of the examined centrelines, the length was appropriate for the geometry but incurred some deviations. To investigate the effect of these deviations, noise was added to a high-quality centreline. Noise (η) was implemented using a random number generator and a range defining

the magnitude of deviations. For each vertex of a high-quality centreline ($N=1000$), random numbers were generated and applied to the x -, y - and z -coordinates. The observed deviations were less numerous than those implemented as noise, but provided a measure of the deviations' magnitude and how it affected mesh-morphing performance. The maximum value of noise was calculated based on the minimum amount of noise added before the centreline intersected the boundary of the target. This resulted in a maximum value of ± 0.5 mm added to each vertex of a centreline. The maximum noise was equivalent 0.83% of the high-quality centreline's length, which defined the low-quality centreline. This level of deviation exceeded those observed from the preliminary investigation, in order to test the upper limits and consequences of this attribute. The medium-quality centreline was defined with half the amount of maximum noise. The range of noise added to the centrelines for each vertex and the percentage with respect to length is presented in *Table 5.2*.

Table 5.2: Quality and properties of each centreline with respect to noise

Reference	Quality	Noise Added Per Vertex	
		Range (mm)	Percentage of Length (%)
Large-Noise Mesh	Low	$-0.50 \leq \eta \leq +0.50$	$0.00 \leq \eta \leq 0.83$
Medium-Noise Mesh	Medium	$-0.25 \leq \eta \leq +0.25$	$0.00 \leq \eta \leq 0.42$
Baseline	High	0.00	0.00

5.2.2 Template Mesh

The high-resolution hexahedral mesh outlined in *Section 4.3.1* was used throughout this investigation, which was generated by a CAD-based method. This mesh had 9,600 elements, additional characteristics were reported in *Table 4.4* and *Figure 4.3*.

5.2.3 Target Mesh

The target mesh (LM_03) was derived from a lateral meniscus anatomy (*Table A.1*). The geometry was selected for its irregular shape and was deemed one of the more challenging menisci to morph. A challenging geometry was selected to further test the limits of the mesh-morphing strategy using lower quality centreline attributes. This mesh was remeshed to create triangular elements with a more even size and distribution. The remeshing process incorporated several mesh processing techniques, including: decimation, subdivision, smoothing and surface fitting. The dimensions (*Figure 4.2*) of the target and template meshes are presented in *Table 5.3*.

Table 5.3: Dimensions of the target and template meshes

Mesh	Transverse		Height (mm)
	Length (mm)	Width (mm)	
Template	29.97	23.35	5.15
Target	38.72	35.46	11.52

5.2.4 Outputs of Interest

In this investigation, two outputs of interest were used to assess the sensitivity of the mesh-morphing strategy for each centreline attribute: surface error (*Section 4.1*) and mesh-quality (*Section 4.2*). When analysing surface errors, values above 5 mm were combined to maintain clarity.

5.3 Results

In general, centreline length had a strong influence on the surface error and mesh-quality, whereas centreline noise did not (*Figure 5.4*). Shorter centrelines led to worse performance metrics due to causing distinct geometry distortions to the morphed meshes, which were concentrated around the horn-attachments. In contrast, centreline noise did not cause any geometry distortions, and the differences between performance metrics were negligible.

Table 5.4: Summary of results for the performance analyses of sub-cases A.1 and A.2

Case	Quality	Surface Error (mm)		Element Quality		
		RMSE	ϵ_{max}	\tilde{J}_{avg}	\tilde{J}_{min}	$\tilde{J} < 0.5$ (%)
Length	Low	2.193±1.655	7.768	0.83±0.15	0.22	4.91
	Med	0.953±0.610	3.867	0.89±0.11	0.29	0.85
	High	0.758±0.426	2.646	0.87±0.10	0.32	0.41
Noise	Low	0.746±0.418	2.635	0.87±0.10	0.32	0.52
	Med	0.760±0.427	2.724	0.87±0.10	0.32	0.43
	High	0.758±0.426	2.646	0.87±0.10	0.32	0.41

5.3.1 Surface Error Analysis

The mesh derived from the high-quality centreline (baseline) had the same overall geometry as the target. However, the mesh-morphing strategy smoothed a large proportion of the irregular features (*Figure 5.2*). All the morphed meshes exhibited these smoothing effects and were not a complete representation of the target's geometry (*Figure 5.3*).




Figure 5.2: The superior view of (a) the target geometry and (b) baseline

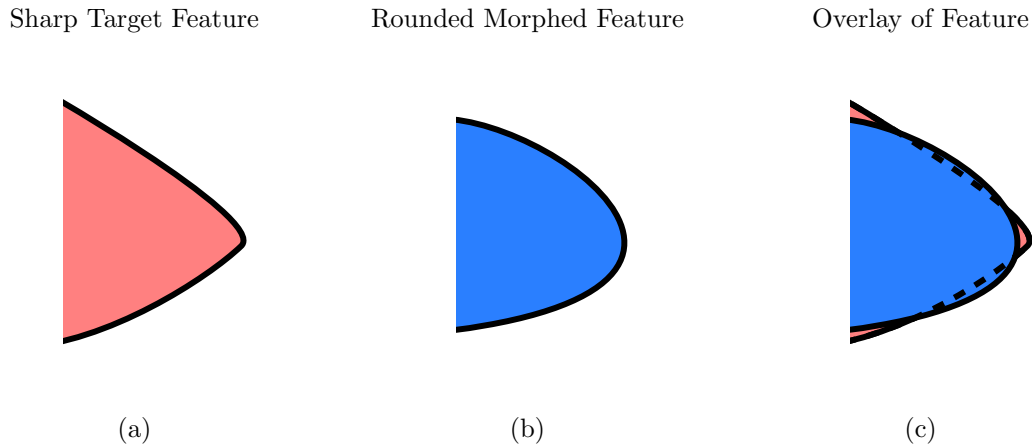


Figure 5.3: The smoothing effect causes (a) sharper features of a target mesh to (b) become rounded in a morphed mesh, (c) which causes regions of over- and under-estimation

A.1: Centreline Length

The length of a centreline had a significant effect on the surface error. Specifically, shorter centrelines led to larger surface errors, evidenced by higher maximum and RMSEs (*Table 5.5*). Additionally, the percentage of larger surface errors (>1 mm) increased with decreasing centreline length (*Figure 5.4a*), which mainly occurred at the horn-attachments (*Figure 5.5*). However, the results showed the increase was non-linear. The increase in the maximum and RMSE for the medium-length mesh compared to the baseline was 26% and 46%, respectively. In contrast, the equivalent values for the small-length mesh both increased by almost 200%. Despite differences at the horn-attachments, the medium-length and baseline meshes displayed similar patterns of surface errors (*Figures 5.5b* and *5.5c*).

A.2: Centreline Noise

The addition of noise to the centrelines had a minimal effect on the surface errors, where all meshes had a maximum and RMSE within a similar range (*Table 5.5*). The largest difference between the maximum and RMSE for these meshes was only 0.089 mm ($<2\%$) and 0.014 mm ($\approx 3\%$), respectively. The comparability was further evidenced by the similar proportions of larger surface errors (*Figure 5.4b*). Consequently, no clear differences could be distinguished from the mesh plots, which looked identical to the baseline (*Figure 5.5c*).

Table 5.5: Comparison of the RMSE and maximum surface error for sub-cases A.1 and A.2

Case	Quality	Surface Error (mm)	
		RMSE	ϵ_{max}
Length	Low	2.193 ± 1.655	7.768
	Med	0.953 ± 0.610	3.867
	High	0.758 ± 0.426	2.646
Noise	Low	0.746 ± 0.418	2.635
	Med	0.760 ± 0.427	2.724
	High	0.758 ± 0.426	2.646




Figure 5.4: Distribution of large surface errors for meshes derived from centrelines with varying (a) length and (b) noise qualities

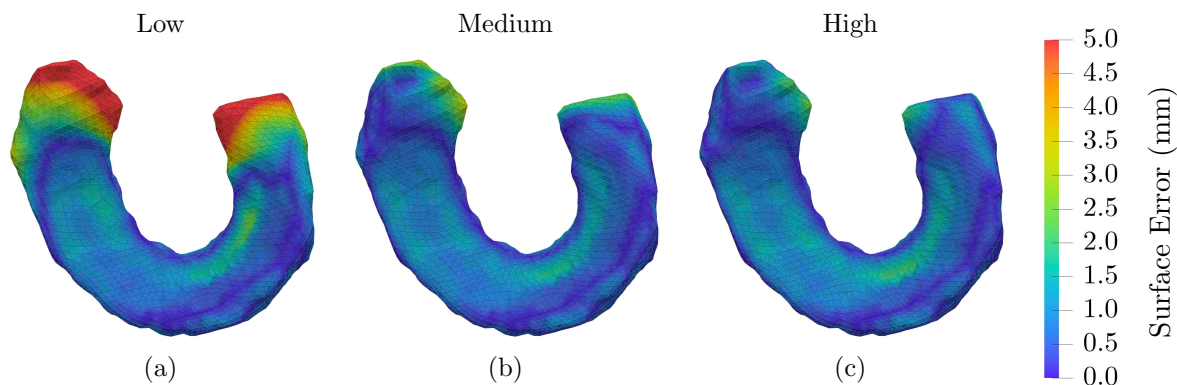


Figure 5.5: Superior view of surface errors (amalgamated above 5 mm) on meshes derived from (a) low-, (b) medium- and (c) high-quality length centrelines (red is worse and blue is better)

5.3.2 Mesh-Quality Analysis

A.1: Centreline Length

Centreline length had a strong effect on the mesh-quality, where shorter centrelines led to a larger percentage of low-quality elements and a smaller minimum scaled-Jacobian (*Table 5.6* and

Figure 5.6a). The increase in low-quality elements occurred at the horn-attachments, which exhibited geometry distortions (*Figure 5.7*). This increase with respect to decreasing centreline length was non-linear. Although there were more than twice the percentage of low-quality elements for the medium-length mesh compared to the baseline, the difference was only 0.44%. In contrast, the equivalent difference between the small-length mesh and the baseline was over a magnitude higher at 4.5%. Overall, the medium-length mesh showed more subtle differences around the horn-attachments. However, the medium-length mesh had the highest proportion of very high-quality elements and the largest average scaled-Jacobian. Additionally, these higher quality elements were observed externally and internally away the horn-attachments (*Figure 5.8*). This was not the case for the small-length mesh, which had the smallest proportion of very-high quality elements, and the smallest average scaled-Jacobian. Also, the cross-section showed significant geometry distortions, and a greater proportion of lower quality elements (*Figure 5.8a*). None of the meshes had any invalid or very low-quality elements.

A.2: Centreline Noise

Centreline noise had only a minor effect on the mesh-quality. All the meshes had the same average and minimum scaled-Jacobian (*Table 5.6*). There were some differences in the proportions of low-quality elements, which increased slightly with greater centreline noise (*Figure 5.6b*). Also, the proportion of very-high quality elements increased slightly with greater centreline noise. However, the largest difference between these meshes with respect to the proportion of low- and very high-quality elements were no greater than 0.11%. As such, no distinguishable differences were observed in either the external or internal elements, which all looked the same as the baseline (*Figures 5.7c* and *5.8c*). There were no invalid elements or very low-quality elements in these meshes.

Table 5.6: Comparison of the average and minimum scaled-Jacobian, and the percentage of very low-, low- and very high-quality elements for sub-cases A.1 and A.2

Case	Quality	Element Quality				
		\tilde{J}_{avg}	\tilde{J}_{min}	$\tilde{J} < 0.1$ (%)	$\tilde{J} < 0.5$ (%)	$\tilde{J} > 0.9$ (%)
Length	Low	0.83±0.15	0.22	0.00	4.91	44.70
	Med	0.89±0.11	0.29	0.00	0.85	60.40
	High	0.87±0.10	0.32	0.00	0.41	46.58
Noise	Low	0.87±0.10	0.32	0.00	0.52	47.50
	Med	0.87±0.10	0.32	0.00	0.43	47.46
	High	0.87±0.10	0.32	0.00	0.41	46.58

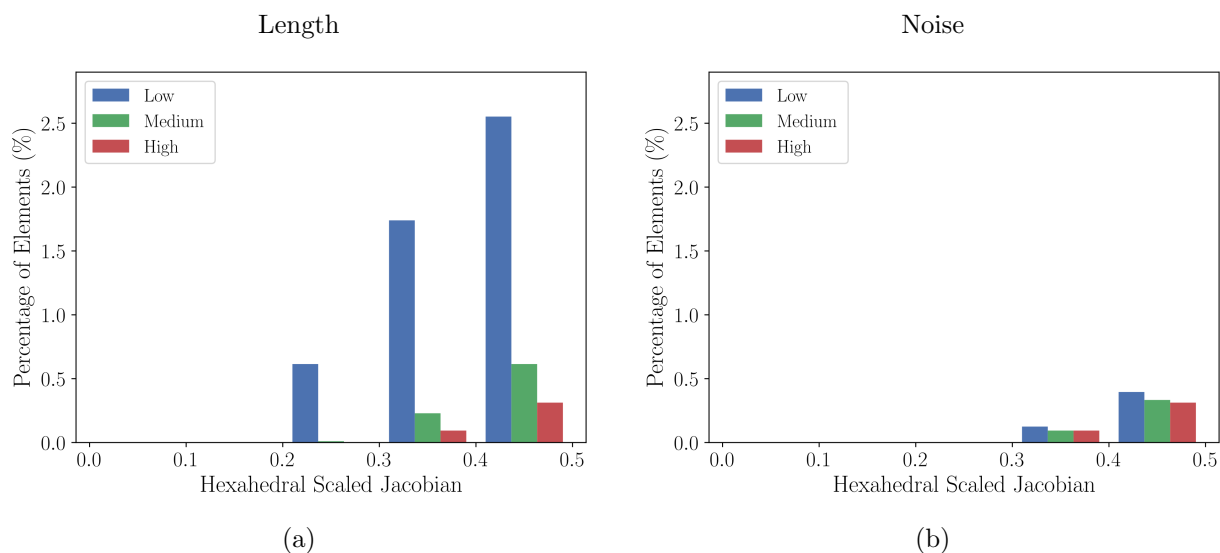


Figure 5.6: Distribution of low-quality elements for meshes derived from centrelines with varying (a) length and (b) noise qualities

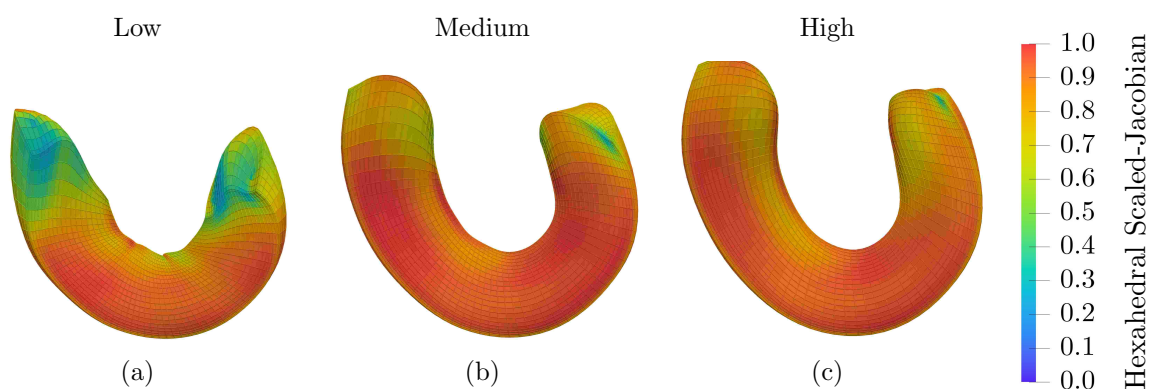


Figure 5.7: Superior view of the hexahedral scaled-Jacobian on meshes derived from (a) low-, (b) medium- and (c) high-quality length centrelines (red is better and blue is worse)

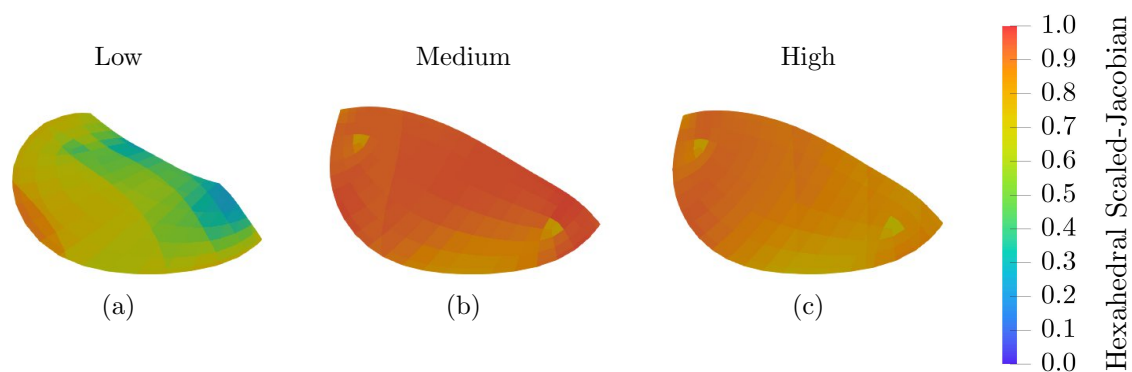


Figure 5.8: Cross-sectional view of the hexahedral scaled-Jacobian on meshes from (a) low-, (b) medium- and (c) high-quality length centrelines (red is better and blue is worse)

5.4 Discussion

This investigation has evaluated the sensitivity of the mesh-morphing strategy with respect to two centreline attributes, and their effect on performance was assessed with respect to surface error and mesh-quality. The use of shorter centrelines with the developed mesh-morphing strategy leads to worse performances and specific geometry changes. In contrast, centreline noise has a minimal effect on the performance. These findings have several implications for the operation of the centreline algorithm and consequences for FE model development.

5.4.1 Geometry Changes

The main effect of shorter centrelines lines is the increase in geometry distortion, concentrated around the horn-attachments. These geometry distortions lead to poorer surface fitting, due to the inaccurate representation of the horn-attachments. Additionally, the global distortions transfer locally into element distortions, which increase the proportion of low-quality elements within these more irregular regions. The geometry distortion is caused by the recession of the mesh around the horn-attachments, which increases with shorter centrelines.

However, the correlation between centreline length and the performance metrics is non-linear. Consequently, further investigation is required to determine a critical cut-off where a particular length is unacceptable and leads to drastically worse results. Ideally, this would include additional centrelines lengths and a range of targets. A critical metric to consider would be the greatest distance from the surface, which in this case would focus on the regions around the horn-attachments. Also, this would represent a normalised metric compared to length, which changes with the size of a geometry. Another finding that warrants further investigation into the optimal centreline length (or distance from target surface) is due to the increase in high-quality elements for the medium-length mesh. Further investigation may find an optimal length which minimises distortion and leads to higher quality meshes.

There are two key differences between the affected regions between the small- and medium-length meshes. The small-length mesh has greater distortion around the horn-attachments and central regions, whereas only the horn-attachments are affected for the medium-length mesh. Surprisingly, the external and internal elements away from this region have a higher quality than the baseline for the medium-length mesh. The regions away from the horn-attachments have

a greater concentration of elements due to the receding effect. The greater density of elements in those regions may cause the increase in higher quality elements. Previously, it has been suggested that performance improvements for mesh-morphing strategies could be achieved for particular regions with increased element densities [46]. In contrast, the higher overall distortion of the small-length mesh likely negates the increase in very high-quality elements.

5.4.2 Effect on Mesh-Morphing Strategy

Centreline Length

The reason centreline length reduces the performance of the mesh-morphing strategy is due to how it affects the subsequent processes in the pipeline. Specifically, shorter centrelines lead to a poorer initial projection (*Section 3.4.1*), defined by greater distortions and correspondence errors. The correspondence errors accumulate around the horn-attachments, where the closest centreline points used to define the projection are further away than for other regions. This causes a lower concentration of vertices around the horn-attachments compared to the baseline, which decreases with shorter centrelines (*Figure 5.9*). This provides an uneven distribution of projected vertices on the target. Essentially, the initial and subsequent projections never fully envelope the target geometry in these regions to the same extent as the remaining geometry. During the surface optimisation algorithm, iterative smoothing and reprojection steps lower the concentration of vertices further, as they recede towards the regions with greater concentrations. This is due to the Laplacian smoothing operator, which in general provides a solution for diffusion processes [422, 423].

The receding effect is exacerbated by the flat horn-attachments on the template mesh, which contains acute and faceted corners (*Figure 5.10*). These acute corners occur at the intersections between the flat horn-attachments and the curved surfaces of the meniscus. This yields two faceted corners at each horn-attachment, which does not resemble the natural aesthetic of an unclipped target geometry. Although a common technique for modelling menisci are to clip the horn-attachments flat, this was not performed for the target geometry as they are reasonably flat. However, the target geometry has horn-attachments with less curvature and greater irregularity. During the morphing operation, these mismatches lead to large surface errors and element distortions, which are observed for the baseline. However, these distortions are amplified for meshes derived from shorter centrelines, which encompass a greater proportion of the

horn-attachments.

Meshes Derived from Shorter Centrelines

Baseline

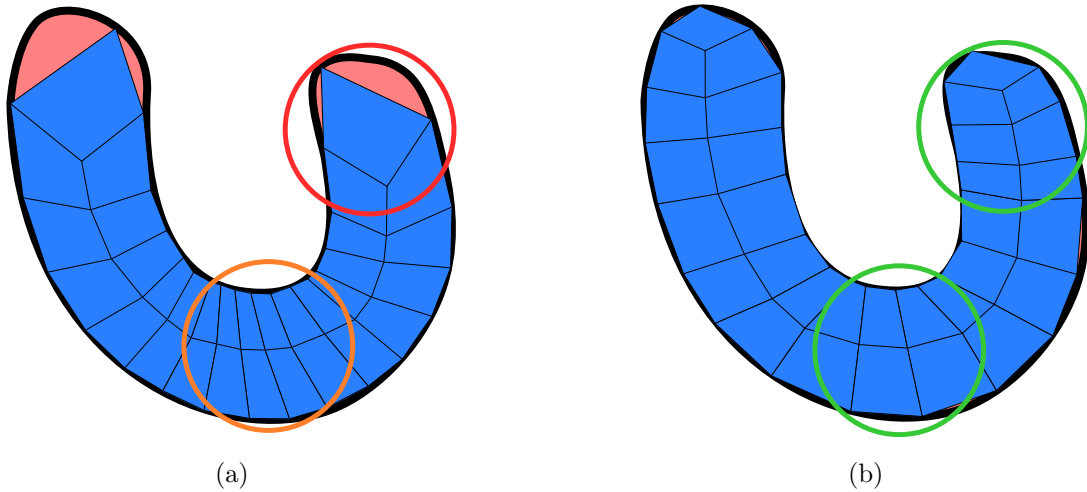


Figure 5.9: The meshes derived from shorter centrelines have lower densities (red circle) around the horn-attachments and higher densities (orange circle) away from this region compared (a) to the baseline that had a more even distribution of densities (green circles) (b)




Figure 5.10: The corners of the template's horn-attachments are sharper (a) than the target's (b), where the large mismatch between angles ($\alpha_M \ll \alpha_T$ and $\beta_M \ll \beta_T$) causes greater correspondence errors and element distortions in the initial projection (c)

Centreline Noise

The initial projections have larger correspondence errors for meshes derived from centrelines with greater noise. However, there are two reasons centreline noise has a minimal effect on the performance and processing of the mesh-morphing strategy. The first reason is due to the appropriate coverage the centrelines provide over the target geometry, which lead to an even distribution of projected vertices. The second is due to the surface optimisation algorithm, where the iterative expansion, smoothing and surface fitting processes do not depend on the centreline. These processes eliminate the correspondence errors and unfold the initial projection, when there are more uniformly distributed vertices. Consequently, given a reasonable initial

projection the strategy is able to negate a significant level of correspondence errors, and produce a bijective map that encompasses the target geometry.

5.4.3 Consequences for Finite Element Modelling

The use of meniscus meshes derived from shorter centrelines would have a negative impact on the automated development and accurate simulation of FE models of the TFJ. First, meshes derived from shorter centrelines would likely cause inaccurate FE models, due to the larger surface errors and geometry distortions. Previously, significant geometry changes have been reported to cause changes in the biomechanics of the meniscus [94].

Secondly, a major advantage of mesh-morphing is the automated transference of model information due to the use of isomorphic meshes [52, 53]. One aspect that could be automated is the definition of horn-attachment springs, as the terminal vertices of the morphed meshes should be the same as the template. However, the meshes derived from shorter centrelines have vertices receded into regions which no longer occupy the horn-attachments. These meshes would lose this aspect of automation and would require manual redefinition of the terminal vertices. If these correspondence errors are not corrected, biomechanical inaccuracies would likely be observed. Consequently, each automatically generated model could require enough manual adjustments to negate some of the benefits this type of strategy can offer. Essentially, this could lead to longer FE model development times than would have been expected through the full automation of these aspects, or using centrelines of an appropriate length.

5.4.4 Improvements and Optimisation of the Centreline Algorithm

The performance of the mesh-morphing strategy is more sensitive to centreline length than noise, and the consequences are worse. These findings have implications on the choice of the centreline algorithm's control parameters, which should be adjusted to minimise the occurrence of shorter centrelines. The step size parameter provides the means to control if a centreline has a longer length or deviations. A small step size can cause deviations in some meshes, but they have an appropriate length. The inverse is true for a larger step size, therefore a smaller step size is chosen for the use of the centreline algorithm for the subsequent chapters in this thesis.

To achieve greater robustness, future work for the algorithm could involve additional modifications or the optimisation of all the parameters (*Section 3.2*). A simple modification could

calculate the volume and area of the minimising meshes, where a volume close to zero would indicate the algorithm derived either a centreline or centreplane. The area of the minimised mesh would indicate which structure is derived, e.g. an area close to zero would indicate a centreline. Once the structure is detected, parameters could be automatically adjusted (i.e. increase step size) to derive centrelines that are less deviated. Deviated centrelines are caused from meshes that tend towards centreplanes. Therefore, this type of change would provide greater robustness for geometries which naturally tend towards centreplanes, as parameter changes can force the generation of centrelines instead of centreplanes.

5.4.5 Limitations

There were four limitations of this investigation: (1) number of targets, (2) number of centrelines, (3) implementation of deviations and (4) use of a non-clipped target. A more robust analysis of the sensitivity to centreline attributes would have included more target geometries and centrelines. However, the target geometry was selected for its challenging shape and was assumed to provide an upper limit for the difficulty of mesh-morphing with lower quality centreline attributes. Also, other geometries would have likely shown the same patterns, as the focus was on the sensitivity of the mesh-morphing strategy to centreline attributes, and not with respect to different geometries. Only three centrelines were chosen to determine the effects of the upper limits of the centreline attributes and the mid-point in severity, which provided evidence for the main patterns and associated changes. Additional centrelines could have provided greater insight for critical centrelines lengths, consequently this has been suggested as an idea for future work (described above).

Although the implementation of noise did not exactly match the presence of deviations that arise from the centreline algorithm, they provided a measure for the magnitude of deviations and represented the same problem. This problem was defined by the positional differences compared to an ideal centreline, which caused irregularities in the projection vectors from the corresponding template centreline. Therefore, it was deemed a sufficient approximation to control the magnitude of deviations using noise.

The use of a non-clipped geometry could be seen as a limitation, as they are commonly clipped for FE models. However, this was not done to negate confounding factors and use the geometry in the shape it was acquired. Also, the horn-attachments were reasonably flat compared to the

other meniscus geometries, as mentioned previously (*Table A.1*). Therefore, it could not be assumed the horn-attachments of this geometry would be clipped for FE analysis.

5.5 Conclusions

In conclusion, the results show that the performance of the mesh-morphing strategy is more sensitive to centreline length than noise. Specifically, the horn-attachments are captured inaccurately and have lower quality elements in meshes derived from shorter centrelines. The use of deviated centrelines decreases the performance slightly, but these differences have a negligible effect on the meshes. The performance losses attributed to the meshes derived from shorter centrelines would likely lead to the development of inaccurate finite element models of the tibio-femoral joint. Consequently, the parameters of the centreline algorithm are adjusted to select for centrelines of an appropriate length with possible deviations, over shorter centrelines. Additionally, these results indicate improvements of the centreline algorithm may be required to achieve greater robustness. The mesh-morphing strategy can accommodate a range of different centrelines and derives valid meshes. However, the critical finding is that the strategy is capable of correcting correspondences errors in the initial projection, providing the vertices are uniformly projected over a target. This is shown for the meshes derived from centrelines with noise but is not the case for those derived from shorter centrelines.

6 Case B: Geometric Robustness

6.1 Introduction

The meniscus has a crescent shape in transverse plane and a wedge shape in the cross-section (*Section 2.1*). In a given population, there exists a significant range of variation between these shapes. The primary variations are present within the transverse and cross-sectional planes, which have been reported in an investigation of a statistical shape model (SSM) [94]. Specifically, the length and width of the meniscus in the transverse plane, and the width and height of the cross-sections between the anterior and posterior horns.

In this chapter, several artificial target meniscus geometries are created to investigate the sensitivity of the mesh-morphing strategy to common geometry variations. The motivation for this lies in the necessity for the developed mesh-morphing strategy to provide high-quality and accurate meniscus meshes for a large proportion of the population. Only the transverse and cross-sectional variations have been investigated, as they represent the majority of the variation. The artificial target geometries represent simplified shapes to minimise confounding factors and understand the effects of the isolated variations, which approximate the common variations determined from anatomical measurements.

Aim: To determine the sensitivity of the mesh-morphing strategy to common geometry variations of the meniscus

Hypothesis: Variations which increase the curvature of the target geometries have a negative effect on the performance of the mesh-morphing strategy

Research Questions:

1. Can a template mesh with an average meniscus geometry be morphed to a significant proportion of the population, represented by common variations?
2. How do common variations and associated aspects affect the performance of the mesh-morphing strategy?

6.2 Methodology

Two sub-cases have been developed to understand how key variations of meniscus geometries affect the performance of the mesh-morphing strategy. The first sub-case investigated the transverse variation of the meniscus (B.1), whereas the second examined the cross-sectional variation (B.2). There were two parameters that defined the transverse (length L_T and width W_T) and cross-sectional (width W_C and height H_C) variations. Each sub-case contained four geometry variations, as two upper limits for each parameter were explored, which contained an isolated variation added or subtracted from an average meniscus geometry. The isolation of individual variations was used to understand which aspects caused significant performance differences using surface error (*Section 4.1*) and mesh-quality metrics (*Section 4.2*).

6.2.1 Template Mesh

The template hexahedral mesh has been created to exhibit the appropriate dimensions of a population average. The dimensions were based on averaged MRI measurements acquired from in-vivo sources [424]. The average meniscus (medial and lateral) values were derived from a cohort of 174 human specimens (99 female, 75 male) who had no underlying pathologies. The average age was 29 (18-60) years, the average weight was 72 kg and the average height was 168.8 cm. The cross-sectional parameters incorporated variations from the anterior, middle and posterior regions of the meniscus. The average transverse and cross-sectional parameters for each region are presented in *Table 6.1*.

Table 6.1: Transverse and cross-sectional dimensions for the template mesh

Region	Transverse		Cross-sectional	
	L_T (mm)	W_T (mm)	W_C (mm)	H_C (mm)
Anterior Horn			8.80	5.30
Middle Body	30.20	20.30	8.30	4.94
Posterior Horn			9.70	4.30

The parametric CAD-based method outlined in *Section 4.3.1* was used to create a high-density template mesh from the transverse and cross-sectional parameters. However, the template mesh generated for this investigation had non-uniform cross-sectional dimensions, to more closely replicate the average geometry of a meniscus. Linear interpolation was used to determine the parameters for the intermediate cross-sections between the anterior horn, middle body and posterior horn. This provided a continuous structure which smoothly transitioned between the

three regions. The dimensional parameters and interpolated values were fit to a B-Spline, which represented a slice of a cross-section (*Figure 4.2a*). A quadrilateral mesh was derived from each slice, which was rotated to the appropriate sweeping angle and structured to achieve a hexahedral mesh (*Figure 4.2b*). This mesh had 9,600 elements and 11,285 vertices with a minimum and average scaled-Jacobian of 0.52 and 0.92 ± 0.06 , respectively.

6.2.2 Target Meshes

Four target meshes for each sub-case have been artificially created to represent key differences that exists between individual meniscus shapes. The template mesh was used as the average geometry, where the isolated parameter changes were applied using the parametric CAD-based method (*Table 6.2*). The triangular target meshes were derived by extracting the quadrilateral surface meshes and subdividing the elements into triangles.

Table 6.2: Dimensional changes applied to an average geometry to create the target meshes





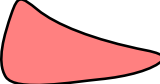


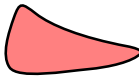

Reference	Transverse		Cross-sectional		Ratios	
	L_T (%)	W_T (%)	W_C (%)	H_C (%)	L_T/W_T	W_C/H_C
\mathbf{T}_L^+	+20	0	0	0	1.7	1.9
\mathbf{T}_L^-	-20	0	0	0	1.3	1.9
\mathbf{T}_W^+	0	+20	0	0	1.3	1.9
\mathbf{T}_W^-	0	-20	0	0	1.6	1.9
\mathbf{C}_W^+	0	0	+20	0	1.5	2.2
\mathbf{C}_W^-	0	0	-20	0	1.5	1.5
\mathbf{C}_H^+	0	0	0	+20	1.5	1.6
\mathbf{C}_H^-	0	0	0	-20	1.5	2.3
Template	0	0	0	0	1.5	1.9

The design of targets was based on simplifications of the primary variations, derived from a SSM [94]. Consequently, the targets encompassed a large proportion of the meniscus shape variation for the selected cohort. Four parameters have been used to define the simplified variations, two for each of the transverse (L_T and W_T) and cross-sectional dimensions (W_C and H_C). The magnitude of difference was set to $\pm 20\%$, as the majority of parameters were closer to $\pm 20\%$ than $\pm 40\%$. This was selected to ensure the geometries incorporated a large proportion of the cohort's variation, and still represent realistic meniscus shapes. Parameter changes of $\pm 40\%$ created meniscus geometries that were unrealistic, due to the adjustment of individual parameters. Due to the nature of the parameter changes, each case had two sets of geometries with difference ratios, i.e. transverse length-to-width (L_T/W_T) ratio and cross-sectional width-to-height (W_C/H_C) ratio.

6.3 Results

Overall, the performance of the mesh-morphing strategy was most sensitive to the cross-sectional variations, where increases to the W_C/H_C ratio led to performances losses (**Table 6.3**). There were some performance differences observed for cross-sectional variations which had equivalent W_C/H_C ratios but a smaller size, and changes to the L_T/W_T ratio. However, their effect was less pronounced compared to changes in the W_C/H_C ratio. Additionally, all morphed meshes had no low-quality or invalid elements, and generally led to mesh-quality improvements over the template mesh.

Table 6.3: Summary of results for the performance analyses of sub-cases B.1 and B.2

Reference	Surface Error		Element Quality			Schematic
	RMSE (mm)	$\epsilon > 1$ mm (%)	\tilde{J}_{avg}	\tilde{J}_{min}	$\tilde{J} < 0.8$ (%)	
\mathbf{T}_L^+	0.845±0.397	27.39	0.93±0.06	0.54	3.77	
\mathbf{T}_L^-	0.841±0.427	26.42	0.93±0.06	0.61	3.29	
\mathbf{T}_W^+	0.836±0.409	26.26	0.92±0.06	0.59	3.38	
\mathbf{T}_W^-	0.819±0.376	23.83	0.92±0.06	0.50	5.20	
\mathbf{C}_W^+	0.854±0.430	31.79	0.92±0.06	0.54	4.72	
\mathbf{C}_W^-	0.875±0.442	26.14	0.94±0.06	0.56	3.88	
\mathbf{C}_H^+	0.798±0.392	23.61	0.94±0.05	0.58	2.59	
\mathbf{C}_H^-	0.889±0.442	34.26	0.92±0.06	0.57	4.51	
Template	N/A	N/A	0.92±0.06	0.52	4.42	

6.3.1 Surface Error Analysis

In general, the geometry of the morphed meshes represented a close approximation to the target surfaces (*Figure 6.1*). The key differences were found towards features with high-curvatures, which became smoothed and rounded. An example of features that have been smoothed were found towards the corners of the horn-attachments (*Figure 6.1b*).

Furthermore, the highest surface errors occurred where the variations caused an increase in curvature with respect to the original template. The highest curvatures were located around the wedge and upper rim (superior-periphery). As the cross-sectional variations had the greatest impact on the curvatures in these regions, more significant differences were observed. Overall, the highest and lowest surface errors were measured from the cross-sectional variations. In contrast, the transverse variations had less of an effect on curvature, which led to greater similarities in the surface errors.




Figure 6.1: The superior view of (a) the target geometry \mathbf{T}_L^+ and (b) the morphed mesh with smoothed features at the horn-attachment corners circled

B.1: Transverse Variation

The transverse variations (\mathbf{T}_L^- and \mathbf{T}_W^+) which decreased the L_T/W_T ratio led to an increase in maximum surfaces errors, which on average were 0.62 mm larger (*Table 6.4*). These variations showed a very slight increase in the proportion of larger surface errors (>1 mm), which on average were 0.73% greater. Overall, the average difference between the RMSEs was very similar ($\downarrow 0.007$ mm). The transverse variations which decreased the L_T/W_T ratio showed their larger surface errors and maximum towards the peripheral corners of the horn-attachments on the inferior surface (*Figures 6.2a-6.2d*). Overall, \mathbf{T}_L^- had the worst surface error with the largest maximum, as well as second highest RMSE and proportion of larger errors. \mathbf{T}_W^- had

the lowest error, evidenced by the smallest maximum, RMSE and proportion of larger errors.

B.2: Cross-sectional Variation

The percentage of larger surface errors (>1 mm) increased with an increase to the W_C/H_C ratio, which on average were 8.15% greater (**Table 6.4**). These errors were mostly concentrated around the wedge of the menisci on the inferior surface (**Figures 6.3a-6.3d**). Between cross-sectional variations with the same W_C/H_C ratio, those with a smaller size (C_W^- and C_H^-) had a slightly greater proportion of larger surface errors, which on average were 2.5% greater. Also, there was a slight increase to the RMSE by 0.056 mm. These larger surface errors were observed on the superior surface towards the corners of the horn-attachments between corresponding ratio pairs, e.g. between C_W^- and C_H^+ (**Figures 6.3e-6.3h**). For C_W^- , this region had the highest maximum surface error of all the cross-sectional variations, which was 0.48 mm larger than the second highest maximum (C_W^+). A cluster containing this maximum was observed near the posterior horn on the superior surface (**Figure 6.3f**). Overall, C_H^+ had the best surface fit and C_H^- had the worst.

Table 6.4: Comparison of the RMSE, maximum surface error and percentage of errors above 1 mm for sub-cases B.1 and B.2

Variation	Surface Error		
	RMSE (mm)	ϵ_{\max} (mm)	$\epsilon > 1$ mm (%)
T_L^+	0.845 ± 0.397	1.681	27.39
T_L^-	0.841 ± 0.427	2.377	26.42
T_W^+	0.836 ± 0.409	2.137	26.26
T_W^-	0.819 ± 0.376	1.600	23.83
C_W^+	0.854 ± 0.430	2.106	31.79
C_W^-	0.875 ± 0.442	2.587	26.14
C_H^+	0.798 ± 0.392	2.089	23.61
C_H^-	0.889 ± 0.442	2.012	34.26

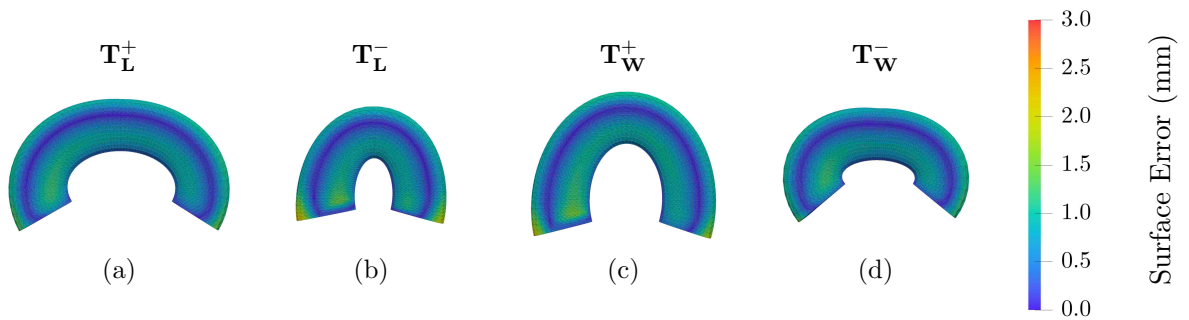


Figure 6.2: Inferior view of surface errors on meshes derived from transverse length changes at (a) +20% and (b) -20%, and width changes at (c) +20% and (d) +20% (red is worse and blue is better)

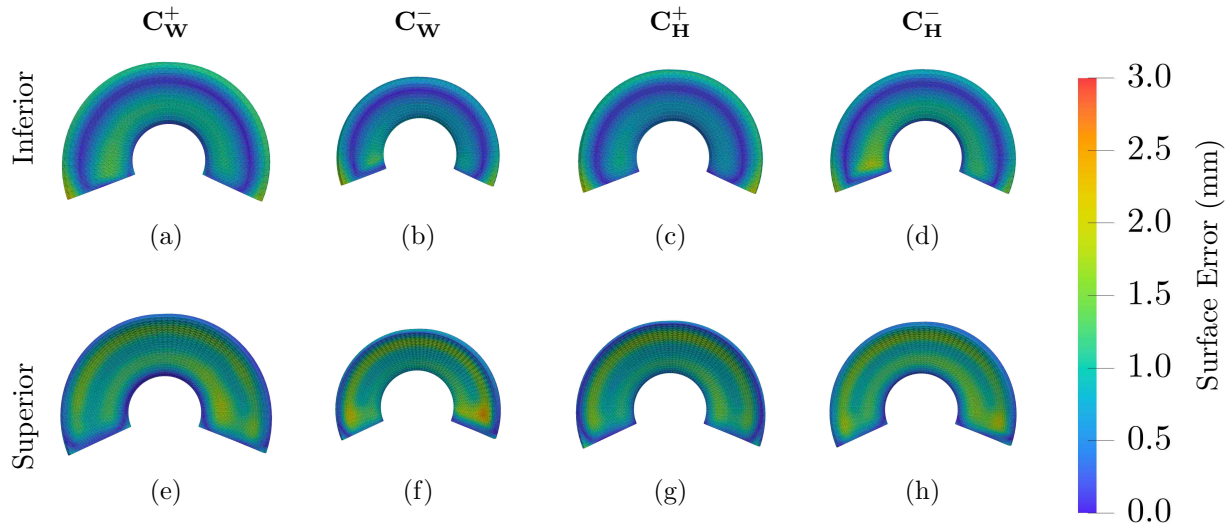


Figure 6.3: Inferior (a-d) and superior (e-h) view of surface errors on meshes derived from cross-sectional width changes at (a,e) +20% and (b,f) -20%, and height changes at (c,g) +20% and (d,h) +20% (red is worse and blue is better)

6.3.2 Mesh-Quality Analysis

The greatest mesh-quality differences occurred between the cross-sectional variations, where a slightly lower mesh-quality was observed for variations with a larger W_C/H_C ratio. The differences between the transverse variations were more subtle, however there was a slight decrease in mesh-quality for variations with a larger L_T/W_T ratio. Overall, the highest quality meshes for all the variations were observed for meshes with the smallest W_C/H_C ratio.

In general, the results between all variations were similar and each derived high-quality meshes. Additionally, improvements to the mesh-quality were observed compared to the template mesh, evidenced by a greater proportion of very high-quality elements ($\tilde{J} > 0.9$) (**Table 6.5**). Only one mesh had both a smaller \tilde{J}_{\min} and a greater proportion of lower quality elements ($\tilde{J} < 0.8$) than the template (\mathbf{T}_W^-), which had the worst quality of all the morphed meshes. Also, the meshes which had a larger W_C/H_C ratio (\mathbf{C}_W^+ and \mathbf{C}_H^-) had a slightly greater proportion of lower quality elements than the template mesh, which on average were 0.2% higher.

B.1: Transverse Variation

In general, there were two trends observed: (1) the proportions of lower quality elements ($\tilde{J} < 0.8$) correlated with \tilde{J}_{\min} , and (2) the proportions of very-high quality elements ($\tilde{J} > 0.9$) correlated with \tilde{J}_{\max} (**Table 6.5**). However, only the first trend could be attributed to specific geometry differences. Consequently, the observations were more subtle for the transverse

variations.

The transverse variations (\mathbf{T}_L^+ and \mathbf{T}_W^-) which increased the L_T/W_T ratio led to a slight decrease in mesh-quality, evidenced by a smaller \tilde{J}_{\min} ($\downarrow 0.08$) and an increase in the proportion of lower quality elements ($\uparrow 1.15\%$). However, this trend was less significant for the \tilde{J}_{avg} ($\downarrow 0.001$) and the proportion of very high-quality elements ($\downarrow 0.14\%$). Consequently, the mesh-quality distributions and patterns for the external and internal elements were similar (**Figures 6.4a-6.4d** and **6.5a-6.5d**). However, a slightly greater proportion of lower quality elements were observed towards the corners of the horn-attachments for variations with a larger L_T/W_T ratio. Overall, \mathbf{T}_L^- had the best mesh-quality and \mathbf{T}_W^- had the worst.

B.2: Cross-sectional Variation

The cross-sectional variations (\mathbf{C}_W^+ and \mathbf{C}_H^-) which increased the W_C/H_C ratio had a slightly lower mesh-quality. This was evidenced by a greater percentage of lower quality elements ($\uparrow 1.38\%$), as well as reductions in the \tilde{J}_{avg} ($\downarrow 0.02$) and the proportions of very high-quality elements ($\downarrow 12\%$) (**Table 6.5**). These lower quality elements were observed towards the wedge in the external and internal elements (**Figures 6.4e-6.4h** and **6.5e-6.5h**). Overall, \mathbf{C}_H^+ had the best mesh-quality and \mathbf{C}_W^+ had the worst.

Table 6.5: Comparison of the average and minimum scaled-Jacobian, and the percentage of lower and very high-quality elements for sub-cases B.1 and B.2

Variation	Element Quality			
	\tilde{J}_{avg}	\tilde{J}_{\min}	$\tilde{J} < 0.8$ (%)	$\tilde{J} > 0.9$ (%)
\mathbf{T}_L^+	0.93±0.06	0.54	3.77	73.80
\mathbf{T}_L^-	0.93±0.06	0.61	3.29	71.89
\mathbf{T}_W^+	0.92±0.06	0.59	3.38	69.85
\mathbf{T}_W^-	0.92±0.06	0.50	5.20	67.66
\mathbf{C}_W^+	0.92±0.06	0.54	4.72	67.68
\mathbf{C}_W^-	0.94±0.06	0.56	3.88	77.74
\mathbf{C}_H^+	0.94±0.05	0.58	2.59	81.21
\mathbf{C}_H^-	0.92±0.06	0.57	4.51	67.27
Template	0.92±0.06	0.52	4.42	63.52

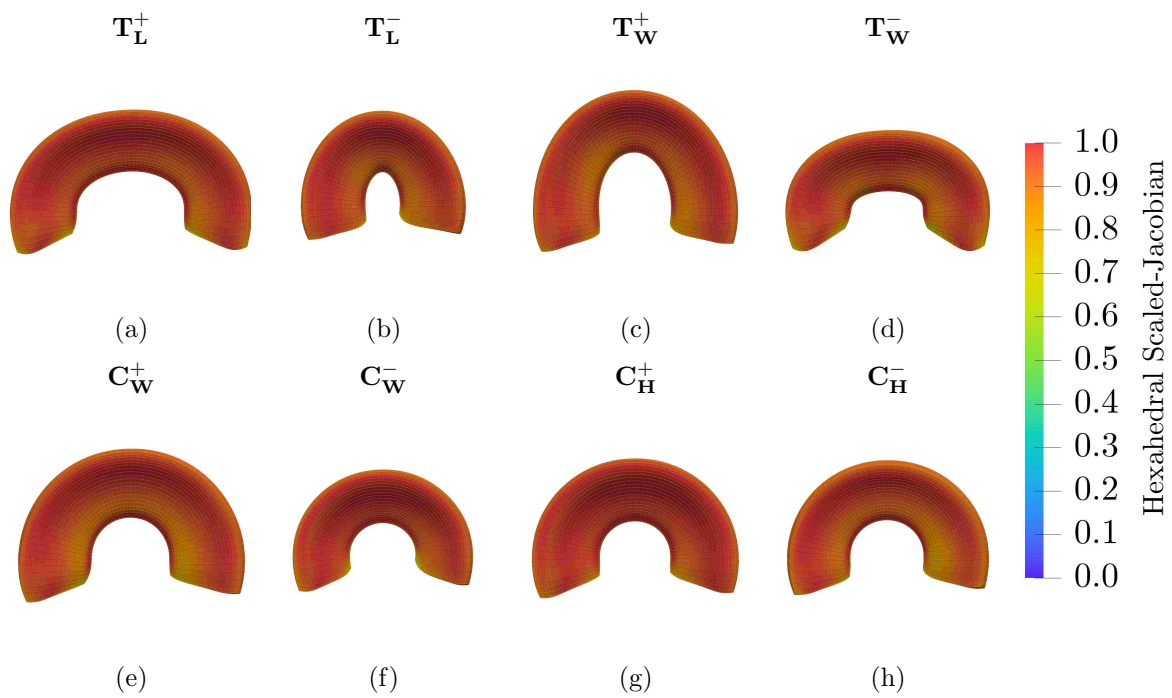


Figure 6.4: Superior view of the hexahedral scaled-Jacobian on meshes derived from transverse length changes at (a) +20% and (b) -20%, and width changes at (c) +20% and (d) -20%. Cross-sectional width changes at (e) +20% and (f) -20%, and height changes at (g) +20% and (h) -20% (red is better and blue is worse)

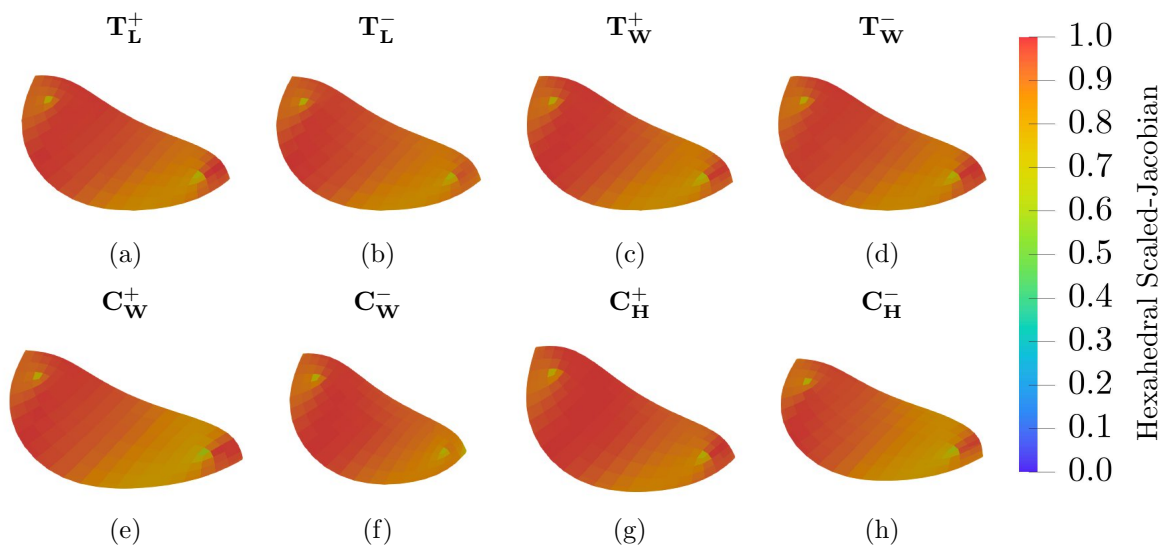


Figure 6.5: Cross-sectional view of the hexahedral scaled-Jacobian on meshes derived from transverse length changes at (a) +20% and (b) -20%, and width changes at (c) +20% and (d) -20%. Cross-sectional width changes at (e) +20% and (f) -20%, and height changes at (g) +20% and (h) -20% (red is better and blue is worse)

6.4 Discussion

The aims of this study were to understand if the mesh-morphing strategy can operate over a range of meniscus shapes with known population variations, and which incur performance losses. Overall, the strategy is able to produce high-quality meshes for a range of artificial menisci, which have an improved mesh-quality compared to the template mesh. The main finding is that the mesh-morphing strategy is most sensitive to the cross-sectional variations of the meniscus, yielding arguably the best ($\mathbf{C}_{\mathbf{H}}^+$) and worst ($\mathbf{C}_{\mathbf{H}}^-$) overall performances. Additionally, the greatest differences are observed when cross-sectional curvature differences exist between the template and target meshes, which are related to the width-to-height ratio. Some subtle differences in performance are associated with size reductions in the cross-sectional variations and differences in the transverse length-to-width ratio.

6.4.1 Template Quality Improvements

Interestingly, all morphed meshes have an improved quality compared to the template hexahedral mesh. This finding contradicts many reports of the mesh-degradation that occurs during a typical mesh-morphing operation [13, 28, 48, 58, 425]. However, there are examples of mesh-morphing strategies maintaining [425] and improving mesh-quality [241].

As the target meshes have been produced using the same CAD-based algorithm, the geometric features are fairly similar to the template mesh except for the key dimensional variations. Additionally, the design of the template and target geometries are idealised and do not contain any anatomical or surface irregularities. These two factors combined with the surface optimisation leads to morphed meshes with a higher mesh-quality, as they are less geometrically constrained as the template due to particular aspects of the mesh-morphing strategy.

The smoothing operations performed during mesh-morphing preferentially improves the mesh-quality over the accuracy of the surface fit. Following the initial projection (*Sections 3.4.1*), there are no surface errors but a high-degree of distortion in the quadrilateral elements of the template mesh. These distortions occur due to geometry differences between the template and target meshes, which can induce invalid elements during the volumetric transformation step. The purpose of the surface optimisation step is to eliminate these distortions, which introduces surface errors and the rounding effect on the geometries. Overall, the surface fit and mesh-

quality compete with each other, where an improvement in one leads to a deterioration in the other. Previously, this trade-off has been commonly observed for other mesh-morphing strategies [45, 53, 58, 307, 403]. Another cause of surface errors is due to the final hexahedral smoothing step, which improves the mesh-quality and ensures there are no invalid elements. This causes a rounding and shrinking effect which is most pronounced at the corners of the horn-attachments.

6.4.2 Effect of Geometric Variations

Overall, the cross-sectional variations show greater differences than the transverse variations, which could be due to the larger change in the width-to-height ratios over the length-to-width ratios, respectively. Compared to the template mesh, the cross-sectional variations had differences in their width-to-height ratios between 0.3 and 0.4. In contrast, the transverse variations had differences in their length-to-width ratios between 0.1 and 0.2. Consequently, this supports findings that larger differences between target and template meshes can cause greater performance losses [46–48]. This indicates that the dimensions of a meniscus are not the critical factor, but the changes respect to a template’s dimensions. There are three dimensional changes that affect the performance of the mesh-morphing strategy, in order of significance:

1. Cross-sectional width-to-height ratio
2. Cross-sectional size
3. Transverse length-to-width ratio

Cross-sectional Curvature Changes

The reason target geometries ($\mathbf{C}_{\mathbf{W}}^+$ and $\mathbf{C}_{\mathbf{H}}^-$) with a larger width-to-height ratio have worse performance metrics is due to the decrease in the wedge angle, which represents the highest curvatures. Also, the wedge of these menisci are the regions where losses are observed in the surface fit and mesh-quality. An increase to the width (W_C) and a decrease to the height (H_C) both cause the wedge angle to become more acute, which increases the curvature around these regions (*Figure 6.6*). These regions of higher curvature cause greater element distortion and reduce the mesh-quality, due to a mismatch between the shape of the regions and an ideal hexahedron. Additionally, during smoothing operations surface losses are more pronounced in regions with higher curvatures, similar to what is generally observed for the horn-attachment

corners. In contrast, target geometries ($\mathbf{C}_{\mathbf{W}}^-$ and $\mathbf{C}_{\mathbf{H}}^+$) that have a smaller width-to-height ratio (larger wedge angle) can provide greater accommodation for the shapes of the hexahedral elements. This reduces the number of lower quality elements required to represent those geometries, and the curvature that causes the larger surface errors. Therefore, it could be inferred that anatomical geometries with a small wedge angle could lead to poorer mesh-qualities. This supports findings from literature, where performance losses have been observed for mesh-morphing strategies operating over thin features [45,426]. Similarly, anatomical geometries with a smaller wedge angle than the template would likely cause higher proportions of lower quality elements and larger surface errors.

Future work to fully understand and appreciate this phenomenon would involve the parametrisation of the wedge angle directly, to study the influence of increasing and decreasing that variable. Although, features with higher curvature are known to cause difficulties for a wide range of surface parametrisation techniques [6,295,347,369,427,428], whereas flat or low-curvature surfaces are easier to parametrise to unit geometries [6,279,295,429]. In the developed mesh-morphing strategy, two surfaces with non-zero global curvatures are parametrised to each other. Historically, the original surface parametrisation problem was to parametrise a curved surface to a flat domain (cartography) [6,429,430]. As such, the majority of surface parametrisation algorithms and early developments are based around these principles [6,347,369,427,431]. Therefore, the developed strategy may encounter several difficulties arising from this aspect of attempting to parametrise a surface without a base domain, such as a unit square, circle or sphere [6,292].

Cross-sectional Size Changes

Cross-sectional size reductions affected the surface error with respect to two attributes: (1) larger errors for equivalent width-to-height ratios, and (2) higher average losses in general. As the pairs for corresponding width-to-heights ratios have essentially the same shape at different scales, it is not clear why this happens for the smaller meshes and not the larger (*Figure 6.7*). However, there is one confounding factor that could be attributed to this effect. The factor is responsible for positional irregularities in the centrelines' shape between target and template meshes, which are caused by the different centreline generation methods (*Figure 6.7a*). This factor is suspected to have an influence due to its occurrence in the peripheral regions of the horn-attachments, where greater centreline irregularities occur and the higher surface errors are observed. This is most clearly observed for a larger cluster of surface errors on the mesh

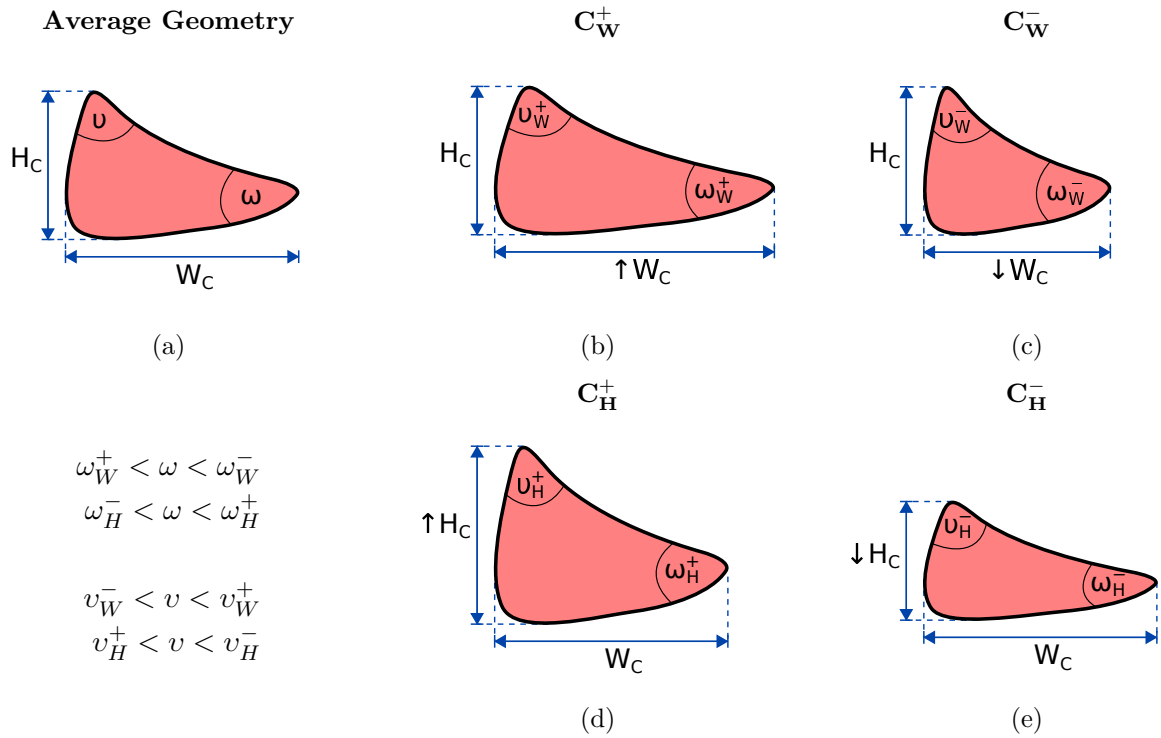


Figure 6.6: Influence of cross-sectional width and height variations on the angle of the wedge ω and upper rim v (superior-periphery): (a) the average geometry of the template mesh, with the (b) increase and (c) decrease of width, and the (d) increase and (e) decrease of height

with reduced cross-sectional width (C_W^-). The positional differences change the cross-sectional origin of the centrelines relative to its position in the template mesh (*Figures 6.7b-6.7c*). The difference in origin causes correspondence errors between the targets and template meshes.




Figure 6.7: Cause of correspondence errors for the cross-sectional variations: (a) different centreline shapes between the target (blue) and template (yellow) meshes, (b-c) similarity of correspondence errors between variations with equivalent width-to-height ratios

Transverse Curvature Changes

The transverse length-to-width ratio causes subtle differences in the performance metrics. Overall, an increase in the ratio leads to smaller surface errors and a higher mesh-quality, and vice

versa. One explanation could be due to the difference between the angle of the horn-attachments compared the template mesh combined with irregular centreline differences, as the performance differences occur in the horn-attachment corners (**Figure 6.8**). However, it is not clear what mechanism initiates the performance differences, and why these occur with an increase and not a decrease in the ratio. Nevertheless, this aspect has a much smaller effect on the metrics compared to the cross-sectional curvature and size changes, and could be an insignificant finding.




Figure 6.8: Correspondence errors associated with different horn-attachment angles (α , β and γ) and centreline shapes: (a) between the average geometry of the template mesh (blue, β) and the variations which (b) decrease (green, α) and (c) increase (red, γ) the transverse length-to-width ratio

6.4.3 Consequences for Finite Element Modelling

Based on the findings of this investigation, there are three possible consequences for the use of the mesh-morphing strategy to develop FE models of the TFJ:

1. Reduced simulation accuracy for menisci with small wedge angles
2. Reduced simulation accuracy for menisci with small cross-sections
3. Improved simulation accuracy for menisci that do not have the aforementioned attributes

Menisci with small wedge angles could cause simulation inaccuracies due to poorer capture of the wedge regions and containing a greater proportion of lower quality elements. Additionally, the lower quality elements exist in regions which would be in contact with both the tibial and femoral cartilage. Ideally, high-quality elements should be situated around interfacing regions during contact analyses [24, 193]. Also, contact analysis is sensitive to geometry changes, and these regions were not accurately captured in the morphed meshes. This is problematic as meniscus

extrusion (reduced wedge angle) is commonly associated with the progression of osteoarthritis within the knee and other pathologies [79,83,432]. Therefore, there would be a higher prevalence of patients with extrusions which require clinical therapies (and for which FE simulations would be useful), compared to healthy pathologies.

Similarly, menisci with small cross-sections could lead to reduced accuracies due to the poorer capture of the superior regions around the horn-attachments. This could cause the inaccurate simulation of the biomechanics between the meniscus-femoral interface, or require the use of additional smaller templates to prevent the increase in larger surface errors. However, the shape of the horn-attachments have previously not been found to have a significant effect on the biomechanics of the meniscus [410].

In general, the morphed meshes show greater smoothness and rounded features compared to the target geometries. Although this leads to surface errors, it results in a large proportion high-quality elements. As such, depending on the degree of surface errors, these attributes could be advantageous. In general, contact analyses favour smoother surfaces and higher quality elements, which can lead to performance and accuracy improvements for a FE model [24, 193, 433]. Further research is required involving simulations with morphed meshes derived from image-based geometries containing these features, to determine if these speculations (c.f. consequences 1-3) are valid.

6.4.4 Limitations

The main limitation of this investigation is the use of CAD-based geometries. As described, each of the artificial geometries has morphed successfully, with mesh-quality improvements compared to the template mesh. However, as the target geometries are derived from a CAD-based method and not derived from image-based procedures, they are not true representations of in-vivo meniscus shapes. In general, image-based meniscus geometries have highly irregular surface features and are not smooth, following image-processing (*Table A.1*). Although, these attributes could be due to limitations of the image-processing techniques and the presence of image artifacts. Regardless, the design of the targets are based on a swept geometry that have smooth contours. Consequently, this generates surfaces that are much smoother and an oversimplification of what would be acquired from image-based procedures. The 3D structure of in-vivo geometries has additional variations, not just through the 2D cross-sections. For example, the horn-attachments

have raised protrusions and are more geometrically different than the cross-sections towards the middle of the anatomy [434]. Additionally, in-vivo menisci can contain several irregularities due to pathologies, such as tears and extrusions (*Figure 2.2*). As such, the results reported here do not indicate that morphing to anatomically derived meniscus geometries would lead to mesh-quality improvements. The greater irregularity and curvature differences of image-based meshes would cause greater surface errors and lower quality meshes, due to larger differences to the template. Overall, these geometries are idealised therefore care should be taken when extrapolating these findings to geometries derived from segmented medical images. Despite the fact that the geometries may not exist in nature to these exact specifications, for the purpose of this study they are sufficient.

6.5 Conclusions

In conclusion, all morphed meshes have an improved mesh-quality compared to the original template but some did not improve to the same extent as others. Although, they do not represent the target geometries with 100% accuracy, due to smoothing and rounding effects. The mesh-morphing strategy has the most difficulty with variations that have a smaller wedge angle than the template geometry. Additionally, the strategy is more sensitive to cross-sectional variations than transverse. Overall, the transverse variations provide reasonably similar results. In contrast, the best and worst performances are from cross-sectional variations, due to an increase or decrease in the wedge angle, respectively. Consequently, morphed meshes derived from meniscus geometries with smaller wedge angles could cause performance and accuracy losses to FE simulations of the tibio-femoral joint. Finally, it seems reasonable to infer that the mesh-morphing strategy can morph a generic template mesh to a large proportion of meniscus variations inherent in a given population. As such, the developed mesh-morphing strategy appears to be suitable for meniscus geometries.

7 Case C: Mesh Sensitivity

7.1 Introduction

In the previous cases (*Chapters 5-6*), a template mesh with the same characteristics has been used, which includes: mesh-resolution, -quality, -topology and -shape. However, these characteristics are variable, and will affect the performance of a mesh-morphing strategy. Consequently, this investigation aims to determine the sensitivity of the mesh-morphing strategy to these template mesh characteristics. Understanding the effects of these characteristics will allow a more informed and appropriate choice of template meshes.

The mesh-resolution defines the quantity of elements and vertices a mesh is composed [122]. This property gains particular interest due to its role at influencing the accuracy of a FE model, which increases with resolution and is generally dictated by convergence studies [18, 36, 63, 76, 134]. Additionally, there are indications that the performance of mesh-morphing strategies can improve with increasing resolution [46, 47]. However, a large resolution can impose superfluous computational requirements [134, 435, 436].

A high-quality template mesh is often used to limit the generation of lower quality elements, as morphing operations generally induce mesh-degradation [13, 28, 48, 58, 425]. Ideally, a mesh used in FE simulations will have a large proportion of high-quality elements and very few, if any, low-quality elements. This is because low-quality elements and their locations have a strong impact on the convergence and accuracy of a FE simulation [54, 170].

Mesh-topology is defined by the connections between the vertices that define the elements [172, 437, 438], which dictates the regularity (*Section 2.4.1: Regularity*). This characteristic is highly dependent on the mesh generation procedure, which can either be designed manually or be subject to the operations of a given method and choice of parameters [38, 40, 192, 212]. Additionally, mesh-topologies can define the number and valence of singularities, or irregular vertices, which are sometimes referred to as defects [167, 192, 439]. Some of these features can act like point constraints and stress concentrators depending on the particular application of the mesh [247, 440]. Consequently, they can have an effect on mesh-processing algorithms, such as Laplacian smoothing. For hexahedral meshes, vertex singularities are defined by an element

connectivity (i.e. valence or neighbourhood) that does not equal four on the boundary or eight within the volume [167].

The shape of a template mesh is typically based on an existing image-based mesh (**Table 2.8**). Additionally, template mesh choices are often selected for their greater similarity to target geometries [46–48, 52, 441]. This is because large geometry differences between a template and target mesh reduce the performance of a mesh-morphing strategy [46–48].

Aim: To determine the sensitivity of the mesh-morphing strategy to four template mesh characteristics: mesh-resolution, -quality, -topology and -shape

Hypothesis: Template meshes with a higher resolution, higher mesh-quality, meniscus-specific topology and greater shape similarity have a positive effect on the performance of the mesh-morphing strategy

Research Questions:

1. (a) How sensitive is the mesh-morphing strategy to template mesh resolution?
(b) Do cross-sectional or central-axis resolution changes affect the strategy?
2. (a) How sensitive is the mesh-morphing strategy to template mesh quality?
(b) Does the strategy improve or retain low-quality elements?
3. (a) How sensitive is the mesh-morphing strategy to template mesh topology?
(b) Does a meniscus-specific topology improve the performance of the strategy?
4. (a) How sensitive is the mesh-morphing strategy to template mesh shape?
(b) Does shape similarity improve the performance of the strategy?

7.2 Methodology

Four sub-cases were developed to determine the sensitivity of the mesh-morphing strategy's performance to mesh-resolution (C.1), -quality (C.2), -topology (C.3) and -shape (C.4). The performance has been quantified using surface error and mesh-quality analyses. Also, one target mesh was used throughout all four investigations, which was the same employed in *Chapter 5*. In the previous investigations (*Chapters 5-6*), a high-resolution and high-quality template mesh with a meniscus-specific topology and an idealised meniscus shape had been used.

The first sub-case was designed to understand the influence of mesh-resolution with a series low-, medium- and high-resolution (baseline) template meshes. Similarly, to understand the influence of mesh-quality, the second sub-case used low-, medium- and high-quality (baseline) template meshes. The third sub-case had two template meshes generated from different methods, to understand the effect of mesh-topology. The topologies differed with respect to their regularity and shape-specificity. To determine the influence of mesh-shape, the fourth investigation had two template meshes that differed with respect to their similarity to an image-based target meniscus geometry.

7.2.1 Template Meshes

C.1: Resolution

Three template meshes with different resolution were created with the CAD-based method outlined in *Section 4.3.1* by controlling two parameters: (1) number of sweeps and (2) number of cross-sectional subdivisions. Through the manipulation of these parameters, the resolutions increased exponentially from the low- to high-resolution meshes. Although these parameters could have been defined to create a linear increase in resolution, the elements would have been too elongated and not representative of a typical hexahedral mesh. However, convergence studies are often conducted with similar non-linear differences, therefore is not outside the typical use for different mesh-resolutions [18, 36, 63, 76, 134]. Consequently, the chosen resolutions were based on common ranges explored during convergences studies [134, 150, 442].

From the low- to medium-resolution meshes, the resolution along the central-axis was doubled. In contrast, the cross-sectional resolution was quadrupled between the medium- to high-resolution meshes. Although this could be seen as a confounding factor, this enabled the deter-

mination of how dimension-specific resolution increases affected the performance. The control parameters and the associated resolution for each template mesh are presented in **Table 7.1**:

Table 7.1: Control parameters and the associated template mesh resolution

Reference	Control Parameters		Mesh-Resolution	
	Sweeps	Subdivisions	N_e	N_v
Low-resolution [7.1a, 7.1d.]	30	1	1,200	1,643
Medium-resolution [7.1b, 7.1e.]	60	1	2,400	3,233
High-resolution [7.1c, 7.1f.]	60	2	9,600	11,285

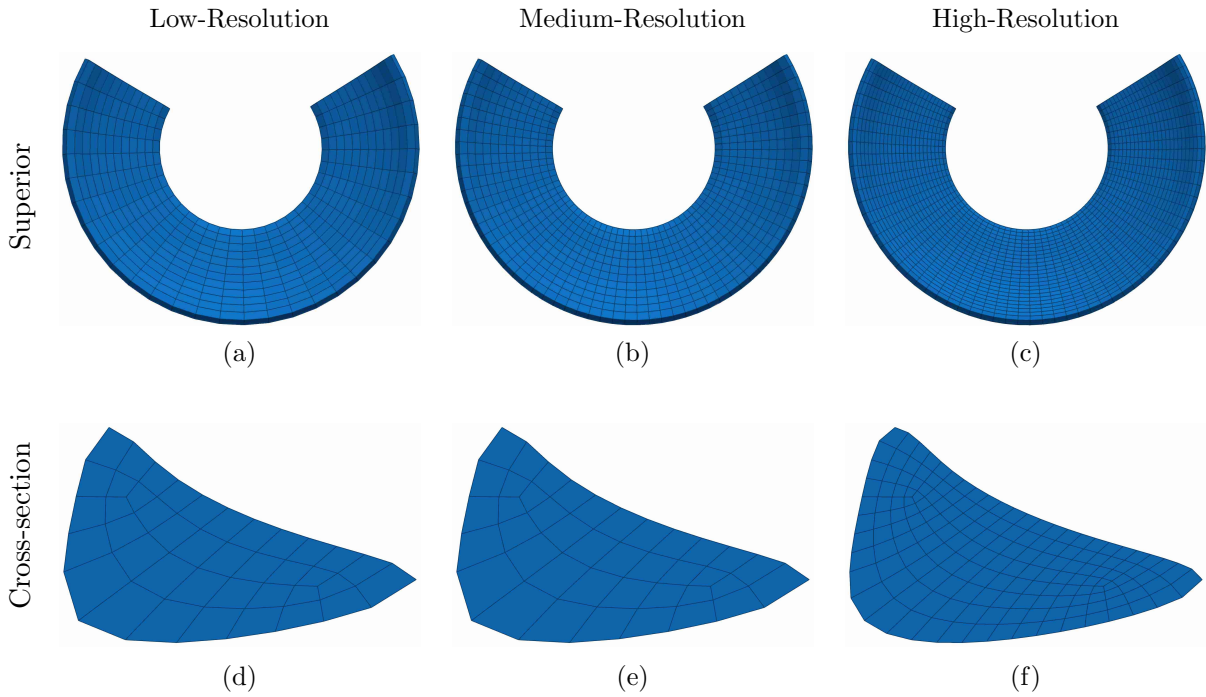


Figure 7.1: Superior (a-c) and cross-sectional (d-f) view of the (a,d) low-, (b,e) medium- and (c,f) high-resolution meshes

C.2: Quality

Three template meshes with variable mesh-quality were induced artificially by modifying the quadrilateral cross-section that was swept in the CAD-based method (**Section 4.3.1**). To control the degree of quality, noise was added to each of the internal vertices of the quadrilateral mesh. Only the internal vertices of the 2D cross-sectional quadrilateral mesh were adjusted, to ensure the template meshes had the same shape. Consequently, the majority of vertex adjustments were situated within the hexahedral mesh, with the only external vertex adjustments situated at the ends of the horn-attachments. However, the quality of the hexahedral elements on the boundary were still reduced, but to a lesser degree compared to elements whose vertices were all internal. Also, as one cross-section was adjusted and swept to create each template

mesh, the mesh-quality differences were the same along the swept axis.

The amount of noise added was proportional to the loss of quality. The noise (η) was created using a random number generator, which had an even distribution above and below zero. This was the same technique applied to induce centreline deviations (noise) in *Chapter 5* (sub-case A.2). To create the low-quality mesh, a maximum amount of noise was empirically determined. First, a significant amount of noise was added, which was steadily reduced until a valid mesh was achieved with very low-quality elements. This was the maximum amount that could be added without initiating invalid elements, which was ± 0.1 mm. The medium-quality mesh was then created using half of this variable, to provide a mid-point between the low- and high-resolution meshes. Additionally, this mesh was verified to ensure the quality was between the low- and high-quality template meshes, and contained low-quality elements but no very low-quality elements. No noise was added to the vertices defining the high-quality template mesh, which was identical to the high-resolution template mesh (baseline). All meshes had the same resolution and topology as the high-resolution mesh. The degree of noise added to each vertex and associated mesh-quality metrics for each template mesh are presented in the *Table 7.2*:

Table 7.2: Degree of noise added and the associated template mesh quality

Reference	Noise Range (mm)	Element Quality			
		\tilde{J}_{avg}	\tilde{J}_{min}	$\tilde{J} < 0.1$ (%)	$\tilde{J} < 0.5$ (%)
Low-quality [7.2a]	$-0.10 \leq \eta \leq +0.10$	0.75 ± 0.21	0.01	1.25	14.37
Medium-quality [7.2b]	$-0.05 \leq \eta \leq +0.05$	0.84 ± 0.13	0.21	0.00	2.50
High-quality [7.2c]	0.00	0.89 ± 0.08	0.55	0.00	0.00

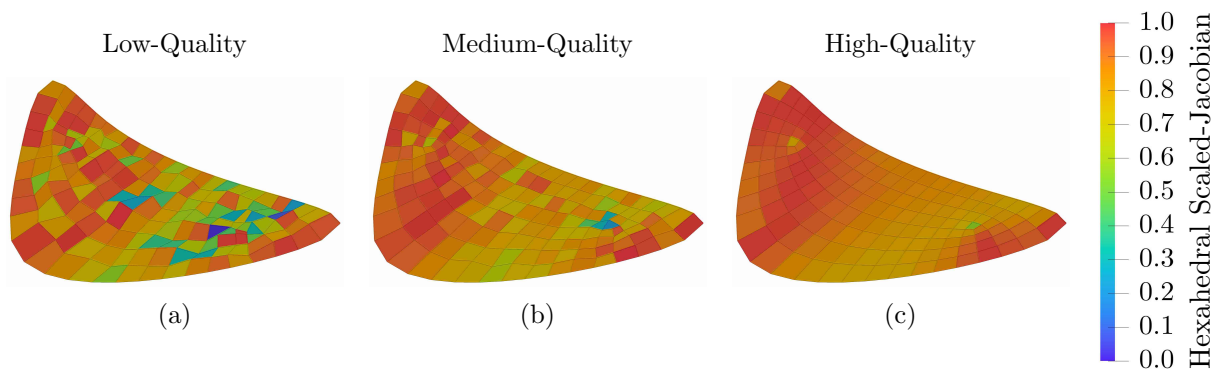


Figure 7.2: Cross-sectional view of the (a) low-, (b) medium- and (c) high-quality meshes

C.3: Topology

Two templates meshes with different topologies were created using the CAD-based (sweeping) method and a multi-block method, which were defined as CAD-Swept and CAD-MB, respec-

tively. Sweeping methods are known for producing high-quality hexahedral meshes of simple shapes with either a structured or semi-structured topology, in a robust and predictable fashion [443]. In contrast, the multi-block methods are state-of-the-art procedures for creating high-quality hexahedral meshes of irregular geometries and anatomies [28, 36, 38]. However, they are less robust and can require greater user-interactivity than sweeping methods, depending on the target geometry. Although it is possible to create semi-structured topologies using the multi-block method, structured topologies are typically derived. All multi-block meshes were created with IA-FEMesh (MIMIX, University of Iowa, Iowa, USA) [38]. Both template meshes were based on the same CAD-based geometry, which was derived from the low-resolution template mesh (n.b. the low-resolution mesh is the same as CAD-Swept). Additionally, both meshes had 1,200 elements, as it was difficult to reach higher resolutions using the multi-block method. This appeared to be towards the limit for the number of elements they could represent without inducing invalid elements.

The CAD-Swept mesh had a meniscus-specific semi-structured topology, whereas the CAD-MB mesh had a generic structured topology. The topological differences were emphasised by the number and position of singularities. Overall, both meshes had 168 singularities, but their locations were different with respect to their positions on the meniscus shape and proportions on the boundary. The CAD-based mesh had 110 external (boundary) and 58 internal (non-boundary) singularities. The cross-section had an external singularity at the wedge and upper rim (superior-periphery), which each had one internally adjacent (*Figure 7.3a*). In contrast, the multi-block meshes contained all 168 singularities on the boundary (i.e. no internal singularities). The cross-section had two singularities around the wedge, and one at the upper and lower rim (inferior-periphery) (*Figure 7.3b*). A potential confounding factor was that the mesh-quality was poorer for the CAD-MB mesh. The characteristics of each template mesh are presented in *Table 7.3*:

Table 7.3: Mesh generation attributes and the associated template mesh topology

Reference	Mesh Generation		Mesh-Resolution (#)		Singularities (#)	
	Geometry	Method	Elements	Vertices	Internal	External
CAD-Swept	CAD	Swept	1,200	1,643	58	110
CAD-MB	CAD	Multi-Block	1,200	1,674	0	168




Figure 7.3: Internal (dashed-line) and external (solid-line) vertex singularity differences between the (a) CAD-Swept and (b) CAD-MB meshes

C.4: Shape

Two template meshes with different shapes were generated using the multi-block method. One of the geometries had an idealised CAD-based meniscus geometry (CAD-MB), whereas the other had an anatomical image-based meniscus geometry (anatomical-MB). The CAD-MB mesh was the same used in the mesh-topology sub-case (C.3). The anatomical-MB mesh was generated from an in-vivo meniscus geometry, which was acquired from the Institute of Medical and Biological Engineering (iMBE, University of Leeds, UK). The geometry was extracted from MRI data acquired from the Open Knee project [3], which represented the lateral meniscus (right knee) of a patient identified as Generation 2 - Specimen 1. The MRI used fat saturation phases to produce a resolution of $0.35 \times 0.35 \text{ mm}^2$ in the sagittal plane with a sliced thickness of 0.70 mm and no slice gap. Segmentation was performed manually using ScanIP (Simpleware Ltd, Exeter, UK), which was also used to extract a triangular mesh of the geometry. This geometry was selected as it represented a geometry with low-irregularity and high-smoothness (LM_00). Additionally, it had the lowest complexity and greatest similarity to the other lateral meniscus geometries available (*Table A.1*). A summary of the information used to generate the template meshes are presented in *Table 7.4*:

Table 7.4: Mesh generation attributes and the associated template mesh shape

Reference	Mesh Generation	
	Geometry	Method
CAD-MB [7.4a]	CAD	Multi-Block
Anatomical-MB [7.4b]	Anatomical	Multi-Block




Figure 7.4: Superior view of the (a) CAD-MB and (b) anatomical-MB meshes

7.3 Results

Overall, the performance of the mesh-morphing strategy was most sensitive to mesh-resolution, where higher resolutions led to performance improvements (**Table 7.5**). Additionally, the mesh-topology had an effect on the mesh-morphing strategy, where a meniscus-specific semi-structured topology (CAD-Swept) was more favourable than a generic structured topology (CAD-MB). Between generic structured meshes with the same topology (CAD-MB and anatomical-MB), the anatomically shaped mesh had greater performance metrics. In contrast, the performance of the mesh-morphing strategy was the least sensitive to mesh-quality, where the differences were more subtle. For the four sub-cases, improvements to the quality of the template meshes were observed for the low- and medium-quality meshes (C.2) and the generic structured meshes (C.3 and C.4). All the other meshes incurred mesh-quality reductions. None of the meshes had any invalid elements.

Table 7.5: Summary of results for the performance analyses of sub-cases C.1-C.4

Case	Mesh	Surface Error (mm)		Element Quality		
		RMSE	ϵ_{\max}	\tilde{J}_{avg}	\tilde{J}_{min}	$\tilde{J} < 0.5$ (%)
C.1	Low-resolution TM	N/A	N/A	0.85 ± 0.08	0.57	0.00
	Med-resolution TM	N/A	N/A	0.86 ± 0.08	0.57	0.00
	High-resolution TM	N/A	N/A	0.89 ± 0.08	0.55	0.00
	Low-resolution MM	0.895 ± 0.581	2.959	0.80 ± 0.15	0.16	6.33
	Med-resolution MM	0.759 ± 0.425	2.712	0.81 ± 0.15	0.11	5.21
	High-resolution MM	0.758 ± 0.426	2.649	0.87 ± 0.10	0.33	0.40
C.2	Low-quality TM	N/A	N/A	0.75 ± 0.21	0.01	14.37
	Med-quality TM	N/A	N/A	0.84 ± 0.13	0.21	2.50
	High-quality TM	N/A	N/A	0.89 ± 0.08	0.55	0.00
	Low-quality MM	0.747 ± 0.423	2.583	0.86 ± 0.10	0.33	0.44
	Med-quality MM	0.757 ± 0.424	2.645	0.87 ± 0.09	0.33	0.40
	High-quality MM	0.758 ± 0.426	2.649	0.87 ± 0.10	0.33	0.40
C.3	CAD-Swept TM	N/A	N/A	0.85 ± 0.08	0.57	0.00
	CAD-MB TM	N/A	N/A	0.84 ± 0.19	0.04	7.08
	CAD-Swept MM	0.895 ± 0.581	2.959	0.80 ± 0.15	0.16	6.33
	CAD-MB MM	1.149 ± 0.774	4.324	0.78 ± 0.20	0.10	11.67
C.4	CAD-MB TM	N/A	N/A	0.84 ± 0.19	0.04	7.08
	Anatomical-MB TM	N/A	N/A	0.81 ± 0.23	0.03	12.25
	CAD-MB MM	1.149 ± 0.774	4.324	0.78 ± 0.20	0.10	11.67
	Anatomical-MB MM	1.064 ± 0.682	3.789	0.84 ± 0.18	0.15	8.83

7.3.1 Surface Error Analysis

C.1: Resolution

The surface errors worsened with decreasing mesh-resolution, evidenced by a higher maximum, RMSE and percentage of larger surface errors (>1 mm) (**Table 7.6**). The larger surface errors were concentrated around the superior surface and corners of the horn-attachments, however the increase in errors was non-linear (**Figures 7.5a-7.5c**). For the medium-resolution mesh compared to the baseline, the average increase in the maximum, RMSE and percentage of larger errors was 4.22%. In contrast, the average increase between these three metrics for the low-resolution meshes was over $5\times$ larger ($\uparrow 21.13\%$). Consequently, the distribution of surface errors was very similar between the medium-resolution mesh and the baseline (**Figures 7.5b-7.5c**).

C.2: Quality

The quality of the template meshes had a minimal effect on the surface errors, where all metrics had values within a similar range (**Table 7.6**). The largest difference between the maximum, RMSE and percentage of larger surface errors was 0.011 mm ($<2\%$), 0.067 mm ($<3\%$) and 0.66% ($<4\%$), respectively. Consequently, no clear differences were observed, which all looked identical to the baseline (**Figure 7.5c**).

C.3: Topology

Lower surface errors were observed for the CAD-Swept mesh compared to CAD-MB. This was evidenced by a lower maximum, RMSE and percentage of larger surface errors (**Table 7.6**). The CAD-MB mesh showed an average increase of approximately 38% between the three metrics compared to the CAD-Swept mesh. The larger surface errors were concentrated around the middle of the superior surface and the corners of the horn-attachments (**Figures 7.5d-7.5e**). Additionally, there was a cluster of very large surface errors (> 3 mm) towards the middle of the superior surface.

C.4: Shape

Lower surface errors were observed for the anatomical-MB mesh compared to CAD-MB (**Table 7.6**). This was evidenced by a decrease in the maximum, RMSE and percentage of larger surface errors of 0.084 mm ($\downarrow 7.34\%$), 0.535 mm ($\downarrow 12.38\%$) and 0.85% ($\downarrow 7.42\%$), respectively. The

larger surface errors were concentrated on the superior surfaces for both meshes (*Figures 7.5e-7.5f*). However, two key differences were observed between the meshes, where the anatomical-MB mesh had: (1) a smaller cluster of larger errors on superior surface, and (2) fewer large surface errors around the horn-attachments.

Table 7.6: Comparison of the RMSE, maximum surface error and percentage of errors above 1 mm for sub-cases C.1-C.4

Case	Mesh	Surface Error		
		RMSE (mm)	ϵ_{\max} (mm)	$\epsilon > 1$ mm (%)
C.1	Low-resolution MM	0.895±0.581	2.959	23.94
	Med-resolution MM	0.759±0.425	2.712	19.75
	High-resolution MM	0.758±0.426	2.649	17.92
C.2	Low-quality MM	0.747±0.423	2.583	17.26
	Med-quality MM	0.757±0.424	2.645	18.14
	High-quality MM	0.758±0.426	2.649	17.92
C.3	CAD-Swept MM	0.895±0.581	2.959	23.94
	CAD-MB MM	1.149±0.774	4.324	33.47
C.4	CAD-MB MM	1.149±0.774	4.324	33.47
	Anatomical-MB MM	1.064±0.682	3.789	32.62

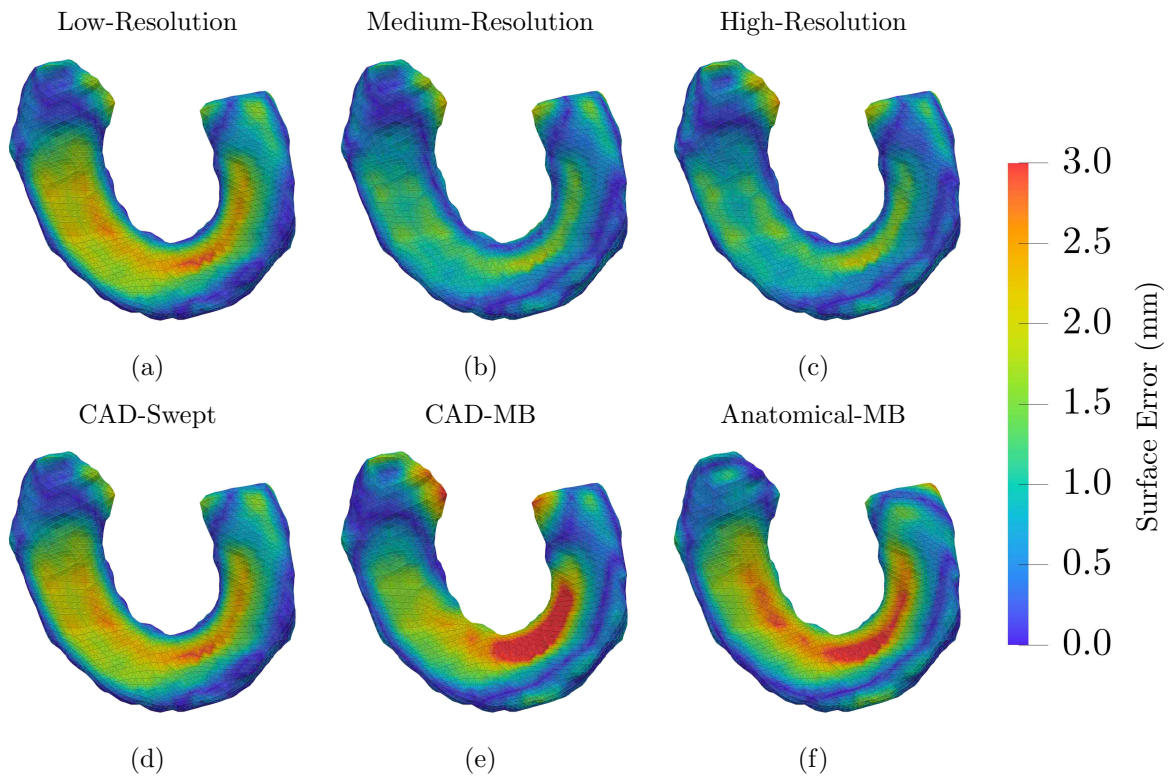


Figure 7.5: Superior view of surface errors (amalgamated above 3 mm) for the (a) low-, (b) medium- and (c) high-resolution meshes, and the (e) CAD-Swept, (f) CAD-MB and (g) anatomical-MB meshes (red is worse and blue is better)

7.3.2 Mesh-Quality Analysis

C.1: Resolution

The mesh-quality improved with increasing mesh-resolution, evidenced by a higher \tilde{J}_{avg} and \tilde{J}_{min} , and a smaller percentage of very low- and low-quality elements (**Table 7.7**). However, the increase was non-linear, where the lower resolution meshes were closer to each other than the baseline. From the low- to medium-resolution meshes, the change in the \tilde{J}_{avg} , \tilde{J}_{min} and percentage of low-quality elements were $\downarrow 0.006$, $\uparrow 0.04$ and $\downarrow 1.13\%$. In contrast, the equivalent values from the medium-resolution mesh to the baseline, were approximately $4\text{-}11\times$ larger ($\downarrow 0.06$, $\downarrow 0.22$ and $\downarrow 4.81\%$, respectively). This was further evidenced by the similarity between the external and internal mesh-quality for the lower resolution meshes compared to the baseline (**Figures 7.6a-7.6c** and **7.7a-7.7c**). The increase in low-quality elements were observed along the external singularities, i.e. the wedge, upper-rim (superior-periphery) and geometric edges of the horn-attachments. Similarly, the internal mesh-quality was lower around the internal singularities, which was more pronounced for the baseline. Overall, the internal mesh-quality was worse and less uniform for the lower resolution meshes. None of the meshes had any very-low quality elements.

C.2: Quality

The quality of the template meshes had a very minor effect on the mesh-quality (**Table 7.7**). The mesh-quality metrics were almost identical between the medium-quality mesh and the baseline, where the average difference between the metrics was 0.09%. Also, the low-quality mesh was very similar, but showed a slight decrease. The low-quality mesh had the same \tilde{J}_{min} with only a decrease to the \tilde{J}_{avg} by 0.002 and an increase in the percentage of low-quality elements by just 0.04%. Consequently, no distinguishable differences were observed for the external and internal elements, which all looked identical to the baseline (**Figures 7.6c** and **7.7c**). As such, the low- and medium-quality morphed meshes showed significant improvements to the mesh-quality with respect to their template meshes. For the medium-quality mesh, the \tilde{J}_{avg} , \tilde{J}_{min} and percentage of low-quality elements improved by $\uparrow 0.03$, $\uparrow 0.12$ and $\downarrow 2.10\%$, respectively. For the low-quality mesh, the equivalent values improved by $\uparrow 0.11$, $\uparrow 0.31$ and $\downarrow 13.94\%$, respectively, and the very low-quality elements were eliminated. Consequently, none of the meshes had any very-low quality elements.

C.3: Topology

A higher mesh-quality was observed for the CAD-Swept mesh compared to CAD-MB (**Table 7.7**). This was evidenced by more favourable mesh-quality metrics, where the \tilde{J}_{avg} and \tilde{J}_{min} were larger by 0.02 and 0.06, respectively. Additionally, the percentage of low-quality elements were 5.3% lower and there were no very-low quality elements, whereas CAD-MB had 0.17%. However, there were some improvements to the mesh-quality metrics for CAD-MB with respect to the template mesh, where the \tilde{J}_{min} and percentage of very low-quality elements improved by $\uparrow 0.06$ and $\downarrow 0.08\%$, respectively. Although, the \tilde{J}_{avg} and percentage of low-quality element worsened by $\downarrow 0.06$ and $\uparrow 4.58\%$, respectively. The low-quality elements for both meshes were observed around the external singularities (**Figures 7.6d-7.6e**). However, the singularities of CAD-MB had a greater proportion of lower quality elements, some of which had a very low-quality criterion. Despite the greater presence of lower quality external elements, the internal elements of CAD-MB had a higher quality than CAD-Swept (**Figures 7.7d-7.7e**). Additionally, CAD-MB had a larger percentage of very high-quality elements by 9.84%, which were located on the inferior surface and internally.

C.4: Shape

The anatomical-MB mesh had a higher mesh-quality than CAD-MB (**Table 7.7**). Each of the mesh-quality metrics were more favourable, where the \tilde{J}_{avg} and \tilde{J}_{min} were both larger by 0.06. Additionally, the percentage of low-quality elements were 2.83% lower and there were no very-low quality elements, whereas CAD-MB had 0.17%. Both meshes had their lower quality elements concentrated around the singularities, but these elements were poorer for CAD-MB (**Figures 7.6e-7.6f** and **7.7e-7.7f**). Additionally, both meshes had a large percentage of very high-quality elements, which were located on the inferior surface and internally. However, anatomical-MB had a greater percentage of these elements by 20.25%.

Table 7.7: Comparison of the average and minimum scaled-Jacobian, and the percentage of very low- and low-quality elements for sub-cases C.1-C.4

Case	Mesh	Element Quality				
		\tilde{J}_{avg}	\tilde{J}_{min}	$\tilde{J} < 0.1$ (%)	$\tilde{J} < 0.5$ (%)	$\tilde{J} > 0.9$ (%)
C.1	Low-resolution TM	0.85±0.08	0.57	0.00	0.00	35.00
	Med-resolution TM	0.86±0.08	0.57	0.00	0.00	35.00
	High-resolution TM	0.89±0.08	0.55	0.00	0.00	46.25
	Low-resolution MM	0.80±0.15	0.16	0.00	6.33	27.58
	Med-resolution MM	0.81±0.15	0.11	0.00	5.21	29.42
	High-resolution MM	0.87±0.10	0.33	0.00	0.40	46.58
C.2	Low-quality TM	0.75±0.21	0.01	1.25	14.37	28.75
	Med-quality TM	0.84±0.13	0.21	0.00	2.50	35.00
	High-quality TM	0.89±0.08	0.55	0.00	0.00	46.25
	Low-quality MM	0.86±0.10	0.33	0.00	0.44	46.90
	Med-quality MM	0.87±0.09	0.33	0.00	0.40	47.03
	High-quality MM	0.87±0.10	0.33	0.00	0.40	46.58
C.3	CAD-Swept TM	0.85±0.08	0.57	0.00	0.00	35.00
	CAD-MB TM	0.84±0.19	0.04	0.25	7.08	58.75
	CAD-Swept MM	0.80±0.15	0.16	0.00	6.33	27.58
	CAD-MB MM	0.78±0.20	0.10	0.17	11.67	37.42
C.4	CAD-MB TM	0.84±0.19	0.04	0.25	7.08	58.75
	Anatomical-MB TM	0.81±0.23	0.03	1.50	12.25	50.58
	CAD-MB MM	0.78±0.20	0.10	0.17	11.67	37.42
	Anatomical-MB MM	0.84±0.18	0.15	0.00	8.83	57.67

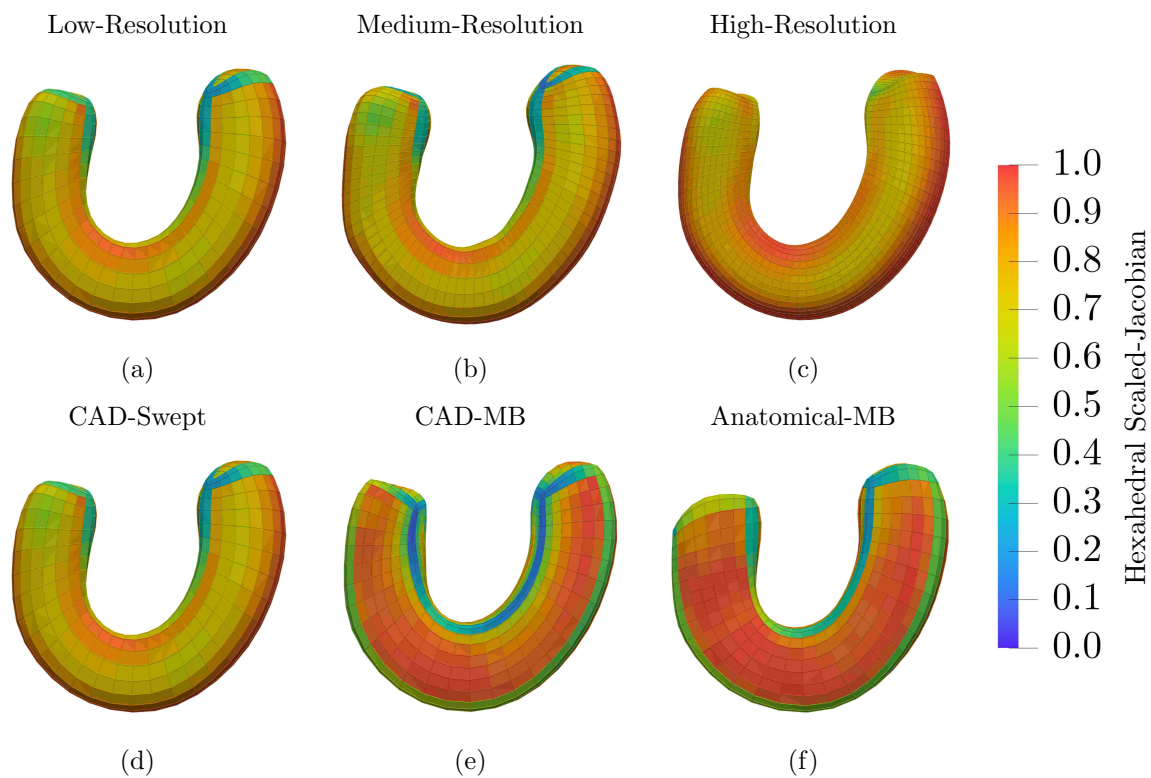


Figure 7.6: Inferior view of the hexahedral scaled-Jacobian on the (a) low-, (b) medium- and (c) high-resolution meshes, and the (e) CAD-Swept, (f) CAD-MB and (g) anatomical-MB meshes (red is better and blue is worse)

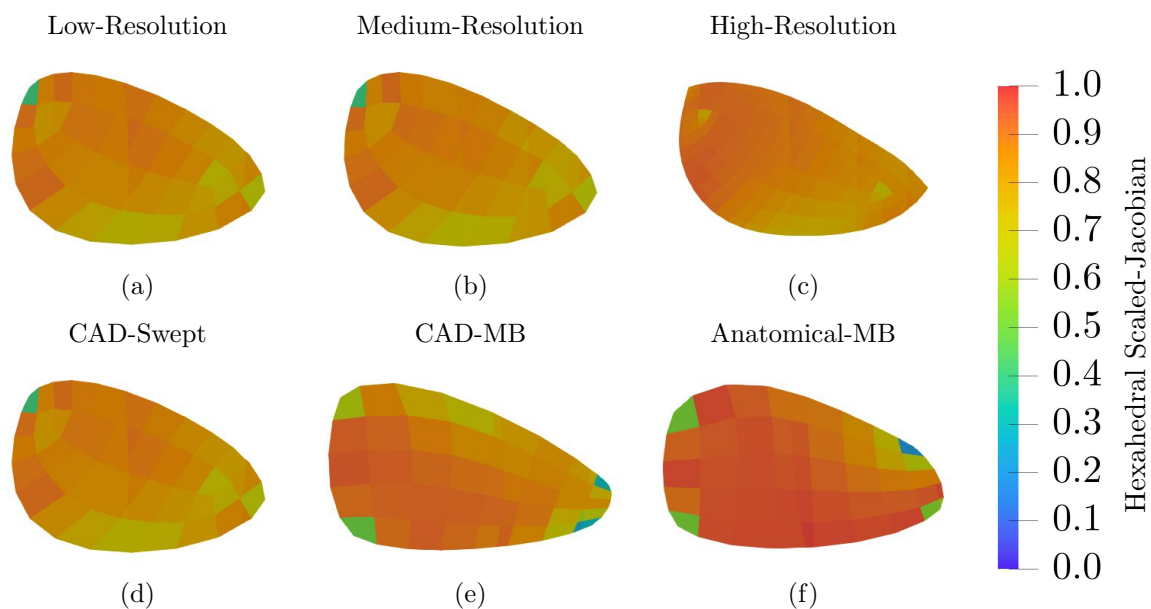


Figure 7.7: Cross-sectional view of the hexahedral scaled-Jacobian on the (a) low-, (b) medium- and (c) high-resolution meshes, and the (e) CAD-Swept, (f) CAD-MB and (g) anatomical-MB meshes (red is better and blue is worse)

7.4 Discussion

This aim of this study was to determine the sensitivity of the mesh-morphing strategy with respect to four template mesh characteristics: mesh-resolution, -quality, -topology and -shape. There are five key findings from this study. First, the mesh-morphing strategy provides performance improvements with increasing resolutions, in both the cross-section and central-axis of the meniscus geometry. Second, the quality of a valid template mesh has a less pronounced effect on the performance of the strategy, which is able to morph meshes to a similar standard of performance. Third, a greater performance is obtained when morphing a template mesh with a meniscus-specific topology, due to the support of internal singularities. Fourth, a template mesh with a shape more similar to a target geometry leads to a better performance, between generic structured topologies. Fifth, the strategy is able to improve the mesh-quality of several template meshes with different mesh-topologies and shapes, but only for those that have low-quality elements. Overall, the best performance was provided by a template mesh with the highest resolution (baseline), which had an idealised geometry and meniscus-specific topology. In contrast, the worst performance was from a template mesh with the lowest resolution, which had an idealised geometry and a generic structured topology (CAD-MB). Additionally, between the template meshes with the lowest resolution (CAD-Swept, CAD-MB and anatomical-MB), the meniscus-specific semi-structured topology provided the best performance (CAD-Swept).

7.4.1 Sensitivity to Template Mesh Resolution

Addressing *research question 1a*, the mesh-morphing strategy is quite sensitive to changes in mesh-resolution. Furthermore, answering *research question 1b*, resolution changes in both the cross-section and along the central-axis affects the strategy. However, they affect the performance metrics differently. Specifically, the surface fit improves the most with increasing central-axis resolution, whereas greater mesh-quality improvements are due to increasing cross-sectional resolutions. This indicates that an increase in resolution alone does not dictate performance improvements, but also the dimension in which they are applied. However, it is not clear which mechanism causes this phenomenon to occur for increasing central-axis resolutions.

The reason a higher resolution mesh provides a lower surface error is due to the increase in projected vertices that sample the target geometry's irregularities. During the mesh-morphing process, there are several Laplacian smoothing operations. These calculate new vertex positions

based on the current vertex position and the neighborhood of vertices surrounding them (*Figure 2.17*). As more vertices are projected onto the target geometry with increasing resolution, this provides more sample points of the surface for a better approximation during smoothing. For the lower resolution meshes, fewer neighborhood vertices are within a close proximity, which causes poorer capture of irregular features as vertices are smoothed away from those regions. In the literature, an increase in vertices provides a better approximation of a given surface, through similar mechanisms like subdivision and projection [38,212]. However, these techniques incrementally increase mesh-resolution by refining a coarse mesh to accurately represent a target surface.

The increase in resolution causing an increase in mesh-quality is a consequence of smaller hexahedral elements fitting a reduced proportion of the irregularities and high-curvature features (*Figure 7.8*). As such, each element represents a fraction of the distortion, where high-curvature features are captured with less compromise to the ideal shape of a hexahedron. This would explain why increases in cross-sectional resolution improve the mesh-quality, as the highest curvature features (e.g. wedge) occur perpendicular to the cross-section. This correlates with suggestions that region-specific density increases would improve the mesh-quality of morphed meshes [46]. However, increases in mesh-resolution have previously led to greater proportions of both higher and lower quality elements, for the automated hexahedral meshing of vascular anatomies [212].

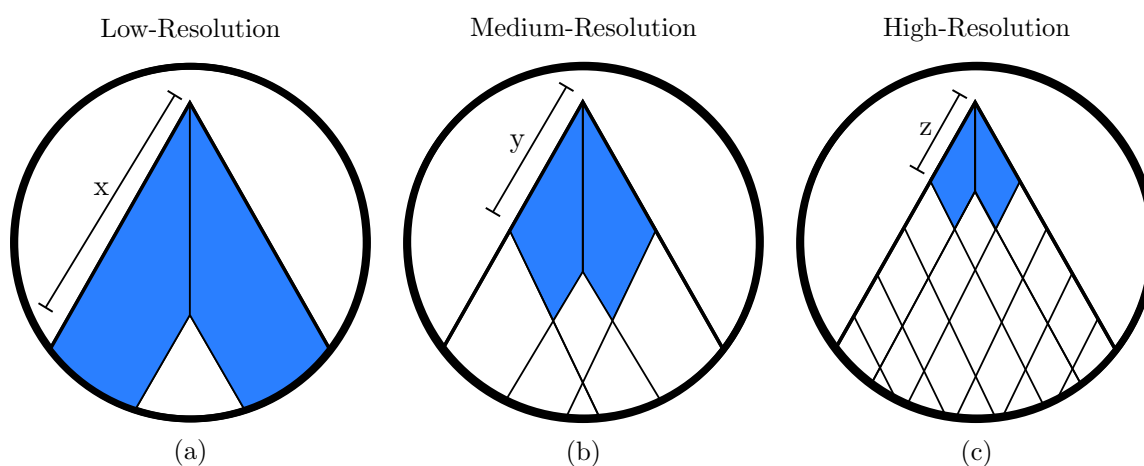


Figure 7.8: The proportion of an element occupying a surface irregularity (x , y and z) and the associated distortion (blue) for (a) low-, (b) medium- and high-resolution meshes, where the irregularities for lower resolution meshes represent greater element distortion ($x > y > z$)

7.4.2 Sensitivity to Template Mesh Quality

With respect to *research question 2a*, the mesh-morphing strategy does not appear to be affected by the presence of low-quality elements in a template mesh. Furthermore, answering *research question 2b*, the strategy is able to improve low-quality elements during the morphing operation. This can be understood as a consequence of three phases of the mesh-morphing strategy: (1) the surface optimisation (*Section 3.5*), (2) linear elastic transformation (*Section 3.6.1*) and (3) hexahedral Laplacian smoothing (*Section 3.6.2*). First, the surface optimisation phase normalises the quality of the meshes' quadrilateral elements, by eliminating the majority of mesh-quality differences. This mechanism is similar to the normalisation of correspondence differences described in *Section 5.4.2: Centreline Noise*. Second, the internal vertices are intrinsically smoothed into their optimal positions following the linear elastic transformation. Third, the final hexahedral mesh smoothing phase further normalises the internal and external vertices, reducing mesh-quality differences. As discussed in *Section 6.4.1*, although it is common for mesh-morphing strategies to cause mesh-degradation [13,28,48,58,425], there are reports this is not always the case [241,425].

7.4.3 Sensitivity to Template Mesh Topology

Regarding *research question 3a*, the mesh-morphing strategy is quite sensitive to the template mesh topology. In response to *research question 3b*, the meniscus-specific semi-structured topology leads to significant performance improvements over the generic structured topology. In general, there are two main effects of the generic structured topology compared to the meniscus-specific semi-structured topology:

1. A greater proportion of low-quality elements and larger surface errors around regions with a high-curvature (i.e. singularities and superior surface)
2. A greater proportion of very high-quality elements around regions with a low-curvature (i.e. inferior surface)

These effects are caused by the differences in the singularities and how they affect Laplacian smoothing (*Figure 7.9*). Specifically, the number of internal and external singularities they contain, their associated valence and where they are positioned with respect to the meniscus geometry. The structured meshes (multi-block) tend towards a rectangular shape with flatter

surfaces, when subject to Laplacian smoothing (*Figures 7.9a-7.9b*). This causes the larger surface errors around the wedge and superior surfaces, and the high-quality elements on the inferior surface. The higher surface error and mesh-quality are due to the mismatching and matching of the associated curvatures, respectively. Additionally, the lack of internal or strategically placed external singularities provide no support for a typical meniscus geometry. This is supported by literature which reports meshes containing only external singularities can induce greater distortions and mismatched features [188].

In contrast, this effect is less pronounced for the semi-structured meshes, as their topology smooths more naturally to the shape of a meniscus (*Figure 7.9c-7.9d*). This is due to the topological design incorporating strategically placed internal and external singularities, to support and constrain the mesh to a meniscus geometry. This is supported by literature, where internal singularities can allow meshes to represent geometries with greater complexity without inducing large distortions [35, 197, 198]. Singularities behave like point constraints [247], as they preserve features of a geometry and regions with a high-curvature [167, 188, 444, 445]. Additionally, they are also described as defects, as they can increase distortion and mesh folding [247, 444, 446]. This explains why the lowest quality elements are observed along the singularities for both the structured and semi-structured meshes. Also, the external singularities have a greater impact on mesh-quality reductions than the internal singularities for the semi-structured meshes. Consequently, the semi-structured meshes have fewer low-quality elements due to containing less than two-thirds of the external singularities of the structured meshes.

Overall, the structured meshes have a much greater proportion of very-high quality elements, for equivalent mesh-resolutions. Additionally, the multi-block mesh derived from an image-based geometry has the highest of all the meshes in this investigation. This is because the Laplacian smoothing operator favours structured meshes with respect to recovering the ideal shape of hexahedral elements. However, this comes at the expense of increasing the surface errors, causing these meshes to have the poorest surface fit of all the meshes. This further supports the evidence of the competing effects between the surface fit and mesh-quality [45, 53, 58, 307, 403] (*Section 6.4.1*). Although, the Laplacian smoothing operator is not able to overcome larger distortions for all the elements, specifically those situated around high-curvature regions. The difficulty of the structured meshes to accommodate high-curvatures causes them to contain the highest proportion of low-quality elements. In contrast, the meniscus-specific semi-structured

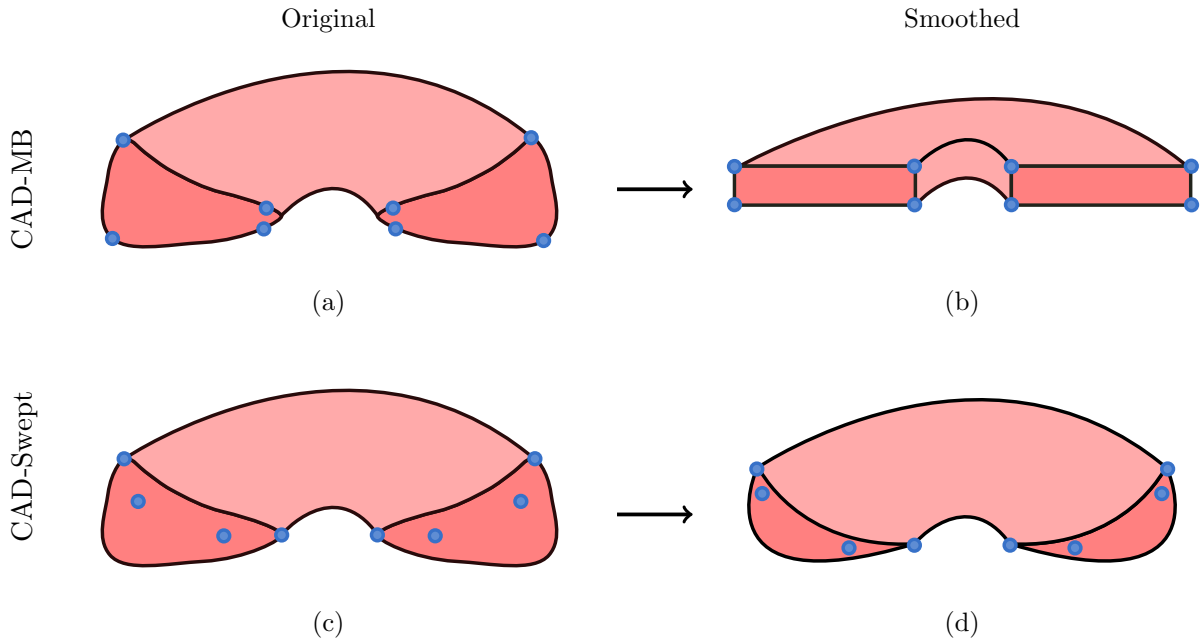


Figure 7.9: Effect of hexahedral Laplacian smoothing on shape change due to singularity locations on the hexahedral mesh derived from the CAD-based (a-b) and multi-block method (c-d)

meshes strike a more acceptable balance between the surface fit, and the percentage of low- and very high-quality elements.

7.4.4 Sensitivity to Template Mesh Geometry

To answer *research question 4a*, the mesh-morphing strategy appears to be sensitive to the template mesh geometry. In response to *research question 4b*, greater shape similarity between a template and target mesh improves the performance of the strategy. The reason for the more favourable performance by the template mesh with an in-vivo meniscus geometry is due to the greater curvature similarities to the in-vivo target geometry. This causes the initial projection to have fewer correspondence errors and mesh-degradations, i.e. tangled quadrilateral elements. Consequently, during the surface optimisation phase, there is less vertex movement as the surface mesh is closer to its global minimum. Additionally, more iterations of the surface optimisation phase are spent improving the shape of the elements, as there are fewer elements to untangle (i.e. removing invalid elements). As such, greater shape similarity causes a morphed mesh to be closer to its global minimum for a given number of iterations, leading to higher performances. This further supports the findings from *Chapter 6* and literature, that greater performance losses result from larger differences between a template and target mesh [46–48].

7.4.5 Consequences for Finite Element Modelling

There are three consequences for the development of FE models of the TFJ with respect to the choice template mesh characteristics. Overall, improved FE simulation accuracy would be expected from morphed meshes derived from templates with the following characteristics:

1. Higher resolutions
2. Meniscus-specific semi-structured topology over a generic structured topology
3. Greater shape similarities to a target geometry

Each of these characteristics would improve FE simulation accuracy as they provide lower surface errors and a greater mesh-quality, between equivalent controls. Similar to the previous discussion in *Section 6.4.3*, lower surface errors provide greater capture of the contacting interfaces. This improves the accuracy of FE simulations as contact analyses are sensitive to geometry changes [24, 193]. Additionally, these characteristics lead to higher quality elements around the wedge and other contacting interfaces, which are more beneficial for contact analyses. Consequently, the opposite of these characteristics would lead to reduced simulation accuracy. Overall, the greatest simulation accuracy would be expected from morphed meshes derived from the high-resolution template mesh, due to providing the most favourable performance metrics. The worst accuracy would be expected from morphed meshes derived from the generic structured template meshes.

7.4.6 Limitations

General

A limitation of the four investigations was that they contain a small number of template meshes, where a similar limitation has been discussed in *Section 5.4.5* (c.f. limitation 2). For the sub-case investigating mesh-resolution, more template meshes could have resolved a critical cut-off. Specifically, to determine the resolution beyond the point of diminishing performance returns. However, the resolution of meshes are often dictated by convergence studies, and therefore the decision on resolution has other influences [13, 28, 48, 58, 425]. A similar argument could be made for the mesh-quality investigation, however there was no strong differences observed from the current implementation of reducing the mesh-quality. For the mesh-topology investigation, no other topologies for meniscus meshes have been observed in the literature. However, a

third option could have been to investigate an unstructured mesh, although this would lose the benefits provided by structured and semi-structured meshes (*Section 2.4.1: Regularity*). For the mesh-shape investigation, an additional geometry with greater dissimilarities could have been used, e.g. a medial meniscus or a more irregular lateral meniscus mesh. Overall, the number of templates used in each investigation were sufficient to show the main patterns and associated differences.

Another limitation was that only one target has been used in each investigation, which has been discussed in *Section 5.4.5* (c.f. limitation 1). For the sub-cases investigating mesh-topology and mesh-shape, one nuance would be that some topologies and shapes could be better suited to particular target geometries, and vice versa.

C.2: Quality

The lower quality template meshes were not a natural artifact of the template meshes, and was induced from the addition of noise to the internal vertices of the cross-sectional geometry. Furthermore, the majority of low-quality elements existed internally, except for those at the ends of the horn-attachments. It could be argued that this is not a true representation of lower quality meshes. However, the purpose of this sub-case was to determine if the mesh-morphing strategy was sensitive to template meshes containing low-quality elements, and if their mesh-quality could be improved.

C.3: Topology

There were four limitations of this sub-case:

1. Only two generation methods were compared
2. Meshes were generated by a proficient but non-expert user.
3. Use of low-resolution meshes
4. Only two topologies were investigated

Several methods are available to create hexahedral meshes of irregular geometries and could have been explored in this investigation, particularly unstructured mesh generators. However, the multi-block method and its derivatives are the most common option used to create high-quality hexahedral meshes for biomedical simulations [28, 38]. Although, another aspect that could

have been explored was the variety of multi-block methods available, including the commercial alternatives that are often used. However, based on the user's technical skill, it was assumed similar results would have been derived.

The template meshes have not been created by a user who has experience using this tool to create state-of-the-art FE models of the meniscus. The technical skill required to reach this level of proficiency would take dedicated time, to gain complete understanding of the tool's intricacies. In addition, the multi-block method is a difficult tool to master and is very sensitive to subtle changes in the position of the multi-blocks used to create the mesh. Although, it does not take a long amount of time to be able to generate meshes to a reasonable standard and quality. As such, the strategy may have been less sensitive if the multi-block meshes were generated by an operator with more experience, which could have addressed some of the deficiencies. These deficiencies had several consequences on the types of multi-block meshes that could have been developed, e.g. mesh-resolution, -topology and -regularity. This could explain the difficulties experienced when attempting to increase the mesh-resolution, which would have been useful to compare against the high-resolution CAD-Swept mesh (best performance). Although it would have been interesting to explore higher resolution multi-block meshes, there was a limit to the number of elements these meshes could contain, as previously mentioned. Additionally, meshes with different topologies and regularities can be created using the multi-block method, however this was beyond the scope of the user's technical skill. However, the purpose of this sub-case was to determine the sensitivity of the mesh-morphing strategy to template meshes with different but commonly used topologies.

C.4: Shape

Four of the five limitations for this sub-case have already been discussed, which were: (1) the number of template meshes, (2) the number of target meshes, (3) use of a multi-block method by a non-expert user, and (4) use of low-resolution meshes. The fifth limitation was due to the poorer mesh-quality these template meshes had compared to the other meshes. However, this was unavoidable and could be attributed to limited technical skill, but this did provide further evidence that the strategy can improve the mesh-quality. Nevertheless, it is most likely that the generic structured topology simply cannot accurately capture irregular geometries without some compromise to the proportions of low-quality elements.

7.5 Conclusions

In conclusion, the mesh-morphing strategy is sensitive to a several template mesh characteristics. The strategy is most sensitive to the resolution of a template mesh, where the performance can be improved by increasing the mesh-resolution. Additionally, the strategy appears to be quite sensitive to the mesh-topology and -shape, which stem from the specificity of the topology and shape with respect to the target geometry. Specifically, a geometry-specific topology and more similar shape provide a greater overall performance. Furthermore, the quantity and positioning of singularities affect several phases of the mesh-morphing strategy. Consequently, the topology of template meshes for irregular geometries should be designed strategically, to maximise the performance of the mesh-morphing strategy. Interestingly, the quality of a template mesh has a negligible effect on the performance. The strategy is reasonably robust to the presence of low-quality elements and can improve template meshes containing such elements to a better standard. Finally, it is expected that more accurate performance improvements would be obtained from template meshes that have a higher resolution, geometry-specific topology and greater shape similarities to a target mesh. However, there is some conflict between combining some of these favourable characteristics in the current setup, e.g. high-resolution and/or geometry-specific topology with an anatomical shape. Overall, the best performances were obtained from meshes with a geometry-specific topology, regardless of mesh-resolution. The main benefit of these meshes are that they do not require large shape similarities to a target geometry and still yield the best performances. This suggests that just one template mesh with a meniscus-specific topology is required, instead of many with a generic structured topology. Consequently, these are selected as the template meshes of choice going forward.

Part III

Performance Evaluation

8 Case D: Performance Benchmark

8.1 Introduction

In part II, the sensitivity of the mesh-morphing strategy was assessed with respect to surface error and mesh-quality metrics, which are the most common methods of assessment for such algorithms. This chapter goes further, in assessing the effect the strategy has on metrics of interest from a FE simulation point of view, comparing its performance to the gold standard in terms of hexahedral meshes (multi-block method).

Aim: To compare the performance of the mesh-morphing strategy against a gold-standard procedure (multi-block method)

Hypothesis: The developed automatic mesh-morphing strategy can generate hexahedral meshes with greater speed than a semi-automatic state-of-the-art procedure, to an equivalent standard with respect to the geometric representation and finite element simulation precision.

Research Questions:

1. Does the strategy produce hexahedral meshes with an equivalent quality and accuracy with respect to a state-of-the-art method?
2. Can the strategy generate hexahedral meshes faster than a state-of-the-art method?
3. Does the biomechanics of a FE model change with the use of the strategy compared to a state-of-the-art method?

Novelty: The novel aspects of the developed mesh-morphing strategy are the ease and speed of the operation relative to the standard of the derived hexahedral mesh. Additionally, this study derives novelty from how an improvement is determined from a state-of-the-art procedure to a newly developed tool, used to generate hexahedral meshes for biomedical simulations. The assessment of a mesh-generation procedure used to create tibio-femoral joint (TFJ) simulations consists of the following:

1. Comparison of time taken to generate meshes
2. Surface-error analysis

3. Mesh-quality analysis
4. Comparison of simulated contact variables

Individually, these analyses are not novel; however, the combined assessments have not been employed to compare mesh generation procedures used to develop biomedical and clinical finite element (FE) simulations, such as the TFJ.

8.2 Methodology

The developed mesh-morphing strategy was compared against the multi-block method for 20 anatomical meniscus geometries. The selected meniscus geometries have been acquired from a variety of sources with the aim of representing a range of anatomical variations and challenging features that would test the performance of a mesh generation procedure. The time required to generate each mesh was recorded to assess the speed advantage of using the new method. The generated meshes were assessed using surface error (*Section 4.1*) and mesh-quality metrics (*Section 4.2*). The purpose of the FE models was to determine if the morphed meshes behave differently in a relevant simulation (*Section 4.5*). The analyses of the contact variables were used to determine if the use of the mesh-morphing strategy would lead to the negative outcome of misleading researchers or clinicians from inaccurate simulation results and cause incorrect conclusions to be drawn.

8.2.1 Target Meshes

Target Meniscus Geometry Selection

A total of 47 meniscus geometries were acquired from a range of sources (*Figure 8.1*). Each geometry was categorised into sub-groups based on the presence of particular features and associated descriptors (*Table 8.1* and *Figure 8.2*):

In this investigation, 20 meniscus geometries (10 lateral; 10 medial) were analysed. The aim was to identify and incorporate the most challenging geometries representative of each group, with a variety of overlap and not to include several geometries that were too similar to each other. Additionally, the meniscus geometries represent anatomies from patients that have degenerative and non-degenerative conditions, which yields particular differences in the shape of the menisci. Furthermore, the geometries have been derived from a range of different MRI and segmentation

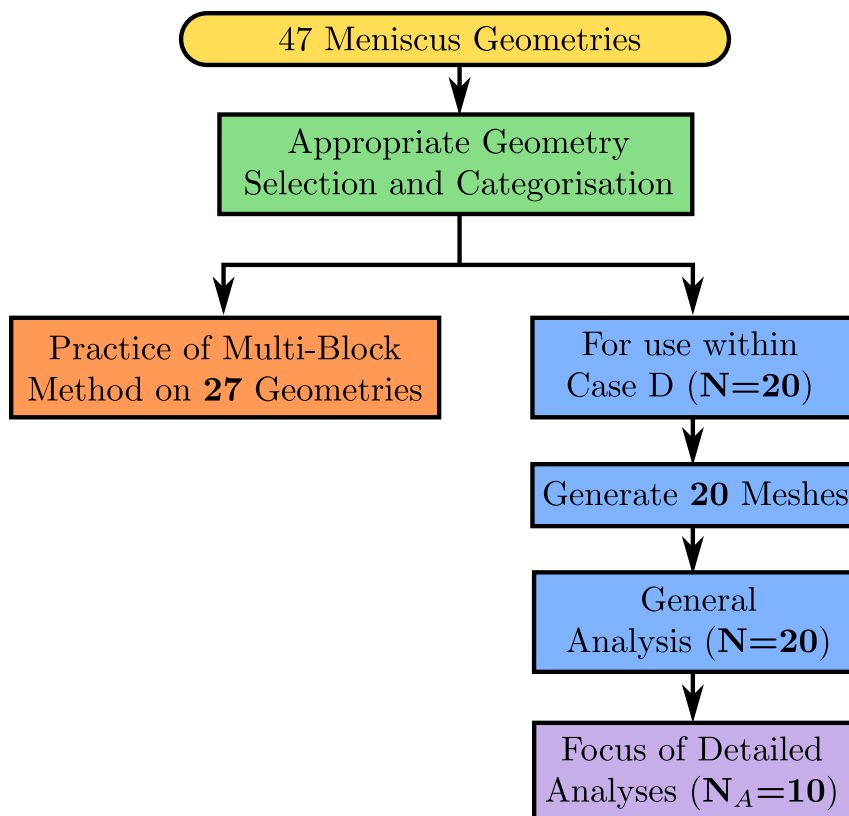


Figure 8.1: Flow-chart of what proportion of the 47 geometries were used for training/practice of the multi-block method and used in Case D

Table 8.1: Features and their descriptors used to classify the difficulty of generating hexahedral meshes from target geometries

Rank	Feature	Descriptor	
		Most Challenging	Least Challenging
1	Semi-Lunar Shape	Non-Semi-Lunar	Semi-Lunar
2	Anatomical Shape	Blocky	Natural
3	Radial/Inferior Bulge	Large	Small
4	Horn-attachments	Small	Large
5	Tapered Shape	Tapered	Non-Tapered
6	Meniscal Tears	Torn	None
7	Surface	Irregular	Smooth
8	Shape	Irregular	Smooth
9	Extrusion	Extruded	Non-extruded
10	Wedge Angle	Small	Large

procedures. These levels of geometric and acquisition variances are important, to ensure the strategy is evaluated against a range of appropriate meniscus geometries that are representative of what different biomedical engineers would derive. The remaining sub-set of 27 geometries were selected for training and practice of the multi-block method to gain a level of proficiency using the tool to create high-quality hexahedral meniscus meshes.

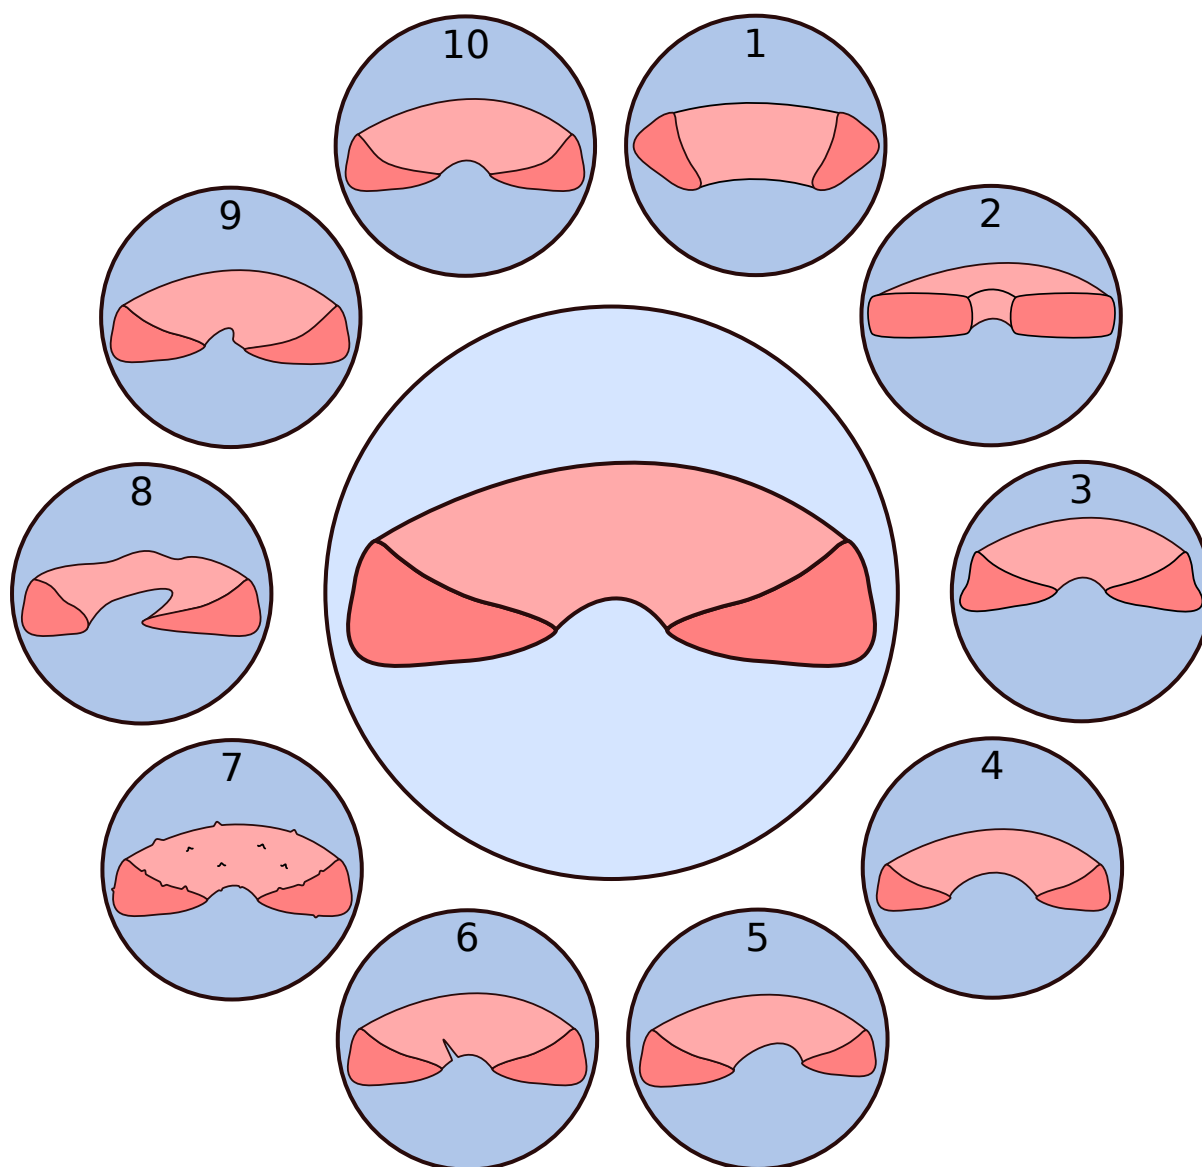


Figure 8.2: The challenging aspects of the ten features used to rank target geometries on their difficulty (the least challenging baseline shown is in the middle)

Identification of Challenging Features

A sub-set of meniscus geometries have been selected to assess the performance of the mesh-morphing strategy. The sub-set consists of menisci which have features that have previously been known to cause the mesh-morphing strategy challenges. The 10 features have been determined based on previous analyses and a review of the known geometric differences of menisci within a given population that may challenge the mesh-morphing strategy (*Table 8.1* and *Figure 8.2*).

In case B (*Chapter 6*), sensitivity analyses of the strategy with respect to shape variation were performed. The findings from this study indicate that there are two key aspects of a target

geometry that can affect the performance of the strategy:

1. Large global curvature differences
2. Small wedge angles

Global curvature changes were found to cause difficulties for the mesh-morphing strategy and could be expanded further to include the following aspects:

1. Irregular Shape
2. Irregular Surface
3. Tapered Geometry
4. Meniscal Tears

Although irregular surface features are localised curvature differences, in case A (*Chapter 5*), surface irregularities were smoothed and caused performance losses. Similarly, the presence of meniscal tears may present similar challenges to irregular surface features.

Small wedge angles were found to lead to poorer performances than target geometries with larger wedge angles and can be thought as a special case of curvature differences. Additionally, this would include meniscus extrusion which is an aspect of degeneration and leads to a reduction in the wedge angle of meniscus structures.

Selection of Geometries Based on Challenging Features

The 47 geometries were ranked to determine how many challenging features they possess. The final set was chosen to represent a diverse range of challenging features and various combinations of those features. The most challenging aspect of each feature was ranked between 1-10, with 10 being the most challenging and 1 being the least challenging (*Table 8.1* and *Figure 8.2*). The least challenging aspect of each feature was ignored. The difficulty of generating a hexahedral mesh from the 47 geometries was then determined from the ratio of challenging features and the degree of difficulty each feature represents.

Therefore, the maximum difficulty a geometry can achieve is 5.5, if all challenging features are present. The lowest difficulty a geometry can achieve is 0, if there are no challenging features. The lateral and medial geometries were then ordered based on their difficulty rating, with the 10 most difficult geometries selected within each group (lateral; medial).

To ensure that an adequate representation of each challenging feature was evaluated, the similarities between each geometry was evaluated and geometries with similar feature sets were replaced.

Finally, the set of geometries were manually reviewed and if any features were not being represented they were manually selected and replaced the lowest difficulty geometries. A sub-set of 10 geometries representing the variety of each group (lateral; medial) were selected for detailed analyses with the same method.

8.2.2 Template Mesh

The low-density mesh used in case C (*Chapter 7*), with 1,200 hexahedral elements, was used throughout this investigation. There are 40 elements in the cross-section of the template mesh. The template mesh was generated using 30 incremental sweeps across a 240° arc. The 240° arc was used to represent an idealised centreline for the geometry of the template mesh. The low-density mesh has been selected to minimise confounding variables, due to the difficulties of generating multi-block meshes of menisci with more than 1,000 hexahedral elements.

8.2.3 Generation of Meniscus Meshes using the Multi-Block Method

The total number of hexahedral elements and the proportion of cross-sectional to transverse hexahedral elements was matched to the template mesh. This was to ensure that the multi-block meshes were the same density as the morphed meshes and to negate the effects of mesh density on the derived contact variables.

To create an equivalent number of elements to the morphed meshes and to match the cross-sectional to transverse element distributions, a particular subdivision scheme was attempted where possible. The aim when generating the multi-block meshes was to create 10 blocks encompassing each meniscus geometry. This would be achievable with the cross-sectional axes being set to have 5 and 8 subdivisions with respect to the height and width of the meniscus and the transverse axis being set to 3 subdivisions.

8.2.4 Generation Times

By comparison to a multi-block method, there are fewer manual interactions required from the user with the use of the developed morphing strategy (*Figure 8.3*). Due to the automatic

nature of the mesh-morphing strategy, a user is only required to set the parameters, ensure input data is valid and run the operation. The first step of the mesh-morphing strategy is to generate centrelines from the template and target, or pre-existing data can be used. However, there is currently no built-in method to check the generated centrelines are valid representations of the geometry, such as being too short. Therefore, they can be generated by the strategy automatically then validated by the operator before executing the remainder of the mesh-morphing strategy. This strategy has been summarised in *Figure 8.3a*. In contrast, the multi-block method requires several stages of decision-making and experience of the intricacies of how the tool creates hexahedral meshes. The most demanding aspects of this tool is understanding how to create multi-block grids, where to manually position their vertices and why to position them at regions close to the surface, to appreciate what effect this will have on a generated hexahedral mesh. The less demanding aspect is the choice of element size or subdivision scheme. However, this can still cause an invalid mesh to be generated and therefore is not always a straightforward aspect of the method. The steps required to use the employed multi-block method has been summarised in *Figure 8.3b*.

Due to the differences of the user-interaction requirements and positioning in their software development cycles, it is difficult to accurately assess and compare the speed of the methods. The multi-block method has a graphical user interface (GUI) which not only allows these interactions but enables a user to view the immediate consequences of these choices. The mesh-morphing strategy is still in the early prototype phases and does not have a GUI. Due to not having undergone end-stage development and finalisation, it would be difficult to navigate for someone either not well-versed in the test applications or not familiar to a command shell environment and the necessary helper-functions required to visualise and intuitively operate the strategy.

The speed of the mesh-morphing algorithm was assessed based on how the finalised version would operate. Centreline computation times were included, however, required adjustments were reported but not accumulated into the final time which can be automated (*Chapter 3*). As described in *Section 3.2*, the finalised version of the centreline generation algorithm will be fully automated and not require manual adjustments to address short length issues. The speed of the multi-block method has been measured from the point geometries have been imported to the point where a valid mesh of adequate element density has been exported. Times acquired from generating the morphed and multi-block meshes were used to determine an estimation

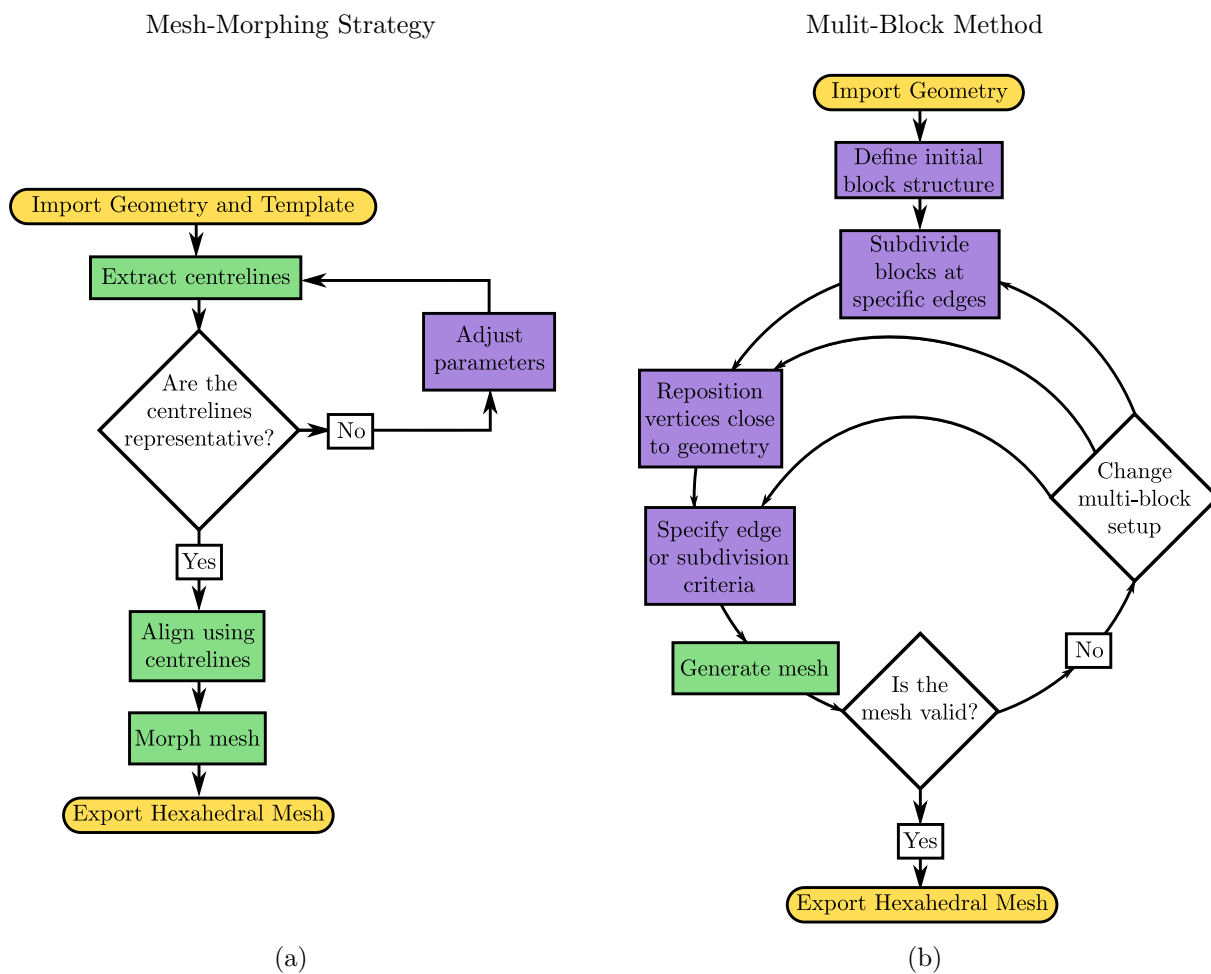


Figure 8.3: Procedural comparisons between (a) an automatic mesh-morphing strategy and (b) a semi-automatic multi-block method. Automatic processes shown in green and operator required processes shown in purple.

of how long each would typically take to operate for meniscus geometries. The times for the mesh-morphing strategy have been measured to the millisecond, and the multi-block method to the second.

8.3 Results

8.3.1 Generation Times

The mesh-morphing strategy generated each mesh with greater speed than the multi-block method (*Table 8.2* and *Figure 8.4*). Additionally, the morphing strategy was able to generate meshes with a more consistent and predictable time-frame, unlike the multi-block method.

Table 8.2: Comparison of the average, standard deviation, minimum and maximum generation times between the mesh-morphing strategy and multi-block method

	Generation Times (seconds)		
	Avg. \pm Std.	Min.	Max.
Morphed Meshes	24.82 \pm 0.56	23.87	25.97
Multi-Block Meshes	1586 \pm 1994	279	8040

All the morphed meshes were generated quite fast, in under 26 seconds. The majority of multi-block meshes (75%) were generated within 23 minutes. However, two meshes took over an hour and 45 minutes with one of them taking 2 hours and 14 minutes. The shortest generation time amongst the multi-block meshes was 4 minutes and 39 seconds. The standard deviation was roughly half a second for the morphing strategy, compared to over half an hour for the multi-block method. Consequently, there was a lower variance between generation times for the mesh-morphing strategy than the multi-block method. This was further supported by the bimodal distribution for the multi-block meshes (*Figure 8.4d*). It was more challenging to produce some geometries in a timely manner than others. Common features for geometries that took over 25 minutes included: surface and shape irregularities, thin-regions (extrusions) and tapered horns. In contrast, the morphed meshes had approximately a perfect normal distribution, where the deviations could be attributed to differences in target mesh vertex densities (*Figure 8.4c*). The number of target vertices had an effect on the time taken to generate each hexahedral mesh (*Figure 8.5a*). Longer computation times were observed when the quantity of target vertices increased. Overall, the mesh-morphing strategy had a computational complexity that was slightly more $\mathcal{O}(\log N)$ than $\mathcal{O}(N)$. Unlike the mesh-morphing strategy, the generation times for the multi-block method were not affected by the number of target vertices but by particular features (*Figure 8.5b*).

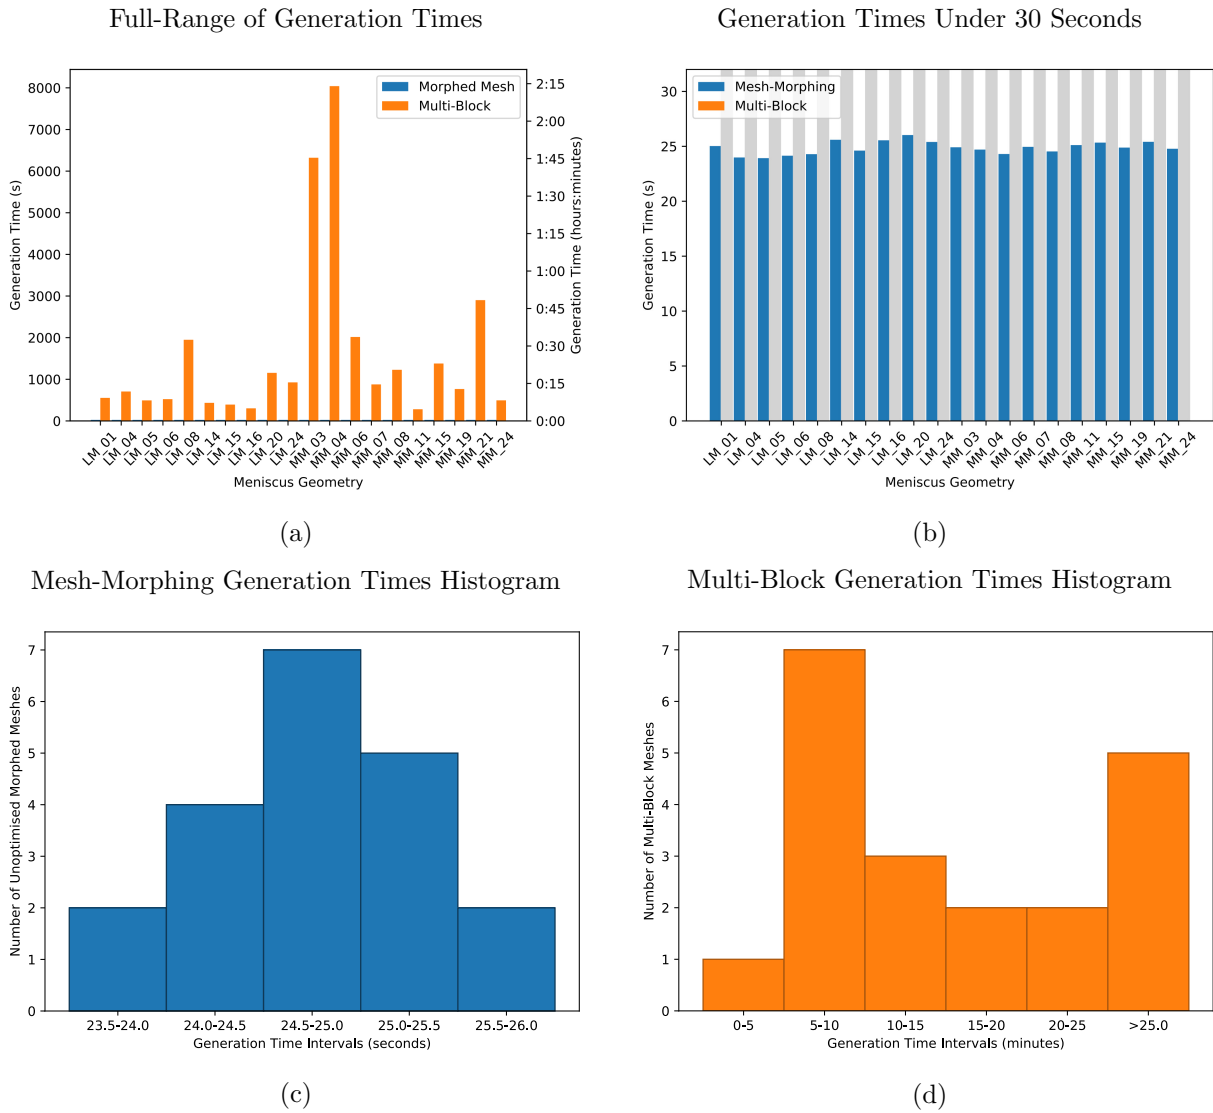


Figure 8.4: Mesh generation times for (a,b) individual geometries and the distribution of times for the (c) morphed and (d) multi-block meshes

8.3.2 Surface Error Analysis

Lower surface errors were achieved using the multi-block method compared to the morphing strategy (*Table 8.3* and *Figure 8.6*). All multi-block meshes had lower RMSEs than the morphed meshes (*Figure 8.6a*). Additionally, the multi-blocks meshes achieved lower maximum surface errors than the morphed meshes (*Figure 8.6b*). Meshes derived from both methods had reasonably consistent surface errors with respect to the RMSE, which showed little variation. This was evident from the small standard deviations and were similar between methods.

The mesh-morphing strategy generated a large proportion of meshes (75%) with RMSEs below 0.96 mm. Additionally, the most common RMSE range was 0.8-0.9 mm, representing 60% of the




Figure 8.5: Comparison of the number of target vertices against the time taken to generate each hexahedral mesh for the (a) mesh-morphing strategy and (b) multi-block method

Table 8.3: Comparison of the average, standard deviation, minimum and maximum surface errors between the mesh-morphing strategy and multi-block method

	RMSE (mm)			Max. Error (mm)		
	Avg. \pm Std.	Min.	Max.	Avg. \pm Std.	Min.	Max.
Morphed Meshes	0.916 \pm 0.073	0.833	1.111	2.887 \pm 0.532	2.154	3.885
Multi-Block Meshes	0.193 \pm 0.069	0.099	0.356	1.318 \pm 0.432	0.524	2.231

geometry cases (**Figure 8.6c**). For the multi-block meshes, 75% had their RMSEs below 0.273 mm. Also, the most frequent RMSE range was between 0.15-0.25 mm, which represented 50% of the geometries (**Figure 8.6d**). The majority of morphed meshes (75%) had a maximum surface error below 3.30 mm. In contrast, 75% of the multi-block meshes had a maximum surface error below 1.60 mm. Additionally, the maximum surface error histogram shows that 75% of the multi-block meshes had a maximum surface error range of 1-2 mm (**Figure 8.6f**).

The RMSE and maximum surface error histograms for the multi-block method were more normally distributed than the corresponding morphed mesh histograms indicating the multi-block method was more robust for a range of meniscus geometries. For the morphing strategy, the histograms showed a similar and strong positive skewness indicating the mesh-morphing strategy was challenged more by particular meniscus geometries. Common features for geometries with high surface errors were: surface and shape irregularities, thin-regions (extrusions) and tapered horns. These are the same features that negatively affected the generation times for the multi-block method.

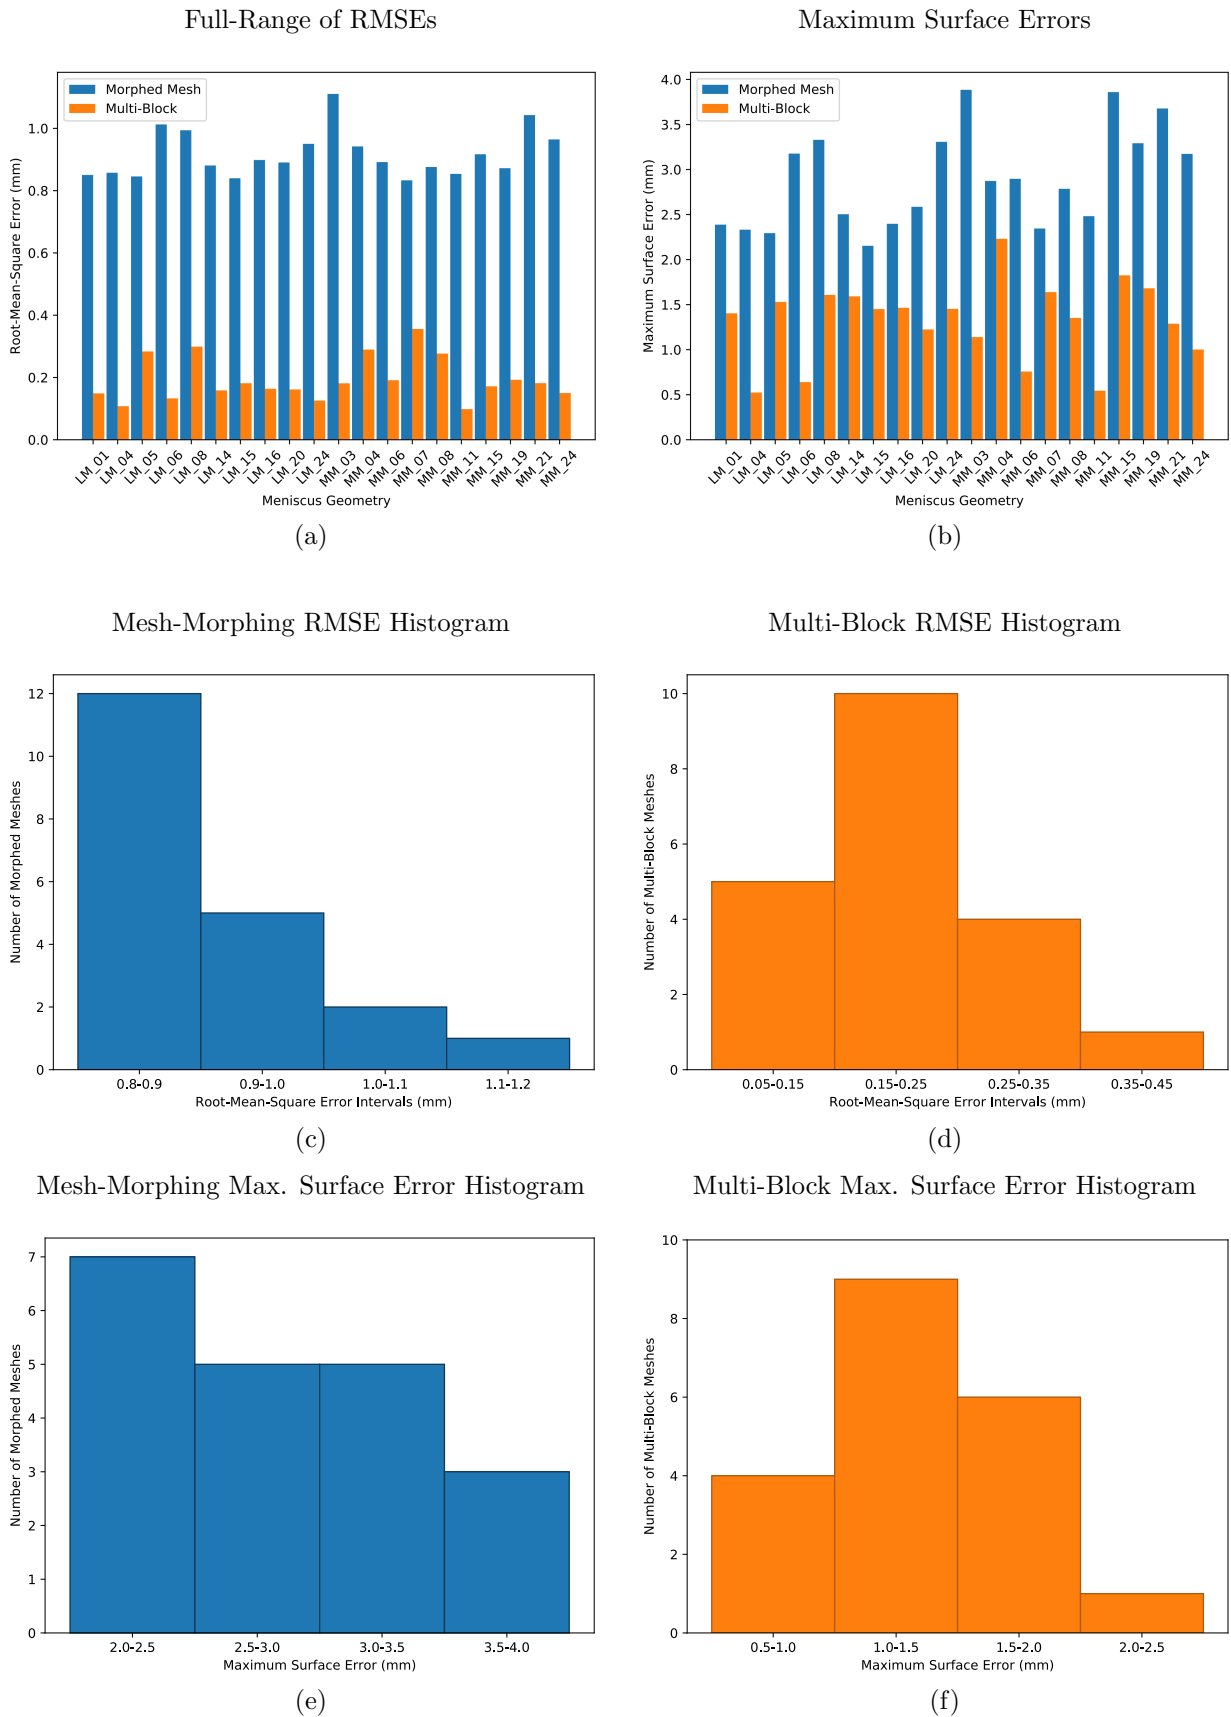


Figure 8.6: Surface errors for (a,b) individual geometries and associated distributions for the (c,e) morphed and (d,f) multi-block meshes

The majority of morphed meshes (80%) had a slightly larger volume than the multi-block meshes, which always underestimated the target mesh volume (*Table 8.4* and *Figure 8.7*).

Table 8.4: Comparison of the average, standard deviation, minimum and maximum absolute volume errors between the mesh-morphing strategy and multi-block method

	Absolute Volume Error (%)		
	Avg. \pm Std.	Min.	Max.
Morphed Meshes	4.31 \pm 3.73	0.04	15.0
Multi-Block Meshes	3.05 \pm 0.87	1.56	4.89




Figure 8.7: (a) Mesh volume comparison between morphed, multi-block and target meshes, (b) target volume difference for each morphed and multi-block mesh

8.3.3 Mesh-Quality Analysis

The mesh-morphing strategy produced higher quality meshes for the majority of meniscus geometries (*Table 8.5* and *Figure 8.8*). Very low-quality elements ($\tilde{J} < 0.1$), which are the most likely to cause problems during the numerical solution of the finite element equations, are far more numerous in the multi-block meshes (*Table 8.5* and *Figure 8.8a*). Only two morphed meshes contained very low-quality elements compared to 19/20 for the multi-block meshes (*Figure 8.8a*). For 75% of the multi-block meshes, the quantity of very low-quality elements were below 1.67%. The one multi-block mesh that had zero very low-quality elements was derived from a meniscus geometry with a very block-like shape (LM_05).

Low-quality elements ($\tilde{J} < 0.5$) are less likely to cause problems for finite-element simulations than very low-quality elements, but are still undesirable and ideally should be limited. The mesh-morphing strategy was able to generate fewer low-quality elements for 80% of the meniscus

Table 8.5: Comparison of the average, standard deviation, minimum and maximum percentage of very low and low-quality elements between the mesh-morphing strategy and multi-block method

	Very Low-Quality Elements (%)			Low-Quality Elements (%)		
	Avg. \pm Std.	Min.	Max.	Avg. \pm Std.	Min.	Max.
Morphed Meshes	0.04 \pm 0.12	0.00	0.50	5.86 \pm 1.16	4.75	9.00
Multi-Block Meshes	1.11 \pm 0.60	0.00	2.08	9.35 \pm 2.57	4.75	12.92

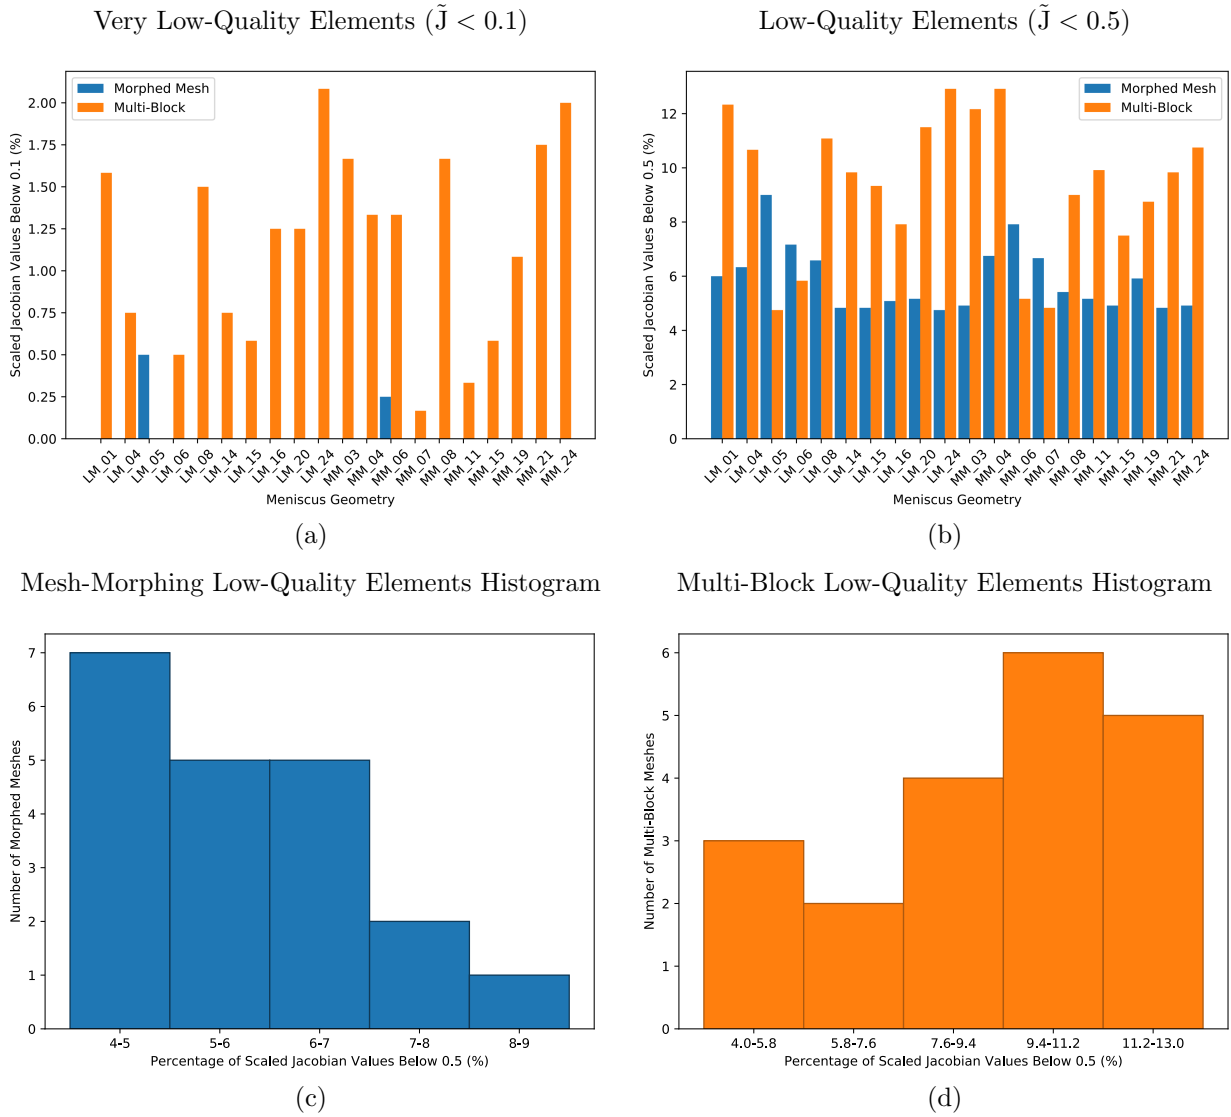


Figure 8.8: The number of (a) very low and (b) low-quality elements for individual geometries and the low-quality element distributions for the (c) morphed and (d) multi-block meshes

geometries (*Figure 8.8b*). The percentage of low-quality elements for 75% of the morphed and multi-block meshes were below 6.67% and 11.17%, respectively. Four multi-block meshes (20%) that contained fewer low-quality elements than the corresponding morphed meshes were derived from block-like shapes (LM_05, LM_06, MM_06 and MM_07). These meshes caused the low-

quality histogram for the multi-block meshes to have a bimodal distribution (*Figure 8.8d*). This indicated that the processing of these four block-like meshes was different to the other geometries.

8.3.4 Finite Element Analysis

Convergence

All 40 simulations derived from this investigation converged. The convergence rates for the morphed mesh simulations were much faster than the multi-block simulations (*Table 8.6* and *Figure 8.9*). Also, the convergence rates were more variable and less predictable for the multi-block simulations. This was evident by a standard deviation of 19 minutes compared to 5 minutes for the morphed mesh simulations.

Table 8.6: Comparison of the average, standard deviation, minimum and maximum convergence times between the mesh-morphing strategy and multi-block method

	Convergence Times (HH:MM)		
	Avg. \pm Std.	Min.	Max.
Morphed Meshes	00:13 \pm 00:05	00:08	00:34
Multi-Block Meshes	00:40 \pm 00:19	00:20	01:27




Figure 8.9: (a) Individual and (b) overall convergence times for simulations derived from the morphed and multi-block meshes

Contact Area

The difference between the contact areas derived from the morphed and multi-block mesh simulations were quite large (*Table 8.7* and *Figure 8.10*). Additionally, the percentage differences

were larger on the tibial cartilage than the femoral cartilage. The majority of percentage differences were less than 33% for both the tibial and femoral cartilage (*Figures 8.10e-8.10f*). Contact area differences were more variable on the tibial cartilage than femoral (*Table 8.7*).

Table 8.7: Comparison of the average, standard deviation, minimum and maximum absolute contact area differences between the mesh-morphing strategy and multi-block method

	Absolute Contact Area Difference (%)			Absolute Contact Area Difference (mm ²)		
	Avg. \pm Std.	Min.	Max.	Avg. \pm Std.	Min.	Max.
Tibial	35.16 \pm 35.66	1.40	128.19	31.65 \pm 25.01	1.41	82.23
Femoral	20.74 \pm 15.65	0.71	61.85	25.43 \pm 16.62	1.01	50.18

The largest contact area difference on the tibial cartilage was from the simulation with LM_24, which also had the largest total difference (128%). For the femoral cartilage, the largest difference was from the simulation with MM_07 (62%). In addition to the simulation with LM_24, another simulation with a total difference over 100% was the simulation with LM_14 (115%). The majority of contact area differences derived from the morphed mesh simulations were larger (positive) than the corresponding multi-block simulations.

8.3.5 Detailed Analysis

Ten geometries were used to provide a more detailed, both qualitative and quantitative, analysis of the strategy performance. All performance metrics from the 20 geometries are available in *Appendix B*.

Generation Times

Variations in geometric features did not appear to have an effect on the processing speed of the mesh-morphing strategy. For multi-block method, the presence of particular geometric features had the most significant impact on the generations times.

Geometries MM_03 and MM_04 had both an irregular surface and thin-regions showing signs of extrusion. These two meshes took the longest time to generate and were much longer than the others for the multi-block method. Irregular surfaces presented difficulties for the multi-block method, where manual placement of vertices were required around surface irregularities and caused the procedure to become error-prone. The simpler multi-block structures that would be sufficient for more regular meniscus geometries were inadequate for geometries with irregular surfaces, tapered-regions and thin-regions.

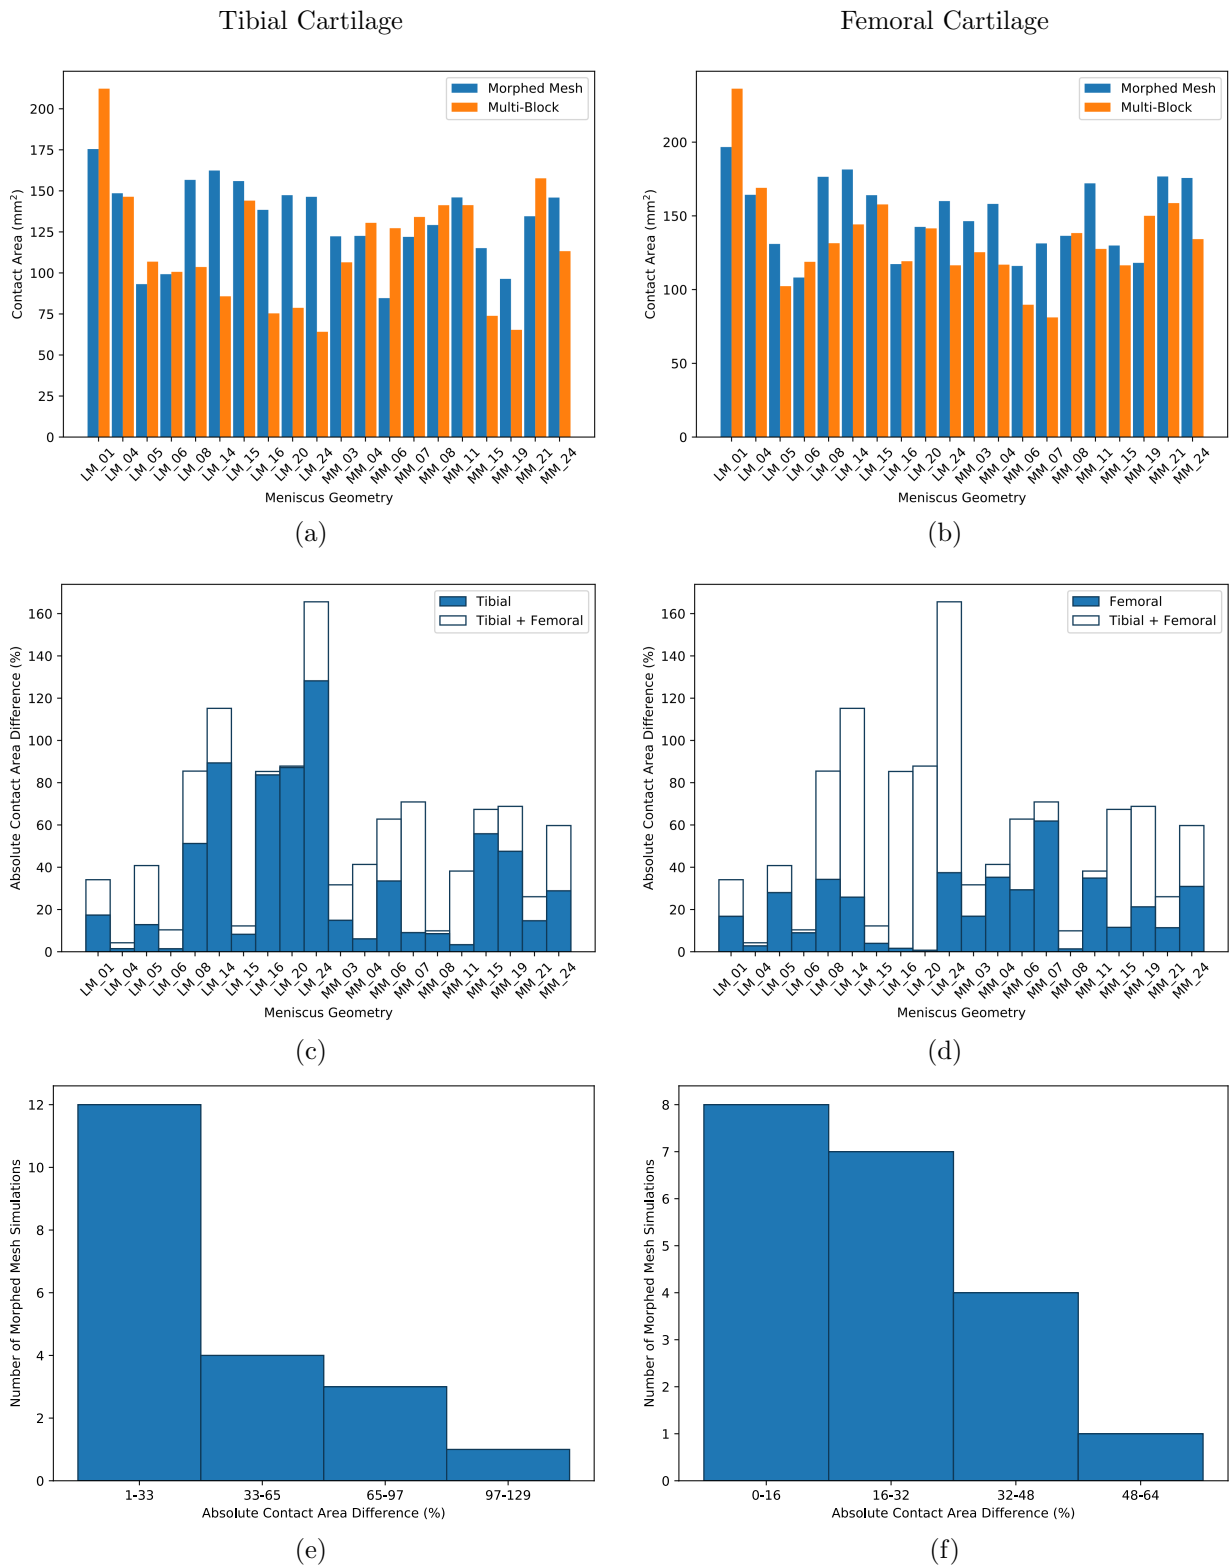


Figure 8.10: The contact areas for individual simulations on the (a) tibial and (b) femoral cartilage. The absolute contact area differences (c-d) and distributions (e-f) for the morphed mesh simulations on the (c,e) tibial and (b,f) femoral cartilage, respectively

Surface Error

For each method, there were particular geometric features that caused higher surface errors than others. The highest surface errors associated with the mesh-morphing strategy were located around larger features that became smoothed and rounded, causing higher errors in large regions of the geometry. For the multi-block method, features that cause higher surface errors were more subtle and smaller due to the greater overall surface-fit.

The mesh-morphing strategy imparted a smoothing and rounding effect to the overall geometry of the menisci and caused features to become less defined, producing higher overall surface errors in the process. The morphed meshes tended to show this effect and high surface errors in three main regions: (1) concave surfaces, (2) angular features and (3) thin features (**Figure 8.11**). Due to the smoothing and rounding effect, the morphed meshes were generally thicker and slightly larger than the target geometry. The highest surface errors were found on the superior surface in the medial region of the meniscus where the geometry was most concave (**Figure 8.11a-8.11b**). Additionally, concave regions of the geometries outside the superior medial area exhibited larger surface errors. Compared to the target geometries, these regions in the morphed meshes were less concave. Angular features were often located at the corners of horn-attachments and geometric irregularities. The corners of the horn-attachment sites often showed higher surface errors, especially where they had been clipped and had a sharp contrast between the boundaries of curved and flat surfaces (**Figure 8.11c**). Additionally, large-scale surface to small-scale geometric irregularities posed a similar difficulty and showed surface errors similar to the corners of the horn-attachment sites (**Figure 8.11d-8.11e**). Another feature that caused higher surface errors were thin-regions and extrusions, which were often more rounded, larger and slightly receded than the target geometry (**Figure 8.11f**).

The multi-block method caused higher surface errors in curved and irregular meniscus geometries, which did not lend themselves to block definitions that enable the creation of more angular and faceted geometries in the derived meshes. The multi-block meshes tended to be more faceted and cropped small rounded and irregular features. Higher surface errors for the multi-block meshes occurred from three aspects: (1) irregular features (e.g. undulations and protrusions), (2) curved geometric edges and (3) multi-block artefacts/defects (**Figure 8.12**). When partially captured, these features became sharper or point-like due to being defined by a small number of elements and lacked the smoothness defined by the target geometry. Where

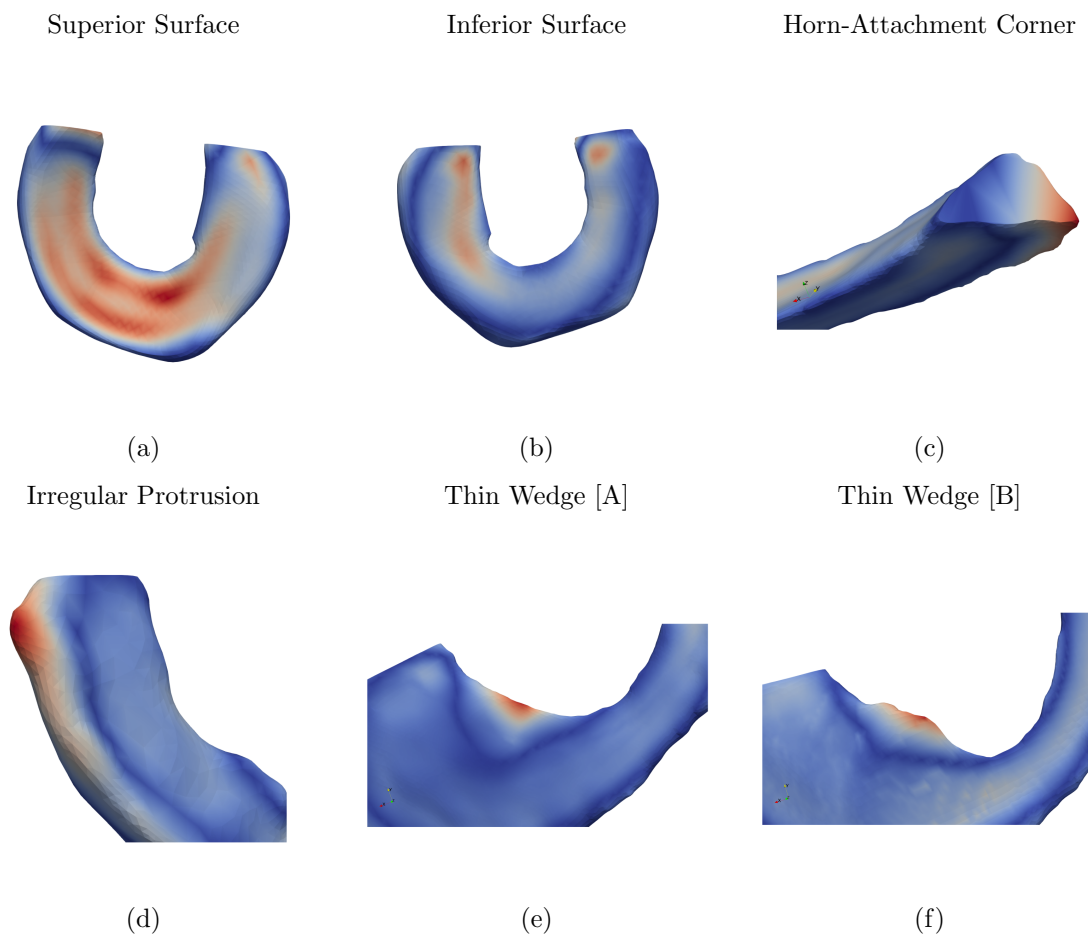


Figure 8.11: Regions and features of the target geometries which caused (a,c-f) higher and (b) lower surface errors for the morphed meshes (blue=lower surface error; red=higher surface error)

both curved edges and irregular features existed, the surface errors were exacerbated and appeared to compound into larger overall errors. The multi-block method sometimes imparted particular artefacts or defects where vertices of a hexahedral mesh, which were automatically generated from the sub-division and meshing process, appeared to slip from the surface. These regions formed dents in the surface of a hexahedral mesh and caused surface errors in these regions. These dents have occurred on the rims of the meniscus geometry between differently curved surfaces and created highly angular concave features in the derived meshes (*Figure 8.12f*).

Mesh-Quality

There were several geometric features where low and high-quality elements commonly occurred for each method. Meshes derived from the mesh-morphing strategy had higher-quality elements towards regions with higher curvature. The multi-block meshes had higher-quality elements in

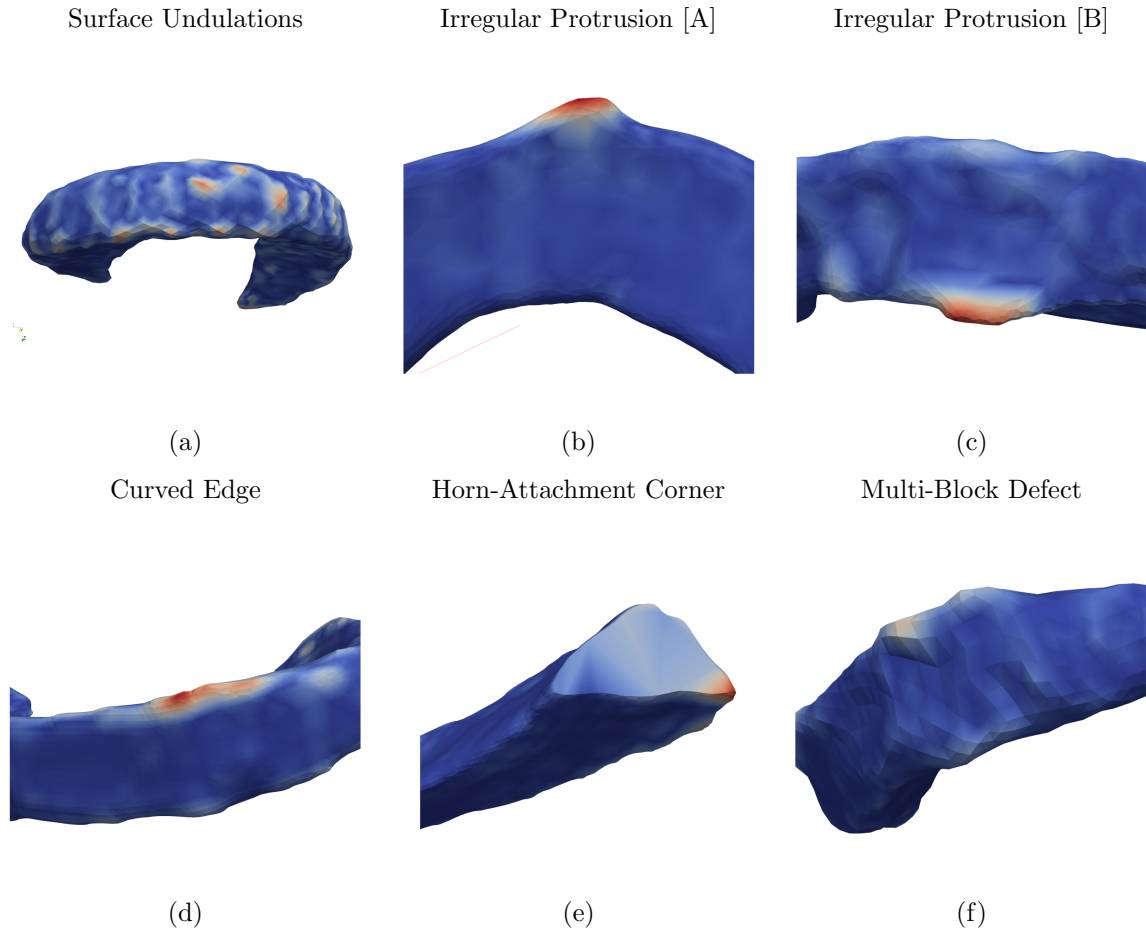


Figure 8.12: Regions and features of the target geometries which caused high surface errors for the multi-block meshes (blue=lower surface error; red=higher surface error)

regions with less curvature. Meshes from both methods had lower-quality elements in areas where there was a sharp change in curvature between different regions. For all the meshes, the singularity and edge regions of the original template or multi-block structure caused the lowest quality regions and affected both methods similarly (*Figure 8.13-8.14*).

The external element-quality of the morphed meshes were generally affected by five geometric aspects: (1) mesh singularities, (2) concave surfaces, (3) flat and convex surfaces, (4) horn-attachments and (5) thin regions (*Figure 8.13*). The highest quality regions were within close proximity to concave surfaces. Regions towards the horn-attachments tended to have less curvature than the medial regions and had lower-quality elements. A mesh-morphing specific element-quality issue occurred near some of the thin sections, where the elements appeared ‘pinched’ or squeezed into a relatively narrow area, causing lower quality elements than the surrounding (*Figure 8.13d-8.13e*). The wedge region of the morphed meshes were not affected by low-quality elements to the same degree as multi-block meshes.

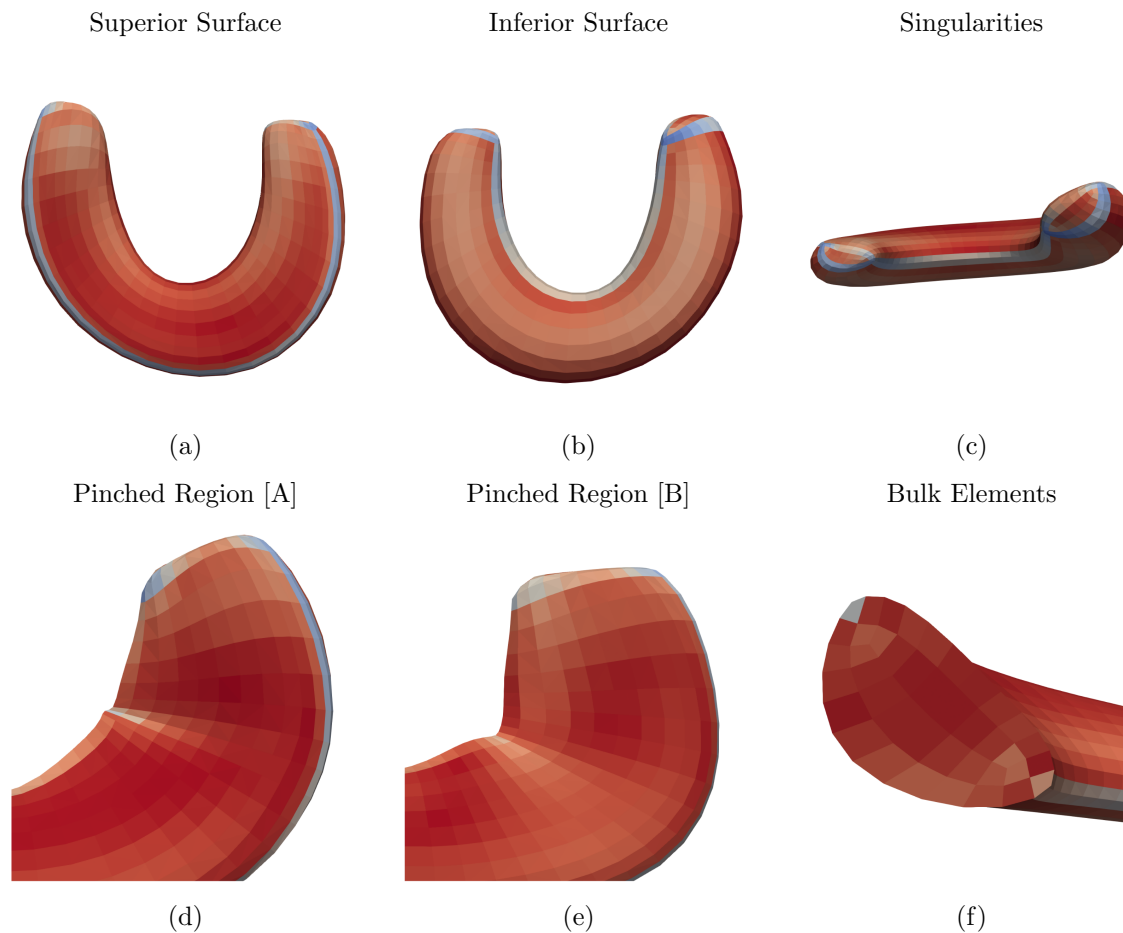


Figure 8.13: Regions and features of the morphed meshes depicting variations of higher and lower quality elements (blue=lower element-quality; red=higher element-quality)

The external element-quality of the multi-block meshes were affected by six aspects: (1) mesh singularities, (2) multi-block structure, (3) concave surfaces, (4) convex surfaces, (5) wedge region and (6) multi-block defects/artefacts (*Figure 8.14*). The inherent multi-block structure used by the method to decompose the target geometry into hexahedral elements imparted a patchy element quality. In comparison, the element-quality for the morphed meshes were continuous and smooth (*Figure 8.13a-8.13b*). The multi-block method tended to favor flat-to-convex surfaces over concave with respect to element-quality, which was the opposite for the morphed meshes. The inferior surfaces had the highest element qualities, whereas the superior surfaces were lower (*Figure 8.14a-8.14b*). The elements between the singularity edges within the wedge regions often had lower-quality elements (*Figure 8.14c*). A multi-block based artefact or defect which had previously been described to cause large surface errors, led to irregularly shaped elements (*Figure 8.14d-8.14e*). These defects were automatically created during the subdivision process and little control could be dictated over them without defining

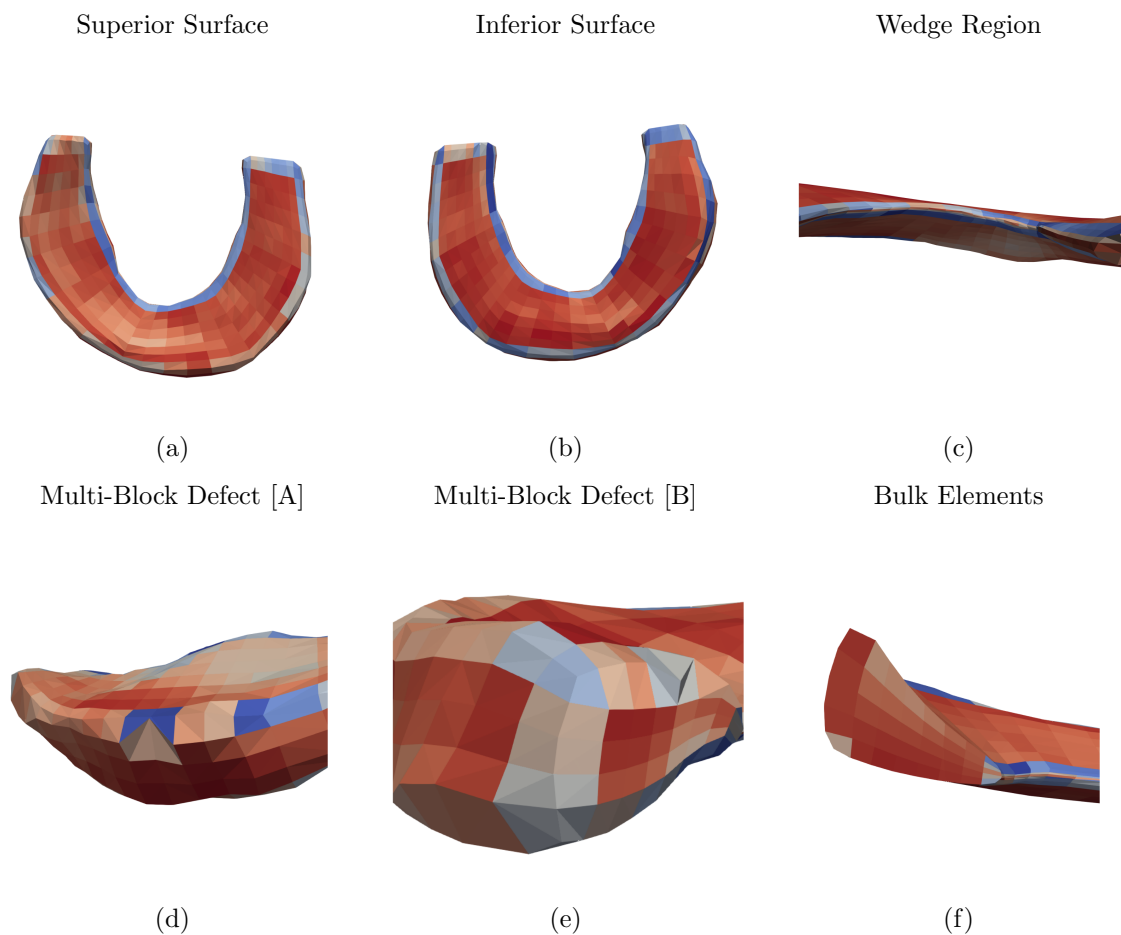


Figure 8.14: Regions and features of the multi-block meshes depicting variations of higher and lower quality elements (blue=lower element-quality; red=higher element-quality)

additional multi-blocks and re-positioning the multi-block vertices.

The quality of the internal (bulk) elements showed similarities to the variabilities observed for the external elements between morphed and multi-block meshes. For both methods, the bulk elements were affected by the quality of the external elements surrounding them. However, the morphed meshes tended to have a higher quality bulk than the multi-block meshes.

Finite Element Analysis

Overall, the general contact pressure distributions between mesh pairs were similar and distinguishable from other geometries (*Figure 8.15*). However, the maximum contact pressure could be dissimilar with respect to the value or location. Although, where differences in the maximum contact pressure occurred, a high-value concentration was often found in the corresponding max-value location.

The simulations with the morphed and multi-block meshes exhibited four common differences:

1. Smoothness and continuity of contact pressure distributions:

The morphed mesh contact pressure distributions were smoother and more continuous than those derived from the corresponding multi-block simulations (*Figures 8.15*). The contact pressures were often higher in the discontinuous patches than the corresponding regions in the morphed mesh simulations, which had lower values in smoother and more continuous distributions.

2. Magnitude of contact pressure towards thin-regions or extrusions:

Morphed meshes derived from target geometries which exhibited thin-regions or extruded features did not show the same contact pressure that were observed in the corresponding multi-block simulations (*Figures 8.15c-8.15d*). In the morphed mesh simulations, the contact pressures were higher and more concentrated within a smaller area for these regions.

3. The shape and position of the horn-attachments:

The horn-attachments - defined between the terminal vertices of the template mesh and location on the tibial plateau - could have a different shape and attachment location relative to the smoother morphed geometry (*Figure 8.16*).

4. Maximum contact pressure locations:

Multi-block meshes often had point-like artefacts that were not present or as distinctive in the corresponding target geometries. For the multi-block simulations, maximum contact pressure values were often located within close proximity to a protrusion or point-like artefact. This was the main cause for location differences between corresponding simulations.

The contact area distributions between pairs were distinguishable from other simulations (*Figure 8.17*). The greater roundness and size of morphed meshes were observable, as well as the greater continuity of the contact areas.

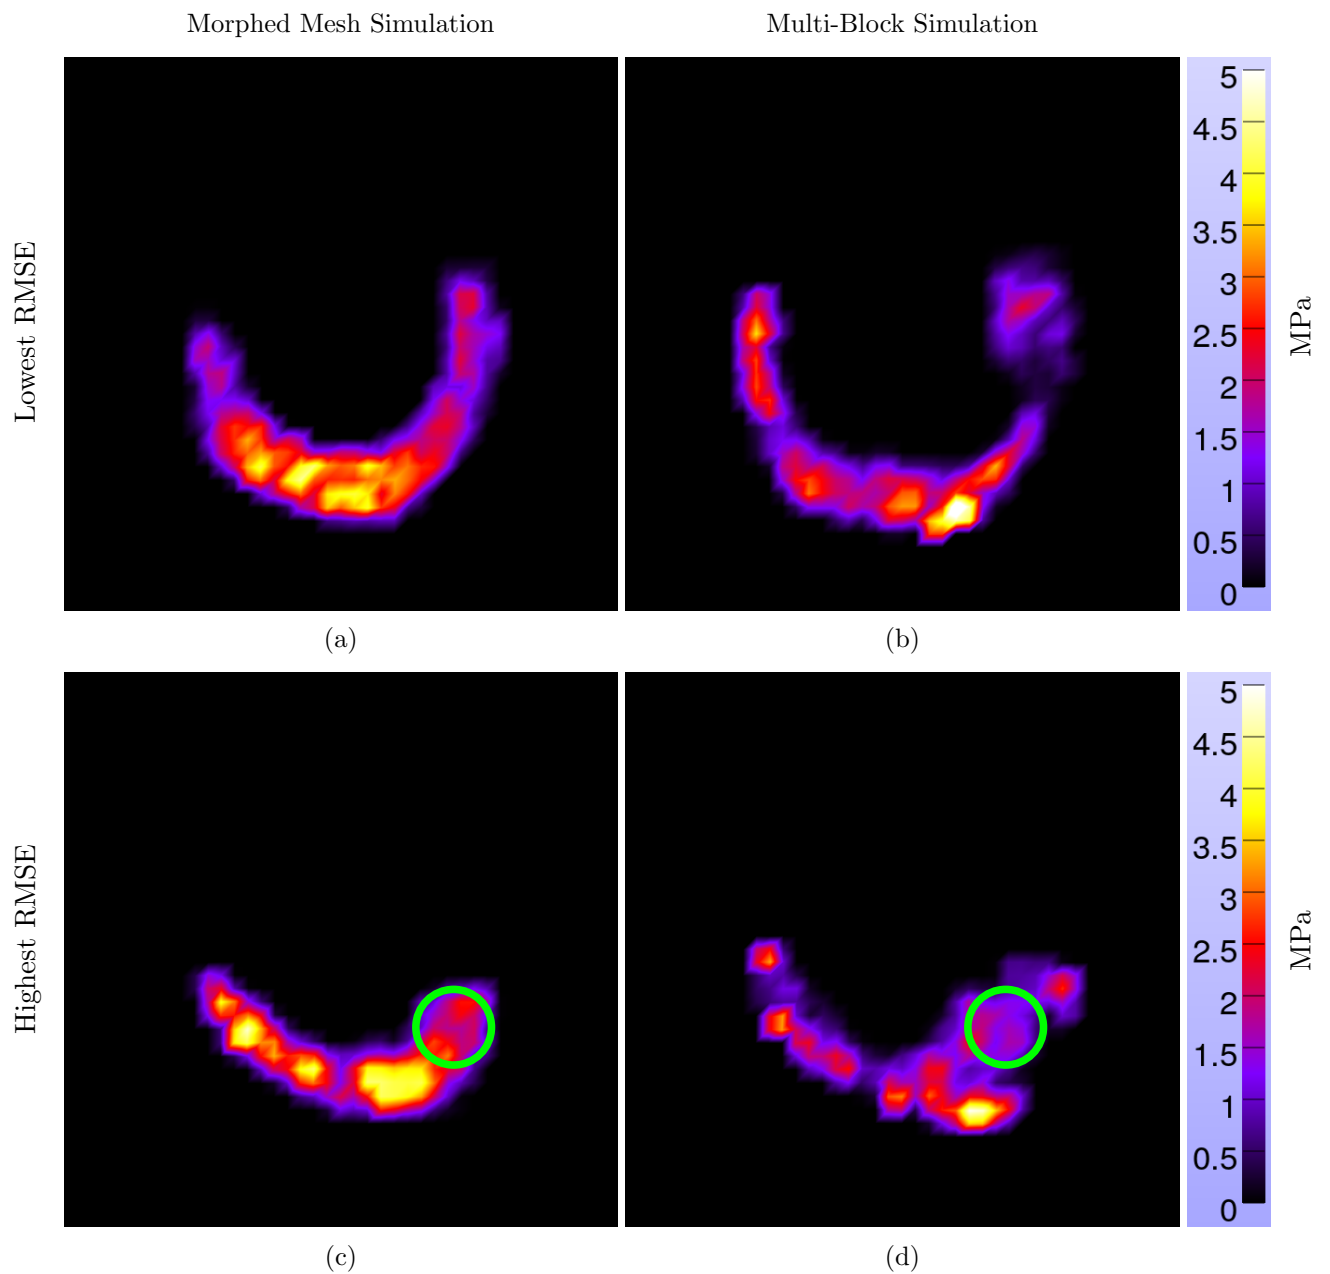


Figure 8.15: Contact pressure distributions of simulation pairs, comparing morphed meshes with the lowest (a-b) and highest (c-d) RMSE. Contact pressure differences around thin-regions highlighted




Figure 8.16: Different shaped horn-attachments between some of the (a) morphed and (b) multi-block meshes

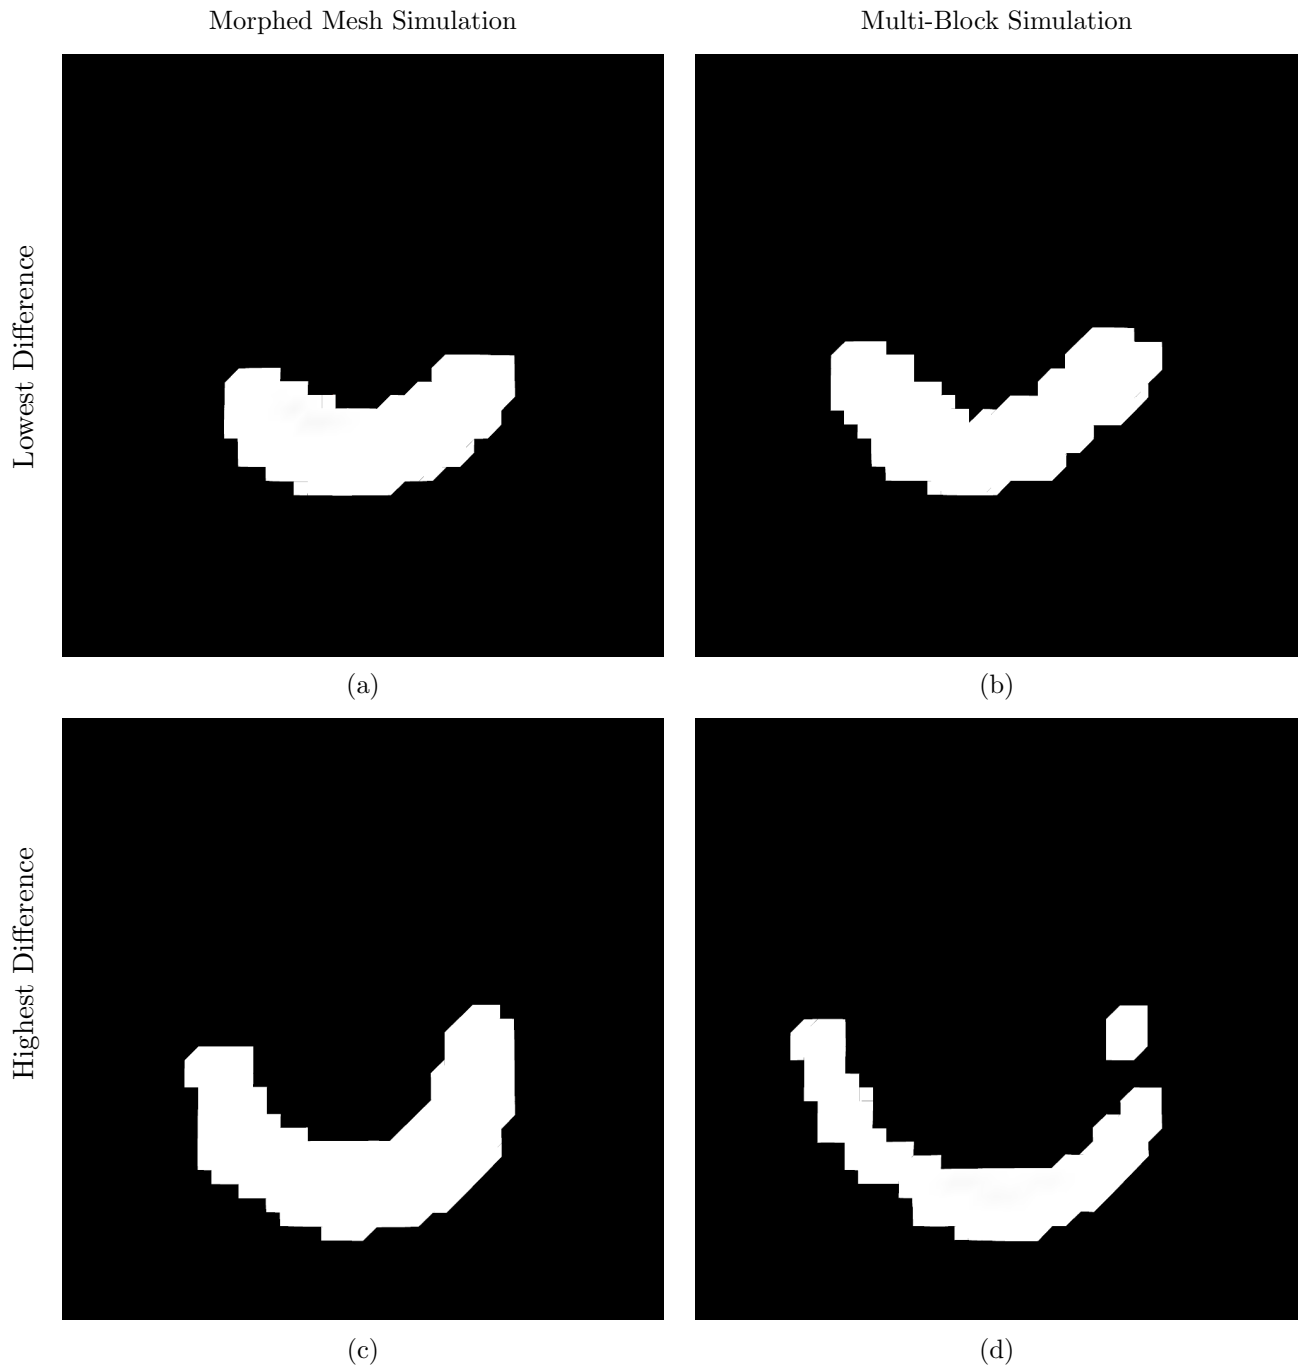


Figure 8.17: Contact area distributions of simulation pairs, comparing the lowest (a-b) and highest (c-d) contact area difference

8.4 Discussion

8.4.1 Generation Times

The mesh-morphing strategy generated hexahedral meshes much faster than the multi-block method, due to being automated and requiring no user-interaction. The speed differences for the strategy were dependent on the composition of algorithms and how they processed target geometries with a different number of vertices. In contrast, the generation times for the multi-block method were strongly dependent on the user-interaction. The differences between the targets' geometric features had the strongest impact on the user-interaction and generation times.

The generation times for the strategy became slightly longer with increasing target vertices (N_t). This was most likely due to the increased computational cost of ray-casting a hexahedral mesh's surfaces vertices with respect to the target geometry. The ray-casting algorithm had a complexity of $\mathcal{O}(N_t)$, when morphing template meshes with the same number of vertices (*Section 3.4.1*) [447]. The majority of computational time was spent in an iteration loop, where this algorithm was used at each iteration and was the most expensive looped algorithm (*Algorithm 3.1*). Consequently, a proportion of this complexity was transferred to the overall processing time for the strategy.

8.4.2 Surface Error

The morphed meshes had a higher surface error than the multi-block meshes. The degree of surface errors were linked to the level of constraints in the meshing processes, which were used to determine target correspondences. The mesh-morphing strategy used an automatic unconstrained technique, whereas, the multi-block method required the manual definition of hard constraints. The automatic unconstrained approach led to a faster process, at the expense of a less well-defined surface. The main surface losses for the strategy were due to the smoothing and rounding of the geometry, which caused a reduced representation of particular features with respect to the target geometry.

The smoothing and rounding effects could be attributed to two aspects of the morphing strategy: (1) use of no constraints and (2) a large number of smoothing operations. The morphing strategy did not constrain any surface vertices and instead relied on an unconstrained technique, defined

by upper and lower energy barriers. Additionally, the surface elements of the morphed meshes (quadrilateral meshes) were continually smoothed during several iterations of the morphing process. The spring energy of the quadrilateral meshes were minimised to estimate surface correspondences, instead of using constraints which could lead to distortions [295]. During the smoothing of the quadrilateral elements, the unconstrained technique would prevent collapse or excessive size reductions when interaction occurred with the target surface. Essentially, the target surface defined a lower energy bound that the quadrilateral mesh could not traverse. Over several iterations, the quadrilateral mesh would untangle, shrink and wrap around the target geometry until the energy of the system was minimised. The aim of this approach was to allow the morphed quadrilateral meshes to unfold between the upper (expanded target) and lower (original target) boundaries representing a thin membrane around the target geometry. Furthermore, the lack of constraints were designed to allow the vertices the freedom to move within this membrane and unfold whilst representing the global structure of the target geometry. However, it was evident from the morphed meshes that the shrinking effect of the quadrilateral mesh occurred globally. Consequently, during numerous iterations a mesh geometry became rounder with some features becoming less defined and concavity reduced for some regions. Also, this process was responsible for the horn-attachment recession observed for some of the geometries, as the vertices were given perhaps too much freedom to move.

Some of these issues likely arose from two aspects of the upper bound energy: (1) enforcement and (2) smoothness. First, there was no upper boundary enforcement preventing the quadrilateral mesh traversing that energy barrier, as it was incorrectly assumed the mesh would shrink everywhere locally. Second, the upper boundary was generated by expanding the original target geometry along the surface normals, instead of simply increasing the size. This would have caused some surface smoothing and rounding, that may have been imparted onto the morphed mesh. One potential way to counteract the first effect would be to apply the same rule for the upper boundary and force vertices to remain within the thin membrane if moved outside the slightly expanded target surface. However, this would have still led to the smoothing of quadrilateral elements into a rounded surface, that was sometimes unfavourable for hexahedral elements during the linear elastic transformation step. Also, this may not have prevented the requirement for final quadrilateral smoothing. The rounding effect was more significant when smoothing a quadrilateral mesh than a hexahedral mesh. For a hexahedral mesh, the internal vertices would act like soft constraints, anchoring the surface to the volume. There

was an example of a similar technique that successfully achieved spherical parameterisations of triangular meshes derived from highly irregular geometries [295]. However, the unconstrained technique in the morphing strategy attempted the opposite, parameterise to a base domain with a highly irregular geometry, a meniscus. Using this type of technique with a highly irregular base domain and without an upper bound being enforced did not provide the same degree of parameterisation.

The pinching effect of the mesh-morphing strategy was caused by the centreline-based projection algorithm. For some geometries with very thin wedge regions where the centreline was relatively far from the surface, the initial projection would cause a high degree of inversion. This was likely due to the projection vectors becoming crossed at the surface of the very thin-regions, which could have been worsened by the curved shape of the centreline projecting these vectors inwards towards these thin sections. The mesh-morphing strategy had difficulties removing these inversions and led to concentrations of vertices within these regions.

8.4.3 Mesh-Quality

The morphed meshes had fewer very low and low-quality elements than the multi-block meshes. The main cause was due to the differences in the topology of the meshes. In particular, the position, quantity and magnitude (valency) of vertex and edge singularities, where the morphed meshes had more strategically placed singularities and fewer corresponding edges. As observed in case C (*Chapter 7*), singularities significantly affected the shape changes of a geometry following smoothing operations. Additionally, this affected the quality of the elements in different geometry configurations. For example, elements neighboring singularities and their edges were highly susceptible to having low-quality values, if they were far from an element ideal. In this investigation, all the lowest quality elements were neighboring these singularities and edges. Due to the morphed meshes having fewer surface singularities, this feature of the topology alone allowed the meshes to have fewer low-quality elements. Singularities can both positively and negatively affect mesh-quality, depending on the proportion, location and valence [39, 192, 197, 448]. If used appropriately, singularities enable a mesh to better capture a geometry [39, 197]. Despite global quality improvements singularities can have on meshes, they can reduce the quality locally to neighboring elements [448]. Also, singularities can have an effect on the mesh smoothing, where poorer quality elements can be induced or exacerbated locally, regardless of global improvements [448]. Another cause for morphed meshes having

fewer poor-quality elements was features being smoothed and rounded that would have caused elements to deviate from their ideal shape. As the multi-block meshes had retained a significant proportion of those features and had a lower surface error, this would have imparted greater distortion to the elements causing them to have lower qualities. Irregularities negatively affect element quality [32, 193, 237], especially without topologies tailored to a particular geometry [448]. Meshes with a more optimal topology provide greater accommodation for geometric irregularities [197, 449]. This is an aspect where the appropriate inclusion of singularities can lead to improved mesh-quality globally [197]. This would also explain the greater quality of the morphed meshes, as the template meshes provided a level of topology tailoring, whereas the multi-block meshes did not.

Another cause for the higher-quality meshes was the higher-quality bulk elements than the multi-block meshes, as well as higher quality elements in the bulk than on the surface. In contrast, the multi-block meshes had a graded quality from higher-quality surfaces (flat to convex) to lower (concave). The higher-quality bulk elements for the morphed meshes were most likely caused from two aspects of the mesh-morphing strategy. First, the surface was optimised to the target geometry separately to the volume elements, undergoing several smoothing iterations in an attempt to both unfold the initial projection and collapse the projection onto the target surface. Following this step, the volume elements were interpolated and then the hexahedral meshes were smoothed. From these two steps in the mesh-morphing strategy, in particular the final smoothing step, the shape of the bulk elements were more adequately optimised than the surface elements.

The morphed meshes had higher quality elements on the concave surface and lower qualities on the flat to convex surfaces. However, this was the opposite situation for the multi-block meshes. The cause for this was again due to the differences in topology, as well as the curvature differences with respect to the target geometries. This further supports observations from case C (*Chapter 7*). Consequently, the differences in the location and quantity of the singularities played a key role in constraining the shape of hexahedral elements to particular regions of the geometries and associated curvatures. The singularity structure of the morphed meshes favoured regions of geometries with higher curvature, which also caused the lower quality elements observed towards the horn-attachments. These regions of the horn-attachments had lower curvature and were relatively flat compared to the middle regions of the meniscus geometries. The cause for

lower quality elements around the horn-attachments in the morphed meshes was due to the smoothing of the quadrilateral mesh in the surface unfolding step. It was observed that the quadrilateral vertices that would become singularity corners in the hexahedral meshes were smoothed away from the idealistic shape of a hexahedral element, causing them to be highly distorted when transferring them to that structure during the linear elastic transformation step. Furthermore, these regions tended to collapse onto the target geometry faster than other regions, therefore these elements did not have the same degree of shape optimisation that would occur elsewhere.

8.4.4 Limitations

There were several limitations of this study that require mention. Regarding the FE simulation, the simplicity of the design may not have allowed the effect of some geometric differences to be revealed. These simplicities could have been overcome by using patient-specific tibial and femoral cartilages, as well as a full condyle imparting flexion. This could have determined the effects of patient-specific geometry mismatches and pushed the simulations to the limits of convergence. However, the FE simulations used are consistent with models published in literature and are able to highlight key biomechanical changes resulting from mesh differences. Another aspect of the FE simulations were that they had a sub-optimal mesh-resolution (*Section 4.5.7*). However, as the goal was to compare simulations with meshes derived from the mesh-morphing strategy and multi-block method, they had to have the same resolution. Additionally, the choice of the idealised femoral and tibial cartilage was designed from a familiarised but not expert developer of simulations. Although, the shapes of the tissues were consistent to those observed in literature and provided suitable appropriations of the structures.

With respect to the development of the hexahedral meshes, only one density was investigated. This was limited by the multi-block method and the difficulties associated with achieving element numbers higher than what was chosen. Additionally, some of the target geometries contained artefacts from segmentations that may have been removed for clinical-grade FE models. Justification for this was to explore a range of geometries, with a variety of features and acquired from different protocols to determine the overall robustness of each method and how particular features affect them.

The final limitation arose from the use of the centreline/centreplane-based surface projection

algorithm to determine the surface errors of the hexahedral meshes. This could cause conflict as a similar algorithm was used in the mesh-morphing strategy. However, closest-point projection was found to be insufficient with respect to accuracy and the centreline/centreplane-based algorithm yielded more realistic results. One issue that may have affected the accuracy of the surface error calculation, unrelated to measurement method, was that some target geometries had their horn-attachment sites clipped. The clipping algorithm filled in the removed space with large triangles, effectively reducing the number of vertices within these regions. Consequently, the number of sample points for these areas were reduced, which potentially could have been overcome by subdividing larger triangles into smaller ones to improve the uniformity of the sample points.

8.4.5 Multi-Block Method as Gold Standard

The multi-block method provided several advantages for the process of developing hexahedral meshes. It enabled a user to generate high-quality hexahedral meshes with low surfaces errors. The surface errors were low enough to retain a significant proportion of patient-specificity. Additionally, operating the tool could be relatively fast depending on the complexity of the geometry and was reasonably straightforward. However, some practice was required to gain a particular level of proficiency and in most cases could not be used or learnt immediately, to create hexahedral meshes of irregular anatomies.

Despite these benefits and being a gold-standard procedure for generating high-quality meshes of irregular anatomies, the multi-block method might not be the most suitable tool for some structures, such as the meniscus. There were a number of disadvantages this tool possessed when considering irregular anatomies like the meniscus. Inherently, there was the inability to change the topology of the mesh to favor the target geometry. Not all geometries tended towards block decompositions on the smaller scales and even less on the larger scales. This was evident when observing the lower quality elements of the multi-block meshes along the wedge regions of the meniscus, where two edges of the block structure had to be condensed onto a thin-region that would be best defined by one edge. It was possible to design this topology for the template of the mesh-morphing strategy, which had comparatively higher-quality elements in these regions. Additionally, there was a resolution limit to the multi-block method, as it was difficult to refine the meshes beyond 1200 elements. Perhaps with further practice and skill this limit could be overcome. However, for some of the meniscus geometries it was a

challenge achieving 1200 elements. This limit has three consequences for the development of FE models using this tool. First, it prevents adequate mesh-sensitivity testing through the inability to enable the determination of mesh suitability for higher resolutions and deriving those meshes, if required. Second, it increases the difficulty of capturing irregular features with higher quality elements, that are observed around thin-regions and the wedge region of the menisci. Third and similarly, the multi-block method can impart ‘sharp’ features from smoother irregular features due to not having an adequate number of vertices to sample such regions accurately. Another disadvantage was the rigidity of the tool, which provided limited control over the general smoothness of the derived hexahedral meshes. The smoothing process could induce negative Jacobians, preventing adequate representation of the meniscus geometries at the chosen resolution. This may be improved with further post-processing of the target geometries prior to operating the multi-block method, such as additional smoothing to ensure the presence of minimal irregularities.

The multi-block method has been designed to be a general purpose tool and most test cases of IA-FEMesh use simpler structures, such as bones. This may indicate that the method is better suited to more regular anatomies. Although, there are examples of use for more complex geometries and studies outside the development team, such as using it to generate hexahedral meshes of meniscus anatomies [36]. Other disadvantages include the time requirements, laborious operations and associated artefacts of the method. The interactive procedures can have large time requirements for some geometries that is not straightforward to predict and requires repeating several tedious steps to achieve a valid hexahedral mesh. Also, there are the issues that the multi-block method may impart particular defects, such as the aforementioned denting effect.

8.4.6 Causes for Key Biomechanical Differences

The differences in contact pressure distributions between mesh pairs were smaller than the inter-specimen differences. Consequently, simulations from both methods could easily be distinguished from those derived from other meniscus geometries. However, the multi-block method did retain a greater degree of patient-specificity, evident from the lower degree of surface errors. Although, this appeared only to have a significant effect for thin-regions along the wedge, or possible extrusions, which is an important meniscus pathology and component of knee osteoarthritis [79]. In general, with the exception of extrusions, the overall shape of the meniscus

geometries were captured well by both methods and their corresponding simulations. Additionally, the size of corresponding mesh pairs were very similar, although, slightly larger for the morphed meshes. This is important as not the shape but size of a meniscus has been found to be the most sensitive factor to the contact variables [63, 450]. Specifically, where a smaller meniscus, such as produced by the multi-block method, leads to reduced contact pressure. However, an adequate geometric representation should still be an objective as the meniscal width was found to impart some influence [63, 451]. Despite the general similarities of the contact pressure distributions between mesh pairs, the degree to how much patient-specificity has been retained was affected by a number of differences associated with each method. At the chosen mesh resolution, the mesh-morphing strategy and multi-block method imparted key geometrical differences onto the derived hexahedral meshes that tended to under or over-represent features of the target geometry, respectively. There were distinct simulation effects associated with some of the hexahedral mesh differences between each method.

The morphed meshes exhibited three geometry variations that caused simulation differences with respect to the multi-block meshes. First, the hexahedral meshes had a smoother surface and had more rounded features, which corresponded to smoother and more continuous contact pressure distributions than the multi-block meshes. The maximum contact pressures were often found in the centres of large concentrated regions. This correlates with other FE studies that have shown the meniscus occupying large continuous regions with the max. pressure location situated near the centre [18, 76, 154, 418, 452]. Additionally, three studies have shown strong differences in contact pressure distributions despite using the same open-knee model under axial compression [18, 150, 452]. The study which had meniscus meshes with greater smoothness and less faceting, due to higher mesh-resolutions, derived more continuous and smoother contact pressure distributions [18]. However, the morphed meshes achieved this with lower mesh-resolutions, due to the greater smoothness imparted from the processing of the morphing strategy. This provides evidence that the simulations with morphed meshes produced more realistic and accurate results, due to greater similarities to the outcomes of higher resolution models.

The second geometric variation was a consequence of the morphing strategy not fully capturing thin-sections or extruded regions of a target geometry. This was caused from the smoothing and rounding of the meniscus geometries. This caused higher contact pressures within a

smaller concentrated area when comparing these regions to the multi-block mesh simulations. Also, this aspect caused the greatest localised contact pressure differences of all the features. In contrast, the most significant global difference was the degree of smoothness and continuousness. Two studies have found slightly different observations regarding extruded menisci with respect to contact pressure and area [453, 454]. Both found that the contact area decreases with extrusions but contradict each other regarding the effect of contact pressure. One study found contact pressure decreases for extrusions [453], which corroborates the results from the multi-block method. This would confirm that the use of morphed meshes for extruded geometries would induce erroneous contact pressure distributions. However, another study found that an extrusion has no effect on the contact pressure, which explored four increasingly more extruded geometries [454], compared to one for the other study [453]. This would indicate that the morphed meshes would not provide incorrect contact pressure distributions. It is clear that a degree of patient-specificity was lost as a result of the rounding effect on extruded regions. However, it is difficult to determine how significant this effect was on the biomechanics, due to the contradictory findings from previous investigations. Additionally, it is also difficult to determine the nature of contact area reduction with respect to capturing extruded features, as the contact areas were higher for the majority of morphed mesh simulations. However, morphed mesh simulations involving potentially extruded menisci showed both increases and decreases in contact area compared to the multi-block simulations. Interestingly, the majority of morphed mesh simulations that showed a decrease in contact area did exhibit signs of meniscus extrusion. This could indicate that some of the extruded features were captured in the morphed mesh simulations. However, the reduced contribution to the total contact area would likely be due to these regions being smoothed, less represented and sometimes ‘pinched’ than purely a consequence of extrusion biomechanics.

The third geometric variation was caused by recession of the horn-attachments for some of the geometries, particularly those with a block-like shape. This affected the general shape of the horn-attachments and where they attached to the meniscus geometry relative to the patient-specific structure. This may affect the motion and dynamics of the affected menisci, which could result in changes to the contact pressure distribution. Although, differences in horn-attachments shape and location reportedly have only a small effect on the contact area and pressure [410]. However, the horn-attachment shape was found to have a significant impact on the degree of meniscus movement [410]. This could lead to losses in patient-specificity and greater differences

in dynamic simulations. Although the degree of movement was not measured, large changes in horn-attachment shape could be alleviated through the manual selection of the terminal vertices that were used to define the spring elements. This would lose an aspect of automation with respect to FE model development, at the expense of accuracy, if meniscus movement was a concern.

Another aspect to consider was that the greater roundness of the morphed meshes sometimes required slightly different initial conditions. The femoral height was sometimes located at a higher distance from the tibial cartilage than what was defined for the multi-block simulations. This was required due to the morphed meshes having a slightly larger and rounder geometry. Differences in the initial conditions and discretisation between identical models (open-knee) have been found to yield dissimilar results with respect to the continuousness of contact pressure distributions [18, 150, 452]. Despite the multi-block and morphed mesh simulations having the same boundary and loading conditions, the differences in tibial-femoral height may have caused some dissimilarities. Consequently, using the morphing strategy to develop patient-specific models could mean differences in congruency that may not occur with the use of the multi-block method. This could cause problems during the assembly of the individual mesh components and lead to differences in conformity between the tibial and femoral cartilage meshes.

There were two noteworthy differences between morphed and multi-block meshes with respect to the simulation process and convergence: (1) simulation time and (2) magnitude of applied force. All the morphed mesh simulations converged faster than the multi-block simulations. The majority of morphed mesh simulations completed in at least half the time, and in some cases almost six times as fast. This was attributed to the greater mesh-quality and smoothness, which improve convergence for deformation [455] and contact [456] heavy simulations. Additionally, the morphed mesh simulations were able to converge when subjected to a greater applied force. During preliminary investigations, the applied force was specified at 1000 N. However, the multi-block simulations failed to converge at this force and was reduced to accommodate them. Additionally, the initial force of 1000 N could be considered quite high for a single compartment compared to other studies [18, 19, 81, 453], therefore was considered a reasonable adjustment.

8.4.7 Consequences of the Mesh-Morphing Strategy

The mesh-morphing strategy poses several positive and negative consequences for FE model development and the clinical interpretation of the simulated biomechanics.

Development of Finite Element Models

For the development of FE models, the strategy offers significant advantages over the multi-block method, particularly with respect to automation. The most profound advantages lie with the ability to easily scale-up model development with respect to mesh quantity and resolution. The mesh-morphing strategy is automated and produces hexahedral meshes much faster and within more consistent time frames than the semi-automatic multi-block method. This can enable large-scale studies with predictable windows of development. The automated nature of the process removes rater-biases and rater-drift, and results in consistently shaped hexahedral meshes. Additionally, it is simpler to increase the mesh resolution of the meniscus meshes and has associated benefits of improving surface errors and element-qualities (*Chapter 7*). Furthermore, this can assist mesh-sensitivity analyses and ensuring the use of convergent meshes and models.

Another advantage that can aid the automation of FE models is the transference of material and model definitions from the template design to the morphed meshes. Generally, the template approach allows greater flexibility and customisation for automating the development of the FE models. For tissues such as the meniscus, which require anisotropic material properties, the fibre-orientations can be defined based on the local element structure which will be identical between template and morphed meshes. For the multi-block method, this has the potential to change and depends on the construction and orientation of the block structure. Also, due to the high one-to-one correspondences between the majority of geometries, zonal properties could be defined for the template and consistently lie in relative regions for the morphed meshes. The one-to-one correspondences can allow the horn-attachment definitions to be automated, as the end regions were often in the correct locations and relied on the same vertices for spring element associations. However, additional checks may be required to prevent inaccurate replication of patients-specific horn-attachment shapes and dynamics, for some meshes. As the condition may occasionally arise where morphed vertices of the horn-attachments recess by a large amount with respect to the original template positions and could require redefinition to lie in locations

representative of defining springs elements. Although, as discussed, this has a minimal effect on the contact pressure and area but does affect the motion of the meniscus [410].

The general operation of the mesh-morphing strategy can be applied to any mesh definition, allowing greater flexibility and control over the mesh topology through template design. This was particularly beneficial when utilising a template mesh which resembles a general meniscus geometry with the cross-section resembling a rounded right-angled triangle. If done sensibly and correctly, appropriate template design can lead to higher quality elements and lower surface errors (*Chapter 7*).

The mesh-morphing strategy imparts a rounding effect to the morphed meshes. This has the tendency of producing less concave curvatures and larger regions in some areas of the geometries. This could lose surface conformity between cartilage meshes and pose difficulties and inaccuracies when building patient-specific models with respect to congruency and pairing these conforming structures. However, the trade-off to the rounding effect is that it leads to higher quality and smoother meshes. This can improve solver convergence, particularly where highly-irregular features might hinder certain FE models from converging. Also, this could be useful for preliminary testing of a FE model design, as well as more general and complex studies of knee joint mechanics which are not entirely focussed on the patient-specificity of menisci.

Biomechanical Interpretation of Finite Element Results

The mesh-morphing strategy has several implications to the accuracy and interpretations of the simulated biomechanics. One advantage of the morphed meshes is that they produce more realistic meniscus properties and therefore biomechanics, than the multi-block meshes. This is evident by the smoother surfaces and more continuous contact pressure distributions derived from the FE simulations. A key function of the meniscus is to evenly distribute contact pressure between the femoral and tibial cartilages [78, 457], which the morphed meshes achieve better than the multi-block meshes. Although the morphed meshes are smoother than their target meshes, these smoother features are likely more representative of actual meniscus geometries. The presence of sharp irregular features are often likely the effect of segmentation artefacts and scanning irregularities [76, 458]. Irregular features are often removed by design during post-processing via smoothing operators [20, 76, 80, 134, 458]. Consequently, use of the mesh-morphing strategy could reduce the time taken to acquire meniscus meshes in two ways. First,

the need to segment, in detail, the naturally smooth contours of the meniscus geometries. Second, reduce the amount of post-processing required from a derived meniscus segmentation and surface representation, when aiming to acquire a smooth structure. Excessive post-processing and smoothing operations could lose potentially important features. Therefore, the smoothing process, which was also used to improve the element-quality in the mesh-morphing strategy, could be consolidated and used once. This would prevent the need for multiple smoothing operations during the segmentation process, which are required to smooth and remove segmentation artefacts. Another benefit of this aspect is that the strategy can robustly create realistic meniscus meshes despite the presence of segmentation errors and artefacts.

Additionally, the transition between the horns of a meniscus and the associated attachments into a tibia show continuous smoothness [459]. The edges do not appear angular and sharp [459], unlike how they are often represented in FE models, notably due to the clipping of these regions and simulation using spring elements [18, 150, 418]. When they have not been clipped [19] or modelled naturally [154], they have shown greater smoothness or the same smooth transitions, respectively. Even when the target geometries have been clipped, the morphed meshes tend to create horn-attachment sites that are curved similar to unclipped meniscus geometries [19]. The incorrect representations of menisci having sharp features could lead to erroneous contact mechanics, particularly when non-zero frictional coefficients are applied. However, contour irregularities can sometimes be observed on MRI scans, where menisci have become torn [460]. It is difficult to determine what degree of irregularity should be considered appropriate for pathological menisci. However, based on the evidence and purpose of the menisci, healthy tissues should be reasonably smooth.

Another concern is that there is a change in some biomechanical aspects when compared to the gold-standard procedure - the multi-block method. Particularly, the location and value of the max. contact pressure may be inaccurate. Although, there is evidence that suggests some of the differences are due to the multi-block method imparting ‘sharp’ defects onto the hexahedral meshes. Additionally, the meshes derived from the mesh-morphing strategy may not adequately represent important irregular features that may have been smoothed or rounded, such as tears and extrusions. Consequently, simulations derived from these meshes may not be able to accurately determine the effects of torn or extruded menisci. Although, no differences for potentially torn regions were observed.

9 Case E: Performance Optimisation

9.1 Introduction

In *Chapter 8*, it became evident that despite the advantages the mesh-morphing strategy provided over the multi-block method (e.g. higher mesh-quality and automation), there were several disadvantages. Specifically, the surface fit was comparatively poor, which led to several differences in the biomechanics compared to the gold-standard procedure. To overcome these shortcomings, the mesh-morphing strategy has been optimised with respect to the surface fit. The optimised mesh-morphing strategy is composed of the same components as the unoptimised strategy, however they are utilised with a different logic and design (*Section 3.8*). The new configuration takes principles from traditional surface parameterisation techniques and applies them to volumetric meshes. The optimised strategy builds on top of the unoptimised strategy by using the derived morphed meshes as templates, as they have a high-quality and reasonably similar geometry to the target meshes. Consequently, the optimised strategy requires a performance evaluation similarly conducted for the unoptimised strategy. In this investigation, the performance of the optimised strategy is compared against the unoptimised strategy and multi-block method for 20 anatomical meniscus geometries. This is to determine if the optimised strategy can improve the surface error whilst maintaining a larger proportion of the mesh-quality, and if this leads to improved FE simulation accuracy.

Aim: To compare the performance of the optimised mesh-morphing strategy against unoptimised strategy and multi-block method

Hypothesis: The optimised mesh-morphing strategy derives meshes with a greater surface fit with some compromise to the mesh-quality, which provides greater accuracy for FE simulations

Research Questions:

1. Does the optimised mesh-morphing strategy improve the surface fit without significant loss to the mesh-quality?
2. Does the optimised mesh-morphing strategy improve the accuracy of the FE simulated biomechanics?

3. Is the performance of the optimised strategy closer to the multi-block method or the unoptimised strategy?

Novelty: This investigation builds on the novelty provided in *Chapter 8*. Specifically, two aspects of the research in this chapter is novel: (1) the design of the optimised mesh-morphing strategy, and (2) the developed methodology used to compare and determine an improvement, with comparisons to a gold-standard procedure. Although the individual analyses are not novel, their use to determine if an optimisation to a mesh-generation procedure leads to performance improvements has not been employed.

9.2 Methodology

The entire methodology from *Chapter 8* has been replicated, with the addition of two metrics: (1) overall performance, and (2) qualitative similarity to the unoptimised strategy and multi-block method. In summary, the performance from morphing the same 20 target meshes were compared to the unoptimised strategy and multi-block method (*Section 8.2.1*). The generation times (*Section 8.2.4*), surface error metrics (*Section 4.1*), mesh-quality metrics (*Section 4.2*), finite element metrics (*Section 4.5*) and descriptive analyses were used to evaluate the performance and key differences to the other methods. Each template mesh had the same topology but different shape, which was more specific to each target geometry, as they were based on the corresponding morphed meshes from the unoptimised strategy.

9.2.1 Overall Performance

To compare the overall performance between the three methods, the percentage of low-quality elements was plot against the RMSE.

Ideally, for low surface errors and few low-quality elements, the points should be as close to zero as possible. These favourable meshes would have a higher performance and be located in the lower-left quadrant (*Figure 9.1*). Meshes that performed the most poorly would be situated in the upper-right quadrant, which would indicate high surface errors and numerous low-quality elements. The upper-left and lower-right quadrants would indicate meshes which only performed well with respect to either the mesh-quality or surface error, respectively. This allowed the determination of whether potential improvements were made by the optimised strategy compared to the unoptimised strategy. Also, to determine how well the optimised

strategy performed compared to the multi-block method. This would be visually clear by the presence of any overlapping regions, as well as the degree of such overlap.

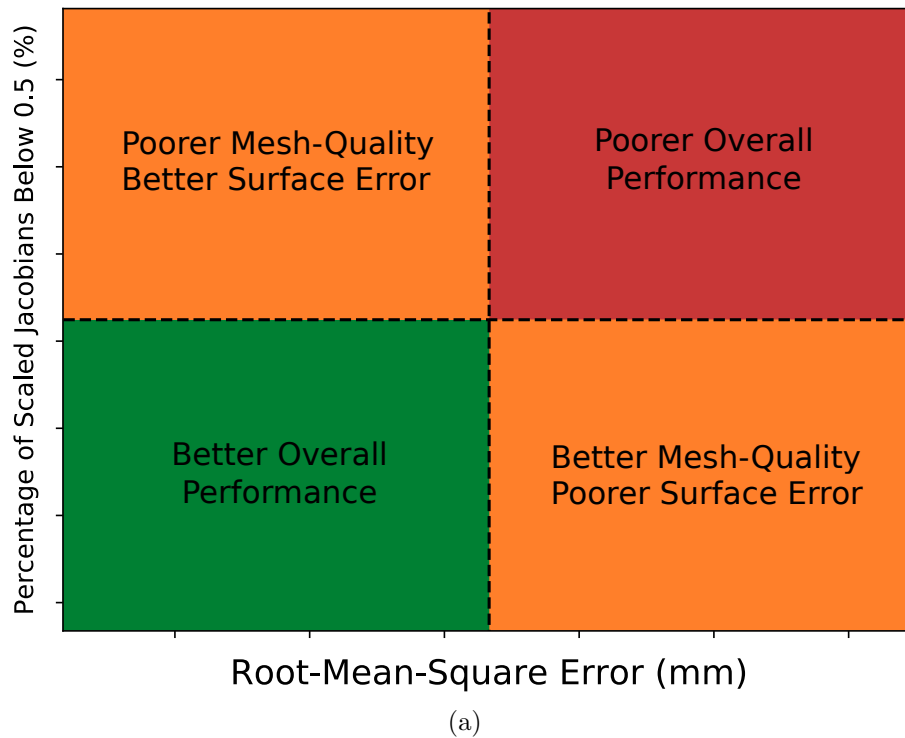


Figure 9.1: An example of a mesh-quality vs. surface error comparison with quadrants highlighting differences in performance

Outliers for the surface error and mesh-quality datasets were identified using two methods, the Z-score and interquartile range (IQR) [461, 462]. An outlier was only selected if both methods identified them.

9.2.2 Qualitative Analysis of Finite Element Outcomes

The contact pressure distribution were characterised with two descriptors: smoothness vs patchiness of the distribution, and connectivity (continuous vs discontinuous distributions).

9.3 Results

9.3.1 Surface Error Analysis

In general, the optimised meshes had lower surface errors than the unoptimised meshes (*Table 9.1* and *Figure 9.2*). The optimisation improved the surface fit for all but one case. However, the optimised surface fit still fell short of what could be achieved using the much more time-consuming multi-block method (*Figures 9.2a-9.2b*). Also, the surface errors for the optimised

strategy were more variable and less consistent compared to both the unoptimised strategy and multi-block method (**Table 9.1**). The mesh that had higher values for the optimised strategy (MM_21) showed the highest degree of shrinking.

Table 9.1: Comparison of the average, standard deviation, minimum and maximum surface errors between the three methods

	RMSE (mm)			Max. Error (mm)		
	Avg. \pm Std.	Min.	Max.	Avg. \pm Std.	Min.	Max.
Unopt. Morphed Meshes	0.916 \pm 0.073	0.833	1.111	2.887 \pm 0.532	2.154	3.885
Multi-Block Meshes	0.193 \pm 0.069	0.099	0.356	1.318 \pm 0.432	0.524	2.231
Opt. Morphed Meshes	0.410 \pm 0.241	0.168	1.239	1.940 \pm 0.728	1.101	4.063




Figure 9.2: RMSE (a) and maximum surface errors (b) for individual geometries

The majority of meshes had reasonably good and competitive surface errors compared to the multi-block method. Large surface errors were observed for geometries which presented common features: (1) large very thin-regions, (2) a narrow or tubular shape and (3) a block-like geometry. However, not all geometries with these features incurred the same degree of surface errors.

All the optimised meshes had a smaller volume than the target geometries, like the multi-block meshes, and had the greatest degree of error overall (**Table 9.2** and **Figure 9.3**). Also, each mesh was smaller than the corresponding unoptimised and multi-block meshes. The optimised meshes had only 15% of meshes closer to the target than the unoptimised meshes and none compared to the multi-block meshes. Also, the medial meniscus meshes (-10.83%) suffered a greater degree of volumetric loss than lateral meniscus meshes (-6.94%).

Table 9.2: Comparison of the average, standard deviation, minimum and maximum absolute volume errors between the three methods

	Absolute Volume Error (%)		
	Avg. \pm Std.	Min.	Max.
Unopt. Morphed Meshes	4.31 \pm 3.73	0.04	15.0
Opt. Morphed Meshes	8.88 \pm 4.96	3.18	22.11
Multi-Block Meshes	3.05 \pm 0.87	1.56	4.89




Figure 9.3: (a) Mesh volume comparison between the unoptimised, optimised, multi-block and target meshes, (b) target volume difference for each mesh

9.3.2 Generation Times

The optimised strategy generated each mesh with a slower speed than the unoptimised strategy (*Figure 9.4b*). On average, the optimised strategy increased the total generation time by 30 s. Also, the length of time to perform the optimisation in isolation took on average 5 s longer than the unoptimised strategy. However, the optimised strategy outperformed the multi-block method for each corresponding mesh and by two orders of magnitude on average (*Figure 9.4a* and *Table 9.3*). There was a clear distinction between the total generation times for the three methods, evident from the lack of overlap between minimum and maximum values. The optimisation of the strategy increased the sensitivity to the number of target vertices (*Figures 9.4c-9.4d*). Consequently, the optimised strategy appeared to approach an $O(N)$ algorithm.

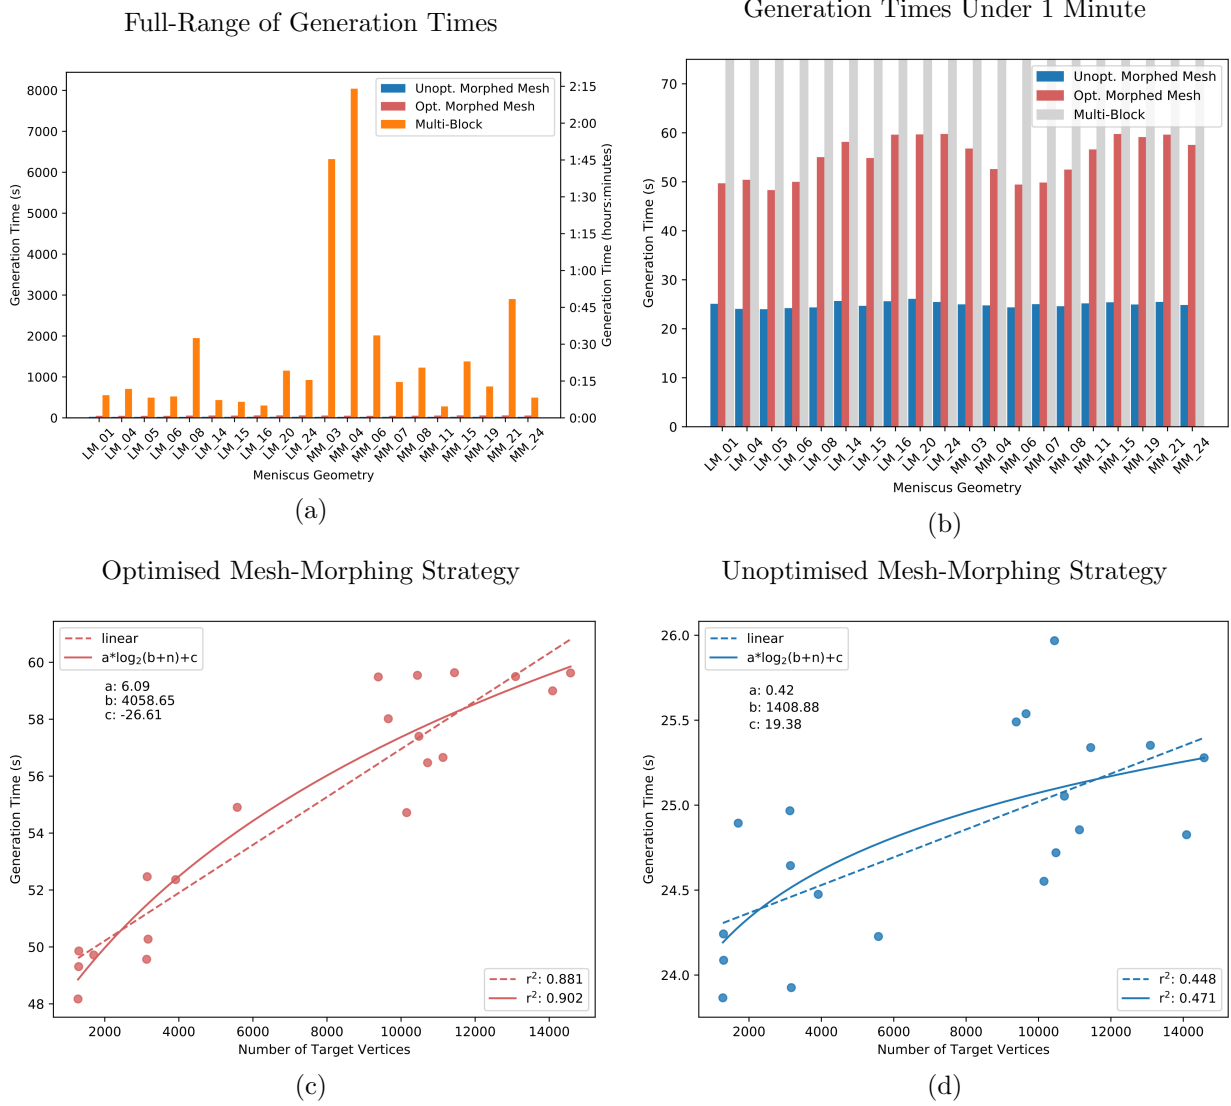


Figure 9.4: Mesh generation times for (a,b) individual geometries. Comparison of the number of target vertices against the time taken to generate each hexahedral mesh for the (c) optimised and (d) unoptimised mesh-morphing strategies

Table 9.3: Comparison of the average, standard deviation, minimum and maximum generation times between the three methods

	Generation Times (seconds)		
	Avg. \pm Std.	Min.	Max.
[Total] Opt. Morphed Meshes	54.83 \pm 4.11	48.17	59.64
[Isolated] Opt. Morphed Meshes	30.02 \pm 3.69	24.31	34.35
Multi-Block Meshes	1586 \pm 1994	279	8040
Unopt. Morphed Meshes	24.82 \pm 0.56	23.87	25.97

9.3.3 Mesh-Quality Analysis

Each optimised mesh had a lower quality than the corresponding unoptimised meshes (*Table 9.4*, and *Figure 9.5*). However, the optimised strategy produced slightly higher quality meshes than the multi-block method. These findings were evident based on the percentage of very low- and low-quality elements on average and between corresponding mesh pairs. Overall, the optimised meshes had a slightly higher proportion of low-quality elements than the multi-block meshes, but generally fewer between corresponding mesh pairs. (*Table 9.4* and *Figures 9.5b*).

Table 9.4: Comparison of the average, standard deviation, minimum and maximum percentage of very low and low-quality elements between the three methods

	Very Low-Quality Elements (%)			Low-Quality Elements (%)		
	Avg. \pm Std.	Min.	Max.	Avg. \pm Std.	Min.	Max.
Unopt. Morphed Meshes	0.04 \pm 0.12	0.00	0.50	5.86 \pm 1.16	4.75	9.00
Multi-Block Meshes	1.11 \pm 0.60	0.00	2.08	9.35 \pm 2.57	4.75	12.92
Opt. Morphed Meshes	0.85 \pm 0.61	0.08	2.17	11.70 \pm 6.19	2.75	32.17

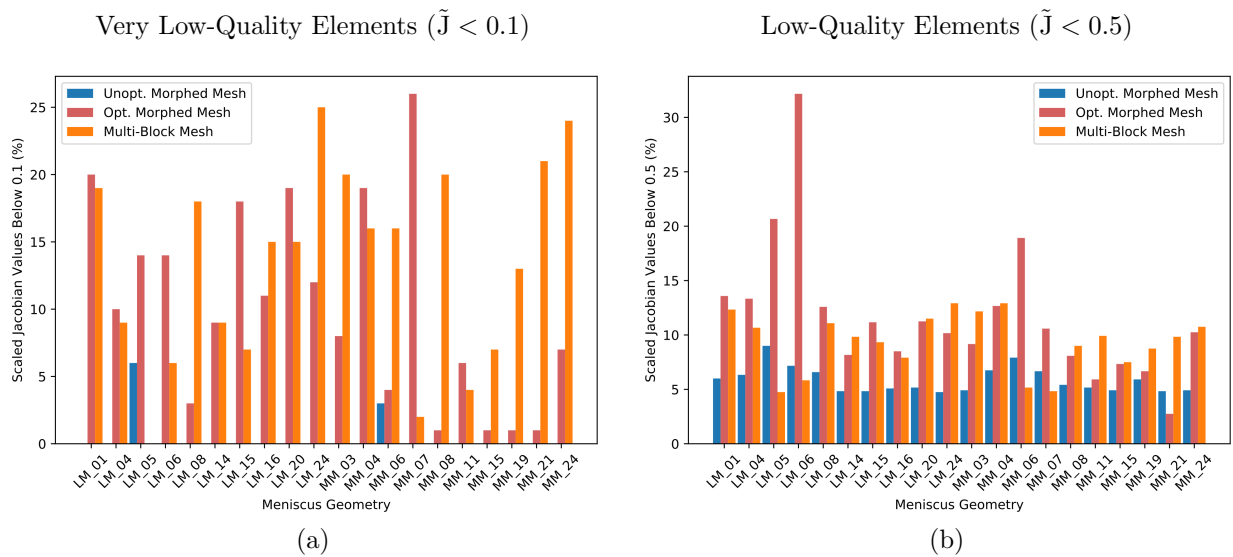


Figure 9.5: The number of (a) very low and (b) low-quality elements for individual geometries

9.3.4 Overall Performance

Overall, the multi-block method consistently performed the best with respect to mesh-quality and surface error, which occupied the favorable lower-left quadrant (*Figure 9.6*). The unoptimised strategy performed the worst, achieving higher surface errors despite consistently achieving smaller proportions of low-quality elements, which caused a dense cluster in the lower-right quadrant. In general, the optimised strategy lied between the two of them, but closer to the multi-block method, situated more broadly in the lower-left quadrant. Additionally, several optimised meshes occupied the region enclosed by the multi-block method (*Figure 9.6b*). This highlighted a degree of overlap between the performance of the optimised strategy and the multi-block method. In contrast, no overlap occurred between the unoptimised strategy and the multi-block method, due to the larger differences in surface errors.

Another aspect these plots highlighted was the robustness of each method at producing a consistent result. The unoptimised strategy was the most robust and occupied the smallest region with the greatest density. This was strongly influenced by the strategy consistently producing small proportions of low-quality elements. This was followed by the multi-block method, which provided a larger region. The range of surface errors was roughly proportional to the range of low-quality elements. The least robust method was the optimised strategy, which had larger ranges for both the surface error and mesh-quality. Consequently, it was the least dense and occupied the largest area. Additionally, two outliers were identified for the optimised strategy, whereas none were for the other two methods.

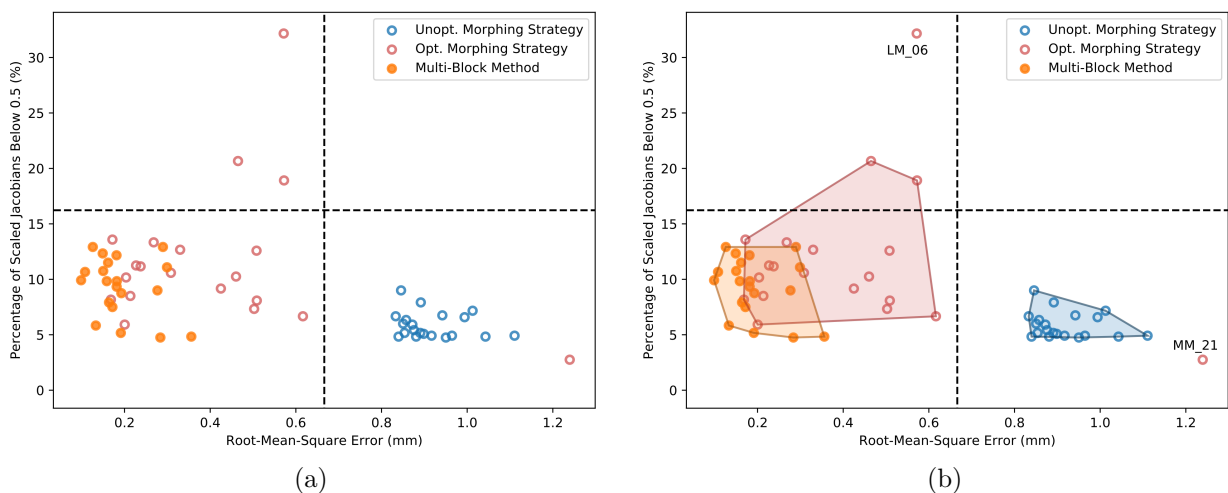


Figure 9.6: Comparison of the root-mean-square error against number of low-quality elements for each method: (a) without enclosing area and (b) with enclosing area, excluding outliers

9.3.5 Finite Element Analysis

Convergence

All 20 simulations with the optimised morphed meshes converged. The convergence rates for the optimised mesh simulations were slightly faster than those with multi-block meshes and slower than those with unoptimised meshes (*Table 9.5* and *Figure 9.7*). Additionally, the convergence times were within a similar range to the multi-block simulations.

Table 9.5: Comparison of the average, standard deviation, minimum and maximum convergence times between the three methods

	Convergence Times (HH:MM)		
	Avg. \pm Std.	Min.	Max.
Unopt. Morphed Meshes	00:13 \pm 00:05	00:08	00:34
Opt. Morphed Meshes	00:38 \pm 00:19	00:16	01:28
Multi-Block Meshes	00:40 \pm 00:19	00:20	01:27




Figure 9.7: (a) Individual and (b) overall convergence times for simulations derived from unoptimised, optimised and multi-block meshes

Contact Area

The contact area differences, with respect to the multi-block mesh simulations, for the optimised mesh simulations were much smaller than those with unoptimised meshes (*Table 9.6* and *Figure 9.8*). Also, the tibial cartilage had higher differences and variability than the femoral cartilage, like the unoptimised simulations.

Table 9.6: Comparison of the average, standard deviation, minimum and maximum absolute contact area differences between the multi-block method and the unoptimised and optimised mesh-morphing strategies

Method	Interface	Absolute Contact Area Difference (%)			Absolute Contact Area Difference (mm ²)		
		Avg. ± Std.	Min.	Max.	Avg. ± Std.	Min.	Max.
Unopt. Strategy	Tibial	35.16±35.66	1.40	128.19	31.65±25.01	1.41	82.23
	Femoral	20.74±15.65	0.71	61.85	25.43±16.62	1.01	50.18
Opt. Strategy	Tibial	17.94±15.04	0.01	56.63	17.48±12.22	0.01	47.01
	Femoral	12.31±10.37	1.35	35.59	16.19±13.76	2.14	46.80

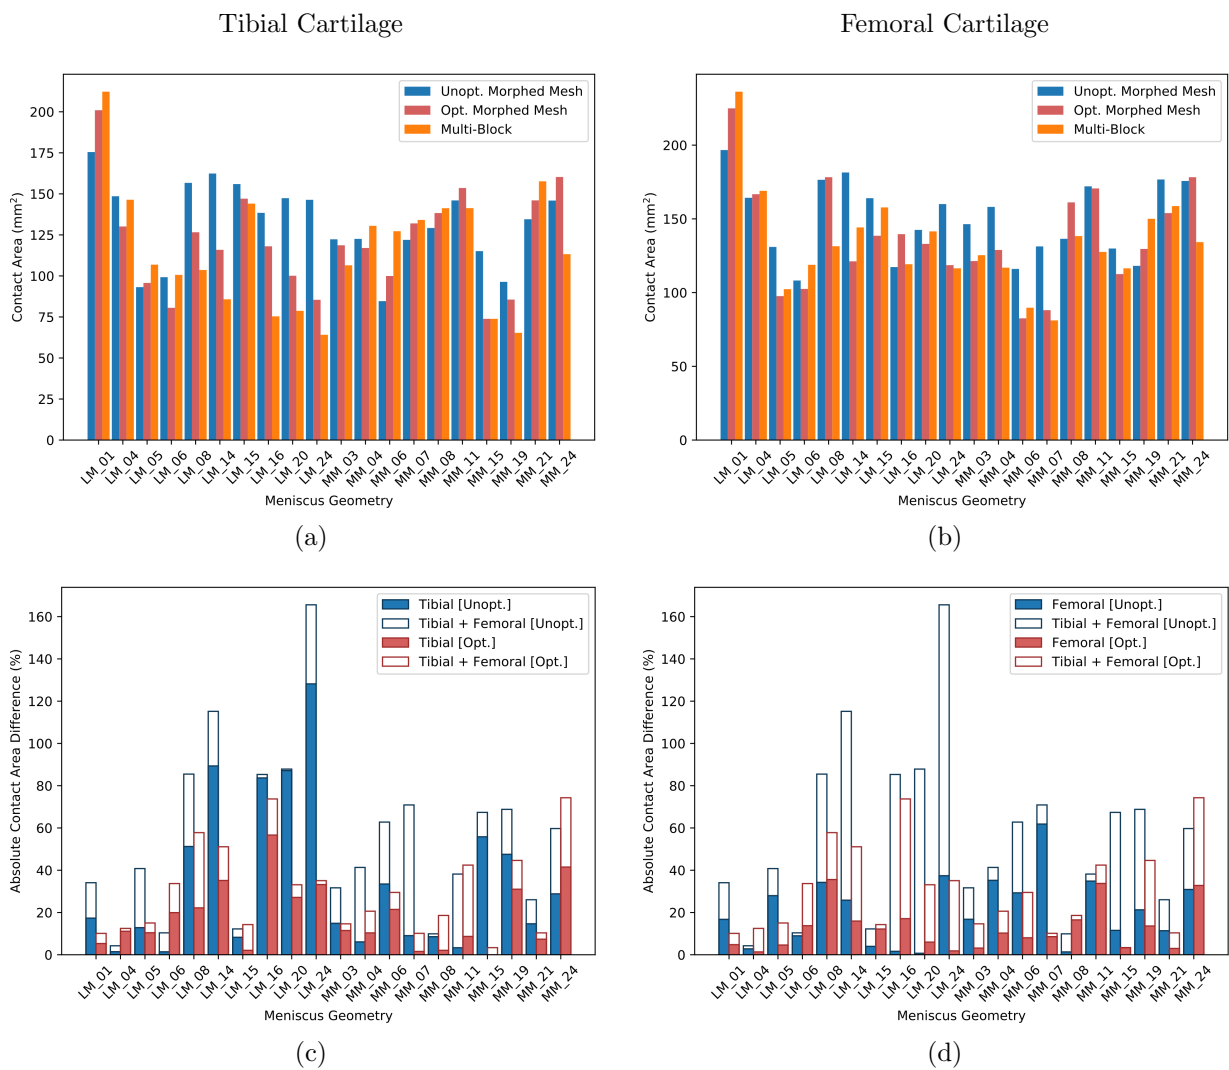


Figure 9.8: The contact areas for individual simulations on the (a) tibial and (b) femoral cartilage. The absolute contact area differences (c-d) for the unoptimised and optimised simulations

9.3.6 Detailed Analysis

Surface Error

In general, the optimised meshes were more patient-specific than the unoptimised meshes. They retained a greater proportion of geometric features and irregularities. Consequently, they appeared more similar to the multi-block meshes and target geometries than the unoptimised meshes. Also, the rounding and smoothing effect was much smaller for the optimised strategy than the unoptimised strategy. However, the optimised meshes appeared smoother than the multi-block meshes. The optimised strategy did not discriminate concave surfaces, evident by the similar degree of errors on both the superior and inferior surfaces, unlike the unoptimised strategy (*Figures 9.9a-9.9b*). Also, the pinching effect did not occur for any of the derived meshes. However, a shrinking effect was observed for the optimised strategy, that was not present in the unoptimised strategy. Some regions which negatively affected the optimised meshes showed similar surface errors in the unoptimised morphed but to a much lower degree.

The highest surface errors associated with the optimised mesh-morphing strategy were due to a shrinking effect. The shrinking effect caused some optimised meshes to have a smaller overall size than the corresponding target geometries. These errors occurred most frequently and severely around the horn-attachment sites and the peripheral rims of the meniscus geometries (*Figures 9.9b-9.9c*). Also, as protrusions often occurred around the peripheral rims, surface errors could be caused in these locations (*Figures 9.9d-9.9e*). In some cases, thin-regions along the wedge were also affected (*Figure 9.9f*). The rounding effect of the unoptimised morphed meshes affected similar regions.

Regardless of the degree of shrinking, the main surface losses occurred for three aspects: (1) horn-attachment corners, (2) surface irregularities and (3) thin-regions (when present). The highest surface errors for the optimised meshes were from horn-attachment corners, which represented the most angular regions of a meniscus geometry, particularly if the geometry was clipped at these regions (*Figure 9.9c*). To a lesser degree than the horn-attachment corners, surface irregularities caused noticeable surface errors (*Figure 9.9a*). Previously, horn-attachment corners and geometric irregularities caused a similar degree of error and were categorised under angular features. The horn-attachment corners were still the most dominant cause of higher surface errors but generally lower than the unoptimised meshes. Geometric irregularities were

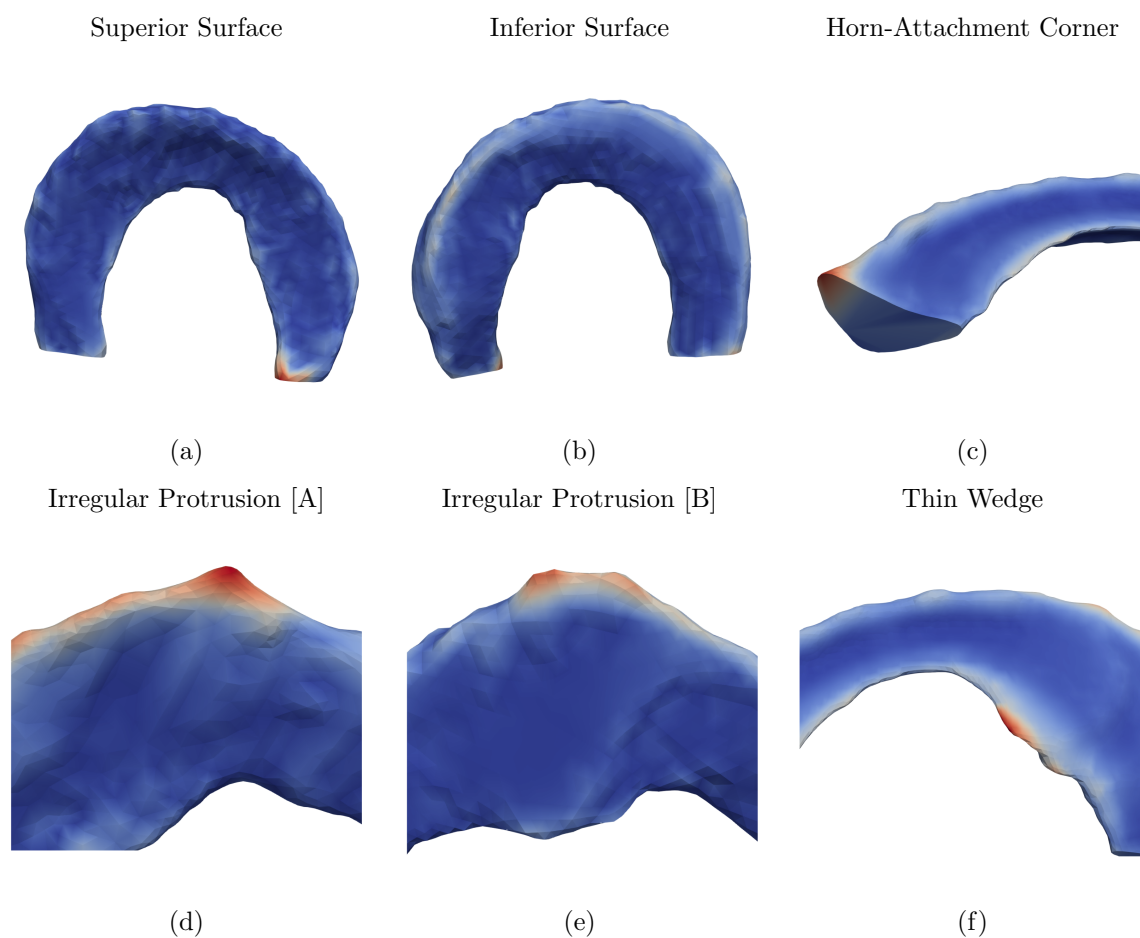


Figure 9.9: Regions and features of the target geometries which caused surface error variations for the optimised meshes (blue=lower surface error; red=higher surface error)

less of a problem and more localised to surface irregularities, similar to the multi-block method. As a result, geometric irregularities were reasonably well-captured but small-scale surface irregularities, such as protrusions, became slightly rounded (small-scale rounding effect) (*Figures 9.9d-9.9e*). For meshes with a lower degree of shrinking, these features were less rounded with the larger surface errors concentrated towards the tip of the features. Where a higher degree of shrinking occurred, the rounding effect was greater, which flattened small features and caused the errors to encompass a larger proportion of the features.

Thin-regions and possible extrusions were reasonably well-captured in the optimised meshes and generally had lower surface errors than the unoptimised meshes. Although, for some geometries these regions caused noticeable surface errors, particularly MM_21 and MM_24. The higher surface errors in these regions were associated with two aspects. These regions became (1) slightly rounded and (2) sometimes slightly distorted. The local rounding effect of thin-regions was much smaller than what was observed for the unoptimised morphed meshes, but still evident.

Mesh-Quality

The same aspects that affected the mesh-quality for the unoptimised meshes contributed to similar variations for the optimised meshes. However, the negative effects of these aspects were more detrimental to the mesh-quality of the optimised meshes. The external elements located within concave regions commonly represented some of the highest quality elements. The mesh singularities tended to be the most common cause of lower quality elements, both internally and externally. Also, the external elements had a slightly lower quality than the internal elements. One noticeable difference between the two mesh-morphing strategies was their response to the degree of surface concavity on a target geometry. Additionally, there appeared to be a proximity effect, that caused lower quality elements to impart their lower quality to neighboring regions. The shrinking effect tended to have a positive effect on the mesh-quality.

The external element-quality for the optimised morphed meshes were generally affected by five aspects: (1) concave regions, (2) singularities, (3) convex regions, (4) horn-attachment sites and (5) thin-regions along a meniscus wedge. The highest quality external elements were found within the concave regions, which was the same for the unoptimised morphed meshes (*Figures 9.10a-9.10b*). The higher quality elements were also found towards the medial regions of the superior surfaces, where the highest concavity was observed amongst the target geometries. Common regions with lower quality elements were often found along mesh singularities. However, lower quality elements were not always found continuously along the singularities (*Figure 9.10d*). Some of the lowest quality elements have been commonly observed at horn-attachment corners, where the valence of the singularities were lowest.

Convex regions have been found to have slightly lower quality elements. This was observed for each mesh, where the inferior (convex) surface had a slightly lower quality than the superior (concave) surface. However, lower quality elements were typically found towards the peripheral rims of both these surfaces (*Figure 9.10e*). Overall, the optimised strategy discriminated mesh-quality differences between concave and convex surfaces less than the unoptimised strategy. Additionally, lower quality elements were found towards horn-attachments, which also had higher relative convexities than other regions. The higher convexities were due to curvature changes at the terminal regions of the horn-attachment sites.

The quality of elements towards the horn-attachments tended to be lower than other global re-

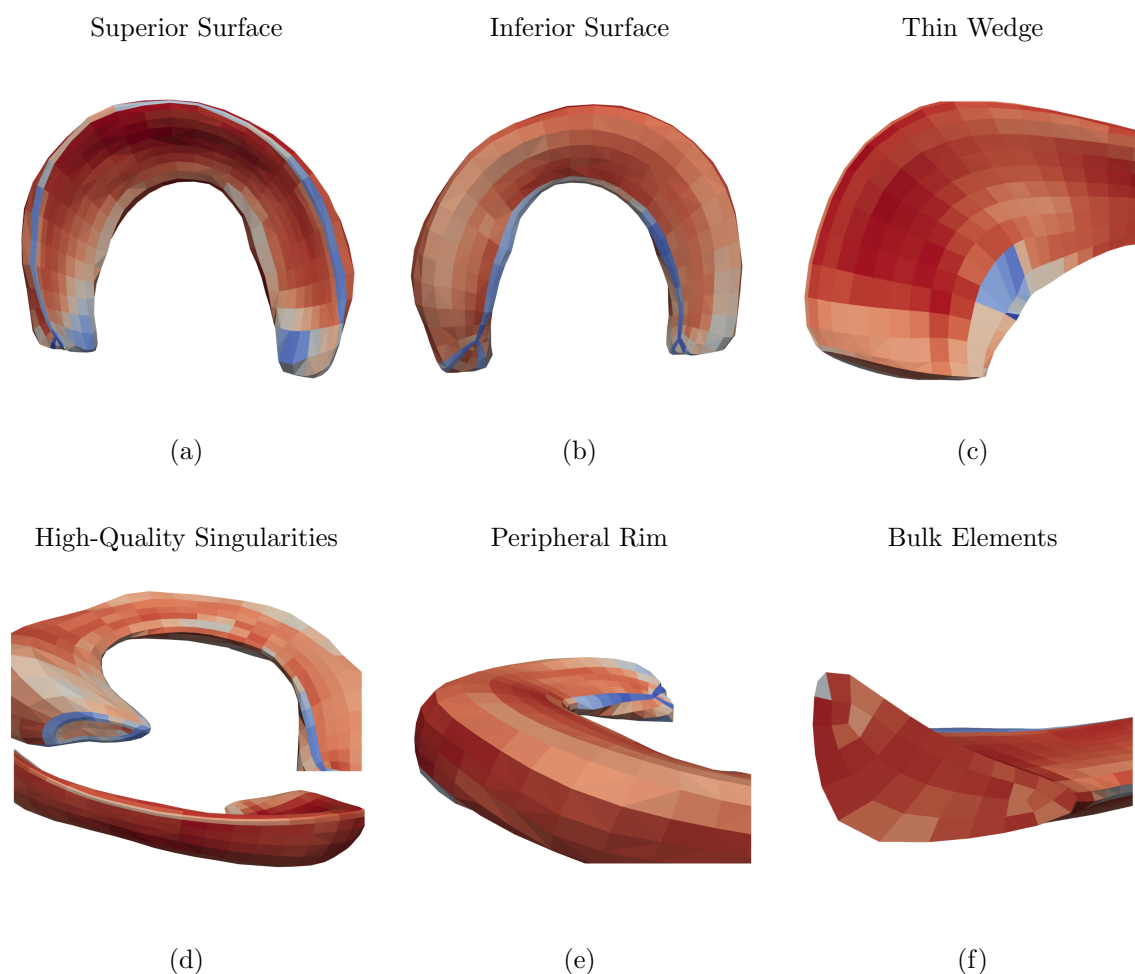


Figure 9.10: Regions and features of the optimised meshes depicting variations of higher and lower quality elements (blue=lower element-quality; red=higher element-quality)

gions (**Figure 9.10a**). The horn-attachment sites had a high concentration of singularities and convexity. Additionally, the lowest valence singularities occupied these regions at the corners. Also, they exhibited a combination of the four aforementioned aspects detrimental to element quality: (1) singularities, (2) convex regions, (3) proximity effects and (4) residual effects.

Thin-regions along the wedge, or possible extrusions, also tended to have lower quality elements, when present (**Figure 9.10c**). These often occurred in regions of the wedge that were thin and had a jagged contour.

In general, the internal elements had a slightly higher quality than the external elements (**Figure 9.10f**). Although, the internal element quality was susceptible to proximity effects, similar to the external elements. Specifically, internal elements adjacent to lower-quality external elements could have a slightly lower-quality compared to those adjacent to higher-quality external elements. This could cause a slightly graded appearance within cross-sections. Due to the in-

ferior surface having a slightly lower quality than the superior, the gradations existed between these to surfaces, transitioning from lower to higher element qualities in the inferior-superior direction. Additionally, internal elements adjacent to internal mesh singularities tended to have lower-quality elements than other internal elements within a given inferior-superior slice. The effect of internal mesh singularities was also observed for the unoptimised meshes. However, the unoptimised meshes were not as affected by the proximity effect of external elements. Consequently, the internal elements had a higher quality than the optimised meshes, with a more pronounced contrast to the external elements.

Finite Element Analysis

The optimised mesh simulations appeared to lie between the differences observed for the multi-block and unoptimised mesh simulations (*Figures 9.11, 9.12 and 9.14*). However, the general contact pressure distributions were more similar between the simulations with the multi-block and optimised meshes (*Figure 9.12*). The differences were less distinct, and these pairs were more distinguishable than the simulations derived from the unoptimised morphed meshes. The higher overall contact pressure was due to larger continuous concentrations of contact pressures that had lower values or discontinuities in the multi-block simulations. These regions were smoother, more continuous and occupied a larger surface area in the optimised mesh simulations, but less than the unoptimised simulations.

There were three notable differences between the simulations with optimised meshes to the multi-block and unoptimised meshes: (1) contact pressure distributions, (2) capture of thin-regions and (3) horn-attachment geometry.

The optimised mesh simulations captured features more similar to the multi-block simulations than the unoptimised mesh simulations. This was evident by similar discontinuities in the contact pressure distributions (*Figures 9.12a-9.12c*). However, there were fewer discontinuities and less patchiness than the multi-block simulations but more than the unoptimised mesh simulations, which were more continuous and smooth (*Figure 9.11*). There was a slight degree of smoothness within concentrations but generally leaned closer to the multi-block simulations than those with unoptimised meshes. Also, there were higher value concentrations within the discontinuities than observed for the multi-block simulations, which were more diffuse.

There were strong similarities between the multi-block and optimised mesh simulations with

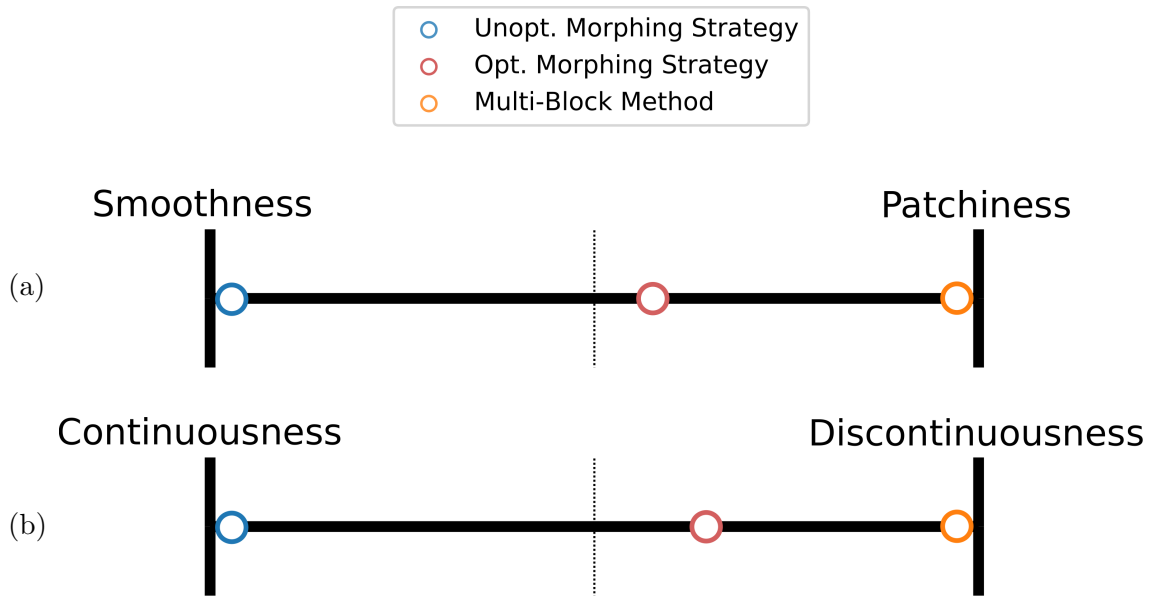


Figure 9.11: Contact pressure descriptor ranges for the (a) smoothness-to-patchiness and (b) continuousness-to-discontinuousness qualitative measures, comparing where the optimised mesh simulations lie compared to those derived from the unoptimised and multi-block meshes

respect to capturing thin-regions of a meniscus (*Figures 9.12e-9.12f*). The thin-regions along the wedge, or possible signs of extrusion, tended to have higher contact pressure values in the unoptimised mesh simulations. However, the contact pressure values were slightly higher compared to the multi-block simulations, where they were slightly more diffuse and less concentrated.

The shape of the horn-attachments were generally more similar to the multi-block meshes than the unoptimised meshes. However, the optimised meshes differed in two ways, the shape and size. The horn-attachment shape of the multi-block meshes approximated a rectangle. In contrast, the optimised meshes had oval-shaped horn-attachment sites, like the unoptimised meshes. Also, there was a smaller degree of recession that occurred for the majority of optimised meshes when compared to the unoptimised meshes. Although, not the shape but terminal vertices tended to be more narrow than those defined on the multi-block meshes (*Figure 9.13*).

Between simulation triplets, the differences observed for the contact pressure distributions were comparable to the contact area distributions (*Figure 9.14*). Quantitative analysis of the contact area differences showed the optimised simulations were closer to the multi-block simulations (*Section 9.3.5*). The qualitative analysis indicated these smaller differences translated into contact area distributions with more similar shaped meniscus impressions.

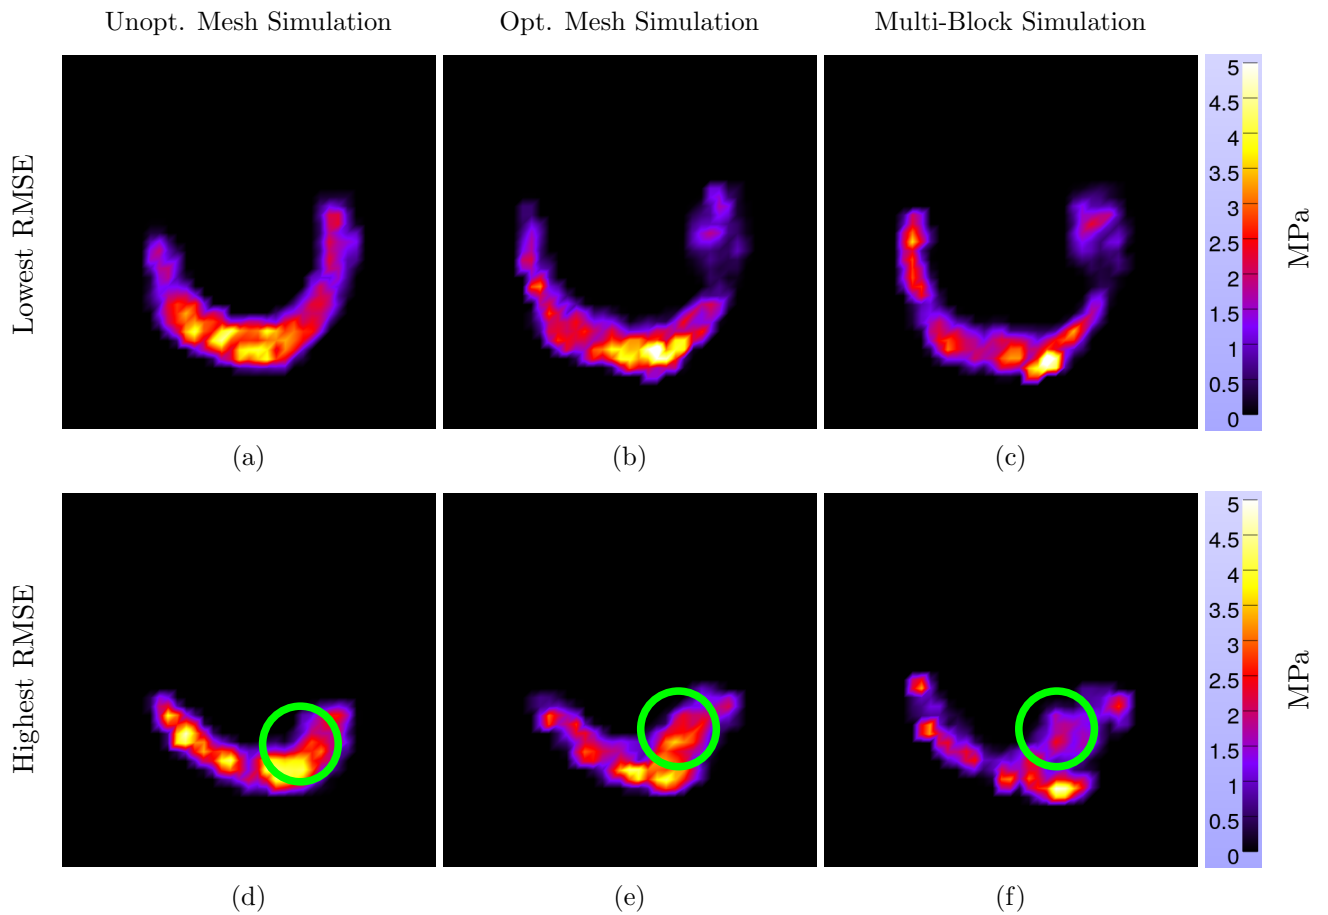


Figure 9.12: Contact pressure distributions of simulation triplets, comparing optimised meshes with the lowest (a-c) and highest (d-f) RMSE. Contact pressure differences around thin-regions highlighted




Figure 9.13: Different shaped horn-attachments between some of the (a) optimised and (b) multi-block meshes

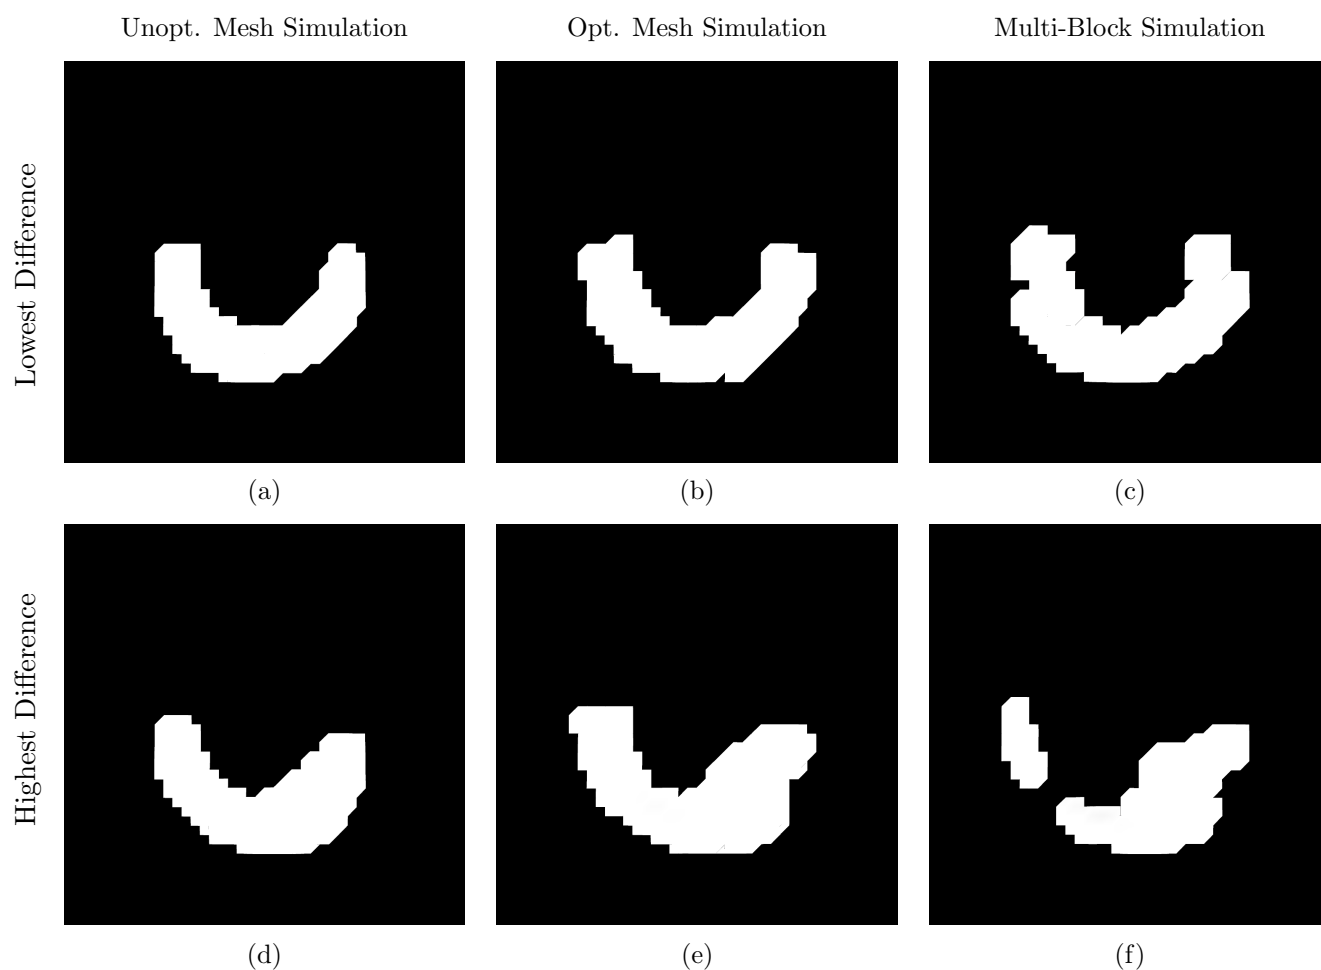


Figure 9.14: Contact area distributions of simulation triplets, comparing optimised mesh simulations with the lowest (a-c) and highest (d-f) contact area difference

9.4 Discussion

9.4.1 Surface Error

The optimisation process improved several aspects of the surface fit. Improvements to the patient-specificity of the meniscus meshes was one of the main benefits. This was the result of greater feature preservation, in particular: thin-regions, convex/concave surfaces and the horn-attachments. Also, the optimisation was able to eliminate the pinching effect observed for some of the potentially extruded anatomies. Despite the greater patient-specificity, the optimised strategy still imparted an adequate level of smoothing and rounding (e.g. removed surface irregularities), which prevented angular and faceted meshes. The main cause for these improvements were due to the use of soft-constraints instead of none. Also, there were two effects that had an impact on the surface fit: (1) a reduced smoothing and rounding effect and

(2) shrinking effect.

The fact that the smoothing and rounding effect was reduced provided several benefits compared to the unoptimised and multi-block meshes. Overall, it allowed the optimised meshes to represent the best geometric attributes from the unoptimised and multi-block meshes. Compared to the multi-block meshes, they had smoother and more meniscus-like geometries with fewer surface irregularities and no faceted or angular features. In addition, they had greater patient-specificity, feature capture and overall shape similarity with reduced losses around horn-attachments and thin-regions compared to the unoptimised meshes.

The shrinking effect was the most detrimental with respect to the degree of surface and volumetric errors. The cause for the shrinking effect was simple and likely preventable with a modification to the strategy. It was caused as the hexahedral meshes were smoothed until there were no non-positive Jacobians. Consequently, if meshes had numerous non-positive Jacobians following the morphing optimisation, they shrank proportional to the difficulty of removing them, as more smoothing iterations were required. One potential solution to this problem would be to improve the global mesh-quality by removing distorted elements individually. In the literature, there was an example of such an algorithm that was used to improve the quality of hexahedral meshes generated for tissues of the tibio-femoral joint, including the meniscus [40]. The algorithm removed the distortion of individual elements by changing the shape and angles to more closely resemble an idealised hexahedral element. Essentially, the vertices of an element were adjusted to yield angles that were closer to 90° . Another cause for the greater shrinking effect observed for some geometries was likely due to the use of a sub-optimal template mesh. In particular, this appeared to have the strongest effect on meshes with block-like shapes, tubular geometries and those with very thin features. An optimally designed topology has been observed to provide a greater surface fit for irregular geometries [39, 197, 449]. Also, providing a template mesh with a closer initial shape has been shown to minimise the distortion induced during a morphing operation [46, 463, 464]. However, there was no correlation between the shrinking effects and the degree of change observed for the contact variables. For example, the medial meshes had greater surface and volume errors but more precise predictions of max. contact pressure locations and contact area overall. The volume errors correlated with the qualitative observations of meshes becoming smaller and shrunken. However, the volume errors were quite different to those caused by the unoptimised strategy, evident by an overall smaller RMSE for

the optimisation. Specifically, the optimised meshes were smaller versions of the target geometries, whereas the unoptimised meshes were rounder versions. Based on this knowledge, the problem could potentially be solved quite easily, by simply increasing the size of the optimised meshes to match the volume of the target meshes. A fairly straightforward algorithm could be developed using an optimisation loop between the use of a linear scale-operator and the volume difference. The derived meshes would have the same mesh-quality, but the surface errors would likely be reduced and the volume errors would be zero.

9.4.2 Generation Times

The generation times were around 30 s slower for almost a 50% reduction to the surface errors. The slower times were inevitable due to the optimised strategy relying on the generation of a more accurate (surface fit) template mesh from the unoptimised strategy. However, the optimisation was still orders of magnitude faster than the multi-block method, with a more comparable surface fit. The complexity of the iteration loop was increased the most by the presence of the linear elastic transformation operator, which only occurred once at the end of the unoptimised strategy. The most demanding aspect of the transformation was associated with the solution to the stiffness matrix of the finite element method [465]. The second cause for increase was due to smoothing more vertices from a volumetric mesh each iteration, instead of the less numerous surface vertices.

There were two key consequences of the complexity increase, besides an increase to generation times. First, the optimisation was more sensitive to the number of target vertices. One potential solution to overcome this issue for larger meshes would be to decimate the target mesh to reduce the number of vertices. The majority of the optimisation could be performed on the decimated mesh and then followed by a few iterations on the original target mesh at the end, if necessary (i.e. using a hierarchical approach). Hierarchical-based morphing strategies have led to reduced computational complexity (e.g. faster generation times) [237] and reduced element distortion [292]. However, this would be more of an issue if morphing other anatomies, particularly those with a larger geometry and a high degree of irregularity. The second consequence would be associated with any parameter adjustments of the optimised strategy. Each optimised iteration required more time than an unoptimised iteration. The unoptimised strategy could perform 20 iterations in less than time than 15 optimised iterations (*Figure 3.1*). Therefore, if a greater global minimisation of a template mesh onto a target geometry was required, an increase to

the number of iterations would need to be specified. This would lead to increasingly longer computations compared to similar modifications for the unoptimised strategy.

9.4.3 Mesh-Quality

The optimisation process improved surface fitting but at the expense of reduced mesh-quality. However, this was not by design but as a consequence of the changes made to reduce the surface errors. Typically, the improved surface fit of the irregular geometries caused greater distortion and element irregularity. These findings correlate with observations from the literature [32, 193, 237]. Additionally, the combination of proximity and residual effects played a role in reducing the mesh-quality of the optimised meshes. Consequently, the optimised mesh-quality was only slightly better than the multi-block meshes. Despite the overall reduction to the mesh-quality, some aspects have improved locally. Most notably, the pinching effect has now been alleviated, as well as some of the lower quality elements that were associated with those regions. However, when present in the unoptimised meshes, they did impart a stronger residual effect than other regions with lower quality elements in the corresponding optimised meshes. Another improvement was the presence of higher quality singularities, due to the greater global minimisation with respect to the target geometries. For the unoptimised and multi-block meshes, the singularities consistently represented lower quality elements, whereas some optimised meshes had singularities with a high mesh-quality.

There were three effects that caused key differences to the mesh-quality: (1) the proximity effect, (2) the minimisation effect and (3) the residual effect. The proximity effect was caused by incremental smoothing, which followed each surface projection and volumetric transformation iteration. This gradually propagated higher and lower element qualities to their neighbours. There were several consequences of this effect. Most critically, this effect lowered the quality of elements near low and very low-quality elements, but also initiated the opposite effect for higher quality elements. This led to smooth and continuous transitions between lower and higher quality elements. Consequently, the proximity effect led to a greater number of lower quality elements, as these regions would encompass a larger proportion of the geometry than observed for the unoptimised meshes. In particular, there were more numerous lower quality elements around singularities and horn-attachments. Additionally, it led to the reduced quality of internal elements, as the lower qualities propagated internally. This mechanism was likely more sensitive to the low-resolution meshes. Therefore, the effect could have been smaller for

higher resolution meshes, as their would have been more layers to propagate through.

The minimisation effect defined the difference between the degree of singularity alignment with corresponding anatomical locations and the associated element-quality improvements. Specifically, the energy minimisation process could align the singularities with the natural edges of the meniscus geometries. This led to higher quality singularities than observed for the unoptimised and multi-block meshes. Some geometries likely required more iterations to reach their global minima, which contributed to the larger mesh-quality variance. In addition, the lack of iterations required to reach a global minimum would explain why singularities around horn-attachments, that had not lined-up, had a lower quality. This knowledge could be used to design templates more optimal for particular geometries. For example, lower quality peripheral edges could potentially be improved with an additional edge singularity along the rim, to capture higher curvature edges.

Despite several smoothing operations during each iteration and at the end of the optimisation, lower quality regions present in the unoptimised meshes remained or became worst. This observation was defined as the residual effect. This made it difficult to determine the pure effects from the optimisation, as some aspects were likely carried over from the unoptimised strategy. The main cause of the residual effect was due to the presence of lower quality elements in the unoptimised meshes. These were most frequently observed around convex surfaces, horn-attachments and pinched regions. Another cause for this effect was due to the optimised strategy inducing greater distortion and irregularities both globally and locally from the improved surface fit. In addition, the initial projection step in the optimisation imparted a large amount of distortion, which worsened the quality of most elements. The residual effect not only led to a global mesh-quality reduction but exacerbated lower quality elements locally. This highlighted an important factor, that the optimised mesh-quality was dependent on the unoptimised mesh-quality. Moreover, if an unoptimised mesh suffered flaws, such as the pinching effect, the overall performance of the optimised mesh would not be as favourable. Furthermore, these flaws could potentially induce a greater shrinking effect, which was observed for MM_21. The sensitivity of the optimised strategy to the quality of the unoptimised meshes increased performance variabilities of the derived meshes. Specifically, the variable template quality combined with a potentially sub-optimal number of minimisation steps led to a greater mesh-quality variance between optimised meshes. Also, as discussed previously, this in turn impacted the surface and volume errors as

it influenced the magnitude of the shrinking effect.

The optimised mesh derived from LM_06 had not only the lowest mesh-quality but was a statistical outlier. There were several reasons this mesh performed worse than the others. The strongest impact was due to the fact there were no non-positive Jacobians, following the iteration loop. The last three negative Jacobians were removed during the 15th (last) iteration. Therefore, no final smoothing step was required to complete the optimisation. Consequently, there was minimal improvement to the element quality, as there was no additional smoothing, other than what was performed at the end of each iteration. Also, there was a large proportion of lower quality elements but none were low enough to be classified as inverted. This could easily be solved by specifying a different criterion for the final smoothing step. The criterion could be defined based on a maximum desired number of low and/or very low-quality elements, but would likely lead to a greater shrinking effect for some meshes. Another aspect that contributed to the lower quality was the block-like shape of the target. All four block-like geometries caused performance issues for both the unoptimised and optimised strategies. Therefore, the residual effect would have been a dependent factor for the performance of these meshes. Additionally, LM_06 had a combination of stronger residual effects and thinner regions than the other meshes. The residual effects for LM_06 were most pronounced at the horn-attachments and singularities. The lower mesh-quality associated with block-like meshes were likely due to the use of a template mesh with a sub-optimal topology and singularity design for their shape, as discussed in relation to their surface error. Specifically, the template mesh was designed for more meniscus-shaped geometries, whereas the multi-block structure (deformed box) was more appropriate to match their 12 edges and 6 faces. In line with this, previous research has shown that the mesh-quality and surface error for geometrically irregular meshes could be improved with a more appropriate topological design [197, 448, 449]. Overall, the block-like meshes contributed to the larger variance associated with the proportions of low and very low-quality elements.

9.4.4 Consequences of the Optimised Strategy

There were several benefits and potential challenges associated with the use of the optimised strategy for developing finite element models of the tibio-femoral joint.

Challenges

There are three challenges associated with the use of the optimised meshes: (1) geometry shrinkage, (2) horn-attachment shape and (3) reduced convergence. The most impactful challenge could be the shrinking of the geometries. However, this could be overcome through a size-optimisation algorithm, as discussed previously. This could either ensure the original volume matches a target mesh, or ensure it is only slightly smaller to establish an adequate pairing between interfacing anatomies, e.g. the tibial and femoral cartilage. For some geometries, a large degree of shrinking would cause a loss of conformity between tibial and femoral cartilage meshes, if not resized. This would be similar to the use of unoptimised meshes but on the other end of the spectrum and through a different mechanism. Although the unoptimised meshes were sometimes too large, they would not be as easily solved with resizing due to the rounded geometry. In fact, resizing an unoptimised mesh to match the volumes would likely cause greater surface errors. Consequently, despite this aspect representing an important challenge, in all likelihood it could easily be resolved and lead to greater conformity. For meshes with a high degree of shrinkage, simulations derived directly from them without adjustment could cause an underestimation of contact pressures [63, 450]. Observations of lower contact pressures could lead to clinicians misdiagnosing knee pathologies, as increased contact pressures often indicate particular problems, such as osteoarthritis and degeneration [22, 154, 456]. Two investigations have identified a trend between decreasing meniscus size and a decrease to the associated contact pressure [63, 450]. Also, one of those studies found the shape of a meniscus to have less of an effect, when explored using statistical shape modelling (SSM) [63]. However, these findings contradict the observed results for the optimised meshes. The results from this chapter instead showed decreased sizes did not affect the contact variables. In contrast, the shape changes had more of an impact, as the smaller optimised meshes were more precise than the rounder shapes of the unoptimised meshes. However, there was a key difference between the shape changes for the current investigation and the one based on a statistical shape model. The unoptimised meshes were simply smoother and rounder versions of a target, whereas in the SSM study the variations were due to isolated population differences, like meniscal width [63]. Additionally, the size differences for those investigations were based on a geometry with the same shape [63, 450]. However, this was not the case for the optimised and unoptimised meshes, which had both different shapes and sizes. Also, there was a subtle difference between the initial conditions

of the optimised and multi-block mesh simulations compared to those derived from the unoptimised meshes. Simulation adjustments were not required for the optimised meshes, despite some shrinking, unlike the unoptimised meshes. Specifically, the optimised simulations used the same initial alignment between the tibial and femoral cartilage meshes as the multi-block simulations.

In general, the horn-attachments were better captured by the optimised strategy than the unoptimised strategy, but were still slightly different to multi-block meshes. Overall, the horn-attachment shapes showed less recession and were more similar to the multi-block meshes. However, the terminal vertices did show some contraction and recession. The shape of horn-attachments has been shown not to be a majorly sensitive factor [410]. However, if concerned, a user could solve this problem by simply re-selecting the terminal vertices of interest. This would lose an automatable aspect inherent from a template-based development approach. Nevertheless, this would be required for the multi-block meshes. Therefore, this would not lead to slower FE model development times than current practices.

The last challenge was related to the loss of convergence benefits that were associated with the unoptimised meshes. The optimised simulations were solved only slightly faster on average than the multi-block meshes. This was most likely due to the reduced geometric smoothness of the meshes and a greater proportion of lower quality elements. From the results, it could be inferred that the unoptimised meshes would be able to withstand more demanding (complex) simulations but provide lower precision. On the other hand, the opposite could be inferred for the optimised and multi-block meshes. Additionally, the optimised meshes could potentially be more convergent than the multi-block meshes, based on the evidence presented in this thesis. Specifically, the optimised meshes had a slightly greater mesh-quality and smoothness, as well as slightly lower convergence times. In the literature, these two factors have been shown to improve the convergence for demanding simulations, which endure large deformations [455] and contact mechanics [456].

Precision Improvements for Finite Element Analysis

There were four benefits the optimised strategy provided with respect to precision improvements for finite element analyses. First, the lower surface errors and overall shape similarities have contributed to more precise simulations, despite some meshes having relatively larger volumetric

losses. Interestingly, there was no correlation between higher or lower surface and volumetric errors with respect to the contact variable differences between the optimised meshes. However, the contact variable differences were larger for the unoptimised simulations, which did have higher surface errors. For the precision of the simulations, shape similarity was a more important factor than size similarity, between unoptimised and optimised simulations. The lack of correlation between simulation precision against the surface and volume errors, compared to the multi-block meshes, was highlighted by the following examples:

1. **[Example A] LM_01 (very close RMSE and max. error):**

- One of the most comparable optimised meshes to a corresponding multi-block mesh
- Had a contact area difference of $\sim 5\%$ (both interfaces)
- Lower difference than the corresponding unoptimised simulation

2. **[Example B] LM_14 (very close RMSE and a lower max. error):**

- Another comparably close optimised mesh
- Had a contact area difference of $\sim 20\text{-}40\%$ (between interfaces)
- Still lower than the corresponding unoptimised simulation

3. **[Example C] MM_21 (highest RMSE and max. error):**

- The most shrunk mesh and was the least comparable
- Had a contact area difference of $\sim 2\text{-}8\%$ (between interfaces)
- Lower difference than the corresponding unoptimised simulation

The first two examples highlighted that lower surface errors for the optimised meshes did not always translate to lower contact variable differences. The third example highlighted that despite large errors, the contact variables were still closer than the corresponding unoptimised simulation. In fact, the combined tibial and femoral contact area difference was less than both the unoptimised tibial and femoral differences in isolation. Additionally, the contact variable differences were remarkably smaller for MM_21 than other optimised meshes with much lower surface errors (**example C vs. example B**). Essentially, large surface errors and shrinking did not cause a greater loss of contact area or typically indicate larger contact variable differences.

The second improvement was that the contact variables showed greater patient-specificity and overall closer similarities to the multi-block simulations. This was evident by greater contact area precision and higher similarities between max. contact pressure locations. In summary, the contact area difference was almost 50% lower compared to the unoptimised simulations. Additionally, a much larger proportion of simulations had max. contact pressure values in either the same or similar locations.

The third improvement was the greater capture of the anatomical and geometric features. Notably, the simulations were closer representations of the original target geometries than the unoptimised meshes. This was most evident by the shape of the impressions, which were more similar to the target geometry. Also, the contact pressure distributions were more similar to the multi-block simulations. Not only were the global patterns more similar but the local features as well. This was most apparent by the capture of thin-features, which were generally more similar to the multi-block simulations than those derived from the unoptimised meshes.

The fourth aspect was an overall benefit compared to both the unoptimised and multi-block simulations. Specifically, the optimised simulations struck a good balance between the best simulation attributes derived from the unoptimised and multi-block meshes. Compared to the unoptimised simulations, the more patient-specific geometries led to more precise contact variables. When compared to the multi-block simulations, the smoother and rounder geometries with fewer irregularities led to more meniscus-like behaviours. Most notably, there were fewer discontinuities and less patchiness in contact pressure and area distributions.

9.5 Conclusions

In conclusion, there were several benefits the optimised strategy provided for the development of finite element models of a tibio-femoral joint. When compared to the unoptimised strategy, the optimised strategy provided greater conformity, patient-specificity and simulation precision. When considering the multi-block method, the optimised strategy was able to derive comparable meshes with much greater ease and speed. Moreover, all meshes were generated automatically and rapidly, in under 60 s. The greater overall similarities led to improved simulation precision and more comparable contact variables. Finally, the optimised meshes provided more realistic meniscus-like behaviours than the multi-block meshes.

Part IV

Discussion

10 Discussion

10.1 Summary

The development of high-quality patient-specific models can aid in both understanding, diagnosing and treating a large variety of pathologies, which impact society and are expensive to health-care systems. Consequently, there has been widespread development of patient-specific models for a variety of anatomies and joints [4, 466]. However, it is a challenging endeavour to develop high-quality models of this nature, as they require several laborious and time-consuming steps. Additionally, proving the validity of these models are non-trivial tasks.

One approach to solving this problem has been to use mesh-morphing techniques. There have been several strategies developed around the core problem of mesh generation for a range of anatomies (**Table 2.8**). However, it can be difficult to determine the suitability of a mesh-morphing strategy and the consequences of their use in clinical applications. Furthermore, core information pertaining to which areas of a strategy cause performance losses are rarely assessed with respect to simulation accuracy. Therefore, it is not always clear how to improve and optimise a strategy for particular applications.

The aims of this research were to develop novel automatic mesh-morphing strategies and methodologies. The methodologies were used to determine the robustness and performance of the mesh-morphing strategies used to develop biomedical FE models. The strategies were evaluated on meniscus anatomies that provide a range of complexity and challenges, which would be applicable to a wide variety of anatomies. Additionally, the methodologies provide a framework of how to assess general mesh-morphing strategies, and mesh-generation procedures, with a focus on practical considerations and goals. Consequently, the core aims of the thesis were achieved (**Section 1.2**).

The sensitivity assessments in **Chapters 5-7** aided the selection of the most optimised conditions for the validation of the strategy against a state-of-the-art procedure. In **Chapter 5** (case A), it was determined that centrelines should be optimised for acceptable length representation over the presence of deviations. In **Chapter 6** (case B), it was found that the strategy was sensitive to meniscus geometries with small wedge angles and high-curvature features. This

information aided the selection of challenging meniscus geometries and features for exploration in cases D and E. In *Chapter 7* (case C), it was found that the mesh-morphing strategy was sensitive to mesh-resolution and mesh-topology, where higher resolution meshes and meniscus-specific topologies yielded the best results. As such, a template mesh with a meniscus-specific topology was chosen for the subsequent performance analyses.

In *Chapter 8* (case D), the results indicated that the unoptimised strategy produced meshes with a more favourable mesh-quality over surface fit, and the opposite for the gold-standard. However, the unoptimised meshes produced more realistic behaviours during simulations. In *Chapter 9* (case E), the results indicated the optimised strategy achieved the goal of achieving lower surface errors, with some compromise to the mesh-quality and generation time. Overall, the optimised strategy was much more comparable to the gold-standard, whilst producing more realistic meniscus behaviours.

10.2 Development and Future Work of Morphing Strategies

There were several key achievements and findings associated with the development and design of the automatic centreline algorithm and mesh-morphing strategies:

1. Development of a novel and highly efficient centreline generation algorithm
2. Development of a novel mesh-morphing strategy that uses non-linear centrelines of both a target and template mesh
3. Development of a novel centreline-based mesh-morphing strategy which optimises surface fit by extending principles of spherical parameterisation to volumetric meshes

10.2.1 Centreline Algorithm

Following the implementation of the original centreline algorithm in this work [328], it became apparent that several components could be improved with respect to efficiency. The first novel improvement was the derivation of a formula that could minimise of a mesh into a centreline without the need to solve a large matrix (*Equation 3.54*). A novel finding was that the constants of this formula could be manipulated to derive centreplanes. Initially, centreplanes were considered a potential option for the mesh-morphing strategies, but instead enabled more precise surface error measurements of menisci with very thin regions.

The second novel improvement was to employ a decimation algorithm (*Section 3.2.3: Mesh Simplification*). The use of such an algorithm has not previously been used to improve the efficiency of these types of centreline algorithms. However, decimation algorithms have been found favourable in previous studies, among other techniques, for improving the efficiency of solving general RBF systems [467].

Robustness and Future Work

In case D, the improved algorithm was shown to be robust, where it was tested on 20 meniscus anatomies, and successfully derived centrelines for each. After parameter tuning (following *Chapter 5*), only one geometry (LM_07) failed to achieve a centreline without further parameter adjustments. LM_07 failed because it was a heavily extruded anatomy that naturally tended towards a centreplane. However, this was overcome by increasing the step parameter.

Other geometries that had large extrusions tended towards more deviated centrelines, as they also tended towards centreplanes. However, the centreplanes of those geometries had minimised enough for the post-processing algorithms to succeed in extracting a centreline, albeit slightly deviated. Several parameters are involved in the centreline algorithm, which includes the kernel radius of the RBF system and cluster radius. Consequently, complete tuning of the entire system could have yielded more robust results.

Alternatively, a modification that could improve the algorithm would measure the area of the triangle elements that defined the minimised surface. When the total area is below a given threshold (close to zero), it could be assumed a centreline has been derived. Otherwise, if the area threshold has not been met the algorithm could automatically adjust the step size accordingly. Also, the original algorithm required a final step that minimised points in a plane perpendicular to a coarse centreline [328]. Although, the improved algorithm did not require this step, it could be used differently to improve the general robustness of the algorithm. Furthermore, it could be used to control the length of the centrelines, by projecting the end-points onto their original surface. The new points along that curve could then be minimised, which would provide a centreline with end-points connected the surface mesh. A distance threshold would then define the removal of centreline points which are within the specified distance from the surface mesh.

Overall, the generated centrelines were accurate representations (qualitatively). In general, they

maintained equidistant positions with respect to their irregular surfaces, had an appropriate length and were reasonably smooth. However, the robustness of this algorithm has only been assessed on meniscus geometries.

10.2.2 Mesh-Morphing Strategy Improvements and Future Work

Detailed discussion on the novelty and benefits of the unoptimised and optimised strategies have been provided in *Sections 3.7* and *3.8*, respectively. The main problem for mesh-morphing consists of balancing the energies (\mathbf{E}^s and \mathbf{E}^v) represented by *Equation 3.77*, and ideally solving them together. By knowing what type of energies exist and how they effect morphing performance, more favourable or balanced energies could be derived. This would provide greater control over the equilibrium, which could be achieved with more sophisticated energy minimisation techniques. Naturally, several quadratic energies have been defined similar to \mathbf{E} for solving similar problems [402, 468–470]. Additionally, there are a wide range of traditional minimisers that are capable of efficiently solving quadratic energies directly. However, the components would have to accurately reflect both the problems of surface parameterisation and volumetric element distortion.

Currently, \mathbf{E}_s^o is represented reasonably accurately by the optimised strategy, and matches the formula for surface error (*Equation 4.3*). Although, the derivation of $V_i^{*\Omega}$ results from centreline-based projection, which could be replaced with a method that determines correspondences more accurately. A promising option could be to adopt techniques used to determine correspondences for SSMs. Some of these non-rigid registration techniques are quite sophisticated, e.g. particle-based optimisation [471]. Also, the use of an existing SSM might be useful for target approximation and identification, but has several aforementioned disadvantages. These techniques could be combined with a force-based volumetric transformation, which has yielded promising results for morphing phalanges [47]. This may eliminate the need for centreline-based projections in the optimisation steps, once a morphed mesh is more geometrically similar to a target.

The main issue with the optimised strategy is the volume energy \mathbf{E}_v^o (Laplacian component), which always induces shrinkage regardless of the degree of element distortion. An improved definition would use a Jacobian-based optimiser [343], which would directly and accurately represent \mathbf{E}_v . Alternatively, other quality metrics and combinations could be used to measure

volumetric distortion. Fortunately, there are a large number of these optimisers available, which have undergone extensive research [340, 343]. Also, these optimisers can reduce element distortion and improve their shape without inducing shrinking effects. Another approach would be to replace the contractive Laplacian energies with more sophisticated smoothing-based energies. There are a vast array of choices, with the majority requiring a straightforward modification to the constants defined by α_{ij} [260, 295, 472] (**Table 2.4**). However, as most applications involve triangular meshes, the constants would have to be suitable for hexahedral meshes and prevent mesh shrinkage [472]. Advanced and novel morphing energies (\mathbf{E}) could be composed by combining these smoothing- and metric-based energies (\mathbf{E}_v) with surface error components (\mathbf{E}_s). This would allow minimisers to solve a morphing operation in one process, instead of two that are distinct and opposing.

However, other approaches could focus on improving the components separately. There are several techniques that can reduce element distortions for volumetric transformations. This could be achieved by defining non-linear interfacial energies, which transfers boundary changes to the volumetric elements less uniformly and more controlled. For FE-based approaches, this can be achieved through the definition of non-linear or inhomogeneous material properties. Non-linear elastic material properties have been used to reduce distortion for large deformations [315]. Similarly, inhomogeneous properties can be achieved by iteratively increasing the stiffness of distorted elements, which have been found to reduce further degradation [54, 58]. Previously, a robust volumetric transformation has been proposed combining non-linear material properties and iterative stiffening [315]. For other methods, the log barrier technique has been found effective at reducing distortion and untangling elements [473]. Also, RBF-based transformations can be used to control the distribution and concentration of distortions [53]. A simpler approach could be achieved through volumetric smoothing, based on the order of point-wise processing. Specifically, this would smooth vertices from the outside-in, as the existing definitions would likely be ordered differently or randomly. This would gradually propagate the boundary modifications to the core of a mesh to minimise distortions.

Consideration should be given for more esoteric approaches which define energy functions based on existing systems. Previously, a novel surface parameterisation approach for vertebral mesh-morphing had taken inspiration from molecular dynamics to minimise the surface energy [52]. Similarly, concepts and techniques used to balance surface and volume energies could be taken

from other fields, e.g. chemistry, physics and metallurgy (alloying). However, a limitation of these methods are the requirement for proficient expertise [45]. Also, a hierarchical approach would reduce the computational complexity of any mesh-morphing strategy [47]. Furthermore, this technique has also been shown to reduce distortion for surface parameterisations [292]. Simple techniques which can decrease computational cost should generally be considered, as they will allow a greater turnover of complex systems and larger anatomies.

10.3 Development and Impact of Methodologies

There were six key findings and achievements associated with the development, analysis and design of the methodologies used to assess sensitivity, robustness and performance; which were the development of a methodology to:

1. measure surface error based on centrelines and centreplanes, which provided greater accuracy for capturing errors around thin regions
2. determine the sensitivity of a centreline-based mesh-morphing strategy to centreline attributes, comparing differences in length and deviations
3. determine target geometry sensitivities, using a parameterised meniscus geometry based on a SSM which captured extreme population differences
4. determine the sensitivity of a mesh-morphing strategy to template attributes, comparing different qualities, resolutions, topologies and shapes
5. determine and prove the robustness of a mesh-morphing strategy to a gold standard procedure, with comprehensive comparisons of the: surface error, volume error, mesh-quality, generation time, simulation precision and convergence
6. determine and prove if a different mesh-morphing strategy yields improvements to simulation precision and performance attributes

The following sections discuss how the developed methodologies could have an impact on evaluating and improving mesh-morphing strategies and mesh-generation procedures. Additionally, these sections describe the main findings resulting from these methodologies and their applicability to other mesh-morphing strategies.

10.3.1 Centreline-based Surface Error Measurement

A method that accurately measures surface error can be defined as a process that derives a bijective piecewise affine map, to calculate the corresponding pointwise distances. Two common methods for determining pointwise correspondences between two surfaces are closest point (vertex-vertex and vertex-plane [189, 474]) and normal projection [475]. However, the geometries must be reasonably similar or lack complex features, such as surface irregularities and thin regions. Additionally, these methods can be inaccurate, resulting in non-bijective maps. Consequently, several problems were encountered when initially using these methods for measuring the surface errors of the unoptimised meshes, as they had large local deviations. The main problem with the closest point methods were the underestimation of errors, due to identifying merely the nearest vertex, which was not always the most appropriate corresponding location. A similar requirement for bijective maps in biomedical engineering are to develop accurate SSMs [291, 476, 477]. For these applications, it has also been reported that closest point methods can lead to large correspondence errors and distortions [478]. As a result, the closest point methods must be utilised with sophisticated techniques to reduce distortions and achieve bijectivity [401, 478]. In contrast, normal projection is often the final step used to determine correspondences during SSM development [401]. For the surface error measurements of the unoptimised meshes, normal projection produced more realistic results and fewer correspondence errors than the closest point methods. However, there were several shortcomings, as the surface normal vectors could still project to non-corresponding locations, which could both over- and under-estimate errors. This was particularly noticeable in highly irregular and thin regions.

To overcome these shortcomings, a novel application of the centrelines and centreplanes were to use them to guide the projections for surface error measurements. This achieved better results, as they provided a constant or anchored location of correspondence between the target and morphed geometries. Additionally, this was the justification behind why centreline-based projection was found to be favourable over these methods for the mesh-morphing strategy. Centreplane-based projections were used to accurately measure the surface error of meshes with highly extruded regions, where the use of a centreline was less accurate. More accurate measurements could have been achieved using alternative methods. These could have been similar to techniques employed in the mesh-morphing strategies or used to develop SSMs, e.g. freeform deformation [401] or other surface parameterisation techniques [260]. However, to limit

effects from algorithms used in the morphing strategy, a more simplified method was chosen. Due to the inherent simplifications, the centreline- and centreplane-based projection algorithms alone could have produced some correspondence errors. However, after visual inspection they appeared to be more realistic, consistent and accurate than the other methods. Additionally, despite the presence of deviations in some centrelines, they did not appear to affect the results.

10.3.2 Sensitivity Evaluation of a Mesh-Morphing Strategy

The sensitivity analyses were broken down into three cases (A-C), evaluating the attributes of different centrelines, target geometry variations and template meshes. The outlined methodologies could not only be used to identify the sensitivity of different strategies for particular anatomies, but also how to optimise these variables. An aspect that has garnered little attention or reporting in the field of mesh-morphing.

Centreline Attributes

The first methodology was developed to determine the sensitivity of the unoptimised strategy to various centrelines (*Chapter 5*). This was used to determine if less ideal centrelines would be acceptable or whether further parameter refinement was required to derive more ideal centrelines. The most critical finding was that shorter centrelines had a negative effect on performance. It is quite likely that this finding would be applicable to other centreline-based mesh-morphing strategies [45, 369, 479, 480]. The main reason for this is that important information needed to determine the correspondences in regions proximal to the centreline end-points is lost. Furthermore, it would be difficult to imagine a trivial algorithm that could overcome this loss of information and mismatch between lengths. Another important finding was that deviations along the centreline, implemented as noise, did not affect the unoptimised strategy. However, it is likely this is not a general finding for centreline-based mesh-morphing strategies, particularly those that employ only pure projection from centrelines. This is because the unoptimised strategy consists of several algorithms that attempt to improve and unfold the initial projection from the centrelines. The execution of these algorithms occupy the majority of the computation time and do not involve the centreline structures. This was designed to eliminate any projection errors that could arise. In contrast, the unfolding algorithm was not able to overcome centreline length issues. Consequently, the focus of empirical parameter tuning selected against deviations in favour of acceptable centreline lengths.

Geometric Variations

The second methodology evaluated the influence of common and idealised target geometry variations with respect to meniscus anatomies (*Chapter 6*). This was another novel methodology which could aid in determining the suitability of a mesh-morphing strategy to known but simplified differences. Poor performances at this stage could eliminate a potential strategy and the need for further evaluations (i.e. cases D and E), or identify components of a strategy that require improvements. Furthermore, the methodology can be applied to a wide range of anatomies [241, 265, 481–484], where statistical shape modelling information is available to construct an idealised parametric model.

The findings from case B were used to identify challenging aspects about meniscus variations for further investigation in cases D and E. Another option would use a high-quality SSM directly. This would consolidate critical aspects of cases B, D and E and potentially evaluate a mesh-morphing strategy in a reduced time-frame. However, this would increase evaluation time-scales if the strategy could have been eliminated as an option after case B. Also, an evaluation protocol based on a SSM could be devised to contain both anatomically accurate population extremes and majorities. This would give a greater picture of general robustness apposed to robustness against challenging and complex geometries, which was the situation for cases D and E. This is an important component of validation to address, as population extremes and majorities may not always contain challenging features. Therefore, this aspect could speed up development but miss some edge cases, which the methodologies in cases D and E should address. However, a degree of both would be necessary for complete evaluation. Consequently, the mesh-morphing strategies are missing this component from their evaluation. Specifically, how robust they are in the general population and for what percentage they can yield respectable results.

The most general finding from case B was that large geometric differences between a template and target had a negative effect on performance. This correlates with literature, where mesh-morphing strategies perform better with smaller differences between a target and template [46–48]. One mesh-morphing strategy was applied against a range of phalanges with varying geometric differences using an anatomically derived phalanx template [46]. The target phalanx with the largest differences led to a morphed mesh with greater distortion and an invalid element (zero volume). Another general finding from case B was that large curvature differences and thin features caused performance losses. Similarly, thin features have been found to cause problems

for other mesh-morphing strategies [45,426]. Although, findings from cases D and E indicated the unoptimised strategy had greater difficulties capturing sharper, high-curvature and thin features than the other methods. Therefore, this finding from case B may be more exaggerated for the unoptimised strategy than other mesh-morphing strategies.

Template Mesh Attributes

The third novel methodology sought to address and answer the sensitivity of a mesh-morphing strategy with respect to template quality, resolution, topology and shape (*Chapter 7*). This could be applied generally to evaluate all mesh-morphing strategies and associated target anatomies. It allows the identification of the most suitable template attributes with respect to an anatomy of interest, which would likely differ between various structures. Therefore, the key findings from case C strongly depend on the unoptimised strategy in combination with meniscus anatomies, and those that would operate similarly. Most notably, the resolution, topology and shape of the template had the greatest effect on the performance. The effect of template shape and topology would primarily test the sensitivity of a general mesh-morphing strategy. Essentially, these determine how similar a strategy's template geometry must be to a target, and which topologies are more favourable for a particular anatomy. Also, it has several implications for the extension of mesh-morphing strategies to other anatomies. This is because it is less trivial to create a high-quality generic hexahedral mesh using CAD-based principles than through an image-based procedure (anatomical origin) [47,485]. Additionally, it is more straightforward to increase the resolution and control the topology.

Remarkably, template mesh-quality had minimal affects on the unoptimised strategy's performance. This contradicts several studies which emphasise starting with a high-quality template [45,189] due to mesh-morphing imparting element distortions [28,48,58,425]. Although there were quality losses when morphing the high-quality template, there were significant improvements for the lower quality templates. Essentially, the strategy was able to normalise the performance of the lower quality templates to achieve similar results to the high-quality template. Despite quality losses being a common result of mesh-morphing, some studies have found their strategies led to improvements [241] or no reductions [425]. One study attributed the quality improvements to smoothing processes imparted by their mesh-morphing strategy [241], which was the same conclusion drawn for the unoptimised strategy.

As a result, the quality of the morphed meshes derived from high- and low-quality templates were reasonably similar. However, quality losses are likely to be more prominent on higher quality meshes (as observed), as well as more topologically structured meshes (**Section 2.4.1: Regularity**). For the unoptimised strategy, this was attributed to the unfolding algorithm, which improved the shape of the boundary elements, in combination with the volumetric transformation. During linear elastic transformations (LET), the internal vertices play less of a role in the final result than the external vertices, whose positions were optimised during the unfolding algorithm. Essentially, the solution determined the optimal positions for the internal vertices regardless of their current positions, similar to RBF-based transformations. However, a key difference between FE- (i.e. LET) and RBF-based transformations is that the former has awareness of the appropriate DOF and shape for hexahedral elements [58, 189]. Also, without some form of surface optimisation, other mesh-morphing strategies would transfer the lower quality attributes of the boundary elements to the volumetric elements.

Of the four input variables, the effects of template resolution was the most general finding applicable to other mesh-morphing strategies. Specifically, where increased resolutions led to higher quality meshes and lower surface errors. This is supported by observations in literature that have indicated increasing resolutions yield more favourable morphed meshes [46, 47]. However, the resolution of a finite element mesh is often dictated by convergence studies [36, 63, 91, 126].

10.3.3 Validation Methodologies to Benchmark and Assess Performance

To the author's knowledge, no investigation has evaluated the performance of their mesh-morphing strategies to the extent reported in this thesis. Additionally, no automatic strategies have compared their simulation convergence nor evaluated their morphed meshes in multi-component contact-based simulations. However, there are examples of semi-automatic landmark-based strategies that have simulated morphed joints, e.g. the tibio-femoral joint (without meniscus) [27] and the pelvic joints [486]. Although, these did not compare the simulation convergence to a gold-standard model either. Of the reviewed literature, only one study had provided information on the surface error, mesh-quality, simulation validity (N=8) and generation times in comparison to a standard procedure [53]. Furthermore, due to the availability of in-vitro analysis on some of their target geometries, they were able to validate their simulations based on experimental data. Unfortunately, the strategies presented in this thesis could not be experimentally validated, as in-vitro data was not available.

The importance of the developed validation methodologies is that they allow direct and straightforward comparisons of robustness between methods with practical considerations. There were two other studies that each compared two strategies for biomedical applications [46, 52]. One performed simulation comparisons between the two strategies and a gold standard procedure, but with only of one specimen [52]. This provided an assessment based on the practical operation but did not demonstrate the strategies' robustness. In contrast, the other compared the performance of six phalanges between the two strategies, but performed no simulations [46]. Their assessment demonstrated aspects of robustness but did not assess how they would perform in real-world applications. Evidently, these studies failed to provide a complete picture of the consequences and benefits for using their strategies.

A goal of the reported validation methodologies were to establish a standard framework for assessing mesh-morphing strategies and generation methods. A standard framework would allow engineers to easily identify methods that would be preferable to their applications. In other cases, it may be preferable to simply eliminate methods that would be unsuitable. These types of frameworks have been employed in other areas of biomedical engineering to improve quality assurances and general comparisons (e.g. ISO-9001 and associated standard operating procedures [487]). Furthermore, it could allow methods to be ranked based on the four core metrics (*Section 2.7*), other evaluation criteria (*Section 2.4.1: Evaluation Criteria*) or a single metric (*Equation 3.77*). The availability of this information would simplify the selection process when considering the entire application (e.g. post-processing and analysis) and future work, not simply the goal of discretising an anatomy. This could normalise methods and target applications to a standard ranking system, instead of comparing performance metrics directly. Particularly, as different anatomies or applications may naturally impose and restrict certain metric thresholds. For example, large surface errors have been considered acceptable for complex anatomies, such as the spine and pelvis [48]. Similarly, meniscus meshes have been found to yield poorer performance metrics compared to other anatomies of knee [376]. Essentially, a general ranking system would compare where certain methods are situated in a hierarchy. This would operate similar to the h-index, used as a tool to compare general academic performance. Currently, there are several mesh-morphing strategies reported for biomedical applications, but not widespread utilisation of them. This is likely due to the general uncertainty of not knowing the consequences of their use, and if another could provide greater benefits.

10.4 Performance and Impact of Mesh-Morphing Strategies

There were several interesting findings and achievements associated with the application and performance assessments of the mesh-morphing strategies:

1. The strategies showed the same template can be used for both lateral and medial menisci
2. Demonstrated the rapid and automatic generation of comparative simulations using a template design, to transfer definitions of attachments and anisotropic material properties
3. Morphed meshes yielded more meniscus-like behaviours than gold-standard meshes (e.g. more evenly distributed contact pressures)
4. The unoptimised morphed meshes had a distinctly greater mesh-quality and solver convergence than gold-standard meshes
5. Key findings associated with the optimised strategy compared to the unoptimised strategy:
 - (a) Improved representation and capture of geometric features, with more accurate correspondences between template and targets (i.e. corners and wedge)
 - (b) Improved surface error without significant losses to mesh-quality and generation time, which improved the precision of the simulated contact variables
 - (c) Improved overall comparability to the gold-standard with respect to simulation precision and performance attributes

10.4.1 Single Template Mesh-Morphing for Different Anatomies

The performance of the strategies provided similar results when morphing a single generic template to lateral and medial menisci. However, there were key differences in their geometries, in general the medial menisci were larger, thicker and wider. These dimensional differences correlated with those observed in literature [405, 457]. Similarly, a mesh-morphing strategy reported in literature was tested on left and right kidneys using one template [241]. This study identified distinct differences with respect to the shape and volume of their left and right kidney geometries, which correlated with literature. Notably, the right kidneys had on average smaller volumes and thickness-to-length ratios than the left kidneys. Although the mesh-quality results were grouped together into a single plot, in general they appeared to be reasonably

similar. However, no surface error analysis or simulation comparisons were provided in that study. Further analysis may have revealed subtle performance differences between the left and right kidney meshes. However, these findings support the evidence that a one template can be morphed to similar but statistically different anatomies. Another mesh-morphing strategy exploited a bilateral symmetry for femoral geometries [53]. The surface error, mesh-quality and simulation accuracy of the morphed meshes were comparable to a gold standard procedure.

These studies [53, 241] mirrored target geometries to morph a single template to left and right anatomical variations. A similar technique has been performed on menisci to develop a generalised statistical shape model [94]. However, mirroring was not required for the CAD-based meniscus template due to the highly generic design. There were two key differences between the target and template geometries used in these studies compared to those used in this thesis: (1) the magnitude of geometrical differences and (2) template generation method. In relation to the first point, the left and right differences were less distinct than those between the medial and lateral menisci. For the second, they used a template derived from an image-based procedure as apposed to one from a CAD-based method. Due to these two points, there would have been fewer geometrical variations between their templates and targets. Consequently, the generic CAD-based template had to overcome a greater number of challenges than the anatomical image-based templates. An additional minor difference applies to the use of a tetrahedral template mesh that was morphed to femora [53]. Specifically, tetrahedral meshes are better able to accommodate greater geometry variations than hexahedral meshes.

For other mesh-morphing strategies, a single image-based template has been morphed to targets with greater geometry variations, such as vertebrae [48] and phalanges [46, 47]. Phalanges and vertebrae arguably represent greater intra- and inter-specimen variations compared to kidneys and femora. Using a single template for different phalanges (e.g. proximal, middle and distal) resulted in performance losses for anatomies with large differences (e.g. distal index) [46, 47]. In contrast, the results from morphing intra-specimen vertebrae yielded acceptable surface errors and quality meshes, although comparative analyses between the vertebrae were not provided [48]. Similar observations were found for the optimised strategy between lateral and medial menisci. Specifically, the medial menisci suffered greater volume and surface error losses, but this did not negatively affect simulation precision. Also, the analyses showed that both strategies provided more realistic results compared to the multi-block simulations, regardless of differences.

In general, these results are encouraging, as the human body is composed of bilateral symmetries and anatomies with similar shapes. Furthermore, it has been shown on several occasions that one template can yield acceptable results for similar but distinct anatomies. Consequently, fewer generic templates could be used to represent a greater proportion of different anatomies.

10.4.2 Rapid and Automatic Generation of Complex Simulations

The templated design enabled two manual aspects of FE model development to be automated. For the large-scale development of image-based models, these aspects would significantly improve the rate of productivity and turnover. The first aspect was the assignment of horn-attachment springs, which can be a tedious process that increases with greater resolution meshes. The second was the definition of fibre-orientations for anisotropic material properties. Fibre-orientations can either be assigned using local coordinates [20, 81], or the more laborious definition of cylindrical axes [488, 489]. Using local coordinate definitions were the obvious choice, as the topologies of the hexahedral meshes were aligned along fibre directions. For the multi-block models, the topological order was different between most of the meshes. Consequently, the end-points (for spring elements) and local directions had to be determined manually for each mesh. In contrast, the end-points and local directions of the template mesh were known. As such, this information was used to automatically define these properties of the morphed mesh simulations.

To the author's knowledge, the complete automation of these aspects with a templated design have not been demonstrated in literature. However, spring attachments for ligaments of a morphed pelvis have been transferred using a semi-automatic landmark-based strategy [486]. In this thesis, the full automation of these components have been used to evaluate the simulation performance for numerous meshes rapidly in cases D and E. Similarly, this could be used to rapidly evaluate further improvements or other mesh-morphing strategies. Previously, other studies have made conclusions about the potential of template designs to automate the definition of complex boundary conditions [53] and aspects of post-processing [52, 53]. For example, the definition of insertion sites for femoral muscles and tendons is a laborious task, which has been suggested could be overcome through mesh-morphing [53]. The key arguments for their potential lies in the shared characteristics between morphed meshes (i.e. topology and correspondence).

10.4.3 Application and Impact of the Mesh-Morphing Strategies

The developed mesh-morphing strategies could immediately be applied to generate hexahedral meshes for a large proportion of menisci. This would include those that are healthy and the majority of degenerative pathologies (e.g. extrusion, horn-laxity and partial meniscectomy). Additionally, the strategies could be used on any other anatomy that could be defined by a non-branching centreline, such as: long bones (e.g. femur, humerus and phalanges), musculoskeletal soft tissues (e.g. tendons, ligaments and muscles), livers, kidneys, ventricles, atria and non-branching vascular structures. The strategies could also be used on structures that can be represented by a single origin instead of multiple, i.e. a centreline. No strategy modifications would be required, as the geometric centroid would simply define the closest centreline position to all surface vertices. Anatomies that could be described by a single origin include: patellae, short bones (e.g. carpals and tarsals), intervertebral discs and the brain, *inter alia*. Furthermore, trivial modifications would enable the strategies to operate over branching centrelines. This would open the door to morphing more complex anatomical structures, such as: branched vascular structures, complex bones (e.g. pelvis and vertebra) and other branching structures (e.g. bronchi and renal calyces). However, branched centrelines between a template and target would have to be isomorphic.

Another trivial modification would allow the strategies to operate over centreplanes instead of centrelines. This would enable thin anatomies to be morphed, such as: articular cartilage (e.g. tibial and femoral), flat bones (e.g. cranium, scapula and ribs) and tissues of the eye (e.g. sclera and cornea). Although, some of these examples could potentially work with a single origin or centreline, e.g. tibial cartilage or ribs, respectively. However, this would result in large correspondence errors for other anatomies with higher curvatures like the cranium and femoral cartilage. Although trivial, the implementation could be more complex for some anatomies, compared to other trivial modifications. For example, parameterisation of a femoral centreplane using B-splines would require partitioning into more regular regions. However, other parameterisation methods are available, which could enable a modified strategy to not require a partitioned map. Additionally, some of these applications and modifications could require parameter adjustments, particularly if using a higher or lower resolution template mesh.

There are several anatomies where healthy and degenerative conditions would not allow the strategies to work without non-trivial modifications. In particular, there are three features that

would likely cause the strategies to fail:

1. Irregular sharp features
2. The presence of irregular or relatively small holes (tunnel loops of non-genus-0 meshes)
3. Large and irregular geometric deviations

N.B. *Irregular refers to a feature that would not be present in healthy pathologies*

For the meniscus this could be caused by the presence of tears (**Figure 2.2**) (c.f. features 1 and 2). For other anatomies these could be naturally occurring, e.g. the sacrum (sacral foramen), maxilla (infraorbital foramen) and mandible (mental foramen) (c.f. feature 2). However, the majority of issues would arise from complex tissue pathologies and injuries, such as: articular cartilage lesions, septal defects, aneurysms, tumors and bone fractures (c.f. features 1, 2, and 3). Some of these examples could potentially be overcome with a branched centreline or centreplane that could represent the irregular features, e.g. the sacrum and cartilage lesions.

However, there are several non-trivial modifications and methods that could allow some of these cases to succeed. Two modifications which would be the simplest to implement are the following:

1. Define a new template mesh, or topologically separate regions of a standard template mesh, to match a target geometry. Next, define branched centrelines or centreplanes, which captures those features for the target and template meshes
2. Remove the irregular features to match a standard case. Next, use a post-processing technique to insert those features into the morphed mesh

The automation of these modifications would require a tunnel loop detection (c.f. modifications 1 and 2) and filling algorithm (c.f. modification 2) [64,65]. These types of algorithms can detect the presence of holes (tunnel loops) and remove them from the geometry. Essentially, they have the ability to reduce the genus of a geometry. These non-trivial modifications would solve the problem of meniscal tears. However, an unavoidable outcome would be that the meshes would be topologically different to other morphed meshes. Therefore, some benefits of a template design could be lost, e.g. simplified FE post-processing and analysis. Also, the presence of very small holes (tunnel loops) could be difficult to capture with branched centrelines and centreplanes. Overall, the performance of these strategies would be strongly dependent on their ability to represent and define the irregular features. Also, some specimens may require hybrid structures

with centreline and centreplane components, to fully capture all geometric irregularities.

10.5 Conclusion

In conclusion, two novel mesh-morphing strategies have been developed, which were used to automatically generate hexahedral meshes of meniscus anatomies. Additionally, several novel methodologies have been presented that can evaluate the sensitivity and robustness of a mesh-morphing strategy. These were used to benchmark the performance of the strategies to a gold-standard procedure. In general, the strategies and methodologies provide frameworks which could be extended in a number of interesting ways. The strategies have demonstrated several advantages, most notably the automatic development of finite element models of the tibio-femoral joint. Additionally, the optimised strategy has achieved performances comparable to a gold-standard procedure. Furthermore, the optimised strategy could be used to automate the development of high-quality finite element models for a range of complex systems and applications. This could lead to significant reductions to the time and effort required to develop sophisticated simulations. Finally, methods of this nature provide an avenue for large-scale population-based studies and could assist in the development of a digital twin with greater productivity.

Bibliography

- [1] E. A. Makris, P. Hadidi, and K. A. Athanasiou, “The knee meniscus: Structure-function, pathophysiology, current repair techniques, and prospects for regeneration,” *Biomaterials*, vol. 32, no. 30, pp. 7411–7431, 2011.
- [2] J. K. Bryceland, A. J. Powell, and T. Nunn, “Knee Menisci: Structure, Function, and Management of Pathology,” *Cartilage*, vol. 8, no. 2, pp. 99–104, 2017.
- [3] A. Erdemir, “Open Knee: Open Source Modeling and Simulation in Knee Biomechanics,” *Journal of Knee Surgery*, 2014.
- [4] T. A. Burkhart, D. M. Andrews, and C. E. Dunning, “Finite element modeling mesh quality, energy balance and validation methods: A review with recommendations associated with the modeling of bone tissue,” *Journal of Biomechanics*, vol. 46, no. 9, pp. 1477–1488, 2013.
- [5] A. Johnen, J. C. Weill, and J. F. Remacle, “Robust and efficient validation of the linear hexahedral element,” *Procedia Engineering*, vol. 203, pp. 271–283, 2017.
- [6] M. Desbrun, M. Meyer, and P. Alliez, “Intrinsic parameterizations of surface meshes,” *Computer Graphics Forum*, vol. 21, no. 3, pp. 209–218, 2002.
- [7] O. Sorkine, “Laplacian Mesh Processing,” *Eurographics - State of the Art Reports*, no. Section 4, pp. 53–70, 2005.
- [8] B. Mocanu and T. Zaharia, “A complete framework for 3D mesh morphing,” *Proceedings - VRCAI 2012: 11th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and Its Applications in Industry*, vol. 1, no. 212, pp. 161–170, 2012.
- [9] N. Hraiech, F. Taddei, E. Malvesin, M. Rochette, and M. Viceconti, “Fast 3D mesh generation of femur based on planar parameterization and morphing,” *2008 5th IEEE International Symposium on Biomedical Imaging: From Nano to Macro, Proceedings, ISBI*, no. 3, pp. 1561–1564, 2008.
- [10] S. Haker, S. Angenent, A. Tannenbaum, R. Kikinis, G. Sapiro, and M. Halle, “Conformal surface parameterization for texture mapping,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 6, no. 2, pp. 181–189, 2000.

-
- [11] D. Sieger, S. Menzel, and M. Botsch, “High quality mesh morphing using triharmonic radial basis functions,” *Proceedings of the 21st International Meshing Roundtable, IMR 2012*, pp. 1–15, 2013.
- [12] M. L. Staten, S. J. Owen, S. M. Shontz, A. G. Salinger, and T. S. Coffey, “A comparison of mesh morphing methods for 3D shape optimization,” *Proceedings of the 20th International Meshing Roundtable, IMR 2011*, pp. 293–311, 2011.
- [13] D. Sieger, S. Menzel, and M. Botsch, “RBF morphing techniques for simulation-based design optimization,” *Engineering with Computers*, vol. 30, no. 2, pp. 161–174, 2014.
- [14] V. Bayona, N. Flyer, and B. Fornberg, “On the role of polynomials in RBF-FD approximations: III. Behavior near domain boundaries,” *Journal of Computational Physics*, vol. 380, pp. 378–399, 2019.
- [15] E. Lehto, “High Order Local Radial Basis Function Methods for Atmospheric Flow Simulations (Thesis),” Doctoral Thesis, Uppsala University, 2012.
- [16] C. Groth, S. Porziani, and M. E. Biancolini, “The medical digital twin assisted by Reduced Order Models and Mesh Morphing,” *Journal of biomechanical engineering*, vol. 140, no. October, pp. 8–9, 2018.
- [17] M. Khoshgoftar, A. C. T. Vrancken, T. G. van Tienen, P. Buma, D. Janssen, and N. Verdonschot, “The sensitivity of cartilage contact pressures in the knee joint to the size and shape of an anatomically shaped meniscal implant,” *Journal of Biomechanics*, vol. 48, no. 8, pp. 1427–1435, 2015.
- [18] P. Łuczkiwicz, K. Daszkiewicz, W. Witkowski, J. Chróścielewski, and W. Zarzycki, “Influence of meniscus shape in the cross sectional plane on the knee contact mechanics,” *Journal of Biomechanics*, vol. 48, no. 8, pp. 1356–1363, 2015.
- [19] T. M. Guess, G. Thiagarajan, M. Kia, and M. Mishra, “A subject specific multibody model of the knee with menisci,” *Medical Engineering and Physics*, vol. 32, no. 5, pp. 505–515, 2010.
- [20] J. Yao, “Sensitivities of Medial Meniscal Motion and Deformation to Material Properties of Articular Cartilage, Meniscus and Meniscal Attachments Using Design of Experiments Methods,” *Journal of Biomechanical Engineering*, vol. 128, no. 3, p. 399, 2005.

- [21] R. E. Carey, L. Zheng, A. K. Aiyangar, C. D. Harner, and X. Zhang, “Subject-Specific Finite Element Modeling of the Tibiofemoral Joint Based on CT, Magnetic Resonance Imaging and Dynamic Stereo-Radiography Data in Vivo,” *Journal of Biomechanical Engineering*, vol. 136, no. 4, p. 41004, 2014.
- [22] T. L. Donahue, M. L. Hull, M. M. Rashid, and C. R. Jacobs, “A finite element model of the human knee joint for the study of tibio-femoral contact,” *Journal of Biomechanical Engineering*, vol. 124, no. 3, pp. 273–280, 2002.
- [23] E. Peña, B. Calvo, M. A. Martínez, D. Palanca, and M. Doblaré, “Finite element analysis of the effect of meniscal tears and meniscectomies on human knee biomechanics,” *Clinical Biomechanics*, vol. 20, no. 5, pp. 498–507, 2005.
- [24] M. E. Mononen, M. K. Liukkonen, and R. K. Korhonen, “Utilizing Atlas-Based Modeling to Predict Knee Joint Cartilage Degeneration: Data from the Osteoarthritis Initiative,” *Annals of Biomedical Engineering*, vol. 47, no. 3, pp. 813–825, 2019.
- [25] R. Mootanah, C. W. Imhauser, F. Reisse, D. Carpanen, R. W. Walker, M. F. Koff, M. W. Lenhoff, S. R. Rozbruch, A. T. Fragomen, Z. Dewan, Y. M. Kirane, K. Cheah, J. K. Dowell, H. J. Hillstrom, C. W. Imhauser, F. Reisse, D. Carpanen, R. W. Walker, and M. F. Koff, “Development and validation of a computational model of the knee joint for the evaluation of surgical treatments for osteoarthritis,” vol. 5842, no. July 2017, 2014.
- [26] R. J. Cooper, R. K. Wilcox, and A. C. Jones, “Finite element models of the tibiofemoral joint: A review of validation approaches and modelling challenges,” *Medical Engineering and Physics*, vol. 74, pp. 1–12, 2019.
- [27] M. A. Baldwin, J. E. Langenderfer, P. J. Rullkoetter, and P. J. Laz, “Development of subject-specific and statistical shape models of the knee using an efficient segmentation and mesh-morphing approach,” *Computer Methods and Programs in Biomedicine*, vol. 97, no. 3, pp. 232–240, 2010.
- [28] S. Ji, J. C. Ford, R. M. Greenwald, J. G. Beckwith, K. Paulsen, L. A. Flashman, and T. W. McAllister, “Automated subject-specific, hexahedral mesh generation via image registration,” *Nano*, vol. 6, no. 9, pp. 2166–2171, 2008.

-
- [29] S. C. Tadepalli, A. Erdemir, S. Sett, and P. R. Cavanagh, “A Comparison of the Performance of Hexahedral and Tetrahedral,” pp. 1–2, 2017.
- [30] B. Rodriguez-Vila, P. Sánchez-González, I. Oropesa, E. J. Gomez, and D. M. Pierce, “Automated hexahedral meshing of knee cartilage structures—application to data from the osteoarthritis initiative,” *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 20, no. 14, pp. 1543–1553, 2017.
- [31] M. Kazemi, Y. Dabiri, and L. P. Li, “Recent advances in computational mechanics of the human knee joint.” *Computational and mathematical methods in medicine*, vol. 2013, p. 718423, 2013.
- [32] K. H. Shivanna, S. C. Tadepalli, V. A. Magnotta, and N. M. Grosland, “A framework for finite element mesh quality improvement and visualization in orthopaedic biomechanics,” pp. 2–3, 2016.
- [33] S. A. Maas, B. J. Ellis, G. A. Ateshian, and J. A. Weiss, “FEBio: Finite elements for biomechanics,” *Journal of Biomechanical Engineering*, vol. 134, no. 1, pp. 1–10, 2012.
- [34] S. A. Maas, B. J. Ellis, D. S. Rawlins, and J. A. Weiss, “Finite element simulation of articular contact mechanics with quadratic tetrahedral elements,” *Journal of Biomechanics*, vol. 49, no. 5, pp. 659–667, mar 2016.
- [35] M. Li and R. Tong, “All-hexahedral mesh generation via inside-out advancing front based on harmonic fields,” *Visual Computer*, vol. 28, no. 6-8, pp. 839–847, 2012.
- [36] A. Kiapour, A. M. Kiapour, V. Kaul, C. E. Quatman, S. C. Wordeman, and T. E. Hewett, “Finite element model of the knee for investigation of injury mechanisms: Development and validation,” *Journal of Biomechanical Engineering*, vol. 136, no. 1, pp. 1–14, 2014.
- [37] J. C. Gardiner and J. A. Weiss, “Subject-specific finite element analysis of the human medial collateral ligament during valgus knee loading,” *Journal of Orthopaedic Research*, vol. 21, no. 6, pp. 1098–1106, 2003.
- [38] N. M. Grosland, K. H. Shivanna, V. A. Magnotta, N. A. Kallemeyn, N. A. DeVries, S. C. Tadepalli, and C. Lisle, “IA-FEMesh: An open-source, interactive, multiblock approach to anatomic finite element model development,” *Computer Methods and Programs in Biomedicine*, vol. 94, no. 1, pp. 96–107, 2009.

- [39] J. F. Shepherd and C. R. Johnson, “Hexahedral mesh generation constraints,” *Engineering with Computers*, vol. 24, no. 3, pp. 195–213, 2008.
- [40] B. Rodriguez-Vila, P. Sánchez-González, I. Oropesa, E. J. Gomez, and D. M. Pierce, “Automated hexahedral meshing of knee cartilage structures – application to data from the osteoarthritis initiative,” *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 5842, no. December, pp. 1–11, 2017.
- [41] A. Natarajan, V. A. Magnotta, and N. M. Grosland, “Hexahedral meshing of subject-specific anatomic structures using registered building blocks,” *ASME 2010 Summer Bioengineering Conference, SBC 2010*, no. PARTS A AND B, pp. 457–458, 2010.
- [42] T. Blacker, “Automated Conformal Hexahedral Meshing Constraints, Challenges and Opportunities,” *Engineering with Computers*, vol. 17, no. 3, pp. 201–210, 2001.
- [43] E. E. Gassman, S. M. Powell, N. A. Kallemeyn, N. A. DeVries, K. H. Shivanna, V. A. Magnotta, A. J. Ramme, B. D. Adams, and N. M. Grosland, “Automated bony region identification using artificial neural networks: Reliability and validation measurements,” *Skeletal Radiology*, vol. 37, no. 4, pp. 313–319, 2008.
- [44] Y. S. Park, “Rater Drift in Constructed Response Scoring Via Latent Class,” 2011.
- [45] S. Behdadfar, L. Navarro, J. Sundnes, M. Maleckar, S. Ross, h. h. Odland, and S. Avril, “A Centerline Based Model Morphing Algorithm for Patient-Specific Finite Element Modelling of the Left Ventricle,” *IEEE Transactions on Biomedical Engineering*, vol. 9294, no. c, 2017.
- [46] V. A. Magnotta, W. Li, and N. M. Grosland, “Comparison of Displacement-Based and Force-Based Mapped Meshing.” *MIDAS journal*, vol. 2008, no. 1, p. 629, aug 2008.
- [47] N. M. Grosland, R. Bafna, and V. A. Magnotta, “Automated hexahedral meshing of anatomic structures using deformable registration,” *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 12, no. 1, pp. 35–43, 2009.
- [48] P. Hadagali, J. R. Peters, and S. Balasubramanian, “Morphing the feature-based multi-blocks of normative/healthy vertebral geometries to scoliosis vertebral geometries: development of personalized finite element models,” *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 21, no. 4, pp. 297–324, 2018.

- [49] B. Delaunay, “Sur la sphere vide. A la memoire de Georges Voronoi,” *Bulletin de l’Academie des Sciences de l’URSS. Classe des sciences mathematiques et na*, no. 6, pp. 793–800, 1934.
- [50] M. Alexa, “Recent advances in mesh morphing,” *Computer Graphics Forum*, vol. 21, no. 2, pp. 173–196, 2002.
- [51] Z. Li, J. Hu, and J. Zhang, “Comparison of Different Radial Basis Functions in Developing Subject-Specific Infant Head Finite Element Models for Injury Biomechanics Study,” in *ASME 2012 Summer Bioengineering Conference, Parts A and B*, no. 2. American Society of Mechanical Engineers, jun 2012, pp. 49–50.
- [52] I. A. Sigal, M. R. Hardisty, and C. M. Whyne, “Mesh-morphing algorithms for specimen-specific finite element modeling,” *Journal of Biomechanics*, vol. 41, no. 7, pp. 1381–1389, 2008.
- [53] L. Grassi, N. Hraiech, E. Schileo, M. Ansaloni, M. Rochette, and M. Viceconti, “Evaluation of the generality and accuracy of a new mesh morphing procedure for the human femur,” *Medical Engineering and Physics*, vol. 33, no. 1, pp. 112–120, 2011.
- [54] M. T. Bah, P. B. Nair, and M. Browne, “Mesh morphing for finite element analysis of implant positioning in cementless total hip replacements,” *Medical Engineering and Physics*, vol. 31, no. 10, pp. 1235–1243, 2009.
- [55] Z. Salo, M. Beek, and C. M. Whyne, “Evaluation of mesh morphing and mapping techniques in patient specific modeling of the human pelvis,” *International Journal for Numerical Methods in Biomedical Engineering*, vol. 29, no. 1, pp. 104–113, jan 2013.
- [56] F. Xu and S. Kenjereš, “Numerical simulations of flow patterns in the human left ventricle model with a novel dynamic mesh morphing approach based on radial basis function,” *Computers in Biology and Medicine*, vol. 130, no. January, 2021.
- [57] L. Baghdadi, D. A. Steinman, and H. M. Ladak, “Template-based finite-element mesh generation from medical images,” *Computer Methods and Programs in Biomedicine*, vol. 77, no. 1, pp. 11–21, 2005.
- [58] N. Lauzeral, D. Borzacchiello, M. Kugler, D. George, Y. Rémond, A. Hostettler, and F. Chinesta, “Shape parametrization of bio-mechanical finite element models based on

- medical images,” *Computer Methods in Biomechanics and Biomedical Engineering: Imaging and Visualization*, vol. 7, no. 5-6, pp. 480–489, 2019.
- [59] R. R. Upendra, B. J. Wentz, R. Simon, S. M. Shontz, and C. A. Linte, “CNN-based Cardiac Motion Extraction to Myocardial Models from Cine MRI,” vol. 2, pp. 1–9.
- [60] A. L. McNulty and F. Guilak, “Mechanobiology of the meniscus,” *Journal of Biomechanics*, vol. 48, no. 8, pp. 1469–1478, jun 2015.
- [61] G. Du, H. Zhan, D. Ding, S. Wang, X. Wei, F. Wei, J. Zhang, B. Bilgen, A. M. Reginato, B. C. Fleming, J. Deng, and L. Wei, “Abnormal Mechanical Loading Induces Cartilage Degeneration by Accelerating Meniscus Hypertrophy and Mineralization After ACL Injuries In Vivo.” *The American journal of sports medicine*, 2016.
- [62] T. L. Haut Donahue, M. L. Hull, M. M. Rashid, and C. R. Jacobs, “How the stiffness of meniscal attachments and meniscal material properties affect tibio-femoral contact pressure computed using a validated finite element model of the human knee joint,” *Journal of Biomechanics*, vol. 36, no. 1, pp. 19–34, 2003.
- [63] M. Khoshgoftar, A. C. T. Vrancken, T. G. van Tienen, P. Buma, D. Janssen, and N. Verdonchot, “The sensitivity of cartilage contact pressures in the knee joint to the size and shape of an anatomically shaped meniscal implant,” *Journal of Biomechanics*, vol. 48, no. 8, pp. 1427–1435, 2015.
- [64] I. Guskov and Z. J. Wood, “Topological noise removal,” *Proceedings - Graphics Interface*, pp. 19–26, 2001.
- [65] T. K. Dey, K. Li, J. Sun, and D. Cohen-Steinen, “Computing geometry-aware handle and tunnel loops in 3D models,” *SIGGRAPH’08: International Conference on Computer Graphics and Interactive Techniques, ACM SIGGRAPH 2008 Papers 2008*, vol. m, 2008.
- [66] A. J. Fox, F. Wanivenhaus, A. J. Burge, R. F. Warren, and S. A. Rodeo, “The human meniscus: A review of anatomy, function, injury, and advances in treatment,” *Clinical Anatomy*, vol. 28, no. 2, pp. 269–287, 2015.
- [67] S. C. Mordecai, N. Al-Hadithy, H. E. Ware, and C. M. Gupte, “Treatment of meniscal tears: An evidence based approach.” *World journal of orthopedics*, vol. 5, no. 3, pp. 233–41, 2014.

- [68] I. D. McDermott, S. D. Masouros, and A. A. Amis, “Biomechanics of the menisci of the knee,” *Current Orthopaedics*, vol. 22, no. 3, pp. 193–201, 2008.
- [69] V. Vedi, a. Williams, S. J. Tennant, E. Spouse, D. M. Hunt, and W. M. Gedroyc, “Meniscal movement. An in-vivo study using dynamic MRI.” *The Journal of bone and joint surgery. British volume*, vol. 81, no. 1, pp. 37–41, 1999.
- [70] A. E. Kedgley, T. H. Saw, N. A. Segal, U. N. Hansen, A. M. Bull, and S. D. Masouros, “Predicting meniscal tear stability across knee-joint flexion using finite-element analysis,” *Knee Surgery, Sports Traumatology, Arthroscopy*, vol. 27, no. 1, pp. 206–214, 2019.
- [71] T. Kusayama, C. D. Harner, G. J. Carlin, J. W. Xerogeanes, and B. A. Smith, “Anatomical and biomechanical characteristics of human meniscofemoral ligaments,” *Knee Surgery, Sports Traumatology, Arthroscopy*, vol. 2, no. 4, pp. 234–237, 1994.
- [72] D. Kohn and B. Moreno, “Meniscus insertion anatomy as a basis for meniscus replacement: A morphological cadaveric study,” *Arthroscopy: The Journal of Arthroscopic and Related Surgery*, vol. 11, no. 1, pp. 96–103, 1995.
- [73] H. Naghibi Beidokhti, D. Janssen, S. van de Groes, J. Hazrati, T. Van den Boogaard, and N. Verdonschot, “The influence of ligament modelling strategies on the predictive capability of finite element models of the human knee joint,” *Journal of Biomechanics*, vol. 65, pp. 1–11, 2017.
- [74] S. Sathasivam and P. S. Walker, “A computer model with surface friction for the prediction of total knee kinematics,” *Journal of Biomechanics*, vol. 30, no. 2, pp. 177–184, 1997.
- [75] Q. Meng, Z. Jin, R. Wilcox, and J. Fisher, “Computational investigation of the time-dependent contact behaviour of the human tibiofemoral joint under body weight,” *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, vol. 228, no. 11, pp. 1193–1207, 2014.
- [76] R. J. Cooper, A. Liu, G. A. Day, V. N. Wijayathunga, L. M. Jennings, R. K. Wilcox, and A. C. Jones, “Development of robust finite element models of porcine tibiofemoral joints loaded under varied flexion angles and tibial freedoms,” *Journal of the Mechanical Behavior of Biomedical Materials*, vol. 109, no. September 2019, p. 103797, 2020.

- [77] A. Hosseini, S. K. Van de Velde, M. Kozanek, T. J. Gill, A. J. Grodzinsky, H. E. Rubash, and G. Li, “In-vivo time-dependent articular cartilage contact behavior of the tibiofemoral joint,” *Osteoarthritis and Cartilage*, vol. 18, no. 7, pp. 909–916, 2010.
- [78] M. Logerstedt, DS, Snyder-Mackler L, Ritter, RC, Axe, “Knee pain and mobility impairments: meniscal and articular cartilage lesions.” vol. 40, no. 6, pp. 1–58, 2011.
- [79] M. J. Berthiaume, J. P. Raynauld, J. Martel-Pelletier, F. Labonté, G. Beaudoin, D. A. Bloch, D. Choquette, B. Haraoui, R. D. Altman, M. Hochberg, J. M. Meyer, G. A. Cline, and J. P. Pelletier, “Meniscal tear and extrusion are strongly associated with progression of symptomatic knee osteoarthritis as assessed by quantitative magnetic resonance imaging,” *Annals of the Rheumatic Diseases*, vol. 64, no. 4, pp. 556–563, 2005.
- [80] T. M. Guess, S. Razu, H. Jahandar, and A. Stylianou, “Predicted loading on the menisci during gait: The effect of horn laxity,” *Journal of Biomechanics*, vol. 48, no. 8, pp. 1490–1498, jun 2015.
- [81] M. Kazemi, L. P. Li, P. Savard, and M. D. Buschmann, “Creep behavior of the intact and meniscectomy knee joints,” *Journal of the Mechanical Behavior of Biomedical Materials*, vol. 4, no. 7, pp. 1351–1358, 2011.
- [82] K. Zhang, L. Li, L. Yang, J. Shi, L. Zhu, H. Liang, X. Wang, X. Yang, and Q. Jiang, “Effect of degenerative and radial tears of the meniscus and resultant meniscectomy on the knee joint: a finite element analysis,” *Journal of Orthopaedic Translation*, vol. 18, pp. 20–31, 2019.
- [83] H. I. Khan, D. Aitken, C. Ding, L. Blizzard, J. P. Pelletier, J. Martel-Pelletier, F. Cicuttini, and G. Jones, “Natural history and clinical significance of meniscal tears over 8 years in a midlife cohort Orthopedics and biomechanics,” *BMC Musculoskeletal Disorders*, vol. 17, no. 1, pp. 1–9, 2016.
- [84] M. R. Klement and P. F. Sharkey, “The Significance of Osteoarthritis-associated Bone Marrow Lesions in the Knee,” *Journal of the American Academy of Orthopaedic Surgeons*, vol. 27, no. 20, pp. 752–759, 2019.
- [85] P. Łuczkiwicz, K. Daszkiewicz, J. Chróścielewski, W. Witkowski, and P. J. Winklewski, “The Influence of Articular Cartilage Thickness Reduction on Meniscus Biomechanics,”

- Plos One*, pp. 1–13, 2016.
- [86] L. Sharma, F. Eckstein, J. Song, A. Guermazi, P. Prasad, D. Kapoor, S. Cahue, M. Marshall, M. Hudelmaier, and D. Dunlop, “Relationship of meniscal damage, meniscal extrusion, malalignment, and joint laxity to subsequent cartilage loss in osteoarthritic knees,” *Arthritis and Rheumatism*, vol. 58, no. 6, pp. 1716–1726, 2008.
- [87] B. Heidari, “Knee osteoarthritis prevalence, risk factors, pathogenesis and features: Part I,” *Caspian Journal of Internal Medicine*, vol. 2, no. 2, pp. 205–212, 2011.
- [88] M. A. Baldwin, C. W. Clary, C. K. Fitzpatrick, J. S. Deacy, L. P. Maletsky, and P. J. Rullkoetter, “Dynamic finite element knee simulation for evaluation of knee replacement mechanics,” *Journal of Biomechanics*, vol. 45, no. 3, pp. 474–483, 2012.
- [89] K. T. Kang, J. Son, D. S. Suh, S. K. Kwon, O. R. Kwon, and Y. G. Koh, “Patient-specific medial unicompartmental knee arthroplasty has a greater protective effect on articular cartilage in the lateral compartment: A finite element analysis,” *Bone and Joint Research*, vol. 7, no. 1, pp. 20–27, 2018.
- [90] M. Z. Bendjaballah, A. Shirazi-Adl, and D. J. Zukor, “Biomechanics of the human knee joint in compression: reconstruction, mesh generation and finite element analysis,” *The Knee*, vol. 2, no. 2, pp. 69–79, 1995.
- [91] M. Tuncer, J. P. Cobb, U. N. Hansen, and A. A. Amis, “Validation of multiple subject-specific finite element models of unicompartmental knee replacement,” *Medical Engineering and Physics*, vol. 35, no. 10, pp. 1457–1464, 2013.
- [92] K. S. Halonen, M. E. Mononen, J. S. Jurvelin, J. Töyräs, J. Salo, and R. K. Korhonen, “Deformation of articular cartilage during static loading of a knee joint - Experimental and finite element analysis,” *Journal of Biomechanics*, vol. 47, no. 10, pp. 2467–2474, 2014.
- [93] B. Akpinar, E. Thorhauer, S. Tashman, J. J. Irrgang, F. H. Fu, and W. J. Anderst, “Tibiofemoral Cartilage Contact Differences Between Level Walking and Downhill Running,” *Orthopaedic Journal of Sports Medicine*, vol. 7, no. 4, pp. 1–7, 2019.
- [94] A. C. Vrancken, S. P. Crijns, M. J. Ploegmakers, C. O’Kane, T. G. van Tienen, D. Janssen, P. Buma, and N. Verdonschot, “3D geometry analysis of the medial meniscus - A statistical

- shape modeling approach,” *Journal of Anatomy*, vol. 225, no. 4, pp. 395–402, 2014.
- [95] R. Shirazi, A. Shirazi-Adl, and M. Hurtig, “Role of cartilage collagen fibrils networks in knee joint biomechanics under compression,” *Journal of Biomechanics*, vol. 41, no. 16, pp. 3340–3348, 2008.
- [96] F. Galbusera, M. Bashkuev, H. J. Wilke, A. Shirazi-Adl, and H. Schmidt, “Comparison of various contact algorithms for poroelastic tissues,” *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 17, no. 12, pp. 1323–1334, 2014.
- [97] T. L. Haut Donahue, M. L. Hull, M. M. Rashid, and C. R. Jacobs, “The sensitivity of tibiofemoral contact pressure to the size and shape of the lateral and medial menisci,” *Journal of Orthopaedic Research*, vol. 22, no. 4, pp. 807–814, 2004.
- [98] S. Gilbert, T. Chen, I. D. Hutchinson, D. Choi, C. Voigt, R. F. Warren, and S. A. Maher, “Dynamic contact mechanics on the tibial plateau of the human knee during activities of daily living,” *Journal of Biomechanics*, vol. 47, no. 9, pp. 2006–2012, jun 2014.
- [99] E. M. Suero, N. Hawi, R. Westphal, Y. Sabbagh, M. Citak, F. M. Wahl, C. Krettek, and E. Liodakis, “The effect of distal tibial rotation during high tibial osteotomy on the contact pressures in the knee and ankle joints,” *Knee Surgery, Sports Traumatology, Arthroscopy*, vol. 25, no. 1, pp. 299–305, 2017.
- [100] K. T. Danielson, “Fifteen node tetrahedral elements for explicit methods in nonlinear solid dynamics,” *Computer Methods in Applied Mechanics and Engineering*, vol. 272, pp. 160–180, 2014.
- [101] B. L. De Oliveira and J. Sundnes, “Comparison of tetrahedral and hexahedral meshes for finite element simulation of cardiac electro-mechanics,” *ECCOMAS Congress 2016 - Proceedings of the 7th European Congress on Computational Methods in Applied Sciences and Engineering*, vol. 1, no. June, pp. 164–177, 2016.
- [102] O. C. Zienkiewicz, “The Finite Element Method: Its Basis and Fundamentals,” *The Finite Element Method: its Basis and Fundamentals*, p. iii, 2013.
- [103] O. C. Zienkiewicz and R. L. Taylor, “The Finite Element Method for Solid and Structural Mechanics,” in *The Finite Element Method for Solid and Structural Mechanics*, 2014.

-
- [104] T. C. Baudouin, J. F. Remacle, E. Marchandise, F. Henrotte, and C. Geuzaine, “A frontal approach to hex-dominant mesh generation,” *Advanced Modeling and Simulation in Engineering Sciences*, vol. 1, no. 1, pp. 1–30, 2014.
- [105] S. H. Lo and C. K. Lee, “On using meshes of mixed element types in adaptive finite element analysis,” *Finite Elements in Analysis and Design*, vol. 11, no. 4, pp. 307–336, 1992.
- [106] S. Yamakawa and K. Shimada, “Fully-automated hex-dominant mesh generation with directionality control via packing rectangular solid cells,” *International Journal for Numerical Methods in Engineering*, vol. 57, no. 15, pp. 2099–2129, 2003.
- [107] N. Sukumar and A. Tabarraei, “Conforming polygonal finite elements,” *International Journal for Numerical Methods in Engineering*, vol. 61, no. 12, pp. 2045–2066, 2004.
- [108] C. J. Stimpson, C. D. Ernst, P. Knupp, P. P. Pébay, and D. Thompson, “The Verdict Geometric Quality Library,” 2007.
- [109] B. W. Clark, “The Development and Evaluation of the knife Finite Element,” 1996.
- [110] ABAQUS-Documentation, “Three-dimensional solid element library,” 2021.
- [111] S. Maas, D. Rawlins, J. Weiss, and G. Ateshian, “FEBio User’s Manual Version 2.9,” 2019.
- [112] J. F. Remacle, J. Lambrechts, B. Seny, E. Marchandise, A. Johnen, and C. Geuzainet, “Blossom-Quad: A non-uniform quadrilateral mesh generator using a minimum-cost perfect-matching algorithm,” *International Journal for Numerical Methods in Engineering*, vol. 89, no. 9, pp. 1102–1119, 2012.
- [113] S. C. Tadepalli, A. Erdemir, and P. R. Cavanagh, “Comparison of hexahedral and tetrahedral elements in finite element analysis of the foot and footwear,” *Journal of Biomechanics*, vol. 44, no. 12, pp. 2337–2343, 2011.
- [114] S. A. Maas, G. A. Ateshian, and J. A. Weiss, “FEBio: History and Advances,” *Annual Review of Biomedical Engineering*, vol. 19, no. 1, pp. 279–299, 2017.
- [115] ABAQUS-Documentation, “Solid (continuum) elements,” 2021.

- [116] S. E. Benzley, E. Perry, K. Merkley, B. Clark, and G. Sjaardema, "A Comparison of All-Hexahedral and All-Tetrahedral Finite Element Meshes for Elastic and Elasto-Plastic Analysis," *4th International Meshing Roundtable, Sandia National Laboratories*, pp. 179–191, 1995.
- [117] A. Ruggiero, R. D'Amato, and S. Affatato, "Comparison of meshing strategies in THR finite element modelling," *Materials*, vol. 12, no. 14, pp. 1–11, 2019.
- [118] A. Sundin, "Accuracy and Reliability of Plane Hybrid Mixed Elements for Two-Dimensional Elasticity (Thesis)," Master's thesis, Luleå University of Technology, 1991.
- [119] G. Irving, J. Teran, and R. Fedkiw, "Tetrahedral and hexahedral invertible finite elements," *Graphical Models*, vol. 68, no. 2, pp. 66–89, 2006.
- [120] K. J. Bathe, *Finite Element Procedures*, 1996.
- [121] A. Karimi, R. Grytz, S. M. Rahmati, C. A. Girkin, and J. C. Downs, *Analysis of the effects of finite element type within a 3D biomechanical model of a human optic nerve head and posterior pole*. Elsevier B.V., 2021, vol. 198.
- [122] W. Zhao and S. Ji, "Mesh Convergence Behavior and the Effect of Element Integration of a Human Head Injury Model." *Annals of biomedical engineering*, vol. 47, no. 2, pp. 475–486, feb 2019.
- [123] E. A. de Souza Neto, F. M. Andrade Pires, and D. R. Owen, "F-bar-based linear triangles and tetrahedra for finite strain analysis of nearly incompressible solids. Part I: Formulation and benchmarking," *International Journal for Numerical Methods in Engineering*, vol. 62, no. 3, pp. 353–383, 2005.
- [124] K. Polgar, M. Viceconti, and J. J. O'Connor, "A comparison between automatically generated linear and parabolic tetrahedra when used to mesh a human femur," *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, 2001.
- [125] J. F. Remacle, R. Gandham, and T. Warburton, "GPU accelerated spectral finite elements on all-hex meshes," *Journal of Computational Physics*, vol. 324, pp. 246–257, 2016.

- [126] N. A. Devries, K. H. Shivanna, S. C. Tadepalli, V. A. Magnotta, and N. M. Grosland, “Ia-FEMesh: anatomic FE models—a check of mesh accuracy and validity.” *The Iowa orthopaedic journal*, vol. 29, pp. 48–54, 2009.
- [127] A. Ramos and J. A. Simões, “Tetrahedral versus hexahedral finite elements in numerical modelling of the proximal femur,” *Medical Engineering and Physics*, vol. 28, no. 9, pp. 916–924, 2006.
- [128] A. O. Cifuentes and A. Kalbag, “A performance study of tetrahedral and hexahedral elements in 3-D finite element structural analysis,” *Finite Elements in Analysis and Design*, vol. 12, no. 3-4, pp. 313–318, 1992.
- [129] M. Freutel, F. Galbusera, A. Ignatius, and L. Dürselen, “Material properties of individual menisci and their attachments obtained through inverse FE-analysis,” *Journal of Biomechanics*, vol. 48, no. 8, pp. 1343–1349, 2015.
- [130] P. S. Donzelli, R. L. Spilker, G. A. Ateshian, and V. C. Mow, “Contact analysis of biphasic transversely isotropic cartilage layers and correlations with tissue failure,” *Journal of Biomechanics*, vol. 32, no. 10, pp. 1037–1047, 1999.
- [131] L. Yin and D. M. Elliott, “A biphasic and transversely isotropic mechanical model for tendon: Application to mouse tail fascicles in uniaxial tension,” *Journal of Biomechanics*, vol. 37, no. 6, pp. 907–916, 2004.
- [132] J. J. Garcia, N. J. Altiero, and R. C. Haut, “An approach for the stress analysis of transversely isotropic biphasic cartilage under impact load,” *Journal of Biomechanical Engineering*, vol. 120, no. 5, pp. 608–613, 1998.
- [133] R. Maulik, “A transversely isotropic viscoelastic biphasic model of articular cartilage in unconfined compression viscoelastic biphasic model (Thesis),” Master’s thesis, Birla Institute of Technology, 2015.
- [134] R. Mootanah, C. W. Imhauser, F. Risse, D. Carpanen, R. W. Walker, M. F. Koff, M. W. Lenhoff, S. R. Rozbruch, A. T. Fragomen, Z. Dewan, Y. M. Kirane, K. Cheah, J. K. Dowell, and H. J. Hillstrom, “Development and validation of a computational model of the knee joint for the evaluation of surgical treatments for osteoarthritis.” *Computer methods in biomechanics and biomedical engineering*, vol. 17, no. 13, pp. 1502–17, 2014.

- [135] P. S. Donzelli, R. L. Spilker, G. A. Ateshian, and V. C. Mow, “Contact analysis of biphasic transversely isotropic cartilage layers and correlations with tissue failure,” *Journal of Biomechanics*, vol. 32, no. 10, pp. 1037–1047, 1999.
- [136] S. Federico, A. Grillo, G. La Rosa, G. Giaquinta, and W. Herzog, “A transversely isotropic, transversely homogeneous microstructural- statistical model of articular cartilage,” *Journal of Biomechanics*, vol. 38, no. 10, pp. 2008–2018, 2005.
- [137] S. M. M. Elhamian, M. Alizadeh, M. M. Shokrieh, and A. Karimi, “A depth dependent transversely isotropic micromechanic model of articular cartilage,” *Journal of Materials Science: Materials in Medicine*, vol. 26, no. 2, pp. 1–10, 2015.
- [138] S. G. Kulkarni, X. L. Gao, S. E. Horner, R. F. Mortlock, and J. Q. Zheng, “A transversely isotropic visco-hyperelastic constitutive model for soft tissues,” *Mathematics and Mechanics of Solids*, vol. 21, no. 6, pp. 747–770, 2016.
- [139] K. M. Fischenich, J. A. Wahlquist, R. L. Wilmoth, L. Cai, C. P. Neu, and V. L. Ferguson, “Human articular cartilage is orthotropic where microstructure, micromechanics, and chemistry vary with depth and split-line orientation,” *Osteoarthritis and Cartilage*, vol. 28, no. 10, pp. 1362–1372, 2020.
- [140] T. M. Guess, H. Liu, S. Bhashyam, and G. Thiagarajan, “A multibody knee model with discrete cartilage prediction of tibio-femoral contact mechanics,” *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 16, no. 3, pp. 256–270, 2013.
- [141] A. Bilgici, C. Dogan, E. Çil, Sakarya, Ö. Kuru, and M. B. Selcuk, “Relationship between pain severity and magnetic resonance imaging features in patients with osteoarthritis of the Knee,” *Turkish Journal of Rheumatology*, vol. 25, no. 4, pp. 184–190, 2010.
- [142] A. Price, D. Beard, and E. Thienpont, “Uncertainties surrounding the choice of surgical treatment for ‘bone on bone’ medial compartment osteoarthritis of the knee,” *Knee*, vol. 20, no. SUPPL.1, pp. S16–S20, 2013.
- [143] T. Neogi, D. Felson, J. Niu, J. Lynch, M. Nevitt, A. Guermazi, F. Roemer, C. E. Lewis, B. Wallace, and Y. Zhang, “Cartilage loss occurs in the same subregions as subchondral bone attrition: A within-knee subregion-matched approach from the multicenter osteoarthritis study,” *Arthritis Care and Research*, vol. 61, no. 11, pp. 1539–1544, 2009.

- [144] W. Maletius and M. Lundberg, “Refixation of large chondral fragments on the weight-bearing area of the knee joint: A report of two cases,” *Arthroscopy*, vol. 10, no. 6, pp. 630–633, 1994.
- [145] C. Zietz, P. Bergschmidt, R. Lange, W. Mittelmeier, and R. Bader, “Third-body abrasive wear of tibial polyethylene inserts combined with metallic and ceramic femoral components in a knee simulator study,” *International Journal of Artificial Organs*, vol. 36, no. 1, pp. 47–55, 2013.
- [146] S. M. Hauptmann, P. Weber, C. Glaser, C. Birkenmaier, V. Jansson, and P. E. Müller, “Free bone cement fragments after minimally invasive unicompartmental knee arthroplasty: An underappreciated problem,” *Knee Surgery, Sports Traumatology, Arthroscopy*, vol. 16, no. 8, pp. 770–775, 2008.
- [147] J. Tipper, A. Galvin, S. Williams, H. McEwen, M. Stone, E. Ingham, and J. Fisher, “Isolation and characterization of UHMWPE wear particles down to ten nanometers in size from in vitro hip and knee joint simulators,” *Journal of Biomedical Materials Research Part A*, vol. 78A, no. 3, pp. 473–480, sep 2006.
- [148] R. D. Boutin, J. A. Januario, A. H. Newberg, C. R. Gundry, and J. S. Newman, “MR Imaging Features of Osteochondritis Dissecans of the Femoral Sulcus,” *American Journal of Roentgenology*, vol. 180, no. 3, pp. 641–645, mar 2003.
- [149] E. Teeple, K. A. Elsaid, B. C. Fleming, G. D. Jay, K. Aslani, J. J. Crisco, and A. P. Mechrefe, “Coefficients of friction, lubricin, and cartilage damage in the anterior cruciate ligament-deficient guinea pig knee,” *Journal of Orthopaedic Research*, vol. 26, no. 2, pp. 231–237, 2008.
- [150] D. Shriram, G. Praveen Kumar, F. Cui, Y. H. D. Lee, and K. Subburaj, “Evaluating the effects of material properties of artificial meniscal implant in the human knee joint using finite element analysis,” *Scientific Reports*, vol. 7, no. 1, pp. 1–11, 2017.
- [151] D. Warnecke, M. Meßmer, L. de Roy, S. Stein, C. Gentilini, R. Walker, N. Skaer, A. Ignatius, and L. Dürselen, “Articular cartilage and meniscus reveal higher friction in swing phase than in stance phase under dynamic gait conditions,” *Scientific Reports*, vol. 9, no. 1, pp. 1–9, 2019.

- [152] E. Peña, B. Calvo, M. A. Martínez, and M. Doblaré, “A three-dimensional finite element analysis of the combined behavior of ligaments and menisci in the healthy human knee joint,” *Journal of Biomechanics*, vol. 39, no. 9, pp. 1686–1701, 2006.
- [153] P. Raut, “Impact of mesh quality parameters on elements such as beam, shell and 3D solid in structural analysis,” *International Journal of Engineering Research and Applications*, vol. 2, no. 6, pp. 99–103, 2012.
- [154] H. R. Bao, D. Zhu, H. Gong, and G. S. Gu, “The effect of complete radial lateral meniscus posterior root tear on the knee contact mechanics: A finite element analysis,” *Journal of Orthopaedic Science*, vol. 18, no. 2, pp. 256–263, 2013.
- [155] R. Kandadai, “Finite Element Modeling of Human Knee Joint -,” no. 3, pp. 38–42, 2015.
- [156] L. Li, X. Yang, L. Yang, K. Zhang, J. Shi, L. Zhu, H. Liang, X. Wang, and Q. Jiang, “Biomechanical analysis of the effect of medial meniscus degenerative and traumatic lesions on the knee joint,” *American Journal of Translational Research*, vol. 11, no. 2, pp. 542–556, 2019.
- [157] L. Y. Zhu, L. Li, Z. A. Li, J. P. Shi, W. L. Tang, J. Q. Yang, and Q. Jiang, “Design and biomechanical characteristics of porous meniscal implant structures using triply periodic minimal surfaces,” *Journal of Translational Medicine*, vol. 17, no. 1, pp. 1–10, 2019.
- [158] A. A. M. Kiapour, V. Kaul, A. A. M. Kiapour, C. E. Quatman, S. C. Wordeman, T. E. Hewett, C. K. Demetropoulos, V. K. Goel, and E. T. Al, “The Effect of Ligament Modeling Technique on Knee Joint Kinematics : A Finite Element Study,” *Applied mathematics*, vol. 4, no. 5A, pp. 91–97, 2013.
- [159] P. Lamata, I. Roy, B. Blazevic, A. Crozier, S. Land, S. A. Niederer, D. Rod Hose, and N. P. Smith, “Quality metrics for high order meshes: Analysis of the mechanical simulation of the heart beat,” *IEEE Transactions on Medical Imaging*, vol. 32, no. 1, pp. 130–138, 2013.
- [160] N. M. Grosland, R. Bafna, and V. a. Magnotta, “Automated hexahedral meshing of anatomic structures using deformable registration.” *Computer methods in biomechanics and biomedical engineering*, vol. 12, no. 1, pp. 35–43, 2009.
- [161] P. M. Knupp, “Remarks on mesh quality,” in *46th AIAA Aerospace Sciences Meeting and Exhibit*, 2008.

-
- [162] M. Batdorf, L. A. Freitag, and C. Ollivier-Gooch, “Computational study of the effect of unstructured mesh quality on solution efficiency,” *13th Computational Fluid Dynamics Conference*, 1997.
- [163] P. M. Knupp, “Algebraic Mesh Quality Metrics,” *SIAM Journal on Scientific Computing*, vol. 23, no. 1, pp. 193–218, jan 2001.
- [164] N. M. Grosland, C. Lisle, K. H. Shivanna, S. Pieper, and V. A. Magnotta, “A Check of Mesh Quality.”
- [165] M. Ray, M. Mongiardini, A. Atahan, C. Plaxico, and M. Anghileri, “Procedures for Verification and Validation of Computer Simulations Used for Roadside Safety Applications,” *Procedures for Verification and Validation of Computer Simulations Used for Roadside Safety Applications*, no. March, 2011.
- [166] X. Gao, J. Huang, K. Xu, Z. Pan, Z. Deng, and G. Chen, “Evaluating Hex-mesh Quality Metrics via Correlation Analysis,” *Computer Graphics Forum*, vol. 36, no. 5, pp. 105–116, 2017.
- [167] G. Xu, R. Ling, J. Zhang, Z. Xiao, Z. Ji, and T. Rabczuk, “Singularity Structure Simplification of Hexahedral Mesh via Weighted Ranking,” 2019.
- [168] T. J. Baker, “Mesh movement and metamorphosis,” *Engineering with Computers*, vol. 18, no. 3, pp. 188–198, 2002.
- [169] G. De Santis, M. De Beule, K. Van Canneyt, P. Segers, P. Verdonck, and B. Verhegghe, “Full-hexahedral structured meshing for image-based computational vascular modeling,” *Medical Engineering and Physics*, vol. 33, no. 10, pp. 1318–1325, 2011.
- [170] M. Batdorf, L. A. Freitag, and C. Ollivier-Gooch, “Computational study of the effect of unstructured mesh quality on solution efficiency,” *13th Computational Fluid Dynamics Conference*, 1997.
- [171] P. M. Knupp, “On the invertibility of the isoparametric map,” *Computer Methods in Applied Mechanics and Engineering*, vol. 78, no. 3, pp. 313–329, 1990.
- [172] M. Mengoni, J.-P. Ponthot, and R. Boman, “Mesh management methods in finite element simulations of orthodontic tooth movement,” *Medical Engineering Physics*, vol. 38, no. 2, pp. 140–147, 2016.

- [173] A. E. Frey, C. A. Hall, and T. A. Porsching, “Some Results on the Global Inversion of Bilinear and Quadratic Isoparametric Finite Element Transformations,” *Mathematics of Computation*, vol. 32, no. 143, p. 725, 1978.
- [174] D. Holmes and D. Holmes, “Generalized Method of Decomposing Solid Geometry into Hexahedron Finite Elements,” *Proc. 4th International Meshing Roundtable*, pp. 141–152, 1995.
- [175] G. R. Liu, K. Y. Dai, and T. T. Nguyen, “A smoothed finite element method for mechanics problems,” *Computational Mechanics*, vol. 39, no. 6, pp. 859–877, 2007.
- [176] M. A. Price, C. G. Armstrong, and M. A. Sabin, “Hexahedral mesh generation by medial surface subdivision: Part I. Solids with convex edges,” *International Journal for Numerical Methods in Engineering*, vol. 38, no. 19, pp. 3335–3359, oct 1995.
- [177] J. Sarrate, E. Ruiz-Gironés, and X. Roca, “Unstructured and Semi-Structured Hexahedral Mesh Generation Methods,” *Computational Technology Reviews*, vol. 10, pp. 35–64, 2014.
- [178] R. Schneiders, “Algorithms for Quadrilateral and Hexahedral Mesh Generation,” *Proceedings of the VKI Lecture Series on . . .*, vol. i, no. November, p. 56, 2000.
- [179] S. J. Owen, “A survey of unstructured mesh generation technology,” *7th International Meshing Roundtable*, no. October, p. 25, 1998.
- [180] W. R. Quadros, “LayTracks3D: A new approach for meshing general solids using medial axis transform,” *CAD Computer Aided Design*, vol. 72, pp. 102–117, 2016.
- [181] M. Callahan, M. J. Cole, J. F. Shepherd, J. G. Stinstra, and C. R. Johnson, “A meshing pipeline for biomedical computing,” *Engineering with Computers*, vol. 25, no. 1, pp. 115–130, 2009.
- [182] M. L. Staten, S. J. Owen, and T. D. Blacker, “Unconstrained paving plastering: A new idea for all hexahedral mesh generation,” *Proceedings of the 14th International Meshing Roundtable, IMR 2005*, pp. 399–416, 2005.
- [183] Y. Yu, J. G. Liu, and Y. J. Zhang, “HexDom: Polycube-Based Hexahedral-Dominant Mesh Generation,” pp. 1–19, 2021.

-
- [184] M. Sabin, “Criteria for comparison of automatic mesh generation methods,” *Advances in Engineering Software and Workstations*, vol. 13, no. 5-6, pp. 220–225, sep 1991.
- [185] T. J. Tautges, “The generation of hexahedral meshes for assembly geometry: Survey and progress,” *International Journal for Numerical Methods in Engineering*, vol. 50, no. 12, pp. 2617–2642, 2001.
- [186] R. Schneiders, “A grid-based algorithm for the generation of hexahedral element meshes,” *Engineering with Computers*, vol. 12, no. 3-4, pp. 168–177, sep 1996.
- [187] S. Canann, “Plastering - A new approach to automated, 3-D hexahedral mesh generation,” no. 54, 1992.
- [188] Y. Li, Y. Liu, W. Xu, W. Wang, and B. Guo, “All-hex meshing using singularity-restricted field,” *ACM Transactions on Graphics*, vol. 31, no. 6, 2012.
- [189] P. Lamata, S. Niederer, D. Nordsletten, D. C. Barber, I. Roy, D. R. Hose, and N. Smith, “An accurate, fast and robust method to generate patient-specific cubic Hermite meshes,” *Medical Image Analysis*, vol. 15, no. 6, pp. 801–813, 2011.
- [190] S. G. Zachariah, J. E. Sanders, and G. M. Turkiyyah, “Automated hexahedral mesh generation from biomedical image data: Applications in limb prosthetics,” *IEEE Transactions on Rehabilitation Engineering*, vol. 4, no. 2, pp. 91–102, 1996.
- [191] W. M. Chen, Y. H. Cai, Y. Yu, X. Geng, and X. Ma, “Optimal mesh criteria in finite element modeling of human foot: The dependence for multiple model outputs on mesh density and loading boundary conditions,” *Journal of Mechanics in Medicine and Biology*, vol. 21, no. 9, pp. 1–14, 2021.
- [192] H. Liu, P. Zhang, E. Chien, J. Solomon, and D. Bommers, “Singularity-constrained octahedral fields for hexahedral meshing,” *ACM Transactions on Graphics*, vol. 37, no. 4, 2018.
- [193] K. H. Shivanna, N. M. Grosland, M. E. Russell, and D. R. Pedersen, “Diarthrodial joint contact models: Finite element model development of the human hip,” *Engineering with Computers*, vol. 24, no. 2, pp. 155–163, 2008.
- [194] T. J. Tautges, “MOAB-SD: Integrated structured and unstructured mesh representation,” *Engineering with Computers*, vol. 20, no. 3, pp. 286–293, 2004.

- [195] J. Daniels, C. T. Silva, J. Shepherd, and E. Cohen, “Quadrilateral mesh simplification,” *ACM Transactions on Graphics*, vol. 27, no. 5, 2008.
- [196] J. E. Makem, C. G. Armstrong, and T. T. Robinson, “Automatic decomposition and efficient semi-structured meshing of complex solids,” *Engineering with Computers*, vol. 30, no. 3, pp. 345–361, 2014.
- [197] H. J. Fogg, C. G. Armstrong, and T. T. Robinson, “Automatic generation of multiblock decompositions of surfaces,” *International Journal for Numerical Methods in Engineering*, vol. 101, no. 13, pp. 965–991, 2015.
- [198] W. Yu, K. Zhang, S. Wan, and X. Li, “Optimizing polycube domain construction for hexahedral remeshing,” *CAD Computer Aided Design*, vol. 46, no. 1, pp. 58–68, 2014.
- [199] M. Livesu, A. Muntoni, E. Puppo, and R. Scateni, “Skeleton-driven Adaptive Hexahedral Meshing of Tubular Shapes,” *Computer Graphics Forum*, vol. 35, no. 7, pp. 237–246, 2016.
- [200] C. Liu, W. Yu, Z. Chen, and X. Li, “Distributed poly-square mapping for large-scale semi-structured quad mesh generation,” *CAD Computer Aided Design*, vol. 90, no. May, pp. 5–17, 2017.
- [201] E. Makris, P. Neofytou, S. Tsangaris, and C. Housiadas, “A novel method for the generation of multi-block computational structured grids from medical imaging of arterial bifurcations,” *Medical Engineering and Physics*, vol. 34, no. 8, pp. 1157–1166, 2012.
- [202] G. Vemaganti and A. Wieting, “Application of a finite element algorithm for high speed viscous flows using structured and unstructured meshes,” in *21st Fluid Dynamics, Plasma Dynamics and Lasers Conference*. Reston, Virginia: American Institute of Aeronautics and Astronautics, jun 1990.
- [203] P. Sampl, “Semi-Structured Mesh Generation Based on Medial Axis,” in *9th International Meshing Roundtable*, 2000, pp. 21–32.
- [204] M. Taylor, R. Bryan, and F. Galloway, “Accounting for patient variability in finite element analysis of the intact and implanted hip and knee: A review,” *International Journal for Numerical Methods in Biomedical Engineering*, vol. 29, no. 2, pp. 273–292, feb 2013.
- [205] T. J. Tautges, T. Blacker, and S. A. Mitchell, “The whisker weaving algorithm: A connectivity-based method for constructing all-hexahedral finite element meshes,” *Inter-*

- national Journal for Numerical Methods in Engineering*, vol. 39, no. 19, pp. 3327–3349, 1996.
- [206] X. Roca, E. Ruiz-Gironés, and J. Sarrate, “Receding front method: A new approach applied to generate hexahedral meshes of outer domains,” *Proceedings of the 19th International Meshing Roundtable, IMR 2010*, pp. 209–225, 2010.
- [207] P. M. Knupp, “Hexahedral and tetrahedral mesh untangling,” *Engineering with Computers*, vol. 17, no. 3, pp. 261–268, 2001.
- [208] S. J. Owen and S. Saigal, “H-Morph: An indirect approach to advancing front hex meshing,” *International Journal for Numerical Methods in Engineering*, vol. 49, no. 1-2, pp. 289–312, 2000.
- [209] D. White, L. Mingwu, and S. Benzley, “Automated Hexahedral Mesh Generation by Virtual Decomposition,” *Proceedings 5th International Meshing Roundtable*, vol. D, no. May, 1996.
- [210] B. Ren and C. Wang, “Mapping-based 3D hexahedral finite element mesh generation method,” *Wuhan University Journal of Natural Sciences*, vol. 12, no. 2, pp. 255–259, 2007.
- [211] H. Mao, H. Gao, L. Cao, V. V. Genthikatti, and K. H. Yang, “Development of high-quality hexahedral human brain meshes using feature-based multi-block approach,” *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 16, no. 3, pp. 271–279, 2013.
- [212] G. Xiong, S. Musuvathy, and T. Fang, “Automated structured all-quadrilateral and hexahedral meshing of tubular surfaces,” *Proceedings of the 21st International Meshing Roundtable, IMR 2012*, pp. 103–120, 2013.
- [213] J. Gregson, A. Sheffer, and E. Zhang, “All-hex mesh generation via volumetric polycube deformation,” *Eurographics Symposium on Geometry Processing*, vol. 30, no. 5, pp. 1407–1416, 2011.
- [214] L. Liu, Y. Zhang, Y. Liu, and W. Wang, “Feature-preserving T-mesh construction using skeleton-based polycubes,” *CAD Computer Aided Design*, vol. 58, pp. 162–172, 2015.
- [215] G. F. Carey, “Hexing the Tet,” *Communications in Numerical Methods in Engineering*, vol. 18, no. 3, pp. 223–227, 2002.

- [216] R. W. Leland, D. J. Melander, R. W. Meyers, S. A. Mitchell, and T. J. Tautges, “The geode algorithm: combining hex/tet plastering, dicing and transition elements for automatic, all-hex mesh generation,” *Proceedings of the 7th International Meshing Roundtable*, pp. 515–521, 1998.
- [217] W. A. Cook and W. R. Oakes, “Mapping method for generating three-dimensional meshes: past and present,” *Proceedings of the ASME Second International Computer Engineering Conference*, 1982.
- [218] F. Ledoux and J.-C. Weill, “An Extension of the Reliable Whisker Weaving Algorithm,” in *Proceedings of the 16th International Meshing Roundtable*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, no. June, pp. 215–232.
- [219] Y. Lu, R. Gadh, and T. J. Tautges, “Volume decomposition and feature recognition for hexahedral mesh generation,” in *8th International Meshing Roundtable, SAND99-2288, Sandia National Laboratories*, 1999, pp. 269–280.
- [220] —, “Feature based hex meshing methodology: Feature recognition and volume decomposition,” *CAD Computer Aided Design*, vol. 33, no. 3, pp. 221–232, 2001.
- [221] J. F. Shepherd and C. R. Johnson, “Hexahedral mesh generation for biomedical models in SCIRun,” *Engineering with Computers*, vol. 25, no. 1, pp. 97–114, 2009.
- [222] M. Whiteley, D. White, S. Benzley, and T. Blacker, “Two and three-quarter dimensional meshing facilitators,” *Engineering with Computers*, vol. 12, no. 3-4, pp. 144–154, 1996.
- [223] T. D. Blacker and R. J. Meyers, “Seams and wedges in plastering: A 3-D hexahedral mesh generation algorithm,” *Engineering with Computers*, vol. 9, no. 2, pp. 83–93, 1993.
- [224] Y. Ito, A. M. Shih, and B. K. Soni, “Octree-based reasonable-quality hexahedral mesh generation using a new set of refinement templates,” *International Journal for Numerical Methods in Engineering*, vol. 77, no. 13, pp. 1809–1833, mar 2009.
- [225] X. Gao, H. Shen, and D. Panozzo, “Feature Preserving Octree-Based Hexahedral Meshing,” *Computer Graphics Forum*, vol. 38, no. 5, pp. 135–149, 2019.
- [226] M. L. Staten, S. A. Canann, and S. J. Owen, “BMSweep: Locating interior nodes during sweeping,” *Engineering with Computers*, vol. 15, no. 3, pp. 212–218, 1999.

-
- [227] M. Nieser, U. Reitebuch, and K. Polthier, “CubeCover- Parameterization of 3D Volumes,” *Eurographics Symposium on Geometry Processing*, vol. 30, no. 5, pp. 1397–1406, aug 2011.
- [228] J. H. Keyak, J. M. Meagher, H. B. Skinner, and C. D. Mote, “Automated three-dimensional finite element modelling of bone: a new method,” *Journal of Biomedical Engineering*, vol. 12, no. 5, pp. 389–397, 1990.
- [229] M. Viceconti, L. Bellingeri, L. Cristofolini, and A. Toni, “A comparative study on different methods of automatic mesh generation of human femurs,” *Medical Engineering and Physics*, vol. 20, no. 1, pp. 1–10, 1998.
- [230] E. Makris, V. Gkanis, S. Tsangaris, and C. Housiadas, “A methodology to generate structured computational grids from DICOM data: Application to a patient-specific abdominal aortic aneurysm (AAA) model,” *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 15, no. 2, pp. 173–183, 2012.
- [231] B. K. Lahkar, P.-Y. Rohan, H. Pillet, P. Thoreux, and W. Skalli, “Fast subject specific finite element mesh generation of knee joint from biplanar x-ray images,” *Cmbbe*, 2018.
- [232] J. Q. Campbell and A. J. Petrella, “An automated method for landmark identification and finite-element modeling of the lumbar spine,” *IEEE Transactions on Biomedical Engineering*, vol. 62, no. 11, pp. 2709–2716, 2015.
- [233] X. Jia, S. Liao, X. Duan, W. Zheng, and B. Zou, “Anisotropic Finite Element Modeling Based on a Harmonic Field for Patient-Specific Sclera,” *BioMed Research International*, vol. 2017, 2017.
- [234] B. Maier, “Scalable Biophysical Simulations of the Neuromuscular System (Thesis),” Doctoral Thesis, Stuttgart University, 2021.
- [235] I. A. Sigal, H. Yang, M. D. Roberts, and J. C. Downs, “Morphing methods to parameterize specimen-specific finite element model geometries,” *Journal of Biomechanics*, vol. 43, no. 2, pp. 254–262, 2010.
- [236] R. Schneiders and R. Bünthen, “Automatic generation of hexahedral finite element meshes,” *Computer Aided Geometric Design*, vol. 12, no. 7, pp. 693–707, 1995.

- [237] N. M. Grosland, R. Bafna, and V. a. Magnotta, “Automated hexahedral meshing of anatomic structures using deformable registration.” *Computer methods in biomechanics and biomedical engineering*, vol. 12, no. 1, pp. 35–43, 2009.
- [238] N. A. Kallemeyn, S. C. Tadepalli, K. H. Shivanna, and N. M. Grosland, “An interactive multiblock approach to meshing the spine,” *Computer Methods and Programs in Biomedicine*, vol. 95, no. 3, pp. 227–235, 2009.
- [239] K. H. Shivanna, S. C. Tadepalli, and N. M. Grosland, “Feature-based multiblock finite element mesh generation,” *Computer-Aided Design*, vol. 42, no. 12, pp. 1108–1116, dec 2010.
- [240] S. C. Tadepalli, K. H. Shivanna, V. A. Magnotta, and N. M. Grosland, “Semi-Automated Patient Specific Hexahedral Mesh Generation of Articular Cartilage,” in *ASME 2009 Summer Bioengineering Conference, Parts A and B*. American Society of Mechanical Engineers, jun 2009, pp. 1007–1008.
- [241] K. M. Yates and C. D. Untaroiu, “Finite element modeling of the human kidney for probabilistic occupant models: Statistical shape analysis and mesh morphing,” *Journal of Biomechanics*, vol. 74, pp. 50–56, 2018.
- [242] E. K. Main, “Constitutive mechanical properties of carpal tunnel soft tissue (Thesis),” Master’s thesis, The University of Iowa, 2011.
- [243] M. F. Li, S. H. Liao, and R. F. Tong, “Facial hexahedral mesh transferring by volumetric mapping based on harmonic fields,” *Computers and Graphics (Pergamon)*, vol. 35, no. 1, pp. 92–98, 2011.
- [244] M. Tada, H. Yoshida, and M. Mochimaru, “Geometric Modeling of Living Tissue for Subject-Specific Finite Element Analysis,” vol. d, no. 1, pp. 6639–6642, 2006.
- [245] J. D. John, M. W. Arun, N. Yoganandan, G. Saravanakumar, and S. N. Kurpad, “Mapping block-based morphing for subject-specific spine finite element models,” *54th Annual Rocky Mountain Bioengineering Symposium, RMBS 2017 and 54th International ISA Biomedical Sciences Instrumentation Symposium 2017*, vol. 2017-March, pp. 6–11, 2017.
- [246] Y. Kawamura, M. S. Islam, and Y. Sumi, “A strategy of automatic hexahedral mesh generation by using an improved whisker-weaving method with a surface mesh modification

- procedure,” *Engineering with Computers*, vol. 24, no. 3, pp. 215–229, 2008.
- [247] D. Bommès, L. Bruno, N. Pietroni, B. Lvy, N. Pietroni, E. Puppo, C. S. A, M. Tarini, and D. Zorin, “State of the Art in Quad Meshing,” *Eurographics STARS*, vol. xx, pp. 1–24, 2012.
- [248] Y. Yu, X. Wei, A. Li, J. G. Liu, J. He, and Y. J. Zhang, “HexGen and Hex2Spline: Polycube-based Hexahedral Mesh Generation and Spline Modeling for Isogeometric Analysis Applications in LS-DYNA,” pp. 1–31, 2020.
- [249] Y. He, H. Wang, C. W. Fu, and H. Qin, “A divide-and-conquer approach for automatic polycube map construction,” *Computers and Graphics (Pergamon)*, vol. 33, no. 3, pp. 369–380, 2009.
- [250] J. Huang, T. Jiang, Y. Wang, Y. Tong, and H. Bao, “Automatic Frame Field Guided Hexahedral Mesh Generation,” pp. 1–7, 2012.
- [251] S. Han, J. Xia, and Y. He, “Hexahedral shell mesh construction via volumetric polycube map,” *Proceedings - 14th ACM Symposium on Solid and Physical Modeling, SPM'10*, pp. 127–136, 2010.
- [252] N. Ray, D. Sokolov, M. Reberol, F. Ledoux, B. Lévy, N. Ray, D. Sokolov, M. Reberol, F. Ledoux, B. Lévy, and H. Meshing, “Hexahedral Meshing : Mind the Gap !” 2017.
- [253] P. E. Bernard, J. F. Remacle, N. Kowalski, and C. Geuzaine, “Hex-dominant meshing approach based on frame field smoothness,” *Procedia Engineering*, vol. 82, no. December, pp. 175–186, 2014.
- [254] Y. Zhang, C. Bajaj, and B.-S. Sohn, “3D finite element meshing from imaging data,” *Computer Methods in Applied Mechanics and Engineering*, vol. 194, no. 48-49, pp. 5083–5106, nov 2005.
- [255] Y. Zhang, T. J. Hughes, and C. L. Bajaj, “An automatic 3D mesh generation method for domains with multiple materials,” *Computer Methods in Applied Mechanics and Engineering*, vol. 199, no. 5-8, pp. 405–415, 2010.
- [256] P. Kraft, “Automatic Remeshing With Hexahedral Elements : Problems , Solutions and Applications,” *Proceedings of 8th international meshing roundtable*, 1999.

- [257] A. Leichner, H. Andrä, and B. Simeon, “A contact algorithm for voxel-based meshes using an implicit boundary representation,” *Computer Methods in Applied Mechanics and Engineering*, vol. 352, pp. 276–299, 2019.
- [258] M. Viceconti, M. Davinelli, F. Taddei, and A. Cappello, “Automatic generation of accurate subject-specific bone finite element models to be used in clinical studies,” vol. 37, pp. 1597–1605, 2004.
- [259] O. Defaux, “Mathematical Geography and Cartography in Antiquity,” *eTopoi. Journal for Ancient Studies*, vol. 9, pp. 54–70, 2021.
- [260] M. S. Floater and K. Hormann, “Surface Parameterization: a Tutorial and Survey,” in *Advances in Multiresolution for Geometric Modelling*, N. A. Dodgson, M. S. Floater, and M. A. Sabin, Eds. Berlin/Heidelberg: Springer-Verlag, 2005, pp. 157–186.
- [261] D. Aneja, S. R. Vora, E. D. Camci, L. G. Shapiro, and T. C. Cox, “Automated Detection of 3D Landmarks for the Elimination of Non-Biological Variation in Geometric Morphometric Analyses,” in *2015 IEEE 28th International Symposium on Computer-Based Medical Systems*, vol. 344, no. 6188. IEEE, jun 2015, pp. 78–83.
- [262] S. Zeng, “An object-tracking algorithm for 3-D range data using motion and surface estimation,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 3, pp. 1109–1118, 2013.
- [263] J. Guevara, J. Gené-Mola, E. Gregorio, M. Torres-Torriti, G. Reina, and F. A. A. Cheein, “Comparison of 3D scan matching techniques for autonomous robot navigation in urban and agricultural environments,” *Journal of Applied Remote Sensing*, vol. 15, no. 02, 2021.
- [264] D. C. Barber and D. R. Hose, “Automatic segmentation of medical images using image registration: Diagnostic and simulation applications,” *Journal of Medical Engineering and Technology*, vol. 29, no. 2, pp. 53–63, 2005.
- [265] T. N. Nguyen, V. D. Tran, H. Q. Nguyen, and T. T. Dao, “A statistical shape modeling approach for predicting subject-specific human skull from head surface,” *Medical and Biological Engineering and Computing*, vol. 58, no. 10, pp. 2355–2373, 2020.
- [266] S. I. Buchaillard, S. H. Ong, Y. Payan, and K. Foong, “3D statistical models for tooth surface reconstruction,” *Computers in Biology and Medicine*, vol. 37, no. 10, pp. 1461–

- 1471, 2007.
- [267] D. Kozinska, F. Carducci, and K. Nowinski, “Automatic alignment of EEG/MEG and MRI data sets,” *Clinical Neurophysiology*, vol. 112, no. 8, pp. 1553–1561, aug 2001.
- [268] R. B. Rusu, N. Blodow, and M. Beetz, “Fast Point Feature Histograms (FPFH) for 3D registration,” pp. 3212–3217, 2009.
- [269] A. G. Buch, D. Kraft, J. K. Kamarainen, H. G. Petersen, and N. Kruger, “Pose estimation using local structure-specific shape and appearance context,” *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 2080–2087, 2013.
- [270] Y. Guo, Y. Gu, and Y. Zhang, “Invariant Feature Point based ICP with the RANSAC for 3D Registration,” *Information Technology Journal*, vol. 10, no. 2, pp. 276–284, jan 2011.
- [271] F. Pomerleau, F. Colas, F. Ferland, and F. Michaud, “Relative Motion Threshold for Rejection in ICP Registration,” 2010, pp. 229–238.
- [272] N. Gelfand, N. J. Mitra, L. J. Guibas, and H. Pottmann, “Robust Global Registration,” in *Proceedings of the Third Eurographics Symposium on Geometry Processing*, ser. SGP ’05. Goslar, DEU: Eurographics Association, 2005, pp. 197–es.
- [273] J. S. Athertya and G. Saravana Kumar, “Automatic segmentation of vertebral contours from CT images using fuzzy corners,” *Computers in Biology and Medicine*, vol. 72, pp. 75–89, 2016.
- [274] H. K. Sangappa and K. R. Ramakrishnan, “A probabilistic analysis of a common RANSAC heuristic,” *Machine Vision and Applications*, vol. 30, no. 1, pp. 71–89, 2019.
- [275] Z. Hossein-Nejad and M. Nasri, “An adaptive image registration method based on SIFT features and RANSAC transform,” *Computers and Electrical Engineering*, vol. 62, pp. 524–537, 2017.
- [276] M. Bennani Dosse, H. A. Kiers, and J. M. Ten Berge, “Anisotropic generalized Procrustes analysis,” *Computational Statistics and Data Analysis*, vol. 55, no. 5, pp. 1961–1968, 2011.
- [277] J. C. Gower, “Procrustes methods,” *Wiley Interdisciplinary Reviews: Computational Statistics*, vol. 2, no. 4, pp. 503–508, 2010.

- [278] J. Yang, Z. Huang, S. Quan, Q. Zhang, Y. Zhang, and Z. Cao, "Toward Efficient and Robust Metrics for RANSAC Hypotheses and 3D Rigid Registration," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 2, pp. 893–906, 2022.
- [279] Y. Li, Z. Yang, and J. Deng, "Spherical Parameterization of Genus-Zero Meshes Using the Lagrange-Newton Method," in *2007 10th IEEE International Conference on Computer-Aided Design and Computer Graphics*. IEEE, oct 2007, pp. 32–32.
- [280] M. Ben-Chen, C. Gotsman, and G. Bunin, "Conformal flattening by curvature prescription and metric scaling," *Computer Graphics Forum*, vol. 27, no. 2, pp. 449–458, 2008.
- [281] A. Sheffer, E. Praun, and K. Rose, "Mesh parameterization methods and their applications," *Foundations and Trends in Computer Graphics and Vision*, vol. 2, no. 2, pp. 105–171, 2006.
- [282] E. Praun and H. Hoppe, "Spherical parametrization and remeshing," *ACM Transactions on Graphics*, vol. 22, no. 3, p. 340, 2003.
- [283] J.-F. Remacle, C. Geuzaine, G. Compère, and E. Marchandise, "High Quality Surface Remeshing Using Harmonic Maps," *International Journal for Numerical Methods in Engineering*, vol. 00, pp. 1–6, 2009.
- [284] R. Karim, Y. L. Ma, M. Jang, R. J. Housden, S. E. Williams, Z. Chen, A. Ataollahi, K. Althoefer, C. A. Rinaldi, R. Razavi, M. D. O'Neill, T. Schaeftter, and K. S. Rhode, "Surface flattening of the human left atrium and proof-of-concept clinical applications," *Computerized Medical Imaging and Graphics*, vol. 38, no. 4, pp. 251–266, 2014.
- [285] K. Fujiwara, "Eigenvalues of Laplacians on a Closed Riemannian Manifold and Its Nets," *Proceedings of the American Mathematical Society*, vol. 123, no. 8, p. 2585, 1995.
- [286] Y. Xiong, G. Li, and G. Han, "Mean Laplace–Beltrami Operator for Quadrilateral Meshes," Berlin, Heidelberg, 2011, vol. 6530, no. April, pp. 189–201.
- [287] C. Geuzaine and J. F. Remacle, "Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities," *International Journal for Numerical Methods in Engineering*, 2009.
- [288] E. Marchandise, G. Compère, M. Willemet, G. Bricteux, C. Geuzaine, and J. F. Remacle, "Quality meshing based on STL triangulations for biomedical simulations," *International*

- Journal for Numerical Methods in Biomedical Engineering*, vol. 26, no. 7, pp. 876–889, 2010.
- [289] T. Kanai, H. Suzuki, and F. Kimura, “Metamorphosis of arbitrary triangular meshes,” *IEEE Computer Graphics and Applications*, vol. 20, no. 2, pp. 62–75, 2000.
- [290] E. Marchandise, C. C. de Wiart, W. G. Vos, C. Geuzaine, and J.-F. Remacle, “High-quality surface remeshing using harmonic maps-Part II: Surfaces with high genus and of large aspect ratio,” *International Journal for Numerical Methods in Engineering*, vol. 86, no. 11, pp. 1303–1321, jun 2011.
- [291] X. H. Wu and Z. X. Hao, “Evaluation of spherical parameterization methods for three-dimensional reconstruction of medical images,” *Applied Mechanics and Materials*, vol. 365-366, pp. 1342–1349, 2013.
- [292] C. Peng and S. Timalena, “Fast mapping and morphing for genus-zero meshes with cross spherical parameterization,” *Computers and Graphics (Pergamon)*, vol. 59, pp. 107–118, 2016.
- [293] S. Nadeem, Z. Su, W. Zeng, A. Kaufman, and X. Gu, “Spherical parameterization balancing angle and area distortions,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 6, pp. 1663–1676, 2017.
- [294] T. Athanasiadis, I. Fudos, C. Nikou, and V. Stamati, “Feature-based 3D morphing based on geometrically constrained spherical parameterization,” *Computer Aided Geometric Design*, vol. 29, no. 1, pp. 2–17, 2012.
- [295] I. Friedel, P. Schröder, and M. Desbrun, “Unconstrained spherical parameterization,” in *ACM SIGGRAPH 2005 Sketches on - SIGGRAPH '05*. New York, New York, USA: ACM Press, 2005, p. 134.
- [296] G. Biros and O. Ghattas, “Parallel Lagrange–Newton–Krylov–Schur Methods for PDE-Constrained Optimization. Part II: The Lagrange–Newton Solver and Its Application to Optimal Control of Steady Viscous Flows,” *SIAM Journal on Scientific Computing*, vol. 27, no. 2, pp. 714–739, 2005.
- [297] S. Abhyankar, J. Brown, E. M. Constantinescu, D. Ghosh, B. F. Smith, and H. Zhang, “PETSc/TS: A Modern Scalable ODE/DAE Solver Library,” vol. V, no. 212, pp. 1–29,

- 2018.
- [298] S. Saba, I. Yavneh, C. Gotsman, and A. Sheffer, “Practical spherical embedding of manifold triangle meshes,” *Proceedings - International Conference on Shape Modeling and Applications, SMI’05*, vol. 2005, no. Siggraph 2003, pp. 258–267, 2005.
- [299] R. H. Davies, C. J. Twining, and C. J. Taylor, “Consistent spherical parameterisation for statistical shape modelling,” *2006 3rd IEEE International Symposium on Biomedical Imaging: From Nano to Macro - Proceedings*, vol. 2006, pp. 1388–1391, 2006.
- [300] S. M. Shontz, “Numerical Methods for Problems with Moving Meshes (Thesis),” Doctoral Thesis, Cornell University, 2005.
- [301] P. Knupp, “Updating meshes on deforming domains: An application of the target-matrix paradigm,” *Communications in Numerical Methods in Engineering*, vol. 24, no. 6, pp. 467–476, may 2007.
- [302] M. A. Heroux and J. M. Willenbring, “A new overview of the Trilinos project,” *Scientific Programming*, vol. 20, no. 2, pp. 83–88, 2012.
- [303] T. Schneider, Y. Hu, X. Gao, J. Dumas, D. Zorin, and D. Panozzo, “A large-scale comparison of tetrahedral and hexahedral elements for finite element analysis,” *arXiv*, pp. 1–27, 2019.
- [304] M. Ammar, F. Muhammad, N. Jamil, N. Nurul, H. Muhmed, and A. R. Yusoff, “Finite element analysis of hybrid energy harvesting of piezoelectric and electromagnetic,” vol. 01039, pp. 1–10, 2017.
- [305] S. Sivaiah, G. Murali, and M. C. K. Reddy, “Finite Element Analysis of Chemical Reaction and Radiation Effects on Isothermal Vertical Oscillating Plate with Variable Mass Diffusion,” *ISRIN Mathematical Physics*, vol. 2012, pp. 1–14, 2012.
- [306] M. Mengoni, A. C. Jones, and R. K. Wilcox, “Modelling the failure precursor mechanism of lamellar fibrous tissues, example of the annulus fibrosus,” *Journal of the Mechanical Behavior of Biomedical Materials*, vol. 63, pp. 265–272, 2016.
- [307] M. Bucki, C. Lobos, and Y. Payan, “A fast and robust patient specific Finite Element mesh registration technique: Application to 60 clinical cases,” *Medical Image Analysis*, vol. 14, no. 3, pp. 303–317, 2010.

-
- [308] J. Park, S. M. Shontz, and C. S. Drapaca, “A combined level set/mesh warping algorithm for tracking brain and cerebrospinal fluid evolution in hydrocephalic patients,” *Lecture Notes in Computational Vision and Biomechanics*, vol. 3, pp. 107–141, 2013.
- [309] M. Ferrant, S. K. Warfield, C. R. Guttman, R. V. Mulkern, F. A. Jolesz, and R. Kikinis, “3D image matching using a finite element based elastic deformation model,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 1679, pp. 202–210, 1999.
- [310] A. E. Bowden, R. D. Rabbitt, and J. A. Weiss, “Anatomical registration and segmentation by warping template finite element models,” *Laser-Tissue Interaction IX*, vol. 3254, p. 469, 1998.
- [311] S. M. Shontz and S. A. Vavasis, “Analysis of and workarounds for element reversal for a finite element-based algorithm for warping triangular and tetrahedral meshes,” *BIT Numerical Mathematics*, vol. 50, no. 4, pp. 863–884, 2010.
- [312] B. T. Helenbrook, “Mesh deformation using the biharmonic operator,” *International Journal for Numerical Methods in Engineering*, vol. 56, no. 7, pp. 1007–1021, 2003.
- [313] A. Masud and T. J. Hughes, “A space-time Galerkin/least-squares finite element formulation of the Navier-Stokes equations for moving domain problems,” *Computer Methods in Applied Mechanics and Engineering*, vol. 146, no. 1-2, pp. 91–126, 1997.
- [314] K. Stein, T. E. Tezduyar, and R. Benney, “Automatic mesh update with the solid-extension mesh moving technique,” *Computer Methods in Applied Mechanics and Engineering*, vol. 193, no. 21-22, pp. 2019–2032, 2004.
- [315] S. M. Shontz and S. A. Vavasis, “A robust solution procedure for hyperelastic solids with large boundary deformation,” *Engineering with Computers*, vol. 28, no. 2, pp. 135–147, 2012.
- [316] S. Maas, D. Rawlins, J. Weiss, and G. Ateshian, “FEBio Theory Manual, Version 2.9,” 2019.
- [317] V. Bonnaillie-Noël, M. Dambrine, F. Héreau, and G. Vial, “Artificial conditions for the linear elasticity equations,” *Mathematics of Computation*, vol. 84, no. 294, pp. 1599–1632, nov 2014.

- [318] M. González, “Stabilized dual-mixed method for the problem of linear elasticity with mixed boundary conditions,” *Applied Mathematics Letters*, vol. 30, pp. 1–5, 2014.
- [319] A. Hagemann, K. Rohr, H. S. Stiehl, U. Spetzger, and J. M. Gilsbach, “Biomechanical modeling of the human head for physically based, nonrigid image registration,” *IEEE Transactions on Medical Imaging*, vol. 18, no. 10, pp. 875–884, 1999.
- [320] N. Cusimano, F. D. Teso, L. Gerardo-Giorda, and G. Pagnini, “Discretizations of the spectral fractional laplacian on general domains with dirichlet, neumann, and robin boundary conditions,” *SIAM Journal on Numerical Analysis*, vol. 56, no. 3, pp. 1243–1272, 2018.
- [321] G. Galeati and G. Gambolati, “On boundary conditions and point sources in the finite element integration of the transport equation,” *Water Resources Research*, vol. 25, no. 5, pp. 847–856, 1989.
- [322] Q. Meng, J. Fisher, and R. Wilcox, “The effects of geometric uncertainties on computational modelling of knee biomechanics,” *Royal Society Open Science*, vol. 4, no. 8, 2017.
- [323] M. Mengoni, K. Vasiljeva, A. C. Jones, S. M. Tarsuslugil, and R. K. Wilcox, “Subject-specific multi-validation of a finite element model of ovine cervical functional spinal units,” *Journal of Biomechanics*, vol. 49, no. 2, pp. 259–266, 2016.
- [324] A. C. Jones and R. K. Wilcox, “Assessment of factors influencing finite element vertebral model predictions,” *Journal of Biomechanical Engineering*, vol. 129, no. 6, pp. 898–903, 2007.
- [325] K. Ljungkvist, “Techniques for Finite Element Methods on Modern Processors (Thesis),” Doctoral Thesis, Uppsala University, 2015.
- [326] L. Ma and T. Li, “An rbf-based h-adaptive cartesian grid refinement method for arbitrary single/multi-body hull modeling and reconstruction,” *Symmetry*, vol. 13, no. 1, pp. 1–34, 2021.
- [327] T. R. Carr, J. C., Beatson, R. K., Cherrie, J. B., Mitchell, T. J., Fright, W. R., McCallum, B. C., Evans, “Reconstruction and Representation of 3D Objects with Radial Basis Functions,” *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, vol. 28, pp. 67 – 76, 2001.

- [328] Wan-Chun Ma, Fu-Che Wu, and Ming Ouhyoung, “Skeleton extraction of 3D objects with radial basis functions,” in *2003 Shape Modeling International*. IEEE Comput. Soc, pp. 207–215.
- [329] A. Knoll, I. Wald, P. Navratil, A. Bowen, K. Reda, M. E. Papka, and K. Gaither, “RBF volume ray casting on multicore and manycore CPUs,” *Computer Graphics Forum*, vol. 33, no. 3, pp. 71–80, 2014.
- [330] D. Zhong, J. Zhang, and L. Wang, “Fast implicit surface reconstruction for the radial basis functions interpolant,” *Applied Sciences (Switzerland)*, vol. 9, no. 24, 2019.
- [331] W. J. Schroeder and K. M. Martin, “The Visualization Toolkit,” in *Visualization Handbook*, 2005.
- [332] M. Tezzele, N. Demo, A. Mola, and G. Rozza, “PyGeM: Python Geometrical Morphing,” *Software Impacts*, vol. 7, no. November 2020, p. 100047, 2021.
- [333] M. M. Zhang and S. A. Williamson, “Embarrassingly parallel inference for Gaussian processes,” *Journal of Machine Learning Research*, vol. 20, pp. 1–26, 2019.
- [334] Y. Yin, W. Dong, Z. Zhan, and J. Li, “A Peak-Selection RBF Mesh Morphing Method for Subject-Specific Child Occupant Modeling,” in *Volume 13: Design, Reliability, Safety, and Risk*, vol. 13. American Society of Mechanical Engineers, nov 2018, pp. 1–9.
- [335] L. Geronzi, E. Gasparotti, K. Capellini, U. Cella, C. Groth, S. Porziani, A. Chiappa, S. Celi, and M. E. Biancolini, “Advanced radial basis functions mesh morphing for high fidelity fluid-structure interaction with known movement of the walls: Simulation of an aortic valve,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 12142 LNCS, pp. 280–293, 2020.
- [336] M. Falsafioon, S. Arabi, R. Camarero, and F. Guibault, “Comparison of Two Mesh Smoothing Techniques for Unstructured Grids,” *IOP Conference Series: Earth and Environmental Science*, vol. 22, 2013.
- [337] Y. Yu, K. Zhou, D. Xuz, X. Shi, H. Bao, B. Guo, and H. Y. Shum, “Mesh editing with poisson-based gradient field manipulation,” *ACM SIGGRAPH 2004 Papers, SIGGRAPH 2004*, pp. 644–651, 2004.

- [338] J. Vollmer, R. Mencl, and H. Müller, “Improved laplacian smoothing of noisy surface meshes,” *Computer Graphics Forum*, vol. 18, no. 3, pp. 131–138, 1999.
- [339] S. M. Shontz and S. A. Vavasis, “A linear weighted laplacian smoothing framework for warping tetrahedral meshes,” *BIT Numerical Mathematics*, oct 2004.
- [340] M. L. Brewer, L. Freitag Diachin, P. M. Knupp, T. Leurent, and D. J. Melander, “The Mesquite Mesh Quality Improvement Toolkit.” in *Proceedings of the 12th International Meshing Roundtable*, 2003.
- [341] D. S. Thompson, S. Chalasani, and B. K. Soni, “Generation of Volume Meshes By Extrusion From Surface Meshes of Arbitrary Topology,” *Imr*, vol. M, no. May, 2000.
- [342] P. M. Knupp, “Winslow smoothing on two-dimensional unstructured meshes,” *Engineering with Computers*, vol. 15, no. 3, pp. 263–268, 1999.
- [343] —, “Achieving finite element mesh quality via optimization of the Jacobian matrix norm and associated quantities. Part II - A framework for volume mesh optimization and the condition number of the Jacobian matrix,” *International Journal for Numerical Methods in Engineering*, vol. 48, no. 8, pp. 1165–1185, jul 2000.
- [344] J. E. Castillo, S. Steinberg, and P. J. Roache, “On the folding of numerically generated grids: Use of a reference grid,” *Communications in Applied Numerical Methods*, vol. 4, no. 4, pp. 471–481, jul 1988.
- [345] L. A. Freitag, “On combining laplacian and optimization-based mesh smoothing techniques,” *American Society of Mechanical Engineers, Applied Mechanics Division, AMD*, vol. 220, pp. 37–43, 1997.
- [346] M. Meyer, M. Desbrun, P. Schröder, and A. H. Barr, “Discrete Differential-Geometry Operators for Triangulated 2-Manifolds,” pp. 35–57, 2003.
- [347] G. Xu, “Discrete Laplace–Beltrami Operator on Sphere and Optimal Spherical Triangulations,” *International Journal of Computational Geometry Applications*, vol. 16, no. 01, pp. 75–93, 2006.
- [348] G. Taubin, “A signal processing approach to fair surface design,” in *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques - SIGGRAPH '95*. New York, New York, USA: ACM Press, 1995, pp. 351–358.

-
- [349] W. R. Buell and B. A. Bush, “Mesh Generation—A Survey,” *Journal of Engineering for Industry*, vol. 95, no. 1, pp. 332–338, feb 1973.
- [350] D. A. Field, “Laplacian smoothing and Delaunay triangulations,” *Communications in Applied Numerical Methods*, vol. 4, no. 6, pp. 709–712, 1988.
- [351] J. Vollmer, R. Mencl, and H. Müller, “Improved laplacian smoothing of noisy surface meshes,” *Computer Graphics Forum*, vol. 18, no. 3, pp. 131–138, 1999.
- [352] S. M. Shontz and S. A. Vavasis, “A mesh warping algorithm based on weighted Laplacian smoothing,” *Proceedings 12th Int. Meshing Roundtable*, pp. 147–158, 2003.
- [353] A. Nealen, T. Igarashi, O. Sorkine, and M. Alexa, “Laplacian mesh optimization,” *Proceedings - GRAPHITE 2006: 4th International Conference on Computer Graphics and Interactive Techniques in Australasia and Southeast Asia*, pp. 381–389, 2006.
- [354] V. Surazhsky and C. Gotsman, “Intrinsic morphing of compatible triangulations,” *International Journal of Shape Modeling*, vol. 9, no. 2, pp. 191–201, 2003.
- [355] D. Anisimov, D. Panozzo, and K. Hormann, “Blended barycentric coordinates,” *Computer Aided Geometric Design*, vol. 52-53, pp. 205–216, 2017.
- [356] A. M. Winslow, “Numerical Solution of the Quasilinear Triangle,” *Journal of Computational Physics*, vol. 2, pp. 149–172, 1967.
- [357] P. M. Knupp, “Applications of mesh smoothing: Copy, morph, and sweep on unstructured quadrilateral meshes,” *International Journal for Numerical Methods in Engineering*, vol. 45, no. 1, pp. 37–45, 1999.
- [358] B. A. Szabo, “Some recent developments in finite element analysis,” *Computers and Mathematics with Applications*, vol. 5, no. 2, pp. 99–115, 1979.
- [359] H. Runnemalm and S. Hyun, “Three-dimensional welding analysis using an adaptive mesh scheme,” *Computer Methods in Applied Mechanics and Engineering*, vol. 189, no. 2, pp. 515–523, 2000.
- [360] T. Ju, S. Schaefer, and J. Warren, “Mean value coordinates for closed triangular meshes,” *ACM Transactions on Graphics*, vol. 24, no. 3, p. 561, 2005.

- [361] M. S. Floater, *Generalized barycentric coordinates and applications*, 2015, vol. 24, no. May.
- [362] ———, “Mean value coordinates,” *Computer Aided Geometric Design*, vol. 20, no. 1, pp. 19–27, 2003.
- [363] D. Baldwin and M. Weber, “Fast Ray-Triangle Intersections by Coordinate Transformation,” *Journal of Computer Graphics Techniques (JCGT)*, vol. 5, no. 3, pp. 39–49, 2016.
- [364] M. S. Floater, G. Kós, and M. Reimers, “Mean value coordinates in 3D,” *Computer Aided Geometric Design*, vol. 22, no. 7 SPEC. ISS., pp. 623–631, 2005.
- [365] L. Chen, J. Huang, H. Sun, and H. Bao, “Cage-based deformation transfer,” *Computers and Graphics (Pergamon)*, vol. 34, no. 2, pp. 107–118, 2010.
- [366] Á. Tóth and R. Kunkli, “An automated method for generating customizable cages using barycentric coordinates,” *Annales Mathematicae et Informaticae*, vol. 49, pp. 167–180, 2018.
- [367] A. Jacobson, D. Panozzo, C. Schüller, O. Diamanti, Q. Zhou, N. Pietroni, and Others, “libigl: A simple C++ geometry processing library,” 2018.
- [368] E. Fogel and M. Teillaud, “The Computational Geometry Algorithms Library CGAL,” *ACM Communications in Computer Algebra*, vol. 49, no. 1, pp. 10–12, 2015.
- [369] R. Zhang, “3D mesh metamorphosis from spherical parameterization for conceptual design (Thesis),” Doctoral Thesis, Iowa State University, 2011.
- [370] A. Sheffer and V. Kraevoy, “Pyramid coordinates for morphing and deformation,” *Proceedings - 2nd International Symposium on 3D Data Processing, Visualization, and Transmission. 3DPVT 2004*, pp. 68–75, 2004.
- [371] J. Kim, B. J. Miller, and S. M. Shontz, “A hybrid mesh deformation algorithm using anisotropic PDEs and multiobjective mesh optimization,” *Computers and Mathematics with Applications*, vol. 70, no. 8, pp. 1830–1851, 2015.
- [372] A. D. Boer, M. S. V. D. Schoot, and H. Bijl, “New Method for Mesh Moving Based on Radial Basis Function interpolation,” *European Conference on Computational Fluid Dynamics ECCOMAS*, pp. 1–16, 2006.

- [373] A. Bijar, P.-y. Rohan, P. Perrier, Y. Payan, A. Bijar, P.-y. Rohan, P. Perrier, and Y. Payan, “Atlas-Based Automatic Generation of Subject-Specific Finite Element Tongue Meshes,” 2016.
- [374] Z. Salo, H. Kreder, and C. M. Whyne, “Influence of pelvic shape on strain patterns: A computational analysis using finite element mesh morphing techniques,” *Journal of Biomechanics*, vol. 116, p. 110207, 2021.
- [375] V. B. Shim, J. W. Fernandez, P. B. Gamage, C. Regnery, D. W. Smith, B. S. Gardiner, D. G. Lloyd, and T. F. Besier, “Subject-specific finite element analysis to characterize the influence of geometry and material properties in Achilles tendon rupture,” *Journal of Biomechanics*, vol. 47, no. 15, pp. 3598–3604, 2014.
- [376] F. P. Nikolopoulos, E. I. Zacharaki, D. Stanev, and K. Moustakas, “Personalized knee geometry modeling based on multi-atlas segmentation and mesh refinement,” *IEEE Access*, vol. 8, pp. 56 766–56 781, 2020.
- [377] S. C. Tadepalli, K. H. Shivanna, V. A. Magnotta, N. A. Kallemeyn, and N. M. Grosland, “Toward the Development of Virtual Surgical Tools to Aid Orthopaedic FE Analyses,” *EURASIP Journal on Advances in Signal Processing*, vol. 2010, no. 1, p. 190293, dec 2009.
- [378] L. Geronzi, E. Gasparotti, K. Capellini, U. Cella, C. Groth, S. Porziani, A. Chiappa, S. Celi, and M. E. Biancolini, *Advanced radial basis functions mesh morphing for high fidelity fluid-structure interaction with known movement of the walls: Simulation of an aortic valve*. Springer International Publishing, 2020, vol. 12142 LNCS.
- [379] N. Hraiech, M. Rochette, F. Taddei, and M. Viceconti, “Finite element femur generation using morphing and spherical parameterization techniques,” *Proceedings of the International Congress on Computational Bioengineering*, p. 27, 2009.
- [380] B. Stroustrup, *the C++ Programming Language 4th Edition*, 2013.
- [381] B. Hoffman and K. Martin, “CMake,” in *The Architecture of Open Source Applications*, 2014.
- [382] C. Sanderson and R. Curtin, “Armadillo: a template-based C++ library for linear algebra,” *The Journal of Open Source Software*, vol. 1, no. 2, p. 26, jun 2016.

- [383] A. Munshi, “The OpenCL specification,” *2009 IEEE Hot Chips 21 Symposium, HCS 2009*, pp. 11–314, 2016.
- [384] C. Sanderson, “Armadillo: An Open Source C++ Linear Algebra Library for Fast Prototyping and Computationally Intensive Experiments,” vol. 1, pp. 1–16, 2011.
- [385] E. W. Dijkstra, “A note on two problems in connexion with graphs,” *Numerische Mathematik*, vol. 1, no. 1, pp. 269–271, dec 1959.
- [386] P. Dierckx, *Curve and Surface Fitting with Splines*. Oxford University Press, 1993.
- [387] T. Möller and B. Trumbore, “Fast, Minimum Storage Ray-Triangle Intersection,” *Journal of Graphics Tools*, vol. 2, no. 1, pp. 21–28, jan 1997.
- [388] R. Anderson, J. Andrej, A. Barker, J. Bramwell, J. S. Camier, J. Cervený, V. Dobrev, Y. Dudouit, A. Fisher, T. Kolev, W. Pazner, M. Stowell, V. Tomov, I. Akkerman, J. Dahm, D. Medina, and S. Zampini, “MFEM: A modular finite element methods library,” *Computers and Mathematics with Applications*, vol. 81, pp. 42–74, 2021.
- [389] L. Greengard and V. Rokhlin, “A fast algorithm for particle simulations,” *Journal of Computational Physics*, vol. 73, no. 2, pp. 325–348, dec 1987.
- [390] M. Smolik, V. Skala, and O. Nedved, “A comparative study of LOWESS and RBF approximations for visualization,” *arXiv*, 2018.
- [391] A. Iske, “Optimal distribution of centers for radial basis function methods,” Tech. Rep. M0004, 2000.
- [392] M.-S. Hartig, “Approximation of Gaussian Curvature by the Angular Defect: An Error Analysis,” *Mathematical and Computational Applications*, vol. 26, no. 1, p. 15, 2021.
- [393] W. Kabsch, “A solution for the best rotation to relate two sets of vectors,” *Acta Crystallographica Section A*, vol. 32, no. 5, pp. 922–923, sep 1976.
- [394] K. S. Arun, T. S. Huang, and S. D. Blostein, “Least-Squares Fitting of Two 3-D Point Sets,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, no. 5, pp. 698–700, sep 1987.
- [395] F. L. Markley and D. Mortari, “How to estimate attitude from vector observations,” *Advances in the Astronautical Sciences*, vol. 103, no. PART III, pp. 1979–1996, 2000.

- [396] G. Wahba, “A Least Squares Estimate of Satellite Attitude,” pp. 409–409, 1965.
- [397] Sven Woop, Carsten Benthin, and Ingo Wald, “Watertight Ray/Triangle Intersection,” *Journal of Computer Graphics Techniques (JCGT)*, vol. 2, no. 1, pp. 65–82, 2013.
- [398] K. Jeffery, “Firefly detection with half buffers,” *Proceedings - DigiPro 2018: Digital Production Symposium*, 2018.
- [399] K. Hormann and A. Agathos, “The point in polygon problem for arbitrary polygons,” vol. 20, pp. 131–144, 2001.
- [400] M. Galetzka and P. Glauner, “A simple and correct even-odd algorithm for the point-in-polygon problem for complex polygons,” *VISIGRAPP 2017 - Proceedings of the 12th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, vol. 1, pp. 175–178, 2017.
- [401] X. Wang and X. Qian, “A statistical atlas based approach to automated subject-specific FE modeling,” *CAD Computer Aided Design*, vol. 70, pp. 67–77, 2016.
- [402] J. Huang, L. Chen, X. Liu, and H. Bao, “Efficient mesh deformation using tetrahedron control mesh,” *Computer Aided Geometric Design*, vol. 26, no. 6, pp. 617–626, 2009.
- [403] J. W. Fernandez, P. Mithraratne, S. F. Thrupp, M. H. Tawhai, and P. J. Hunter, “Anatomically based geometric modelling of the musculo-skeletal system and other organs,” *Biomechanics and Modeling in Mechanobiology*, vol. 2, no. 3, pp. 139–155, 2004.
- [404] “meshplex 0.15.14,” [Accessed: 2022-01-29]. [Online]. Available: <https://web.archive.org/web/20220129131221/https://pypi.org/project/meshplex/0.15.14/>
- [405] C. van Elst, P. Caekebeke, H. Vandenuecker, J. Bellemans, and L. Scheys, “Can a ratio between medial and lateral meniscal volumes be calculated to determine critical meniscal volume in view of post-meniscectomy symptoms?” *Acta Orthopaedica Belgica*, vol. 86, no. 1, pp. 77–81, 2020.
- [406] D. Xu, J. Van Der Voet, N. M. Hansson, S. Klein, E. H. Oei, F. Wagner, S. M. Bierma-Zeinstra, and J. Runhaar, “Association between meniscal volume and development of knee osteoarthritis,” *Rheumatology (United Kingdom)*, vol. 60, no. 3, pp. 1392–1399, 2021.

- [407] J. Hashemi, N. Chandrashekar, B. Gill, B. D. Beynon, J. R. Slauterbeck, R. C. Schutt, H. Mansouri, and E. Dabezies, “The geometry of the tibial plateau and its influence on the biomechanics of the tibiofemoral joint,” *Journal of Bone and Joint Surgery - Series A*, vol. 90, no. 12, pp. 2724–2734, 2008.
- [408] L. Zheng, R. Carey, E. Thorhauer, S. Tashman, C. Harner, and X. Zhang, “In vivo tibiofemoral skeletal kinematics and cartilage contact arthrokinematics during decline walking after isolated meniscectomy,” *Medical Engineering Physics*, vol. 51, no. 1, pp. 41–48, jan 2018.
- [409] P. Yin, J.-S. Li, W. A. Kernkamp, T.-y. Tsai, S.-h. Baek, A. Hosseini, L. Lin, P. Tang, and G. Li, “Analysis of in-vivo articular cartilage contact surface of the knee during a step-up motion,” *Clinical Biomechanics*, vol. 49, pp. 101–106, nov 2017.
- [410] J. Yao, J. Crockett, M. D’Souza, G. Day, R. Wilcox, A. Jones, and M. Mengoni, “Sensitivity Analysis of Meniscus Assumptions in a Validated Tibiofemoral Finite Element Model,” *Proceedings of ESB2021 – 26th Congress of the European Society of Biomechanics*, 2021.
- [411] “Open Knee(s): Virtual Biomechanical Representations of the Knee Joint,” [Accessed: 2022-01-31]. [Online]. Available: <https://web.archive.org/web/20220131073353/https://simtk.org/projects/openknee>
- [412] D. M. Pierce, “UConn - Interdisciplinary Mechanics Laboratory: Software and Downloads,” [Accessed: 2022-01-31]. [Online]. Available: <https://web.archive.org/web/20220131142808/https://im.engr.uconn.edu/downloads.php>
- [413] “The Osteoarthritis Initiative,” [Accessed: 2022-01-31]. [Online]. Available: <https://web.archive.org/web/20220131142021/https://nda.nih.gov/oai/>
- [414] M. C. Nevitt and G. Lester, “Osteoarthritis Initiative,” *JAMA*, vol. 291, no. 16, p. 1951, apr 2004.
- [415] C. G. Peterfy, E. Schneider, and M. Nevitt, “The osteoarthritis initiative: report on the design rationale for the magnetic resonance imaging protocol for the knee,” *Osteoarthritis and Cartilage*, vol. 16, no. 12, pp. 1433–1441, 2008.
- [416] J. A. Weiss, B. N. Maker, and S. Govindjee, “Finite element implementation of incompressible, transversely isotropic hyperelasticity,” *Computer Methods in Applied Mechanics*

- and Engineering*, vol. 135, no. 1-2, pp. 107–128, aug 1996.
- [417] K. Miller, *Biomechanics of the Brain Second Edition*, 2019, vol. 339, no. 24.
- [418] G. Papaioannou, G. Nianios, C. Mitrogiannis, D. Fyhrie, S. Tashman, and K. H. Yang, “Patient-specific knee joint finite element model validation with high-accuracy kinematics from biplane dynamic Roentgen stereogrammetric analysis,” *Journal of Biomechanics*, vol. 41, no. 12, pp. 2633–2638, 2008.
- [419] J. D. Hunter, “Matplotlib: A 2D Graphics Environment,” *Computing in Science Engineering*, vol. 9, no. 3, pp. 90–95, 2007.
- [420] G. Van Rossum and F. L. Drake, *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace, 2009.
- [421] J. Ahrens, B. Geveci, and C. Law, “ParaView: An End-User Tool for Large-Data Visualization,” in *Visualization Handbook*. Elsevier, 2005, pp. 717–731.
- [422] G. Taubin, “Geometric Signal Processing on Polygonal Meshes,” 2000.
- [423] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr, “Implicit fairing of irregular meshes using diffusion and curvature flow,” in *Proceedings of the 26th annual conference on Computer graphics and interactive techniques - SIGGRAPH '99*. New York, New York, USA: ACM Press, 1999, pp. 317–324.
- [424] H. Erbagci, E. Gumusburun, M. Bayram, G. Karakurum, and A. Sirikci, “The normal menisci: In vivo MRI measurements,” *Surgical and Radiologic Anatomy*, vol. 26, no. 1, pp. 28–32, 2004.
- [425] Z. Li, J. Hu, M. P. Reed, J. D. Rupp, C. N. Hoff, J. Zhang, and B. Cheng, “Development, validation, and application of a parametric pediatric head finite element model for impact simulations,” *Annals of Biomedical Engineering*, vol. 39, no. 12, pp. 2984–2997, 2011.
- [426] Z. Salo, M. Beek, D. Wright, A. Maloul, and C. M. Whyne, “Analysis of pelvic strain in different gait configurations in a validated cohort of computed tomography based finite element models,” *Journal of Biomechanics*, vol. 64, pp. 120–130, 2017.
- [427] U. Pinkall and K. Polthier, “Computing discrete minimal surfaces and their conjugates,” *Experimental Mathematics*, 1993.

- [428] Herholz, “General discrete Laplace operators on polygonal meshes,” *Informatik.Hu-Berlin.De*, p. 49, 2012.
- [429] M. Specht, “Spherical surface parameterization and its application to geometric morphometric analysis of the braincase (Thesis),” Doctoral Thesis, University of Zurich, 2007.
- [430] E. Livieratos, “On the Study of the Geometric Properties of Historical Cartographic Representations,” *Cartographica: The International Journal for Geographic Information and Geovisualization*, vol. 41, no. 2, pp. 165–176, jun 2006.
- [431] M. Wardetzky, S. Mathur, F. Kaelberer, and E. Grinspun, “Discrete Laplace operators: No free lunch,” *Eurographics Symposium on Geometry Processing*, pp. 33–37, 2007.
- [432] J. M. Marzo and J. Gurske-DePerio, “Effects of medial meniscus posterior horn avulsion and repair on tibiofemoral contact area and peak contact pressure with clinical implications,” *American Journal of Sports Medicine*, vol. 37, no. 1, pp. 124–129, 2009.
- [433] F. Bachtar, X. Chen, and T. Hisada, “Finite element contact analysis of the hip joint,” *Medical and Biological Engineering and Computing*, vol. 44, no. 8, pp. 643–651, 2006.
- [434] K. Messner and J. Gao, “The menisci of the knee joint. Anatomical and functional characteristics, and a rationale for clinical treatment,” *J Anat*, vol. 193 (Pt 2, pp. 161–178, 1998.
- [435] D. Campbell, “Finite element mesh generation,” *New Zealand Engineering*, vol. 51, no. 4, pp. 12–X7, 1996.
- [436] S. T. More and R. S. Bindu, “Effect of Mesh Size on Finite Element Analysis of Plate Structure,” *International Journal of Engineering Science and Innovative Technology*, vol. 4, no. 3, pp. 181–185, 2015.
- [437] A. Logg, “Efficient representation of computational meshes,” *International Journal of Computational Science and Engineering*, vol. 4, no. 4, pp. 283–295, 2009.
- [438] G. Salomon, A. Leclercq, S. Akkouche, and E. Galin, “Normal control using N-adic subdivision schemes,” *Proceedings - SMI 2002: Shape Modeling International 2002*, vol. 2002, pp. 21–28, 2002.

- [439] K. Xu and G. Chen, “Hexahedral Mesh Structure Visualization and Evaluation,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 1173–1182, 2019.
- [440] E. Soriano, P. Tornabell, D. Naicu, and G. H. Filz, “Topologically-based curvature in thin elastic shell networks,” 2015.
- [441] D. J. Yoo, “Three-dimensional morphing of similar shapes using a template mesh,” *International Journal of Precision Engineering and Manufacturing*, vol. 10, no. 1, pp. 55–66, 2009.
- [442] T. Haut Donahue, M. Hull, M. Rashid, and C. Jacobs, “A Finite Element Model of the Human Knee Joint for the Study of Tibio-Femoral Contact,” *Journal of Biomechanical Engineering*, vol. 124, no. 3, pp. 273–280, 2002.
- [443] X. Roca, J. Sarrate, and A. Huerta, “Surface Mesh Projection for Hexahedral Mesh Generation By Sweeping,” *Imr’04*, vol. m, 2004.
- [444] C. H. Peng, E. Zhang, Y. Kobayashi, and P. Wonka, “Connectivity editing for quadrilateral meshes,” *ACM Transactions on Graphics*, vol. 30, no. 6, pp. 1–12, 2011.
- [445] N. Aghdaii, “5-6-7 Meshes,” *Proceedings of Graphics Interface*, pp. 27–34, 2012.
- [446] T. Jiang, J. Huang, Y. Wang, Y. Tong, and H. Bao, “Frame field singularity correction for automatic hexahedralization,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 8, pp. 1189–1199, 2014.
- [447] B. Walter and P. Shirley, “Cost Analysis of a Ray Tracing Algorithm,” pp. 1–18, 1997.
- [448] L. Branets and G. F. Carey, “A local cell quality metric and variational grid smoothing algorithm,” *Engineering with Computers*, vol. 21, no. 1, pp. 19–28, 2005.
- [449] H. Couteau, Y. Payan, and H. Lavalley, “The mesh-matching algorithm: an automatic 3D mesh generator for finite element structures,” vol. 33, pp. 1005–1009, 2000.
- [450] M. Dienst, P. E. Greis, B. J. Ellis, K. N. Bachus, and R. T. Burks, “Effect of lateral meniscal allograft sizing on contact mechanics of the lateral tibial plateau: An experimental study in human cadaveric knee joints,” *American Journal of Sports Medicine*, 2007.

- [451] A. Huang, M. L. Hull, S. M. Howell, and T. H. Donahue, "Identification of cross-sectional parameters of lateral meniscal allografts that predict tibial contact pressure in human cadaveric knees," *Journal of Biomechanical Engineering*, vol. 124, no. 5, pp. 481–489, 2002.
- [452] H. Naghibi Beidokhti, D. Janssen, M. Khoshgoftar, A. Sprengers, E. S. Perdahcioglu, T. Van den Boogaard, and N. Verdonschot, "A comparison between dynamic implicit and explicit finite element simulations of the native knee joint," *Medical Engineering and Physics*, vol. 38, no. 10, pp. 1123–1130, 2016.
- [453] R. Kubota, H. Koga, N. Ozeki, J. Matsuda, Y. Kohno, M. Mizuno, H. Katano, and I. Sekiya, "The effect of a centralization procedure for extruded lateral meniscus on load distribution in porcine knee joints at different flexion angles," *BMC Musculoskeletal Disorders*, vol. 21, no. 1, pp. 1–8, 2020.
- [454] P. Debieux, A. E. Jimenez, J. V. Novaretti, C. C. Kaleka, D. E. Kriscenski, D. C. Astur, E. Obopilwe, L. M. Tamburini, L. N. Muench, M. P. Cote, M. Cohen, and K. J. Coyner, "Medial meniscal extrusion greater than 4 mm reduces medial tibiofemoral compartment contact area: a biomechanical analysis of tibiofemoral contact area and pressures with varying amounts of meniscal extrusion," *Knee Surgery, Sports Traumatology, Arthroscopy*, no. 0123456789, 2020.
- [455] C. Zhou and R. Willing, "Multiobjective Design Optimization of a Biconcave Mobile-Bearing Lumbar Total Artificial Disk Considering Spinal Kinematics, Facet Joint Loading, and Metal-on-Polyethylene Contact Mechanics," *Journal of Biomechanical Engineering*, vol. 142, no. 4, 2020.
- [456] Y. Bei and B. J. Fregly, "Multibody dynamic simulation of knee contact mechanics," *Medical Engineering Physics*, vol. 26, no. 9, pp. 777–789, nov 2004.
- [457] P. R. P. Braz and W. G. Silva, "Meniscus morphometric study in humans," *Journal of Morphological Sciences*, vol. 27, no. 2, pp. 62–66, 2010.
- [458] I. F. Cengiz, M. Pitikakis, L. Cesario, P. Parascandolo, L. Vosilla, G. Viano, J. M. Oliveira, and R. L. Reis, "Building the basis for patient-specific meniscal scaffolds: From human knee MRI to fabrication of 3D printed scaffolds," *Bioprinting*, vol. 1-2, pp. 1–10, 2016.

- [459] J. Hasan, “Decellularisation and characterisation of porcine (Thesis),” Doctoral Thesis, University of Leeds, 2014.
- [460] R. Kijowski, H. Rosas, A. Williams, and F. Liu, “MRI characteristics of torn and untorn post-operative menisci,” *Skeletal Radiology*, vol. 46, no. 10, pp. 1353–1360, oct 2017.
- [461] J. W. Tukey, *Exploratory Data Analysis*. Addison-Wesley, 1977.
- [462] B. Iglewicz and D. C. Hoaglin, *Volume 16: How to Detect and Handle Outliers*. ASQC/Quality Press, 1993, vol. 36, no. 3.
- [463] W. Zhang, Y. Ma, J. Zheng, and W. J. Allen, “Tetrahedral mesh deformation with positional constraints,” *Computer Aided Geometric Design*, vol. 81, p. 101909, 2020.
- [464] Z. Salo, M. Beek, and C. M. Whyne, “Evaluation of mesh morphing and mapping techniques in patient specific modeling of the human pelvis,” *International Journal for Numerical Methods in Biomedical Engineering*, vol. 29, no. 1, pp. 104–113, jan 2013.
- [465] I. Farmaga, P. Shmigelskyi, P. Spiewak, and L. Ciupinski, “Evaluation of computational complexity of finite element analysis,” *2011 11th International Conference - The Experience of Designing and Application of CAD Systems in Microelectronics, CADSM 2011*, no. 4, pp. 213–214, 2011.
- [466] M. L. Neal and R. Kerckhoffs, “Current progress in patient-specific modeling,” *Briefings in Bioinformatics*, vol. 11, no. 1, pp. 111–126, 2009.
- [467] M. E. Biancolini, A. Chiappa, U. Cella, E. Costa, C. Groth, and S. Porziani, *Radial Basis Functions Mesh Morphing*. Springer International Publishing, 2020, vol. 1.
- [468] S. L. Lee, M. Lerotic, A. Sani, Y. Zhao, J. Keegan, Y. Hao, and G. Z. Yang, “Articulated postures for subject-specific RF simulation,” *2010 International Conference on Body Sensor Networks, BSN 2010*, pp. 96–101, 2010.
- [469] P. Rondao Alface, M. De Craene, and B. B. Macq, “Three-dimensional image quality measurement for the benchmarking of 3D watermarking schemes,” *Security, Steganography, and Watermarking of Multimedia Contents VII*, vol. 5681, no. May 2014, p. 230, 2005.
- [470] Z. Karni and C. Gotsman, “Spectral compression of mesh geometry,” *Proceedings of the ACM SIGGRAPH Conference on Computer Graphics*, pp. 279–286, 2000.

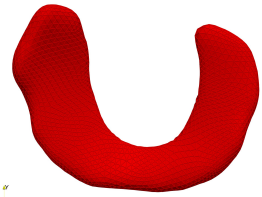



- [471] J. Cates, S. Elhabian, and R. Whitaker, “ShapeWorks: Particle-Based Shape Correspondence and Visualization Software,” *Statistical Shape and Deformation Analysis: Methods, Implementation and Applications*, pp. 257–298, 2017.
- [472] Y. Zhang, C. Bajaj, and G. Xu, “Surface smoothing and quality improvement of quadrilateral/hexahedral meshes with geometric flow,” *Proceedings of the 14th International Meshing Roundtable, IMR 2005*, pp. 449–468, 2005.
- [473] T. Toulorge, C. Geuzaine, J. F. Remacle, and J. Lambrechts, “Robust untangling of curvilinear meshes,” *Journal of Computational Physics*, vol. 254, pp. 8–26, 2013.
- [474] T. P. Salsón, “Simplification, Approximation and Deformation of Large Models (Thesis),” Doctoral Thesis, University of Girona, 2011.
- [475] M. Pereañez, “Enlargement, subdivision and individualization of statistical shape models: Application to 3D medical image segmentation (Thesis),” Doctoral Thesis, University Pompeu Fabra, 2017.
- [476] M. Zhang and P. Golland, “Statistical shape analysis: From landmarks to diffeomorphisms,” *Medical Image Analysis*, vol. 33, no. 3, pp. 155–158, oct 2016.
- [477] L. Zhu and T. Jiang, “Parameterization of 3D brain structures for statistical shape analysis,” *Medical Imaging 2004: Image Processing*, vol. 5370, p. 1254, 2004.
- [478] S. Seshamani, G. Chintalapani, and R. Taylor, “Iterative refinement of point correspondences for 3D statistical shape models,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 6892 LNCS, no. PART 2, pp. 417–425, 2011.
- [479] S. Yoshizawa, A. Belyaev, and H. P. Seidel, “Skeleton-based variational mesh deformations,” *Computer Graphics Forum*, vol. 26, no. 3, pp. 255–264, 2007.
- [480] N. D. Cornea, D. Silver, and P. Min, “Curve-skeleton properties, applications, and algorithms,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 3, pp. 530–548, 2007.
- [481] C. K. Fitzpatrick, M. A. Baldwin, P. J. Laz, D. P. FitzPatrick, A. L. Lerner, and P. J. Rullkoetter, “Development of a statistical shape model of the patellofemoral joint for

- investigating relationships between shape and function,” *Journal of Biomechanics*, vol. 44, no. 13, pp. 2446–2452, 2011.
- [482] D. Kainmueller, H. Lamecker, S. Zachow, and H. C. Hege, “An articulated statistical shape model for accurate hip joint segmentation,” *Proceedings of the 31st Annual International Conference of the IEEE Engineering in Medicine and Biology Society: Engineering the Future of Biomedicine, EMBC 2009*, pp. 6345–6351, 2009.
- [483] S. Ordas, E. Oubel, R. Leta, F. Carreras, and A. F. Frangi, “A statistical shape model of the heart and its application to model-based segmentation,” *Medical Imaging 2007: Physiology, Function, and Structure from Medical Images*, vol. 6511, p. 65111K, 2007.
- [484] P. Soltanmohammadi, J. Elwell, V. Veeraraghavan, G. S. Athwal, and R. Willing, “Investigating the Effects of Demographics on Shoulder Morphology and Density Using Statistical Shape and Density Modeling,” *Journal of Biomechanical Engineering*, vol. 142, no. 12, pp. 1–15, 2020.
- [485] Y. Zhang, “Challenges and advances in image-based geometric modeling and mesh generation,” *Lecture Notes in Computational Vision and Biomechanics*, vol. 3, pp. 1–10, 2013.
- [486] Z. Salo, M. Beek, D. Wright, and C. Marisa Whyne, “Computed tomography landmark-based semi-automated mesh morphing and mapping techniques: Generation of patient specific models of the human pelvis without segmentation,” *Journal of Biomechanics*, vol. 48, no. 6, pp. 1125–1132, 2015.
- [487] S. Zyga, “Application of ISO 13485:2003 in Biomedical Engineering: a Systematic Review,” *International Journal of Caring Sciences*, vol. 4, no. 2, pp. 58–65, 2011.
- [488] M. Freutel, A. M. Seitz, F. Galbusera, A. Bornstedt, V. Rasche, M. L. Tate, A. Ignatius, and L. Dürselen, “Medial meniscal displacement and strain in three dimensions under compressive loads: MR assessment,” *Journal of Magnetic Resonance Imaging*, vol. 40, no. 5, pp. 1181–1188, 2014.
- [489] J. Schwer, M. M. Rahman, K. Stumpf, V. Rasche, A. Ignatius, L. Dürselen, and A. M. Seitz, “Degeneration Affects Three-Dimensional Strains in Human Menisci: In situ MRI Acquisition Combined With Image Registration,” *Frontiers in Bioengineering and Biotechnology*, vol. 8, no. September, pp. 1–12, 2020.

Appendices






A Meniscus Target Geometries

Table A.1: Meniscus target geometries with their feature descriptions and sources (meshes selected for comparative analysis in cases D and E are marked with an asterisk)

Ref.	Superior	Feature Description
LM_00		<p>Smooth Surface</p> <p>Smooth Geometry</p> <p>Small Radial Bulge</p> <p>Large Wedge Angle</p> <p>Large Inferior Bulge (Anterior and Posterior Regions)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>)</p>
LM_01*		<p>Smooth Surface</p> <p>Possible Extrusion</p> <p>Large Wedge Angle</p> <p>Tear (Medial/Anterior Region)</p> <p>Large Inferior Bulge (Anterior and Posterior Regions)</p> <p>Large Radial Bulge (Anterior and Posterior Regions)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [19]</p>
LM_02		<p>Smooth Surface</p> <p>Extrusion*</p> <p>Small-to-Large Wedge Angle (Anterior-Posterior)</p> <p>Large Inferior Bulge</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [80]</p>
LM_03		<p>Irregular Surface</p> <p>Large Horn-Attachments</p> <p>Possible Degeneration or Tears (Wedge Region)</p> <p>Large Inferior Bulge</p> <p>Source: SimTK/Open Knee (<i>uconn.edu</i>) [40]</p>

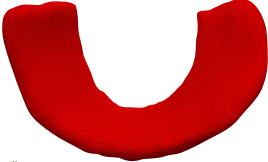
Continued on next page

Table A.1 – *Continued from previous page*

Ref.	Superior	Feature Description
LM_04*		<p>Irregular Surface</p> <p>Tapered Shape</p> <p>Small Horn-Attachments</p> <p>Large Inferior Bulge</p> <p>Source: SimTK/Open Knee (<i>uconn.edu</i>) [40]</p>
LM_05*		<p>Tapered Shape</p> <p>Blocky</p> <p>Smooth Surface</p> <p>Small-to-Large Wedge Angle (Posterior-Anterior)</p> <p>Small Horn-Attachments</p> <p>Source: iMBE (<i>Grant Number: EP/P001076/1</i>)</p>
LM_06*		<p>Blocky</p> <p>Smooth Surface</p> <p>Small Horn-Attachments</p> <p>Small-to-Large Wedge Angle (Anterior-Posterior)</p> <p>Large Inferior Bulge</p> <p>Source: iMBE (<i>Grant Number: EP/P001076/1</i>)</p>
LM_07		<p>Blocky</p> <p>Non-Semi-Lunar Shape</p> <p>Smooth Surface</p> <p>Source: iMBE (<i>Grant Number: EP/P001076/1</i>)</p>
LM_08*		<p>Irregular Shape</p> <p>Irregular Surface</p> <p>Possible Degeneration and/or Tear</p> <p>Large Wedge Angle</p> <p>Large Inferior Bulge (Posterior Region)</p> <p>Large Radial Bulge (Medial-Anterior Region)</p> <p>Source: iMBE (<i>Grant Number: EP/P001076/1</i>)</p>


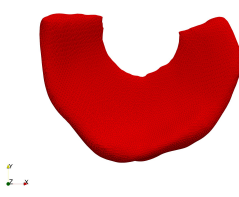
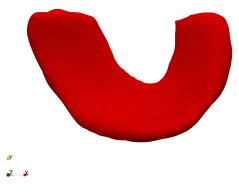


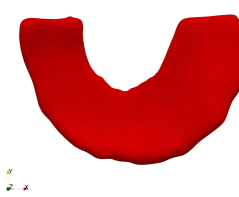
Continued on next page

Table A.1 – *Continued from previous page*

Ref.	Superior	Feature Description
LM_09		<p>Smooth Surface</p> <p>Possible Extrusion</p> <p>Small Wedge Angle</p> <p>Large Radial Bulge (Medial Region)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
LM_10		<p>Irregular Geometry (Medial/Anterior Wedge Region)</p> <p>Possible Extrusion</p> <p>Small Wedge Angle</p> <p>Possible Degeneration and/or Tear</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
LM_11		<p>Smooth Surface</p> <p>Small-to-Large Wedge Angle (Anterior-Posterior)</p> <p>Possible Tear (Medial/Posterior Wedge Region)</p> <p>Possible Degeneration (Medial/Posterior Wedge Region)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
LM_12		<p>Smooth Surface</p> <p>Small-to-Large Wedge Angle (Anterior-Posterior)</p> <p>Possible Degeneration (Medial Wedge Region)</p> <p>Possible Tear (Medial Wedge Region)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
LM_13		<p>Smooth Surface</p> <p>Possible Extrusion</p> <p>Small Wedge Angle</p> <p>Large Radial and Inferior Bulge</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>


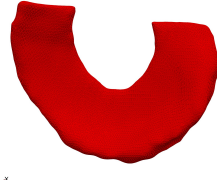



Continued on next page

Table A.1 – Continued from previous page

Ref.	Superior	Feature Description
LM_14*		Smooth Surface Small-to-Large Wedge Angle (Anterior-Posterior) Large Radial and Inferior Bulge Source: SimTK/Open Knee (<i>simtk.org</i>) [40]
LM_15*		Smooth Geometry Large Extrusion Small-to-Large Inferior Bulge (Anterior-Posterior) Large Radial Bulge (Medial) Source: SimTK/Open Knee (<i>simtk.org</i>) [40]
LM_16*		Smooth Geometry Possible Extrusion (Posterior) Small Wedge Angle Large Radial Bulge (Anterior and Posterior) Large Inferior Bulge (Anterior and Posterior) Source: SimTK/Open Knee (<i>simtk.org</i>) [40]
LM_17		Duplicate (Original of LM_04) Source: SimTK/Open Knee (<i>simtk.org</i>) [40]
LM_18		Duplicate (Original of LM_03) Source: SimTK/Open Knee (<i>simtk.org</i>) [40]
LM_19		Smooth Geometry Possible Extrusion (Anterior) Possible Degeneration or Tears (Posterior) Small-to-Large Wedge Angle (Anterior-Posterior) Large Inferior Bulge Source: SimTK/Open Knee (<i>simtk.org</i>) [40]





Continued on next page

Table A.1 – Continued from previous page

Ref.	Superior	Feature Description
LM_20*		<p>Irregular Geometry</p> <p>Irregular Surface</p> <p>Possible Degeneration or Tears (Wedge Region)</p> <p>Small-to-Large Wedge Angle (Anterior-Posterior)</p> <p>Large Inferior Bulge</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
LM_21		<p>Irregular Geometry</p> <p>Irregular Surface</p> <p>Possible Extrusion</p> <p>Large Wedge Angle</p> <p>Large Inferior Bulge</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
LM_22		<p>Irregular Geometry</p> <p>Possible Degeneration or Tears (Wedge Region)</p> <p>Large Wedge Angle</p> <p>Small-to-Large Inferior Bulge (Posterior-Anterior)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
LM_23		<p>Irregular Geometry</p> <p>Possible Degeneration or Tears (Medial-Anterior)</p> <p>Large Wedge Angle</p> <p>Large Inferior Bulge</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
LM_24*		<p>Irregular Geometry</p> <p>Possible Extrusion or Degeneration (Posterior)</p> <p>Small Wedge Angle</p> <p>Large Inferior Bulge</p> <p>Small-to-Large Radial Bulge (Anterior-Posterior)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>



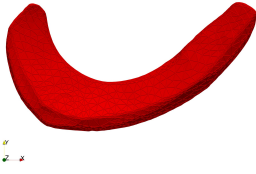

Continued on next page

Table A.1 – *Continued from previous page*

Ref.	Superior	Feature Description
MM_01		<p>Smooth Surface</p> <p>Smooth Geometry</p> <p>Small Horn-Attachments</p> <p>Small Wedge Angle (Medial-Region)</p> <p>Large Inferior Bulge</p> <p>Large Radial Bulge</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [19]</p>
MM_02		<p>Smooth Surface</p> <p>Extrusion*</p> <p>Large Wedge Angle</p> <p>Large Inferior Bulge (Medial-Anterior Region)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [80]</p>
MM_03*		<p>Irregular Surface</p> <p>Possible Extrusion, Degeneration or Tears</p> <p>Tapered Shape</p> <p>Small-to-Large Wedge Angle (Posterior-Anterior)</p> <p>Large Inferior Bulge</p> <p>Large Radial Bulge</p> <p>Source: SimTK/Open Knee (<i>uconn.edu</i>) [40]</p>
MM_04*		<p>Irregular Surface</p> <p>Irregular Shape</p> <p>Tapered Shape</p> <p>Extrusion</p> <p>Small-to-Large Wedge Angle (Posterior-Anterior)</p> <p>Large Inferior Bulge</p> <p>Large Radial Bulge (Anterior Region)</p> <p>Source: SimTK/Open Knee (<i>uconn.edu</i>) [40]</p>

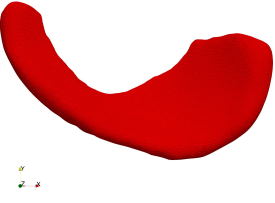
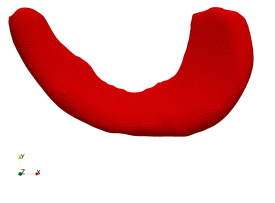
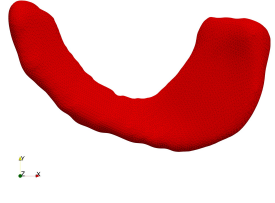
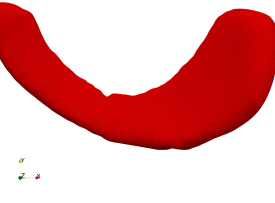
Continued on next page

Table A.1 – *Continued from previous page*

Ref.	Superior	Feature Description
MM_05		<p>Smooth Surface</p> <p>Blocky</p> <p>Tapered Shape</p> <p>Small Radial Bulge</p> <p>Large Wedge angle</p> <p>Large Inferior Bulge (Anterior Region)</p> <p>Source: iMBE (<i>Grant Number:</i> EP/P001076/1)</p>
MM_06*		<p>Smooth Surface</p> <p>Blocky</p> <p>Non-Semi-Lunar</p> <p>Small Horn-Attachments</p> <p>Small Radial Bulge</p> <p>Small-to-Large Wedge Angle (Anterior-Posterior)</p> <p>Large Inferior Bulge (Anterior Region)</p> <p>Source: iMBE (<i>Grant Number:</i> EP/P001076/1)</p>
MM_07*		<p>Smooth Surface</p> <p>Blocky</p> <p>Non-Semi-Lunar</p> <p>Small-to-Large Wedge Angle (Anterior-Posterior)</p> <p>Large Radial Bulge</p> <p>Source: iMBE (<i>Grant Number:</i> EP/P001076/1)</p>
MM_08*		<p>Smooth Surface</p> <p>Irregular Shape</p> <p>Possible Degeneration or Tear</p> <p>Small Horn-Attachments</p> <p>Large Wedge Angle</p> <p>Large Inferior Bulge</p> <p>Large Radial Bulge (Anterior and Posterior Horns)</p> <p>Source: iMBE (<i>Grant Number:</i> EP/P001076/1)</p>

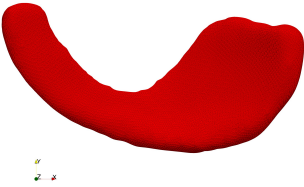
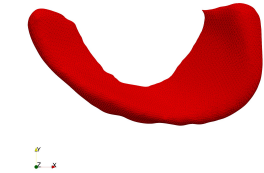
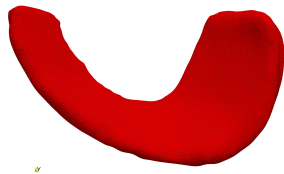
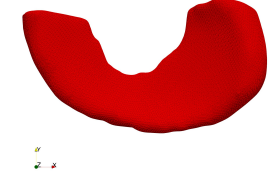
Continued on next page

Table A.1 – *Continued from previous page*

Ref.	Superior	Feature Description
MM_09		<p>Smooth Surface</p> <p>Tapered Shape</p> <p>Large Extrusion (Posterior)</p> <p>Small-to-Large Wedge Angle (Posterior-Anterior)</p> <p>Small Radial and Inferior Bulge</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
MM_10		<p>Smooth Surface</p> <p>Possible Extrusion (Posterior)</p> <p>Tapered Shape</p> <p>Possible Degeneration or Tears (Wedge Region)</p> <p>Large Wedge Angle</p> <p>Large Inferior Bulge</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
MM_11*		<p>Smooth Surface</p> <p>Possible Extrusion (Posterior)</p> <p>Tapered Shape</p> <p>Possible Degeneration or Tears (Medial-Anterior)</p> <p>Large Wedge Angle</p> <p>Small-to-Large Inferior Bulge (Posterior-Anterior)</p> <p>Large Radial Bulge (Medial)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
MM_12		<p>Smooth Surface</p> <p>Possible Extrusion (Posterior)</p> <p>Tapered Shape</p> <p>Possible Degeneration or Tears (Medial)</p> <p>Large Wedge Angle</p> <p>Large Inferior Bulge</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>

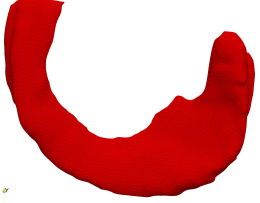
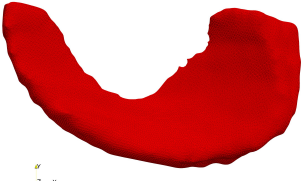


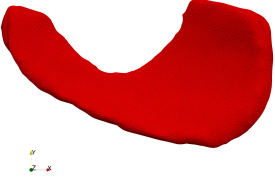
Continued on next page

Table A.1 – *Continued from previous page*

Ref.	Superior	Feature Description
MM_13		<p>Smooth Surface</p> <p>Large Extrusion (Posterior)</p> <p>Tapered Shape</p> <p>Large Wedge Angle</p> <p>Small-to-Very-Large Inferior Angle (Posterior-Anterior)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
MM_14		<p>Smooth Surface</p> <p>Possible Extrusion (Posterior)</p> <p>Tapered Shape</p> <p>Possible Degeneration or Tears (Wedge Region)</p> <p>Small-to-Large Wedge Angle (Posterior-Anterior)</p> <p>Small-to-Large Inferior Bulge (Anterior-Posterior)</p> <p>Large Radial Bulge (Medial)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
MM_15*		<p>Smooth Surface</p> <p>Extrusion (Posterior)</p> <p>Possible Degeneration or Tears (Medial-Anterior)</p> <p>Tapered Shape</p> <p>Small-to-Large Wedge Angle (Anterior-Posterior)</p> <p>Very-Large Inferior Bulge</p> <p>Large Radial Bulge (Posterior Horn)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
MM_16		<p>Smooth Surface</p> <p>Extrusion (Posterior)</p> <p>Possible Degeneration or Tears (Wedge Region)</p> <p>Small-to-Large Wedge Angle (Anterior-Posterior)</p> <p>Large Inferior Bulge</p> <p>Large Radial Bulge (Anterior and Posterior Horns)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>



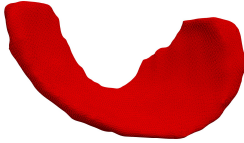
Continued on next page

Table A.1 – Continued from previous page

Ref.	Superior	Feature Description
MM_17		Duplicate (Original of MM_04) Source: SimTK/Open Knee (<i>simtk.org</i>) [40]
MM_18		Duplicate (Original of MM_03) Source: SimTK/Open Knee (<i>simtk.org</i>) [40]
MM_19*		Irregular Surface Possible Degeneration or Tears (Wedge Region) Large Wedge Angle Large Inferior Bulge Large Radial Bulge (Posterior) Source: SimTK/Open Knee (<i>simtk.org</i>) [40]
MM_20		Irregular Surface Possible Extrusion (Posterior) Possible Degeneration or Tears (Wedge Region) Small-to-Large Inferior Bulge (Posterior-Anterior) Large Radial Bulge (Anterior Horn and Posterior) Source: SimTK/Open Knee (<i>simtk.org</i>) [40]
MM_21*		Irregular Surface Extrusion (Posterior) Tapered Shape Possible Degeneration or Tears (Wedge Region) Small-to-Large Inferior Bulge (Posterior-Anterior) Large Radial Bulge (Anterior and Posterior Horns) Source: SimTK/Open Knee (<i>simtk.org</i>) [40]

Continued on next page

Table A.1 – Continued from previous page

Ref.	Superior	Feature Description
MM_22		<p>Irregular Surface</p> <p>Extrusion (Posterior)</p> <p>Possible Degeneration or Tears (Wedge Region)</p> <p>Tapered Shape</p> <p>Large Wedge Angle</p> <p>Small-to-Very-Large Inferior Bulge (Posterior-Anterior)</p> <p>Large Radial Bulge (Posterior Horn)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
MM_23		<p>Irregular Surface</p> <p>Extrusion (Posterior)</p> <p>Possible Degeneration (Medial Wedge Region)</p> <p>Possible Tears (Medial Wedge Region)</p> <p>Tapered Shape</p> <p>Large Wedge Angle</p> <p>Large Inferior Bulge (Anterior Horn)</p> <p>Large Radial Bulge</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>
MM_24*		<p>Irregular Surface</p> <p>Extrusion (Posterior)</p> <p>Possible Degeneration or Tears (Wedge Region)</p> <p>Small Wedge Angle</p> <p>Tapered Shape</p> <p>Small-to-Large Inferior Bulge (Posterior-Anterior)</p> <p>Large Radial Bulge (Anterior and Posterior Horns)</p> <p>Source: SimTK/Open Knee (<i>simtk.org</i>) [40]</p>

B Case D: Summary of Results

Lateral Meniscus Meshes

Table B.1: Case D: Summary of results for the performance analyses from the lateral meniscus meshes

Geometry	Method	Time (s)	Surface Error		Element Quality		Volume (mm ³)	Contact Area	
			RMSE (mm)	ϵ_{\max} (mm)	$\tilde{J} < 0.1$ (%)	$\tilde{J} < 0.5$ (%)		Tibial (mm ²)	Femoral (mm ²)
LM_01	Morphed Mesh	24.97	0.850	2.388	0.00	6.00	1114.81	175.47	196.68
	Multi-Block	554	0.149	1.404	1.58	12.33	999.34	212.26	236.27
LM_04	Morphed Mesh	23.93	0.857	2.332	0.00	6.33	1306.81	148.54	164.29
	Multi-Block	706	0.108	0.524	0.75	10.67	1182.73	146.41	169.02
LM_05	Morphed Mesh	23.87	0.846	2.294	0.50	9.00	775.30	93.17	130.97
	Multi-Block	493	0.284	1.530	0.00	4.75	678.38	106.88	102.34
LM_06	Morphed Mesh	24.09	1.013	3.177	0.00	7.17	693.93	99.24	108.20
	Multi-Block	522	0.133	0.640	0.50	5.83	586.63	100.65	118.85
LM_08	Morphed Mesh	24.23	0.994	3.330	0.00	6.58	1789.43	156.69	176.50
	Multi-Block	1948	0.299	1.608	1.50	11.08	1743.43	103.61	131.47
LM_14	Morphed Mesh	25.54	0.881	2.504	0.00	4.83	1414.55	162.39	181.46
	Multi-Block	434	0.159	1.592	0.75	9.83	1357.65	85.76	144.21
LM_15	Morphed Mesh	24.55	0.840	2.154	0.00	4.83	1515.73	155.96	164.03
	Multi-Block	391	0.182	1.451	0.58	9.33	1504.77	144.06	157.77
LM_16	Morphed Mesh	25.49	0.898	2.397	0.00	5.08	1397.79	138.43	117.30
	Multi-Block	301	0.164	1.465	1.25	7.92	1382.44	75.37	119.25
LM_20	Morphed Mesh	25.97	0.890	2.586	0.00	5.17	1394.03	147.37	142.53
	Multi-Block	1153	0.162	1.225	1.25	11.50	1346.19	78.74	141.52
LM_24	Morphed Mesh	25.34	0.950	3.307	0.00	4.75	1507.74	146.37	160.04
	Multi-Block	925	0.126	1.454	2.08	12.92	1506.35	64.15	116.47

Medial Meniscus Meshes

Table B.2: Case D: Summary of results for the performance analyses from the medial meniscus meshes

Geometry	Method	Time (s)	Surface Error		Element Quality		Volume (mm ³)	Contact Area	
			RMSE (mm)	ϵ_{\max} (mm)	$\tilde{J} < 0.1$ (%)	$\tilde{J} < 0.5$ (%)		Tibial (mm ²)	Femoral (mm ²)
MM_03	Morphed Mesh	24.86	1.111	3.885	0.00	4.92	2066.99	122.32	146.42
	Multi-Block	6320	0.181	1.140	1.67	12.17	2093.76	106.47	125.36
MM_04	Morphed Mesh	24.64	0.942	2.873	0.00	6.75	1749.80	122.56	158.11
	Multi-Block	8040	0.290	2.231	1.33	12.92	1664.86	130.52	116.93
MM_06	Morphed Mesh	24.24	0.892	2.897	0.25	7.92	836.91	84.65	116.04
	Multi-Block	2015	0.191	0.757	1.33	5.17	755.73	127.26	89.75
MM_07	Morphed Mesh	24.89	0.833	2.345	0.00	6.67	1244.02	121.97	131.33
	Multi-Block	876	0.356	1.639	0.17	4.83	1125.51	134.10	81.14
MM_08	Morphed Mesh	24.48	0.876	2.786	0.00	5.42	1732.24	129.19	136.48
	Multi-Block	1227	0.277	1.353	1.67	9.00	1655.29	141.26	138.35
MM_11	Morphed Mesh	25.05	0.854	2.482	0.00	5.17	1570.39	146.00	172.06
	Multi-Block	279	0.099	0.545	0.33	9.92	1503.46	141.30	127.59
MM_15	Morphed Mesh	25.28	0.917	3.860	0.00	4.92	1957.54	115.11	129.90
	Multi-Block	1378	0.172	1.826	0.58	7.50	1996.12	73.87	116.46
MM_19	Morphed Mesh	24.83	0.872	3.292	0.00	5.92	1832.52	96.39	118.14
	Multi-Block	766	0.193	1.681	1.08	8.75	1848.78	65.33	150.05
MM_21	Morphed Mesh	25.35	1.043	3.678	0.00	4.83	1763.39	134.53	176.72
	Multi-Block	2901	0.182	1.289	1.75	9.83	1789.02	157.65	158.66
MM_24	Morphed Mesh	24.72	0.964	3.173	0.00	4.92	1441.10	145.92	175.70
	Multi-Block	494	0.150	1.002	2.00	10.75	1417.73	113.26	134.23

C Case E: Summary of Results

Lateral Meniscus Meshes

Table C.1: Case E: Summary of results for the performance analyses from the lateral meniscus meshes

Geometry	Method	Time (s)	Surface Error		Element Quality		Volume (mm ³)	Contact Area	
			RMSE (mm)	ϵ_{\max} (mm)	$\tilde{J} < 0.1$ (%)	$\tilde{J} < 0.5$ (%)		Tibial (mm ²)	Femoral (mm ²)
LM_01	Unopt. Morph	24.97	0.850	2.388	0.00	6.00	1114.81	175.47	196.68
	Opt. Morph	49.57	0.172	1.101	1.67	13.58	990.87	200.96	224.95
	Multi-Block	554	0.149	1.404	1.58	12.33	999.34	212.26	236.27
LM_04	Unopt. Morph	23.93	0.857	2.332	0.00	6.33	1306.81	148.54	164.29
	Opt. Morph	50.27	0.268	1.550	0.83	13.33	1114.97	130.14	166.73
	Multi-Block	706	0.108	0.524	0.75	10.67	1182.73	146.41	169.02
LM_05	Unopt. Morph	23.87	0.846	2.294	0.50	9.00	775.30	93.17	130.97
	Opt. Morph	48.17	0.465	1.738	1.17	20.67	633.46	95.73	97.63
	Multi-Block	493	0.284	1.530	0.00	4.75	678.38	106.88	102.34
LM_06	Unopt. Morph	24.09	1.013	3.177	0.00	7.17	693.93	99.24	108.20
	Opt. Morph	49.86	0.572	2.611	1.17	32.17	510.50	80.58	102.50
	Multi-Block	522	0.133	0.640	0.50	5.83	586.63	100.65	118.85
LM_08	Unopt. Morph	24.23	0.994	3.330	0.00	6.58	1789.43	156.69	176.50
	Opt. Morph	54.90	0.508	1.791	0.25	12.58	1618.99	126.60	178.26
	Multi-Block	1948	0.299	1.608	1.50	11.08	1743.43	103.61	131.47

Continued on next page

Table C.1 – *Continued from previous page*

Geometry	Method	Time (s)	Surface Error		Element Quality		Volume (mm ³)	Contact Area	
			RMSE (mm)	ϵ_{\max} (mm)	$\tilde{J} < 0.1$ (%)	$\tilde{J} < 0.5$ (%)		Tibial (mm ²)	Femoral (mm ²)
LM_14	Unopt. Morph	25.54	0.881	2.504	0.00	4.83	1414.55	162.39	181.46
	Opt. Morph	58.02	0.168	1.118	0.75	8.17	1339.82	115.91	121.21
	Multi-Block	434	0.159	1.592	0.75	9.83	1357.65	85.76	144.21
LM_15	Unopt. Morph	24.55	0.840	2.154	0.00	4.83	1515.73	155.96	164.03
	Opt. Morph	54.72	0.238	1.551	1.50	11.17	1488.16	147.07	138.55
	Multi-Block	391	0.182	1.451	0.58	9.33	1504.77	144.06	157.77
LM_16	Unopt. Morph	25.49	0.898	2.397	0.00	5.08	1397.79	138.43	117.30
	Opt. Morph	59.49	0.214	1.377	0.92	8.50	1356.56	118.05	139.64
	Multi-Block	301	0.164	1.465	1.25	7.92	1382.44	75.37	119.25
LM_20	Unopt. Morph	25.97	0.890	2.586	0.00	5.17	1394.03	147.37	142.53
	Opt. Morph	59.54	0.227	1.409	1.58	11.25	1323.50	100.11	133.05
	Multi-Block	1153	0.162	1.225	1.25	11.50	1346.19	78.74	141.52
LM_24	Unopt. Morph	25.34	0.950	3.307	0.00	4.75	1507.74	146.37	160.04
	Opt. Morph	59.64	0.204	1.498	1.00	10.17	1465.09	85.46	118.62
	Multi-Block	925	0.126	1.454	2.08	12.92	1506.35	64.15	116.47

Medial Meniscus Meshes

Table C.2: Case E: Summary of results for the performance analyses from the medial meniscus meshes

Geometry	Method	Time (s)	Surface Error		Element Quality		Volume (mm ³)	Contact Area	
			RMSE (mm)	ϵ_{\max} (mm)	$\tilde{J} < 0.1$ (%)	$\tilde{J} < 0.5$ (%)		Tibial (mm ²)	Femoral (mm ²)
MM_03	Unopt. Morph	24.86	1.111	3.885	0.00	4.92	2066.99	122.32	146.42
	Opt. Morph	56.66	0.425	2.252	0.67	9.17	1943.57	118.68	121.39
	Multi-Block	6320	0.181	1.140	1.67	12.17	2093.76	106.47	125.36
MM_04	Unopt. Morph	24.64	0.942	2.873	0.00	6.75	1749.80	122.56	158.11
	Opt. Morph	52.47	0.330	1.768	1.58	12.67	1611.49	117.02	128.93
	Multi-Block	8040	0.290	2.231	1.33	12.92	1664.86	130.52	116.93
MM_06	Unopt. Morph	24.24	0.892	2.897	0.25	7.92	836.91	84.65	116.04
	Opt. Morph	49.31	0.572	2.655	0.33	18.92	687.05	99.95	82.54
	Multi-Block	2015	0.191	0.757	1.33	5.17	755.73	127.26	89.75
MM_07	Unopt. Morph	24.89	0.833	2.345	0.00	6.67	1244.02	121.97	131.33
	Opt. Morph	49.72	0.308	1.184	2.17	10.58	1124.46	131.99	88.09
	Multi-Block	876	0.356	1.639	0.17	4.83	1125.51	134.10	81.14
MM_08	Unopt. Morph	24.48	0.876	2.786	0.00	5.42	1732.24	129.19	136.48
	Opt. Morph	52.36	0.509	1.985	0.08	8.08	1559.69	138.31	161.18
	Multi-Block	1227	0.277	1.353	1.67	9.00	1655.29	141.26	138.35

Continued on next page

Table C.2 – *Continued from previous page*

Geometry	Method	Time (s)	Surface Error		Element Quality		Volume (mm ³)	Contact Area	
			RMSE (mm)	ϵ_{\max} (mm)	$\tilde{J} < 0.1$ (%)	$\tilde{J} < 0.5$ (%)		Tibial (mm ²)	Femoral (mm ²)
MM_11	Unopt. Morph	25.05	0.854	2.482	0.00	5.17	1570.39	146.00	172.06
	Opt. Morph	56.47	0.200	1.373	0.50	5.92	1478.32	153.56	170.64
	Multi-Block	279	0.099	0.545	0.33	9.92	1503.46	141.30	127.59
MM_15	Unopt. Morph	25.28	0.917	3.860	0.00	4.92	1957.54	115.11	129.90
	Opt. Morph	59.63	0.503	2.639	0.08	7.33	1835.74	73.86	112.56
	Multi-Block	1378	0.172	1.826	0.58	7.50	1996.12	73.87	116.46
MM_19	Unopt. Morph	24.83	0.872	3.292	0.00	5.92	1832.52	96.39	118.14
	Opt. Morph	59.00	0.616	2.957	0.08	6.67	1609.01	85.60	129.62
	Multi-Block	766	0.193	1.681	1.08	8.75	1848.78	65.33	150.05
MM_21	Unopt. Morph	25.35	1.043	3.678	0.00	4.83	1763.39	134.53	176.72
	Opt. Morph	59.50	1.239	4.063	0.08	2.75	1432.08	146.03	153.92
	Multi-Block	2901	0.182	1.289	1.75	9.83	1789.02	157.65	158.66
MM_24	Unopt. Morph	24.72	0.964	3.173	0.00	4.92	1441.10	145.92	175.70
	Opt. Morph	57.40	0.460	2.174	0.58	10.25	1270.75	160.28	178.25
	Multi-Block	494	0.150	1.002	2.00	10.75	1417.73	113.26	134.23