Imperial College of Science, Technology and Medicine
Department of Aeronautics

Multiscale structural optimisation with concurrent coupling between scales

Ryan Murphy

Supervisors: Dr. Robert Hewson & Dr. Matthew Santer

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in the Department of Aeronautics of Imperial College London,
October 2021

Abstract

A robust three-dimensional multiscale topology optimisation framework with concurrent coupling between scales is presented. Concurrent coupling ensures that only the microscale data required to evaluate the macroscale model during each iteration of optimisation is collected and results in considerable computational savings. This represents the principal novelty of the framework and permits a previously intractable number of design variables to be used in the parametrisation of the microscale geometry, which in turn enables accessibility to a greater range of mechanical point properties during optimisation. Additionally, the microscale data collected during optimisation is stored in a re-usable database, further reducing the computational expense of subsequent iterations or entirely new optimisation problems. Application of this methodology enables structures with precise functionally-graded mechanical properties over two-scales to be derived, which satisfy one or multiple functional objectives. For all applications of the framework presented within this thesis, only a small fraction of the microstructure database is required to derive the optimised multiscale solutions, which demonstrates a significant reduction in the computational expense of optimisation in comparison to contemporary sequential frameworks.

The derivation and integration of novel additive manufacturing constraints for open-walled microstructures within the concurrently coupled multiscale topology optimisation framework is also presented. Problematic fabrication features are discouraged through the application of an augmented projection filter and two relaxed binary integral constraints, which prohibit the formation of unsupported members, isolated assemblies of overhanging members and slender members during optimisation. Through the application of these constraints, it is possible to derive self-supporting, hierarchical structures with varying topology, suitable for fabrication through additive manufacturing processes.

Acknowledgements

Foremost, I would like to express my sincere gratitude to my supervisors Dr. Robert Hewson and Dr. Matthew Santer for their unwavering support, wealth of technical expertise and perpetual enthusiasm throughout the duration of this research project. Their passion for science is infectious and it has been a genuine pleasure to work with them both.

I would also like to acknowledge Dr. Martin Muir for his continuous support and for sharing his extensive expertise within the field of additive manufacturing, which has helped to the guide the direction of this research project.

I would also like to express my gratitude to Chikwesiri, Dilaksan and Morgan (The Lattice Kids) for their friendship and support. It has been a pleasure to share this experience with you. I will fondly remember our heated debates and our many table football tournaments.

I am deeply indebted to my family (specifically Mum, Dad & Daniela) for their continuous encouragement, support and the many sacrifices that they have made for me throughout my life. Finally, I would like to extend my sincerest gratitude to my partner Em for all her love and support.

Statement of Originality

I hereby declare that except where specific reference is made to the work of others, the contents of this thesis are original and have not been submitted in whole or in part for consideration for any other degree or qualification. This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration, except where specifically indicated in the text. It has also been seen by my supervisors before submission.

Copyright Statement

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that:

- You credit the author and do not use it, or any derivative works, for a commercial purpose.
- When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Nomenclature

Acronyms

AM	Additive Manufacturing
BESO	Bi-directional Evolutionary Structural Optimisation
ESO	Evolutionary Structural Optimisation
KKT	Karush-Kuhn-Tucker
LSF	Level-set Function
LSM	Level-set Method
MTO	Multiscale Topology Optimisation
SIMP	Solid Isotropic Material with Penalisation
TO	Topology Optimisation

Greek Symbols

α	Line search step size
β	Line search scaling factor
γ	Level-set (boundary)
Γ	Surface
δ	Ratio of microscale to macroscale characteristic lengths
δ_{ij}	Kronecker delta function
ϵ_{ij}	Infinitesimal strain tensor
$\bar{\epsilon}_{ij}$	Macroscale infinitesimal strain tensor
ϵ^*_{ij}	Fluctuating infinitesimal strain tensor
ϵ_j	Infinitesimal strain vector (Voigt notation)
E	Young's modulus
E^0	Young's modulus (void phase)
E^1	Young's modulus (solid phase)
E_P	Young's modulus (primary phase)
E_S	Young's modulus (secondary phase)
$\hat{E}$	Normalised Young's modulus

$E(\rho)$	Stiffness interpolant
$\hat{E}(\rho)$	Normalised stiffness interpolant
E_{ijkl}	Fourth-order elasticity tensor
E^H_{ijkl}	Homogenised fourth-order elasticity tensor
E_{ij}	Fourth-order elasticity tensor (Voigt notation)
E^H_{ij}	Homogenised fourth-order elasticity tensor (Voigt notation)
$E^H_{ij,\text{parent}}$	Homogenised fourth-order elasticity tensor (Voigt notation) (parent)
$E^H_{ij,\text{child}}$	Homogenised fourth-order elasticity tensor (Voigt notation) (child)
ζ	Logical expression (unsupported and overhanging members)
Z	Binary integral constraint (unsupported and overhanging members)
η	Heaviside thresholding parameter
θ	Aggregate sensitivity
$\hat{\theta}$	Normalised aggregate sensitivity
κ	Nodal mapping matrix
λ	L-BFGS memory
λ_L	Lamé's first parameter
Λ	Node set of parent microstructures
μ	Barrier parameter
μ_L	Lamé's second parameter
ν	Poisson's ratio
ξ	DOE level
Ξ	Node set (slender members)
ρ	Pseudo-density
σ	L-BFGS scaling parameter
σ_{ij}	Cauchy stress tensor
σ_i	Cauchy stress vector (Voigt notation)
τ	Level-set function
υ	Logical expression (slender members)
U	Binary integral constraint (slender members)

ϕ	Node set for a single enclosing simplex
$\phi_{\alpha\beta}$	Projected variable
$\phi^*_{\alpha\beta}$	Optimal projected variable
$\phi^*_{\alpha\beta,PP}$	Optimal post-processed projected variable
Φ	Node set for all enclosing simplices
X	Design variable matrix
X'	Transformed design variable matrix
$\tilde{X}$	Filtered design variable matrix
$\hat{\tilde{X}}$	Projected filtered design variable matrix
$\hat{\tilde{X}}^*$	Optimal projected filtered design variable matrix
$\hat{\tilde{X}}^*_{PP}$	Optimal post-processed projected filtered design variable matrix
ψ	Mapping functional
Ψ	Node set (unsupported and overhanging members)
ω	Virtual displacement vector
Ω	Design domain

Roman Symbols

A	Post-processing scaling factor
AR_{max}	Maximum aspect-ratio
$c(x)$	Constraint vector
d	KKT Search direction
D	L-BFGS matrix
e	Structural member aspect-ratio
f	Body force vector
$f(x)$	Objective function
F	External force
F_1	Line search merit function
F_2	Line search merit function
g	L-BFGS Lagrangian gradient vector
G	L-BFGS Lagrangian gradient matrix

J	Objective functional
l	Characteristic unit-cell length
l_s	Minimum slender member length
l_b	Minimum bracing assembly length
L	L-BFGS matrix
m	Number of constraints
M	Block transformation matrix
n	Number of design variables
n_{cpu}	Number of computer processing units
n_e	Number of divisions per unit length
n_P	Fraction of nodes assigned primary phase
n_s	Number of microscale simulations
n_S	Fraction of nodes assigned secondary phase
N	Number of nodes
p	Lagrange multiplier vector (lower bound)
P	Projected variable exponent
q	Lagrange multiplier vector (upper bound)
r	Radius
$\hat{r}$	Normalised radius
r_{min}	Minimum manufacturable radius
r_{max}	Maximum permitted radius
$\hat{r}_{min}$	Minimum manufacturable normalised radius
R	Helmholtz filter radius
$R_{x,y,z}$	Rotation matrix
$R_{xy,xz,yz}$	Reflection matrix
s	Slack variable vector (constraints)
S	Piecewise summation of Euclidean norm
t	Surface traction vector
t_p	Pseudo-time

t_1	Two-dimensional thickness parameter
t_2	Two-dimensional thickness parameter
t_s	Time per microscale simulation
T	Transformation matrix
T_s	Total simulation time
u	Displacement vector
v	Slack variable vector (upper bound)
V	Volume-fraction
V_f	Target volume-fraction
w	Slack variable vector (lower bound)
w_b	Bracing radius
x	Design variables
x_L	Lower bound
x_U	Upper bound
y	L-BFGS design variable vector
Y	L-BFGS design variable matrix
z	Lagrange multiplier vector (constraints)

Miscellaneous

$\mathcal{L}$	Lagrangian functional

Contents

Abstract **i**

Acknowledgements **iii**

Nomenclature **v**

1 Introduction **1**

1.1 Thesis Outline 3

1.2 Thesis Objectives 4

1.3 Publications 4

2 Review of Structural Optimisation Literature **6**

2.1 Topology Optimisation 7

2.1.1 Topology Optimisation Frameworks 8

2.1.2 Manufacturing 14

2.2 Metamaterials 16

2.3 Homogenisation 17

2.4 Multiscale Topology Optimisation 19

2.4.1 Concurrent Multiscale Optimisation 20

2.4.2 Univariate Parametrisation Frameworks 21

2.4.3 Multivariate Parametrisation Frameworks 22

2.4.4 Manufacturing 24

2.5 Summary 25

3 Multiscale Topology Optimisation Framework 27

3.1 Linear Elastic Assumptions 27

3.2 Microscale Model 28

3.2.1 Nodal Property Allocation 30

3.2.2 Derivation of Homogenised Mechanical Properties 34

3.3 Macroscale Model 37

3.4 Concurrent Coupling 38

3.4.1 Material Property Transformations 40

3.4.2 Piecewise Continuous Linear Interpolation 45

3.4.3 Microscale Model — Mesh Convergence 48

3.4.4 Mechanical Property Space Visualisation 53

3.4.5 Concurrent Coupling Example (2D) 59

3.5 Optimisation Formulation 61

3.5.1 Helmholtz Filter 61

3.5.2 Optimisation Formulation 61

3.5.3 Sensitivities 62

3.5.4 Interior-point Optimisation Algorithm 65

3.5.5 Addition of Microscale Parameters 75

3.6 Macroscale Geometry Reconstruction 79

3.7 Summary 79

4 Multiscale Topology Optimisation Results **81**

4.1 Base Multiscale Topology Optimisation Framework 81

4.1.1 Cantilever Beam — Compliance Minimisation 82

4.1.2 Unit Cube — Compliance Minimisation 90

4.1.3 Computational Expense Incurred by Microscale Simulations 98

4.1.4 Comparison to IPOPT 98

4.2 Addition of Microscale Parameters 99

4.2.1 Cantilever Beam — Compliance Minimisation 100

4.2.2 Unit Cube — Compliance Minimisation 104

5 Design for Additive Manufacturing **107**

5.1 Manufacturability Requirements 108

5.2 Minimum Length Scale Control 112

5.3 Unsupported and Overhanging Members 113

5.4 Slender Members 120

5.5 Post-processing Algorithm 122

5.6 Updated Optimisation Formulation 124

5.6.1 Sensitivities 125

5.7 Summary 126

6 Design for Additive Manufacturing Results **127**

6.1 Additive Manufacturing Constraints 127

6.1.1 Unit Cube — Compliance Minimisation 129

6.1.2 Cantliever Beam — Compliance Minimisation 134

7 Conclusion **141**

7.1 Summary of Thesis Achievements 144

7.2 Future Work 145

A Derivation of Enclosing Simplex **147**

Bibliography **150**

List of Tables

3.1 Microscale parameter bounds and levels 39

3.2 Resultant RMSE from two-dimensional hyperplanes 58

4.1 Parameter initialisation (Sec. 4.1) 82

4.2 Cantilever beam case compliance comparison to SIMP 89

4.3 Convergence comparison of IPOwDA and IPOPT — Cantilever beam case . . 98

4.4 Convergence comparison of IPOwDA and IPOPT — Unit-cube case 99

4.5 Parameter initialisation (Sec. 4.2) 99

4.6 Normalised aggregate sensitivities for the cantilever beam case 101

4.7 Comparison of compliance functionals — cantilever beam case 101

4.8 Normalised aggregate sensitivities for the unit-cube case 104

4.9 Comparison of compliance functionals — unit-cube case 105

5.1 Microscale parameter bounds and levels used for AM constraints 108

6.1 Parameter initialisation (Sec. 6.1) 128

6.2 Definition of AM constraint parameters 128

6.3 Optimised compliance functionals for the unit-cube case 134

List of Figures

2.1 Illustrations depicting the three sub-classes of structural optimisation 7

2.2 SIMP penalisation parameters 11

2.3 Example level-set function 13

2.4 Overhang constraints in topology optimisation 16

2.5 Homogenisation illustration 19

2.6 Sequential and concurrent coupling 26

3.1 Microscale parametrisation 28

3.2 Aspect-ratio definition 29

3.3 Nodal allocation example (Case 1) 33

3.4 Nodal allocation example (Case 2) 34

3.5 Macroscale finite-element model 38

3.6 Example two-dimensional DOEs 40

3.7 Example material property transformations 43

3.8 Microscale model — mesh convergence (Case 1) 49

3.9 Microscale model — mesh convergence (Case 2) 50

3.10 Microscale model — mesh convergence (Case 3) 51

3.11 Microscale model — convergence time 52

3.12 Simplicial Interpolant (Case 1) 55

3.13 Simplicial Interpolant (Case 2) 57

3.14 Simplicial Interpolant (Case 3) 58

3.15 Two-dimensional example of concurrent coupling 60

3.16 Nested structure of full-factorial DOE 76

4.1 Cantilever beam case (Sec. 4.1.1) 82

4.2 Cantilever beam case — Optimised normalised radii distributions 83

4.3 Cantilever beam case — Aspect-ratio distributions 84

4.4 Convergence plot for the cantilever beam case 84

4.5 Cantilever beam case — Maximum and minimum principal stresses 86

4.6 Reconstructed cantilever beam geometry 88

4.7 Computational expense incurred by the cantilever beam case 89

4.8 Unit-cube case (Sec. 4.1.2) 91

4.9 Unit-cube case — Optimised normalised radii distributions 92

4.10 Unit-cube case — Aspect-ratio distributions 93

4.11 Convergence plot for the unit cube beam 93

4.12 Unit-cube case — Maximum and minimum principal stresses 94

4.13 Reconstructed unit-cube geometry 95

4.14 Computational expense incurred by the unit-cube case 96

4.15 Unit-cube case compliance comparison to SIMP 97

4.16 Convergence plots for the cantilever beam case (dimension addition activated) 100

4.17 Convergence plots for the cantilever beam case (dimension addition disabled) . 102

4.18 Computational expense comparison — cantilever beam case 103

4.19 Cumulative average arc length comparison — cantilever beam case 103

4.20 Convergence plots for the unit-cube case (dimension addition activated) 104

4.21 Convergence plots for the unit-cube case (dimension addition disabled) 105

4.22 Computational expense comparison — unit-cube case 106

4.23 Cumulative average arc length comparison — unit-cube case 106

5.1 Microscale parametrisation for AM constraints 108

5.2 Minimum manufacturable radius definition 109

5.3 Unsupported and overhanging members 110

5.4 Slender members 111

5.5 Augmented projection filter 113

5.6 Supporting members 115

5.7 Intra-cell support 116

5.8 Half-member overhang 117

5.9 Projection filter 117

5.10 Slender member bracing 121

6.1 Unit-cube case (Sec. 6.1.1) 129

6.2 Reconstruction of the unconstrained unit-cube geometry 131

6.3 Reconstruction of the constrained unit-cube geometry 132

6.4 Histogram comparing the unconstrained and constrained normalised radii distributions 133

6.5 Cantilever beam case (Sec. 6.1.2) 135

6.6 Reconstruction of the unconstrained cantilever beam geometry 136

6.6 Reconstruction of the constrained cantilever beam geometry 139

6.7 Histogram comparing the unconstrained and constrained normalised radii distributions 140

6.8 Optimised compliance functionals for the cantilever beam case 140

A.1 Hypercube bisection example 149

Chapter 1

Introduction

The rapidly advancing field of Additive Manufacturing (AM) has provided a revolutionary foundation to develop the next generation of structural optimisation techniques. The design flexibility of AM enables the precise manufacture of intricate multiscale structures, which would otherwise be prohibitively expensive to fabricate (Attaran, 2017; Plocher and Panesar, 2019). As a consequence, the nascent field of Multiscale Topology Optimisation (MTO) is gathering momentum, as researchers seek to leverage the increased geometrical freedom permitted through AM processes to design frameworks capable of efficiently optimising structures over two-scales (Rodrigues et al., 2002; Sivapuram et al., 2016; Zhu et al., 2017; Li et al., 2018; Garner et al., 2019; Imediegwu et al., 2019; Cheng et al., 2019; Colabella et al., 2019; Wang et al., 2020a; Liu et al., 2020a; Thillaithevan et al., 2021; Murphy et al., 2021; Nightingale et al., 2021).

The extremal point properties afforded by bespoke periodic microstructures within MTO frameworks represents an advantage over classical Topology Optimisation (TO) techniques (Mattheck and Burkhardt, 1990; Xie and Steven, 1993; Sethian and Wiegmann, 2000; Bendsøe and Sigmund, 2004; Wang et al., 2003; Allaire et al., 2004), which are typically constrained to isotropic point material properties. The accessibility to extremal point properties should not be understated, as to achieve equivalent local properties using TO requires the heterogeneous assembly of isotropic material, which is heavily dependent on initialisation conditions,

optimisation algorithm, mesh resolution and regularisation.

Many MTO frameworks opt for a parametrisation based strategy, where the design variables controlling the microstructural geometry are optimised to fulfil mechanical objectives at the macroscale (Cheng et al., 2019; Imediegwu et al., 2019; Colabella et al., 2019; Wang et al., 2020a; Thillaithevan et al., 2021; Nightingale et al., 2021). Efficient parametrisations can provide clear accessibility to extremal point properties, as individual parameters can be configured to couple tightly with specific elements of the elasticity tensor, enabling straightforward control of the mechanical point properties during optimisation. It is noted, however, that the parametrisation based MTO frameworks referenced within this paragraph all require a precomputed database of microstructures before optimisation can be performed. These are described as sequentially coupled frameworks. Precomputing a database is prohibitive, as it places a limit on either the number of design variables available to parametrise the microscale geometry or the fidelity of microscale simulations.

One of the primary aims of the work presented within this thesis is to circumvent the limitations imposed by precomputing a database, by coupling the micro and macro scale models concurrently (Weinan et al., 2007). This permits a previously intractable number of design variables to be optimised at the microscale, enabling enhanced control over the point properties during optimisation, by not requiring a precomputed database and instead assembling a partial database on-the-fly. Further to this primary aim, a novel optimisation strategy is also explored, which seeks to further reduce the computational expense of optimisation by initially optimising a subset of the microscale design parameters before activating additional microscale design parameters as the optimisation progresses. It is important to note, that this optimisation strategy is only beneficial through the application of a concurrently coupled MTO framework.

In addition, it is also noted within these parametrisation based MTO frameworks that open-walled (truss-like) microstructures are exclusively used to represent the microscale geometry, as they are straightforward to parametrise and afford both structural and manufacturing based benefits. However, in all cases, a non-zero lower bound is applied to constrain the radii

of the constituent structural members during optimisation to guarantee manufacturability, as a non-zero lower bound constrains the minimum length scale of an optimised geometry and prohibits the formation of unsupported members or isolated assemblies of overhanging members. Consequently, however, structural performance is traded for manufacturability, as accessibility is denied to microstructures with low volume-fractions.

A further aim of the work presented within this thesis is to circumvent the necessary application of prohibitive lower bound constraints to guarantee manufacturability, by integrating novel AM constraints within the present MTO framework. Through the inclusion of these constraints, it is permissible to derive manufacturable structures with substantially increased mechanical performance relative to those derived using prohibitive lower bound constraints.

1.1 Thesis Outline

This thesis is outlined as follows. Chapter 2 contains a review of the literature within the field of structural optimisation which is relevant to the MTO framework presented within this thesis. This includes a survey of classical TO techniques, which serve as the foundation to the present methodology, and a further survey detailing the advantages and disadvantages associated with various MTO frameworks. Chapter 3 contains a detailed breakdown of the methodology applied to construct the concurrently coupled MTO framework and the novel optimisation strategy (which includes the derivation of the foundational interior point optimisation algorithm). Chapter 4 contains the results derived through the application of the methodology presented within Chapter 3, which includes two cases demonstrating the capability of the concurrently coupled MTO framework and two cases demonstrating the capability of the novel optimisation strategy. Chapter 5 contains a detailed breakdown of the methodology applied to derive the AM constraints and the corresponding results derived through the application of these constraints is presented in Chapter 6. Finally, conclusions are drawn in Chapter 7 with respect to both methodology and results chapters.

1.2 Thesis Objectives

The objectives of this thesis are summarised as follows:

- To develop a concurrently coupled parametrisation based MTO framework, which enables optimisation to be performed at a significantly reduced computational cost in comparison to equivalent sequentially coupled frameworks.
- To implement a bespoke optimisation strategy, which further reduces the computational cost of optimisation, by initially optimising only the most sensitive microscale parameters, before introducing the insensitive parameters at a later stage of the optimisation procedure.
- To develop novel AM constraints for open-walled microstructures, which permit self-supporting, hierarchical structures with varying topology, suitable for fabrication through AM processes to be derived.

1.3 Publications

Journal Papers

1. Murphy, R., Imediegwu, C., Hewson, R., Santer M. Multiscale structural optimization with concurrent coupling between scales. *Struct Multidisc Optim 63, 1721 - 1741 (2021).* https://doi.org/10.1007/s00158-020-02773-3
2. Murphy, R., Hewson, R., Santer M. In-loop additive manufacturing constraints for open-walled microstructures. *Additive Manufacturing 48, (2021).* https://doi.org/10.1016/j.addma.2021.102385

Conference Papers

1. Murphy, R., Imediegwu, C., Hewson, R., Santer, M., Muir M. Multiscale Concurrent Optimization Towards Additively Manufactured Structures. *AIAA 2019-1635. AIAA*

Scitech 2019 Forum. January 2019. https://doi.org/10.2514/6.2019-1635

2. Murphy, R., Imediegwu, C., Hewson, R., Santer, M., Muir M. Multiscale Concurrent Multi-Objective Structural Optimization of a Goose Neck Hinge. *AIAA 2020-2260. AIAA Scitech 2020 Forum. January 2020.* https://doi.org/10.2514/6.2020-2260

Author Contributions

In addition to the supervision provided by Dr. Robert Hewson and Dr. Matthew Santer, it is important to highlight two further contributions which represent integral components of the MTO framework presented within this thesis. The code applied to derive the homogenised mechanical properties of singular microstructural geometries has been written by Dr. Chikwesiri Imediegwu. Only the parametrisation sub-routine has been modified for the present work. In addition, the code applied to reconstruct optimised macroscale geometries has been written in isolation by Dr. Robert Hewson.

Chapter 2

Review of Structural Optimisation Literature

Structural optimisation denotes the class of algorithms which combine computational structural mechanics and mathematical optimisation techniques to determine the optimal distribution of material within a set domain, subject to a set of boundary conditions and applied loads, which satisfies one or multiple functional objectives (Christensen and Klarbring, 2008). Typically, objectives such as a maximising stiffness, minimising weight or minimising maximum stress are targeted. Structural optimisation techniques are becoming increasingly integrated within industrial design processes, as they have the potential to yield significant increases in the operational efficiency of both automotive and aerospace based platforms, which will ultimately lead to a reduction in the operational cost of these systems (Keane and Nair, 2005).

Structural optimisation algorithms are typically classified into three distinct sub-classes — size, shape and topology, with each sub-class denoting the mechanism through which material can be distributed within the set domain during optimisation (see Fig. 2.1). This thesis focuses solely on the topology branch of structural optimisation.

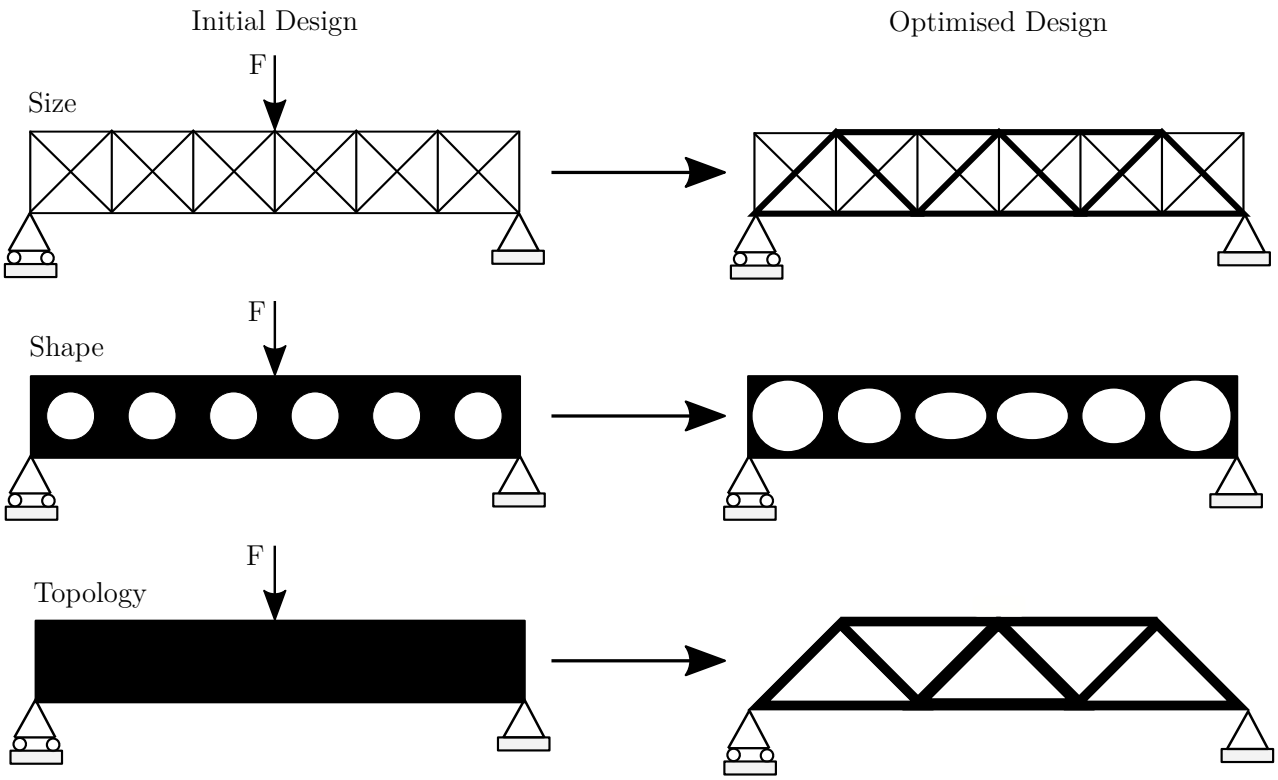


Figure 2.1: Illustrations depicting the three sub-classes of structural optimisation techniques — size, shape and topology. Each sub-class denotes the mechanism through which material can be distributed throughout the set domain. Adapted from Bendsøe and Sigmund (2004).

2.1 Topology Optimisation

Topology optimisation encompasses the sub-class of structural optimisation algorithms which determine the optimum material distribution within a set domain, without a predefined parametrisation model (Bendsøe and Sigmund, 2004). It is necessary in most cases to constrain TO problems to prohibit the derivation of trivial structures, typically by enforcing a volume-fraction constraint. This approach is in contrast to size and shape optimisation based approaches, where the connectivity of a structure remains fixed during optimisation and material can only be distributed using predefined model parameters. In most cases, parametrisation models severely restrict the permissible structural optimality (Christensen and Klarbring, 2008).

This approach was first documented within the 1988 homogenisation paper published by Bendsøe and Kikuchi (1988), which is widely acknowledged to be the seminal paper within

this field. Following this paper, the vast majority of TO research has focussed on mechanical based problems, which is also the sole focus of this thesis. However, due to the inherent design flexibility afforded through TO based approaches, numerous engineering disciplines have since leveraged this approach to derive structures with tailored thermal (Dbouk, 2017), dynamic (Xu et al., 2019), acoustic (Dühring et al., 2008) and magnetic (Bjørk et al., 2017) responses.

2.1.1 Topology Optimisation Frameworks

For the TO of continuum structures targeting mechanical objectives, optimisation problems are typically formulated using the weak variational form (Bendsøe and Sigmund, 2004)

$$\begin{aligned} &\underset{\boldsymbol{x}}{\text{Minimise:}} \quad J(\boldsymbol{u}(\boldsymbol{x}), \boldsymbol{x}) \\ &\text{Subject to:} \quad \int_{\Omega} E_{ijkl}(\boldsymbol{x})\epsilon_{ij}(\boldsymbol{u})\epsilon_{kl}(\boldsymbol{\omega})\, d\Omega = \int_{\Omega} \boldsymbol{f}\boldsymbol{\omega}\, d\Omega + \int_{\Gamma} \boldsymbol{t}\boldsymbol{\omega}\, d\Gamma \\ &\qquad\qquad \int_{\Omega} V(\boldsymbol{x})\, d\Omega - V_f \geq 0 \end{aligned} \tag{2.1}$$

where $\boldsymbol{x}$ denotes a spatially varying vector over the design domain, $\boldsymbol{u}$ denotes the displacement, J denotes the objective function, E_{ijkl} denotes the fourth-order elasticity tensor, ϵ_{ij} denotes the infinitesimal strain tensor, $\boldsymbol{\omega}$ denotes the virtual displacement, Ω denotes the design domain, $\boldsymbol{f}$ denotes body forces, $\boldsymbol{t}$ denotes surface tractions, Γ denotes surfaces, V denotes volume-fraction and V_f denotes target volume-fraction. The infinitesimal strains are calculated as a function of displacement

$$\epsilon_{ij}(\boldsymbol{u}) = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right) \tag{2.2}$$

where x_i denotes a co-ordinate axis. Solutions to the TO problem defined by Eq. 2.1 are non-existent, as it is ill-posed (Murat and Tartar, 1997). However, it is permissible to transform the continuous weak variational form into a well-posed problem by discretising Eq. 2.1 (Bendsøe and Sigmund, 2004), as this ensures the existence of solutions as discretisation fixes the resolution of the TO problem. It is then feasible to solve the discretised representation through the application of finite-element methods (Larson and Bengzon, 2013). Although, it is important to note that the solution is non-convergent, as structures begin to form over

multiple spatial length scales as the characteristic discretisation length tends to zero (Bourdin, 2001). To alleviate this issue, it is necessary to regularise the material distribution to ensure that convergence of the solution is obtained as the discretisation is refined (Sigmund, 2007).

The optimal material distribution for a discretised TO problem is commonly represented as a binary scalar field, with one of two material phases allocated to each node within the discretisation. Therefore, the discretised representation of Eq. 2.1 denotes a binary-integer program, which is computationally intractable to solve due to the vast number of design permutations. As a result, numerous algorithms have been recorded in literature which seek to overcome this computational intractability by reformulating the optimisation problem.

2.1.1.1 Solid Isotropic Material with Penalisation (SIMP)

The most prominent TO algorithm is the Solid Isotropic Material with Penalisation (SIMP) (Bendsøe, 1989; Zhou and Rozvany, 1991). This approach has seen widespread application across the automotive (Yang and Chahande, 1995) and aerospace (Zhu et al., 2016) industries and it is embedded within numerous popular computer-aided engineering packages such as Abaqus (TOSCA) (Smith, 2009), COMSOL Multiphysics (Multiphysics, 1998) and Optistruct (HyperWorks, 2017).

Within this framework, the binary material decision variable is represented as a continuous pseudo-density (ρ), which is defined in the interval $[0, 1]$ — corresponding to void and isotropic material properties at 0 and 1 respectively. This enables the binary integer program to be transformed into a continuous optimisation problem, which is analogous to a sizing optimisation problem (Bendsøe and Sigmund, 2004). In this continuous form, the TO problem can be efficiently solved, significantly reducing the computational expense of optimisation.

The disadvantage of this transformation is that it is permissible to derive structures with intermediate pseudo-densities during the optimisation procedure which do not have a physical representation. Therefore, it is necessary to discourage intermediate pseudo-densities from within the optimised solution by artificially penalising their material properties. For mechan-

ical problems, this is formulated as

$$E_{(}\rho) = E^0 + \rho^p(E^1 - E^0) \tag{2.3}$$

where $E(\rho)$ denotes the stiffness interpolant as a function of pseudo-density, E^1 and E^0 denote the Young's moduli for the solid and void phases respectively and p denotes the penalisation power (Bendsøe and Sigmund, 2004). To demonstrate how the penalisation power affects the stiffness, a visualisation is provided in Fig. 2.2 using multiple exponents. In addition, the Hashin-Shtrikman upper and lower energy bounds for two-phase composite materials (Hashin and Shtrikman, 1963) are also included to highlight the physically permissible values of the penalisation exponent.

Applying a penalisation reduces the stiffness-to-volume ratio of the intermediate pseudo-densities, whilst conserving this ratio at 0 and 1. Therefore, it is less advantageous for the optimiser to include intermediate pseudo-densities (for $p > 1$) within the optimised solution for objectives where maximising the stiffness-to-volume ratio is important. However, whilst penalisation helps to reduce the number of intermediate pseudo-densities, it does not result in a truly binary solution and as a consequence SIMP based TOs typically require post-processing to obtain manufacturable geometries (Liu et al., 2018).

There is a significant debate as to what is the optimal value of the penalisation parameter, as both low and high values afford unique benefits. A high penalisation parameter ($p \geq 5$) penalises the intermediate pseudo-densities more significantly, which results in a more binary optimised solution, however, the optimiser is more susceptible to becoming trapped within a local-minimum due to the high non-linearity of the stiffness interpolant. Conversely, for lower penalisation parameters ($p \leq 3$) the non-linearity is reduced, which reduces the difficulty of optimisation, particularly in the case where $p = 1$ as this represents a convex problem (Stolpe and Svanberg, 2001), however, the optimised material distributions are less binary as a result.

To alleviate the penalisation parameter selection problem, continuation approaches have been proposed (Sigmund and Maute, 2013). Continuation based approaches cast the penalisation

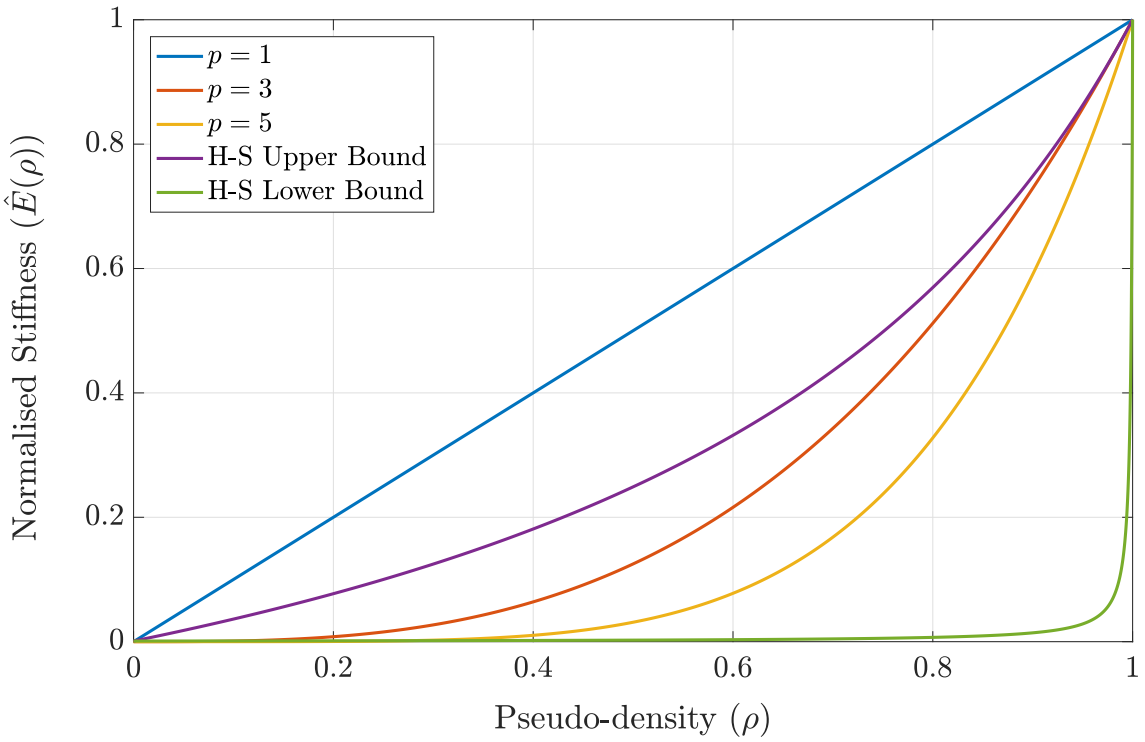


Figure 2.2: How a normalised stiffness interpolant ($\hat{E}(\rho)$) varies as a function of the penalisation parameter. In addition, the Hashin-Shtrikman upper and lower energy bounds for two-phase composite materials (Hashin and Shtrikman, 1963) are also illustrated.

parameter as a function of the iteration number. Initially, the penalisation value is set to a low value and then incremented regularly throughout the optimisation before obtaining a sufficiently high value. This reduces the initial difficulty of optimisation, helping the optimiser to avoid poor local minima, but ensures that the solution is sufficiently binary upon convergence.

2.1.1.2 Level-set Methods

Another well-established framework to solve TO problems is the level-set method (LSM) (Sethian and Wiegmann, 2000; Allaire et al., 2002; Wang et al., 2003; Allaire et al., 2004). In contrast to density based approaches, a boundary embedded in a scalar function of a higher-dimensional space is used to implicitly represent the structure (Wang et al., 2003). This is achieved through the application of a level-set function (LSF), which is defined as

$$\gamma = \{\boldsymbol{x} \mid \tau(\boldsymbol{x}) = 0\},\ \boldsymbol{x} \in R^n \tag{2.4}$$

where $\tau(\boldsymbol{x})$ denotes the LSF and γ denotes the level-set (boundary). By extension, material and void regions of the set domain (Ω) are trivially determined as

$$\tau(\boldsymbol{x}) > 0 \leftrightarrow \boldsymbol{x} \in \Omega \text{ (Interior)} \tag{2.5}$$

$$\tau(\boldsymbol{x}) < 0 \leftrightarrow \boldsymbol{x} \in (D/\Omega) \text{ (Exterior)} \tag{2.6}$$

Consequently, LSMs enable a discrete transition between material and void phases (see Fig. 2.3), which alleviates the intermediate pseudo-density problem detailed for SIMP within Sec. 2.1.1.1. Functional objectives are targeted within LSMs by optimising the LSF, which implicitly parametrises the geometry. One way in which this can be achieved is by initially reformulating Eq. 2.4 as a function of pseudo-time (t_p)

$$\gamma(t) = \{(\boldsymbol{x}(t_p)) \mid \tau(\boldsymbol{x}(t_p), t_p) = 0\},\ \boldsymbol{x}(t) \in R^n \tag{2.7}$$

and then by differentiating both sides of Eq. 2.7 with respect to pseudo-time

$$\frac{\partial \tau(\boldsymbol{x}, t_p)}{\partial t_p} + \nabla \tau(\boldsymbol{x}, t_p) \frac{d\boldsymbol{x}}{dt_p} = 0 \tag{2.8}$$

as this yields an initial value problem, analogous to a Hamilton-Jacobi equation, which can be efficiently solved. The LSF function is then updated iteratively in accordance with the solution of Eq. 2.8 until convergence. Whilst this formulation enables significant changes to the shape of the geometry during optimisation, the maximum genus of the geometry is fixed, and therefore the nucleation of new holes is prohibited during optimisation (Allaire et al., 2002). Consequentially, extensions to this method have been proposed which leverage topological derivatives to nucleate new holes during optimisation (Burger et al., 2004; Fulmanski et al., 2008).

Whilst LSMs afford a number of advantages over density based approaches, they are more susceptible to becoming trapped in local-minima and the final result is more dependent on the initial guess, which is a consequence of the shape-derivatives applied at the core of this method

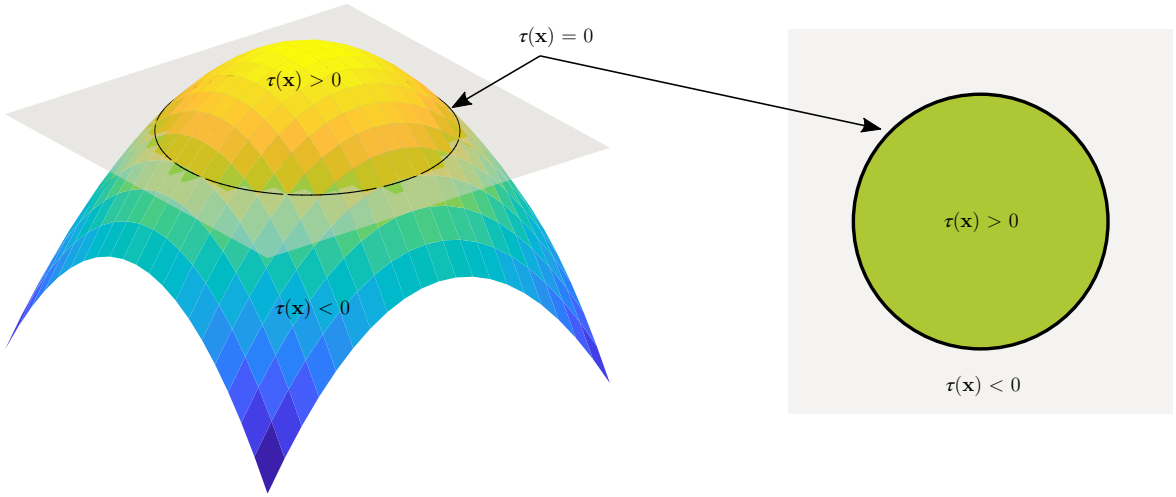


Figure 2.3: An example level-set function and resultant boundary ($\tau(\boldsymbol{x}) = 0$). Adapted from Luo et al. (2008).

(van Dijk et al., 2012, 2013). These limitations may explain why density based methods continue to be the most popular choice of TO algorithm (Sigmund and Maute, 2013).

2.1.1.3 Evolutionary Structural Optimisation

Evolutionary structural optimisation (ESO) algorithms are the third and final branch of algorithms typically applied to solve TO problems (Mattheck and Burkhardt, 1990; Xie and Steven, 1993). In contrast to density and LSM approaches which leverage gradient based optimisation strategies, ESO algorithms apply heuristic methods (Wang and Chen, 2013) to drive towards the optimal structure by removing material from discrete elements inefficiently located within the set-domain. Various criteria can be used to inform the optimiser which elements should be removed, however, the most popular choices are commonly based upon elemental stress. Whilst this methodology enables the efficient optimisation of structures, it is susceptible to becoming trapped in local-minima, as once material is removed from an element, it cannot be reintroduced in later iterations of the optimisation procedure (Edwards et al., 2007).

To combat this local-minima issue, an extension to ESO called Bi-Directional Structural Optimisation (BESO) was introduced (Querin et al., 1998; Young et al., 1999). Within this framework, both the addition and removal of material is permitted, which enables poor local-minima to be avoided. In contrast to the removal of material, material can be reintroduced to void elements if the elemental stresses in neighbouring elements are sufficiently large. Whist efficacious, both ESO and BESO suffer from non-convergent solutions (Bruns and Tortorelli, 2003; Bendsøe and Sigmund, 2004).

2.1.2 Manufacturing

The precise fabrication of topologically optimised structures represents a complex challenge, due to the inherent design flexibility afforded through TO based approaches. As a result, it is typically infeasible to fabricate these designs using traditional manufacturing techniques, such as machining, casting or moulding (Zhu et al., 2021). In most cases, the infeasibility is a result of the inability to locate the fabrication tools in all necessary regions of the set domain. Modifications to TO algorithms have been proposed which couple the design constraints of traditional manufacturing techniques to the optimiser, which permit the derivation of manufacturable topologically optimised designs. However, these methods are intrusive and significantly affect the optimality of the derived structures (Norato et al., 2015; Zhang et al., 2016; Mirzendehdel et al., 2020). As a consequence of these limitations, TO has traditionally only been used in the conceptual design stage, with the optimised solutions used as a foundation to realise more conventional designs (Bici et al., 2017).

A significant number of the challenges associated with manufacturing topologically optimised designs have been alleviated through the introduction of AM techniques. Additive manufacturing enables the fabrication of complex geometries, by building a structure layer-by-layer, as opposed to traditional subtractive manufacturing techniques, which incrementally remove material to realise designs (Gebhardt and Htter, 2016). Consequently, the build costs associated with AM are relatively insensitive to geometrical complexity (Attaran, 2017).

Whilst AM affords vast design flexibility, it is still necessary in most cases to impose manu-

facturing constraints to guarantee manufacturability. For topologically optimised structures, this can include constraints on the minimum feature size and constraints which ensure that the structure is self-supporting and is therefore not composed of any overhanging geometry, as this can lead to the warping of the geometry during fabrication as a consequence of the accumulation of thermal stresses (Muthusamy et al., 2018). In addition, it is necessary to prohibit the formation of enclosed-volumes for structures manufactured through powder bed fusion processes, as they result in trapped powder which leads to undesirable increases in the mass of a fabricated structure. Limitations with respect to powder bed fusion processes are important to highlight, as they represent the preferred AM technique applied to fabricate light-weighted optimised structures within the automotive and aerospace industries (Yap et al., 2015). If these fabrication constraints are not satisfied, invasive post-processing procedures are required, which severely affect the optimality of the derived structures (Zhu et al., 2021).

To eliminate the requirement of post-processing, numerous strategies have been implemented which couple the fabrication constraints of AM to TO based algorithms to guarantee manufacturability . In one example, a combination of both density and projection filters are applied to implicitly control the minimum length scale (Wang et al., 2011). Similarly, in another example, an anisotropic density based projection filter is implemented to discourage overhanging geometry within the optimised solution (Gaynor and Guest, 2016; Johnson and Gaynor, 2018). A further example presented by Langelaar (2016) details a layer-wise filtering procedure to discourage overhanging geometry. Whilst these approaches are effective at eliminating overhanging geometry, it is noted that significant perturbations in the structural layout are required in some instances to satisfy these constraints; illustrated in Fig. 2.4.

Furthermore, multiple techniques have been documented which seek to discourage the formation of enclosed-volumes during the optimisation procedure (Penumakala et al., 2020). In one example, a virtual temperature method is applied by representing the void and solid regions as highly conductive and insulating materials respectively (Liu et al., 2015; Li et al., 2016). Then, by continuously heating the void material and constraining the maximum temperature within the set domain, enclosed-volumes are actively discouraged.

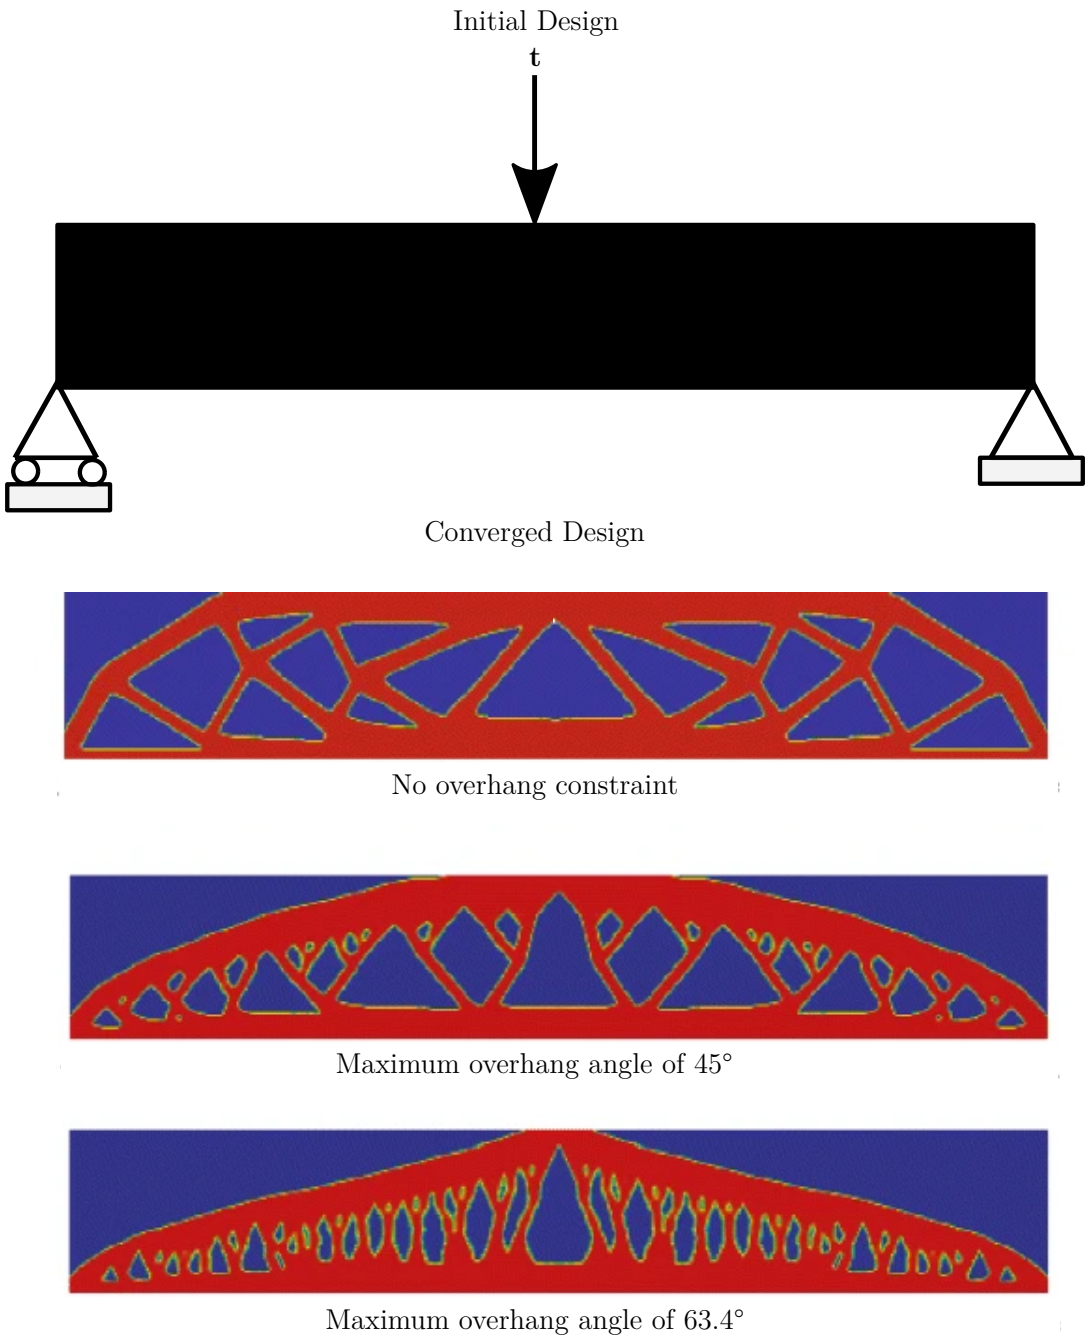


Figure 2.4: An example TO problem which demonstrates how overhang constraints can lead to significant changes in the layout of the optimised geometry. Adapted from Gaynor and Guest (2016).

2.2 Metamaterials

The material properties of the microstructure within all major branches of TO is typically constrained to be isotropic, as this generalises the framework and reduces the number of design variables to be optimised. It is apparent, however, that this places an artificial limit on the

permissible optimality of structures derived through mono-scale (classical) TO techniques (Sigmund et al., 2016; Wu et al., 2021a). As a consequence, in the pursuit of increased structural performance, extensive research has been conducted into the derivation of microstructures with extremal or counter-intuitive material properties (Chen et al., 2018). Typically, these microstructures are referred to as meta-materials, as they exhibit material properties which do not naturally occur in the base material (Kadic et al., 2019).

Meta-materials first rose to prominence within the field of optics, where they were leveraged to control electromagnetic waves (Veselago, 1968). Following their introduction, numerous engineering disciplines have since leveraged the benefits afforded by meta-materials, such as acoustics (Zhang et al., 2009), thermodynamics (Wang et al., 2020b) and mechanics. Applications within the field of mechanics is broad, as many unique meta-materials have been derived, such those exhibiting a negative Poisson's ratio (auxetic)(Lakes, 1987), enhanced energy-absorbing characteristics (Wang et al., 2020b) and extremal bulk moduli (Sigmund, 2000).

Whilst it is apparent that meta-materials provide a foundation to realise structures with enhanced or counter-intuitive material properties, in most applications, it is not clear what the maximum theoretical performance is. This is not the case for the mechanical bulk and shear moduli of an isotropic two-phase composite material, as both the upper and lower bounds have been derived for these quantities using energy-based variational methods (Hashin and Shtrikman, 1963). These bounds, commonly referred to as the H-S bounds, are of particular importance for mechanical compliance based optimisations, as they provide a global lower bound on mechanical compliance (Sigmund et al., 2016).

2.3 Homogenisation

Primarily, TO techniques are applied to derive optimal macroscale topologies, however, through the application of inverse homogenisation it is also permissible to apply these techniques to derive periodic microstructures with prescribed homogenised material tensors (meta-materials)

(Sigmund, 1994). Analogous to TO, inverse homogenisation problems are formulated to determine the optimal distribution of material, which satisfies the constraints imposed on the homogenised material property tensor. As it is possible for multiple topologies to exhibit an arbitrary material property, it is common practice for the lightest structure to be targeted during the optimisation procedure.

Homogenisation in this context refers to a spatial averaging procedure which enables the derivation of an effective material property tensor for a heterogeneous assembly of material (Hassani and Hinton, 1998). Multiple homogenisation techniques have been recorded in literature, with the choice of methodology typically based on a trade-off between computational expense and accuracy (Arabnejad and Pasini, 2013). For finite assemblies, it is permissible to represent the microscale geometry as a continuum with homogenised material properties to evaluate the macroscopic response of a structure, so long as a separation of scales exists between the micro and macroscale characteristic lengths. To demonstrate this, an asymptotic expansion of a two-scale displacement field in δ is considered (Hollister and Kikuchi, 1992)

$$\boldsymbol{u}^{\delta}(\boldsymbol{x},\boldsymbol{y}) = \boldsymbol{u}_0(\boldsymbol{x},\boldsymbol{y}) + \delta\boldsymbol{u}_1(\boldsymbol{x},\boldsymbol{y}) + \delta^2\boldsymbol{u}_2(\boldsymbol{x},\boldsymbol{y}) + O(\delta^3) \tag{2.9}$$

where δ denotes the ratio of the micro to macro scale characteristic lengths, $\boldsymbol{u_\delta}$ denotes the true displacement field, $\boldsymbol{u_0}$ denotes the macroscopic displacement field, $\boldsymbol{u_1}$, $\boldsymbol{u_2}$ and higher-order terms denote perturbations to the macroscopic displacement field, and $\boldsymbol{x}$ and $\boldsymbol{y}$ denote the macro and micro level coordinates respectively. This expansion implies that the microscopic displacement field varies δ^{-1} times quicker than the macroscopic displacement field. As a consequence, so long as the ratio of the characteristic lengths obeys the inequality constraint

$$0 < \delta \ll 1 \tag{2.10}$$

perturbations imposed by microscopic fluctuations on the macroscopic displacement field can be considered negligible. Consequently, it is permissible to derive an inexpensive and accurate evaluation of the macroscopic response of a structure defined over multiple spatial length scales by representing it as a continuum with homogenised material properties; illustrated for clarity

in Fig. 2.5. This is in contrast to direct, full-scale analysis, which is computationally intractable in most cases due to the vast computational resources required to both mesh and simulate the geometry (Otero et al., 2018). It is important to note that the primary disadvantage associated with homogenisation is the inherent loss of information at the microscale as a result of the averaging procedure, however, in most cases, it is only the macroscopic response of a structure which is of interest.

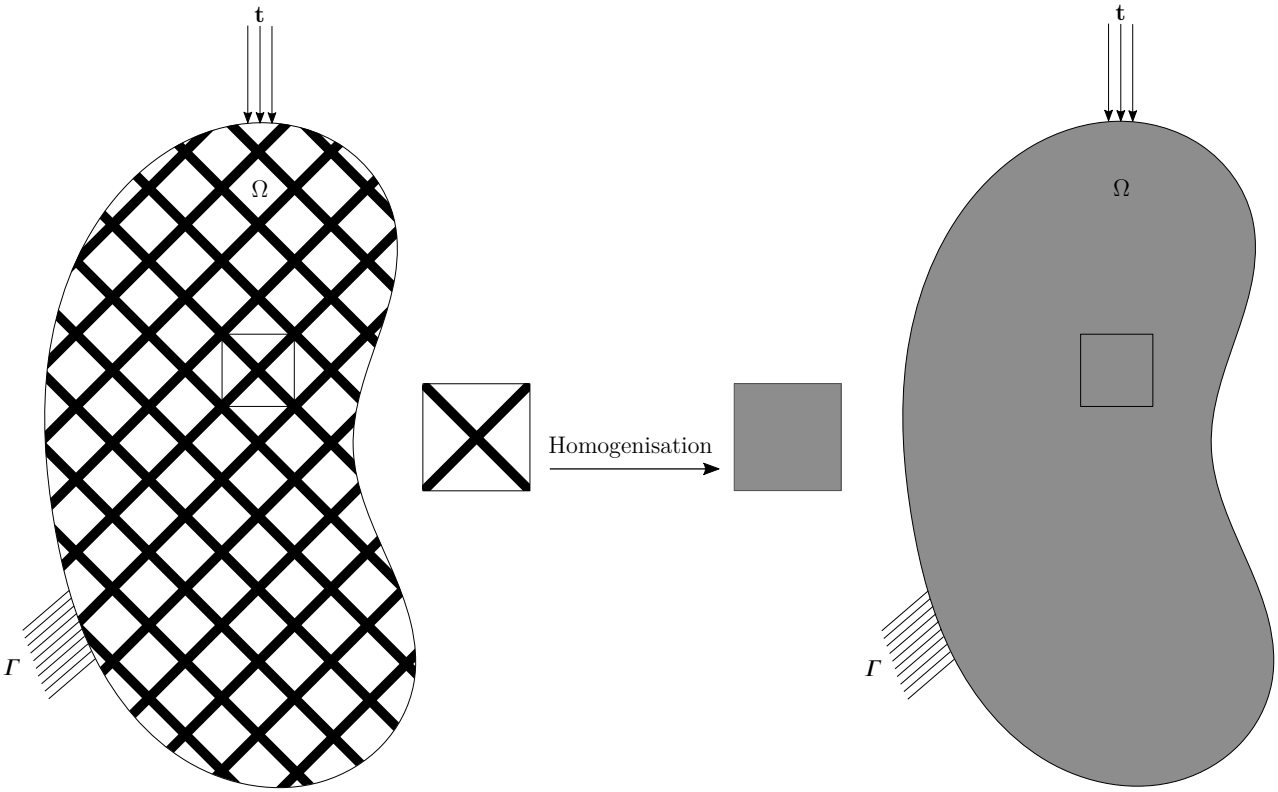


Figure 2.5: An illustration demonstrating how a heterogeneous (periodic) structure can be represented as a homogeneous structural continuum. It is important to note that a sufficient separation of scales should exist between the micro and macro scales for this homogeneous representation to be valid.

2.4 Multiscale Topology Optimisation

In Sec 2.1.1, it is stated that as the discretisation of a mono-scale (classical) TO is increased to a sufficiently high resolution, structures over multiple length scales begin to form during the optimisation procedure with increased structural performance. However, for most practical applications, high-fidelity TOs are computationally intractable (Aage et al., 2017). Further

to this, prior to the introduction and subsequent mass adoption of AM techniques, it was not economically viable to fabricate structures with more than one spatial length scale. This explains why TO research has historically focussed on the development of mono-scale methodologies. However, as detailed in Sec. 2.1.2, recent and rapid advancements within the field of AM have enabled the efficient fabrication of structures defined over multiple scales. As a consequence, the nascent field of MTO is gathering momentum, as researchers seek to combine the benefits afforded by asymptotic homogenisation theory, extremal meta-materials and the design flexibility of AM to design structures with increased structural performance.

2.4.1 Concurrent Multiscale Optimisation

A common and well established approach to MTO is a hierarchical computational procedure, which is typically referred to as concurrent multiscale optimisation. Within this hierarchical framework, both the micro and macroscale description of a structure is optimised simultaneously (Rodrigues et al., 2002; Liu et al., 2008; Garner et al., 2019; Liu et al., 2020b). At the macroscale, a TO problem is solved to determine the optimal spatial distribution of material, and at the microscale, inverse homogenisation techniques are applied to determine optimal microscale topologies as a function of the homogeneous solution to the macroscale problem.

During each iteration of the optimisation procedure, optimisation of the macroscale material distribution precedes optimisation of the microscale. As a consequence, it is permissible to parallelise optimisation of the microscale, to prohibit time-intensive sequential computation. Connectivity between adjacent microstructures is not guaranteed through this approach, as numerous topologies are capable of exhibiting equivalent homogenised material properties. This can lead to significant differences between the predicted and true macroscopic response of a structure (Garner et al., 2019; Liu et al., 2019). To alleviate this issue, Liu et al. (2020b) proposed a modification by introducing pre-defined connected regions within neighbouring microstructures to guarantee connectivity.

By definition, this approach affords vast flexibility over both the micro and macroscale description of a structure. Consequently, assuming local minima are avoided, this method is capable

of deriving theoretically optimal multiscale topologies. As a result, in situations where the optimal microscale topology is not known, concurrent multiscale optimisation provides a means to determine optimal topologies. However, whilst this method is more computationally efficient than equivalent high-fidelity TOs (due to parallelisation architecture), it still requires a vast computational resource, which similarly, in most practical applications is computationally intractable.

To alleviate this issue, a modification to this approach was proposed by Sivapuram et al. (2016), which decouples the linear relationship between the number of inverse homogenisation problems and the number of nodes within the macroscale discretisation. This is achieved by dividing the macroscale domain into a series of sub-domains, where a single inverse homogenisation problems is solved for the microstructure of each sub-domain. The number of sub-domains is then used as a metric to control the computational cost, however, this can severely degrade the optimality of the optimised structure.

2.4.2 Univariate Parametrisation Frameworks

Another common approach to MTO is to leverage the optimised pseudo-densities from a monoscale SIMP based TO to functionally-grade a parametrised microstructure (Brackett et al., 2011; Cheng et al., 2017; Panesar et al., 2018; Liu et al., 2020c). This approach enables the problems associated with intermediate pseudo-densities to be alleviated, as in this instance they correspond to physically manufacturable structures. Importantly, this eliminates the requirement of a binary material distribution, enabling a lower penalisation parameter (see Eq. 2.3) to be used during the optimisation procedure, which reduces the chance of the optimiser from becoming trapped within a local minimum and therefore increases structural performance. In addition, connectivity between adjacent microstructures is guaranteed. It is crucial, however, that the penalisation curve (see Fig. 2.2) is co-aligned to the material property distribution of the parametrised microstructure, as any discrepancy between the two curves has the potential to lead to significant differences between the predicted and true structural response.

To avoid this material property alignment problem, it is common to construct a continuous univariate response surface for the homogenised material properties of the parametrised microstructure as a function of volume-fraction (pseudo-density) (Cheng et al., 2019; Yu et al., 2020). The homogenised material properties used to construct the response surface are derived by fitting a response-surface to the homogenised material properties derived for a series of uniformly spaced microstructures throughout the allowable volume-fraction range. Topology optimisation is then performed at a single-scale using the continuous response surface model in-place of the classical penalisation curve. The optimised multiscale structure is then subsequently derived upon convergence through de-homogenisation of the optimal solution.

It should be noted that analytical models also exists for the material properties of some microstructures (Wu et al., 2021a), which eliminates the requirement of a surrogate model. However, in most cases it is not possible to derive an analytical model for the material properties.

In contrast to the hierarchical approaches described within Sec. 2.4.1, the computational expense of MTO using a univariate response surface is identical to mono-scale TO techniques (excluding negligible pre-computation). Furthermore, it is important to note that once the response surface has been derived it can be exhaustively interrogated within subsequent optimisations for no further computation cost. Whilst efficacious, these methods do not enable a fundamental advantage of MTO to be leveraged — anisotropic extremal microstructures. Consequently, the permissible optimality of a univariate framework is severely restricted in comparison to hierarchical approaches.

2.4.3 Multivariate Parametrisation Frameworks

The natural extension to univariate parametrisation frameworks is to introduce a multivariate parametrisation of the microscale geometry in order to effectively leverage the benefits afforded by anisotropic extremal microstructures. Multivariate parametrisations are capable of providing clear accessibility to extremal mechanical properties, as individual parameters can be configured to couple tightly with specific elements of the material property tensor, enabling

straightforward control of the homogenised properties during optimisation (Imediegwu et al., 2019; Colabella et al., 2019; Thillaithevan et al., 2021; Nightingale et al., 2021).

Typically, multivariate parametrisation frameworks are configured in a sequential fashion to enable the multivariate response surfaces to be derived in advance of optimisation. For the most popular response surface methodologies, such as least-squares polynomial regression (Conte and de Boor, 2017), Kriging (Forrester et al., 2008) and radial basis functions (Buhmann, 2003), a sequential framework is mandatory as all of the microscale data is required prior to the generation of these response surfaces. In contrast to the univariate response surfaces models detailed in Sec. 2.4.3, a computationally expensive pre-computation step is required to collect the microscale data. This is a direct consequence of the exponential increase in the number of microscale simulations required to generate accurate response surfaces as the number of microscale design parameters increases (Bellman et al., 1957) . For example, in the multivariate parametrisation framework presented by Imediegwu et al. (2019), a database containing $823,543$ unique microstructures was required to establish a relationship between the seven geometric parameters controlling the open-walled (truss-like) microscale description and all 21 unique elements of the elasticity tensor. It is prudent to note, however, that through the application of rotational and reflective symmetries of the microscale geometry, only $40,817$ microscale simulations were required in this instance. Similarly, Colabella et al. (2019) recorded a further example of a multivariate parametrisation framework, where a database containing $41,990$ unique microstructures was required to fit multivariate polynomial equations to establish continuous relationships between the seven microscale design parameters and the 9 non-zero elements of the elasticity tensor. Given in both cases each microscale simulation takes approximately four minutes on average to execute, this represents a significant computational expense.

In theory, it is permissible for multivariate parametrisation frameworks to achieve similar structural performance to hierarchical approaches (see Sec. 2.4.1) for a significantly reduced computational expense, so long as the microscale parametrisation is capable of achieving extremal performance for the targeted application. However, as detailed Sec. 2.2, in many

instances the optimal material property tensor for a given application is unknown, therefore, it is impermissible to generate a parametrisation capable of achieving extremal performance in many instances. Consequently, it is prudent to afford as much freedom in the microscale geometry as is possible in these instances. However, for sequential multivariate parametrisation frameworks, this leads to an intractable computational expense as the pre-computation requirement needed to assemble a microstructure database increases exponentially with the number of microscale design parameters (Imediegwu et al., 2019; Colabella et al., 2019).

2.4.4 Manufacturing

The precise and economic fabrication of optimised multiscale structures is only feasible through the application of AM techniques, as it is the only manufacturing process which provides sufficient design flexibility for a modest cost (Plocher and Panesar, 2019). The design constraints imposed by AM processes on MTO frameworks are identical to the constraints imposed on mono-scale TOs (see Sec. 2.1.2). For completeness, this set includes constraints on the minimum length scale and on the formation of overhanging geometry and enclosed-volumes.

For the hierarchical MTO frameworks detailed in Sec. 2.4.1, it is permissible to leverage the projection and density filters applied within mono-scale TO frameworks to control the minimum length scale. Importantly, it would only be necessary to apply these techniques at the microscale, as this is where the smallest features will arise. Similarly, anisotropic filters can be applied to discourage overhanging geometry within hierarchical frameworks, however, in this instance it would be necessary to apply these techniques at both scales of the discretisation. At present, the fabrication constraints of AM have not been directly coupled to hierarchical MTO frameworks.

For parametrisation based MTO frameworks, identical design constraints are imposed by AM processes to those mentioned at the beginning of this section, however, it is not possible to leverage equivalent techniques, as TO is not performed at the microscale. Consequently, it is common practice to explicitly avoid these fabrications constraints by leveraging an open-walled (truss-like) microstructure and placing a non-zero lower bound on the radii of the

constituent structural members, as this explicitly places a limit on the minimum length scale and prohibits the formation of overhanging geometry and enclosed-volumes (Imediegwu et al., 2019; Wang et al., 2020b; Thillaithevan et al., 2021; Nightingale et al., 2021). As a result, however, structural performance is traded for manufacturability, as accessibility is denied to microstructures with low volume-fractions due to the non-zero lower bound on structural member radii. Similarly, at present, the design constraints imposed by AM processes have not been directly coupled to parametrisation based MTO frameworks.

2.5 Summary

The MTO frameworks referenced within Sec. 2.4 correspond to the state-of-the-art methodologies within the field of MTO. It is clear, however, that these frameworks are either computationally intractable to execute for three-dimensional geometries (concurrent multiscale optimisation), severely limited by the permissible optimality of the microscale geometry (univariate parametrisation frameworks) or require an intractable computational resource to precompute a microscale database for a sufficiently general microscale parametrisation (multivariate parametrisation frameworks).

Consequently, the primary aim of the work presented within this thesis is to circumvent the limitations imposed by pre-computation on multivariate parametrisation frameworks, by coupling the micro and macro scale models concurrently (Weinan et al., 2007). Concurrent coupling ensures that only the microscale data required to evaluate the macroscale model during each iteration of optimisation is collected. It is assumed that coupling the micro and macro scale models concurrently will lead to a reduction in the computational expense of optimisation, as only a subset of the microscale database will be required to perform a single MTO. As a consequence, concurrent coupling enables accessibility to a space of designs, unobtainable through sequential methods (see Fig. 2.6). This will permit a previously intractable number of design variables to be used in the parametrisation of the microscale geometry, enabling a step towards the geometric freedom afforded by hierarchical MTO frameworks for a significantly

reduced computational cost.

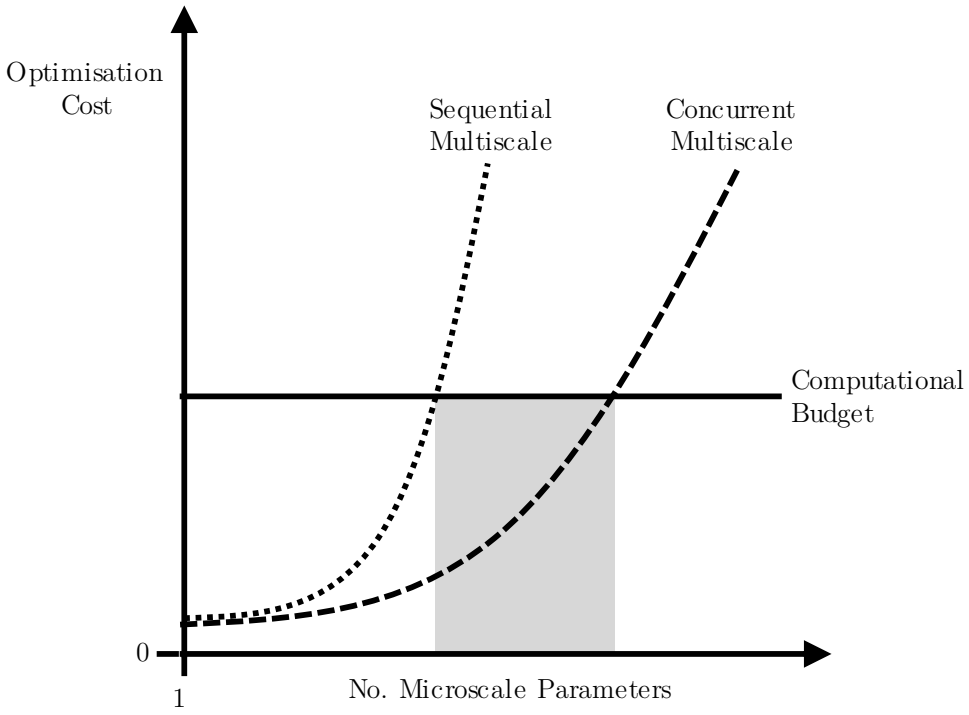


Figure 2.6: The assumed computational costs of optimisation associated with sequentially and concurrently coupled multivariate parametrisation frameworks. The grey region denotes a space of designs only accessible through concurrent methods.

Further to this primary aim, a novel optimisation strategy is also explored which seeks to further reduce the computational expense of optimisation by initially optimising a subset of the microscale design parameters before activating additional microscale design parameters as the optimisation progresses. It is important to note that this optimisation strategy is only permissible through the application of a concurrently coupled MTO framework.

The secondary aim of the work presented within this thesis is to circumvent the necessary application of prohibitive minimum length scale constraints to guarantee manufacturability within multivariate parametrisation frameworks, by integrating novel AM constraints. Through the inclusion of these constraints, it is permissible to derive manufacturable structures with substantially increased mechanical performance relative to contemporary multivariate parametrisation frameworks. Through the application of this methodology, it is permissible to derive self-supporting, hierarchical structures with varying topology, suitable for fabrication through AM processes. In addition, implicit minimum length scale control is also afforded.

Chapter 3

Multiscale Topology Optimisation Framework

Outlined within this section is a detailed description of the methodology underpinning the MTO framework. This includes the derivation of the microscale model, macroscale model, concurrent coupling strategy, optimisation formulation, addition of microscale parameters optimisation strategy and AM constraints.

3.1 Linear Elastic Assumptions

Structural considerations for the remainder of this thesis are limited to the theory of three-dimensional linear elasticity, so small deformations are assumed at both scales. As a consequence, constituent nodes used in the construction of both the microscale and macroscale finite-element models are required to satisfy the constitutive law for linear elasticity

$$\sigma_{ij} = E_{ijkl}\epsilon_{kl} \tag{3.1}$$

where σ_{ij} represents the Cauchy stress tensor, E_{ijkl} denotes the fourth-order elasticity tensor with symmetries

$$E_{ijkl} = E_{jikl} = E_{ijlk} = E_{klij} \tag{3.2}$$

which for a fully-anisotropic material reduces to 21 unique terms due to symmetries, and ϵ_{kl} denotes the infinitesimal strain tensor. Strain is calculated using the strain-displacement equations

$$\epsilon_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right) \tag{3.3}$$

where u_i denotes the displacement in a single coordinate axis and x_i denotes a coordinate axis. Further to these initial assumptions, static equilibrium is also assumed in the evaluation of both structural models

$$\frac{\partial \sigma_{ij}}{\partial x_j} + F_j = 0 \tag{3.4}$$

where F_j denotes the sum of the external forces acting on the set domain.

3.2 Microscale Model

The microscale model serves as a callable function within the present framework, which accepts an array of input parameters describing a single microstructural geometry and after computation returns the corresponding homogenised mechanical properties. The periodic microscale geometry (illustrated in Fig. 3.1) is constructed using 7 parameterised structural members, each with a unique but fixed orientation and located within the boundaries of a unit cell.

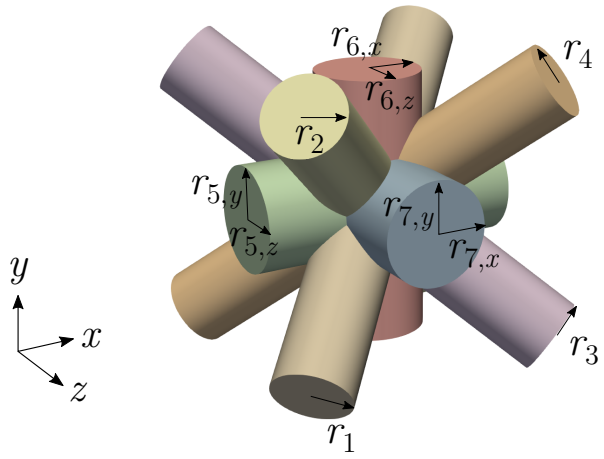


Figure 3.1: The parameterised microscale geometry used within the present framework, with all 10 unique parameters depicted

The 7 members comprise 4 diagonal members, connecting opposing vertices of the unit cell and 3 face members, connecting the centroids of opposing faces. The diagonal members are cylindrical in shape, individually parametrised by a normalised radius $(\hat{r}_1, \hat{r}_2, \hat{r}_3, \hat{r}_4)$. Normalised radii $(\hat{r}_\eta)$ are calculated by dividing the dimensioned radii (r_η) by the characteristic unit cell length (l), and are used as they enable a clearer interpretation of the optimised results. In contrast to the diagonal members, the face members have configurable elliptical cross-sections, parametrised by normalised radius $(\hat{r}_5, \hat{r}_6, \hat{r}_7)$ and aspect-ratio (e_5, e_6, e_7). The normalised radii and aspect-ratios are used in combination to derive the two orthogonal axes $(\hat{r}_{\eta,1}, \hat{r}_{\eta,2})$ defining the elliptical cross-sections, using a piecewise continuous relationship defined $\forall\, \eta \in \{5, 6, 7\}$

$$\hat{r}_{\eta,1}(\hat{r}_\eta, e_\eta) = \begin{cases} \hat{r}_\eta e_\eta, & AR_{max} \leq e_\eta \leq 1 \\ \hat{r}_\eta, & 1 < e_\eta \leq 2 - AR_{max} \end{cases} \tag{3.5}$$

$$\hat{r}_{\eta,2}(\hat{r}_\eta, e_\eta) = \begin{cases} \hat{r}_\eta, & AR_{max} \leq e_\eta \leq 1 \\ \hat{r}_\eta \left(2 - e_\eta\right), & 1 < e_\eta \leq 2 - AR_{max} \end{cases} \tag{3.6}$$

where $AR_{max} \in (0, 1]$ denotes the maximum permissible aspect-ratio of an elliptical cross-section. This formulation permits either orthogonal axis $(r_{5,y}, r_{5,z}, r_{6,x}, r_{6,z}, r_{7,x}, r_{7,y})$ to become the dominant axis; as illustrated in Fig. 3.2. The combination of normalised radius and aspect-ratio is preferred over directly using the two orthogonal axes as parameters, as it enables a limit to be placed on the maximum aspect-ratio of an elliptical cross-section.

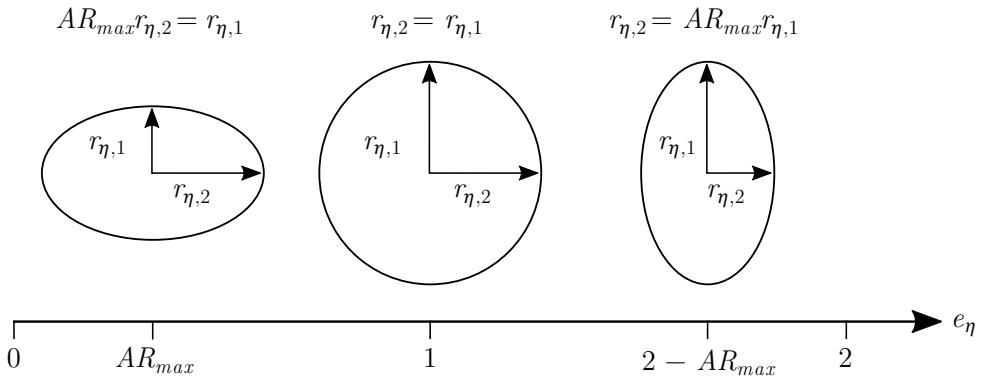


Figure 3.2: The permissible aspect-ratios of an elliptical cross-section, defined as a function of the maximum allowable aspect-ratio (AR_{max})

This parametrisation of the microscale geometry is selected as it provides clear accessibility to a range of mechanical properties, whilst satisfying the requirements of the heterogeneous multiscale method (Weinan et al., 2007), which requires the geometry to be periodic and the mechanical property space to vary continuously as a function of the microscale parameters. In addition, this parametrisation prohibits the formation of problematic fabrication features such as locally enclosed-volumes. Through the precise configuration of the microscale parameters, orthotropic microstructures with efficient stiffness characteristics can be configured to resist a range of load cases, which are ideal for compliance minimisation-based problems. Additionally, the parameters controlling the precisely-orientated diagonal members, enable direct manipulation of the shear-extension terms within the elasticity matrix, providing a mechanism to couple shear and normal forces, and therefore facilitating clear accessibility to anisotropic mechanical properties. In addition to optimising for minimum compliance, this microstructure is well suited at targeting other structural optimisation problems, such as obtaining a target deformation profile. It is important to note however, that the primary target of the work presented within this thesis is to demonstrate the potential of concurrent coupling within a MTO framework, and not to focus on the optimality of the microstructure used with respect to the Hashin-Shtrikman upper bounds for bulk and shear moduli (Hashin and Shtrikman, 1963). Furthermore, it should be also noted that it is permissible to interchange the microstructure used within this framework with another, so long as it is satisfies the requirements of the heterogeneous multiscale method.

3.2.1 Nodal Property Allocation

The microscale model is initialised for all parameter configurations by discretising a cubic domain into a structured mesh using hexahedral elements. A fixed mesh is preferred, as it aids with the implementation of periodic boundary conditions, due to the consistent discretisation on opposing faces. It is important to note that hexahedral elements are preferred over the application of beam finite-elements, as they enable accurate predictions for the homogenised properties of microstructures across the full volume-fraction range to be obtained. This is in contrast to beam finite-elements, which do not accurately capture the structural mechanics of

member intersection (significant in high volume-fraction microstructures) which results in an over prediction of the mechanical properties (Arabnejad and Pasini, 2013). Modifications to improve the accuracy of beam finite-element models within this context have been suggested, however, significant differences are still noted in the prediction of the mechanical properties (Guo et al., 2020).

Microstructures are assembled within the fixed domain by assigning one of two possible material phases to each constituent node. The primary phase is represented by an isotropic material with properties $E = 209$ GPa and $\nu = 0.3$, and the secondary phase is represented by a quasi-void material with isotropic properties $E = 10^{-7}$ GPa and $\nu = 0.3$. The primary phase is assigned to a node if it lies on the interior or boundary of one or more structural members. To determine which nodes are assigned this phase, a node set (N_η) is derived for each member by computing the set of nodes (P) which lie on the interior or boundary, using the perpendicular distance separating the nodes and the member centrelines

$$N_\eta = \left\{ P \in \mathbb{R}^3 \;\middle|\; \frac{\|P \circ \kappa_\eta \times (P \circ \kappa_\eta - \boldsymbol{l}_\eta)\|}{\|\boldsymbol{l}_\eta\|} \leq \hat{r}_\eta \right\}, \qquad \forall\, \eta \in \{1, 2, 3, \ldots, 7\} \tag{3.7}$$

where $\boldsymbol{l}_\eta$ denotes the direction vector of a member ($\hat{r}_\eta$) and

$$\kappa_{\eta k} = \begin{cases} \hat{r}_{5,z}/\hat{r}_{5,y}, & \text{if } \hat{r}_{5,y} < \hat{r}_{5,z} \;\wedge\; \eta = 5 \;\wedge\; x_k = 2 \\ \hat{r}_{5,y}/\hat{r}_{5,z}, & \text{if } \hat{r}_{5,y} \geq \hat{r}_{5,z} \;\wedge\; \eta = 5 \;\wedge\; x_k = 3 \\ \hat{r}_{6,z}/\hat{r}_{6,x}, & \text{if } \hat{r}_{6,x} < \hat{r}_{6,z} \;\wedge\; \eta = 6 \;\wedge\; x_k = 1 \\ \hat{r}_{6,x}/\hat{r}_{6,z}, & \text{if } \hat{r}_{6,x} \geq \hat{r}_{6,z} \;\wedge\; \eta = 6 \;\wedge\; x_k = 3 \\ \hat{r}_{7,x}/\hat{r}_{7,y}, & \text{if } \hat{r}_{7,y} < \hat{r}_{7,x} \;\wedge\; \eta = 7 \;\wedge\; x_k = 2 \\ \hat{r}_{7,y}/\hat{r}_{7,x}, & \text{if } \hat{r}_{7,y} \geq \hat{r}_{7,x} \;\wedge\; \eta = 7 \;\wedge\; x_k = 1 \\ 1, & \text{otherwise} \end{cases} \qquad \forall\, k \in \{1, 2, 3\} \tag{3.8}$$

denotes a nodal mapping matrix which enables the elliptical cross-sections of the face members to be compensated for by applying transformations to the cubic domain, where x_k denotes a coordinate axis. The node sets are combined using a finite union, to establish which nodes are assigned the primary phase

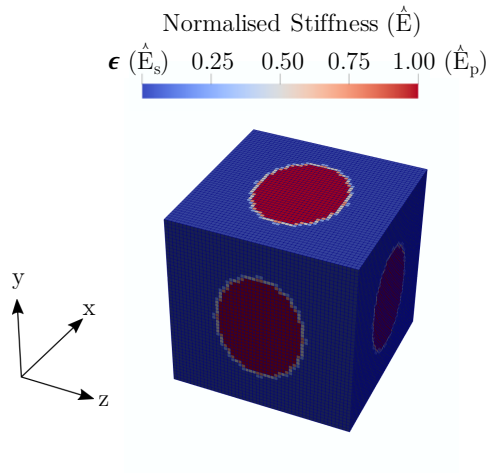
$$N_P = \bigcup_{\eta=1}^{7} N_\eta \tag{3.9}$$

and conversely by symmetric difference, which nodes are assigned the secondary phase (N_S). Stiffness properties are allocated to each element within the cubic domain, as a function of the fraction of nodes assigned each phase

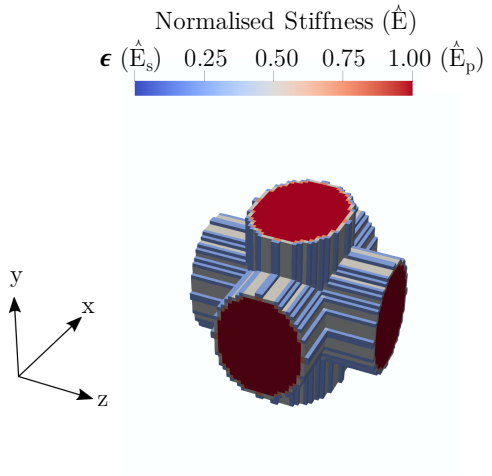
$$E(n_P, n_S) = n_P E_P + n_S E_S \tag{3.10}$$

where E_P and E_S correspond to the stiffness properties of the primary and secondary phases respectively, and n_P and n_S correspond to the fraction of nodes assigned either the primary or secondary phase respectively. As the characteristic discretisation length tends to zero, the analytical representation of the microscale geometry is recovered through nodal property allocation.

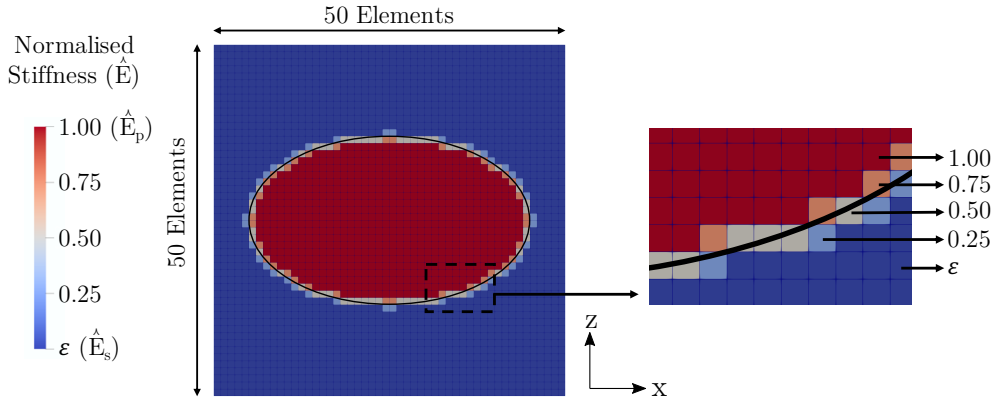
Two examples of nodally allocated meshes are illustrated within Figs. 3.3 and 3.4 for two different microstructures. It is clearly illustrated within Fig. 3.3c how the elements located at the boundary of structural member r_7 are allocated a stiffness proportional to the number of nodes assigned the primary and secondary phases. In addition, it is also important to highlight the additional material allocated near to the vertices of the cubic mesh within Fig. 3.4, which depicts nodal property allocation for a single diagonal member (r_3), as this ensures periodicity of the microstructure (see Fig. 3.4c). Allocation of this additional material is achieved by perturbing the centroid of each diagonal member to the eight vertices of the unit cube and applying Eq. 3.7.



(a) Isometric projection of nodally allocated mesh

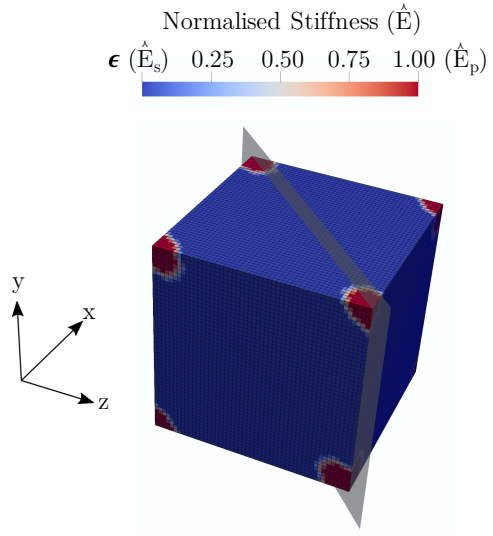


(b) Isometric projection of thresholded ($\hat{E} > \epsilon$) nodally allocated mesh

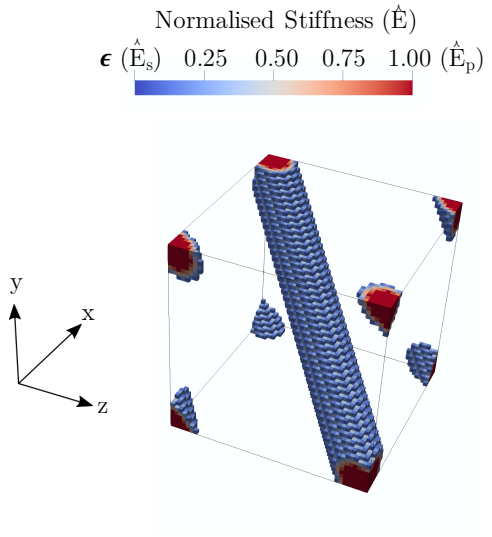


(c) Projection of the nodally allocated mesh onto the $X - Z$ plane

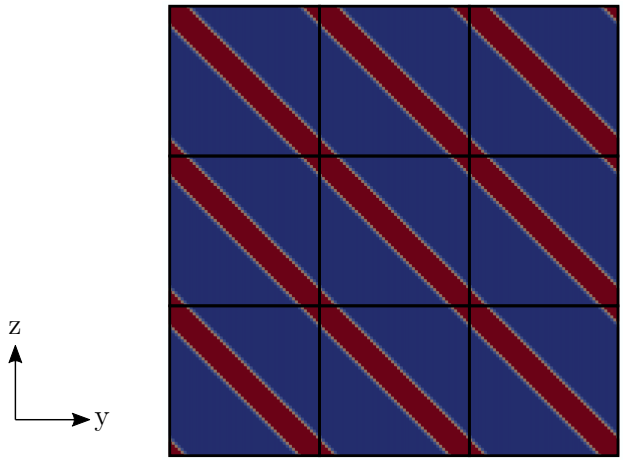
Figure 3.3: Illustration of a nodally allocated hexahedral mesh for three eccentric face members



(a) Isometric projection of nodally allocated mesh



(b) Isometric projection of thresholded ($\hat{E} > \epsilon$) nodally allocated mesh



(c) An assembly of the projection of the nodally alloacted mesh onto the $Y - Z$ plane using the geometry co-aligned to the plane depicted within 3.4a to demonstrate periodicity

Figure 3.4: Illustration of a nodally allocated hexahedral mesh for a single diagonal member (r_3)

3.2.2 Derivation of Homogenised Mechanical Properties

The homogenised mechanical properties of assembled microstructures are derived through periodically constrained finite-element analyses[1], conducted using the open source finite-element module FEniCS (Logg et al., 2012). For a three-dimensional periodic microstructure, a total of six strain deformation analyses are required to provide sufficient information to derive the homogenised mechanical properties.

Strain deformation analyses are conducted using a modified representation of the weak variational form for linear elasticity. In order to derive this modified representation, it is necessary to first consider an asymptotic expansion of a periodic displacement field (Hollister and Kikuchi, 1992)

$$\boldsymbol{u}^{\delta}(\boldsymbol{x},\boldsymbol{y}) = \boldsymbol{u}_0(\boldsymbol{x},\boldsymbol{y}) + \delta\boldsymbol{u}_1(\boldsymbol{x},\boldsymbol{y}) + \delta^2\boldsymbol{u}_2(\boldsymbol{x},\boldsymbol{y}) + O(\delta^3) \tag{3.11}$$

where $\boldsymbol{u^{\delta}}$ denotes the true displacement field, $\boldsymbol{u_0}$ denotes the macroscopic displacement field, $\boldsymbol{u_1}$, $\boldsymbol{u_2}$ and higher-order terms denote perturbations to the macroscopic displacement field, $\boldsymbol{x}$ and $\boldsymbol{y}$ denote the macro and micro level coordinates, and δ denotes the ratio of the micro to macro scale characteristic lengths ($\delta = \frac{\boldsymbol{x}}{\boldsymbol{y}}$). Derivatives of the periodic displacement field are

[1]Code written by Imediegwu (Imediegwu et al., 2019)

then derived through the application of the total derivative

$$\frac{D\boldsymbol{u}(\boldsymbol{x},\boldsymbol{y}(\boldsymbol{x}))}{D\boldsymbol{x}} = \frac{\partial \boldsymbol{u}(\boldsymbol{x},\boldsymbol{y}(\boldsymbol{x}))}{\partial \boldsymbol{x}} + \frac{\partial \boldsymbol{u}(\boldsymbol{x},\boldsymbol{y}(\boldsymbol{x}))}{\partial \boldsymbol{y}}\frac{\partial \boldsymbol{y}}{\partial \boldsymbol{x}} \tag{3.12}$$

which in this instance equals

$$\frac{D\boldsymbol{u}(\boldsymbol{x},\boldsymbol{y}(\boldsymbol{x}))}{D\boldsymbol{x}} = \frac{\partial \boldsymbol{u}(\boldsymbol{x},\boldsymbol{y}(\boldsymbol{x}))}{\partial \boldsymbol{x}} + \frac{1}{\delta}\frac{\partial \boldsymbol{u}(\boldsymbol{x},\boldsymbol{y}(\boldsymbol{x}))}{\partial \boldsymbol{y}} \tag{3.13}$$

as this enables an expression for the infinitesimal strain tensor (see Eq. 3.3) of a periodic displacement field to be derived

$$\begin{aligned}\epsilon_{ij}(\boldsymbol{u}) = \frac{1}{2}\Bigg[&\left(\frac{\partial u_{0i}}{\partial x_j} + \frac{\partial u_{0j}}{\partial x_i}\right) + \left(\frac{\partial u_{1i}}{\partial y_j} + \frac{\partial u_{1j}}{\partial y_i}\right) + \delta\left(\frac{\partial u_{1i}}{\partial x_j} + \frac{\partial u_{1j}}{\partial x_i}\right) + \delta\left(\frac{\partial u_{2i}}{\partial y_j} + \frac{\partial u_{2j}}{\partial y_i}\right) \\ &+\mathcal{O}(\delta^2)\Bigg]\end{aligned} \tag{3.14}$$

which by ignoring terms greater than $\mathcal{O}(\delta)$ can be simplified to

$$\epsilon_{ij} = \bar{\epsilon}_{ij} + \epsilon^*_{ij} \tag{3.15}$$

where $\bar{\epsilon}_{ij}$ denotes the macroscale strain tensor

$$\bar{\epsilon}_{ij} = \frac{1}{2}\left(\frac{\partial u_{0i}}{\partial x_j} + \frac{\partial u_{0j}}{\partial x_i}\right) \tag{3.16}$$

and where ϵ^*_{ij} denotes the fluctuating strain tensor

$$\epsilon^*_{ij} = \frac{1}{2}\left(\frac{\partial u_{1i}}{\partial y_j} + \frac{\partial u_{1j}}{\partial y_i}\right) \tag{3.17}$$

This simplified expression for the infinitesimal strain tensor can then be substituted into the linear elastic equilibrium equation (represented in weak variational form) which is denoted by

$$\int_\Omega E_{ijkl}\epsilon_{ij}(\boldsymbol{u})\epsilon_{kl}(\boldsymbol{\omega})\, d\Omega = \int_\Omega \boldsymbol{f}\boldsymbol{\omega}\, d\Omega + \int_\Gamma \boldsymbol{t}\boldsymbol{\omega}\, d\Gamma \tag{3.18}$$

which after substitution and neglecting body forces is equal to

$$\int_{\Omega^\delta} E_{ijkl}(\epsilon^0_{ij}(\boldsymbol{\omega}) + \epsilon^1_{ij}(\boldsymbol{\omega}))(\bar{\epsilon}_{kl}(\boldsymbol{u}) + \epsilon^*_{kl}(\boldsymbol{u}))\, d\Omega^\delta = \int_\Gamma \boldsymbol{t\omega}\, d\Gamma \tag{3.19}$$

where $\epsilon^0_{ij}(\boldsymbol{\omega})$ and $\epsilon^1_{ij}(\boldsymbol{\omega})$ denote the macroscale and microscale virtual displacements respectively, and Ω^δ represents both the microscale and macroscale domains. It is then permissible to assume a separation of scales ($\delta \to 0$), which enables the following relation to be leveraged (Hollister et al., 1991)

$$\lim_{\delta\to 0} \int_{\Omega^\delta} d\Omega^\delta = \int_\Omega \frac{1}{V} \int_V dV d\Omega \tag{3.20}$$

where V denotes the volume of a single unit cell and in combination with assuming that the virtual displacement varies only at the microscale ($\epsilon^0_{ij}(\boldsymbol{\omega}) = 0$) permits Eq. 3.19 to be reformulated as

$$\int_\Omega \frac{1}{V} \int_V E_{ijkl}\epsilon^1_{ij}(\boldsymbol{\omega})(\bar{\epsilon}_{kl}(\boldsymbol{u}) + \epsilon^*_{kl}(\boldsymbol{u}))\, dV d\Omega = 0 \tag{3.21}$$

where the right hand side is set equal to zero, as there are no tractions applied at the microscale. In order for Eq. 3.21 to be satisfied, it is only necessary for the interior integral to be resolved, which enables a simplification to be derived

$$\int_V E_{ijkl}\epsilon^1_{ij}(\boldsymbol{\omega})\epsilon^*_{kl}(\boldsymbol{u})\, dV = -\int_V E_{ijkl}\epsilon^1_{ij}(\boldsymbol{\omega})\bar{\epsilon}_{kl}(\boldsymbol{u})\, dV \tag{3.22}$$

By solving Eq. 3.22, it is possible to derive the fluctuating strain tensor for any arbitrary macroscale strain tensor. For each microstructure, Eq. 3.22 is solved six times using a unique macroscale unit strain tensor for each case. The microscale stress tensor is then derived for each case

$$\sigma_{ij} = 2\mu_L\epsilon_{ij} + \lambda_L\delta_{ij}\epsilon_{kk} \tag{3.23}$$

where δ_{ij} denotes the Kronecker delta function and where

$$\mu_L = \frac{E}{2(1+\nu)} \tag{3.24}$$

and

$$\lambda_L = \frac{E\nu}{(1+\nu)(1-2\nu)} \tag{3.25}$$

denote the Lamé constants. Subsequently, through the application of a volume-averaged representation of the constitutive law for linear elasticity, it is permissible to use each strain deformation analysis to derive the i^{th} row of the homogenised elasticity tensor

$$E_{ij}^H = \frac{1}{\bar{\epsilon}_j V} \int_V \sigma_i \, dV \tag{3.26}$$

which is represented using Voigt notation in this instance for clarity. Equation 3.22 must be solved in combination with periodic boundary conditions to ensure periodicity of the strain field. Periodic boundary conditions in this context are implemented by coupling the displacement of node sets on opposing faces and corners of the unit cube, which is trivial to implement due to the fixed cubic mesh. Extensive verification of the outputs of the microscale model has been performed and is presented in reference to analytical models within Imediegwu (2020).

3.3 Macroscale Model

The macroscale model is constructed using an assembly of N microstructures, individually defined by a row vector within the design variable matrix

$$X_{\alpha\beta} = \begin{bmatrix} \hat{r}_{1,1} & \hat{r}_{2,1} & \cdots & \hat{r}_{7,1} & e_{5,1} & e_{6,1} & e_{7,1} \\ \hat{r}_{1,2} & \hat{r}_{2,2} & \cdots & \hat{r}_{7,2} & e_{5,2} & e_{6,2} & e_{7,2} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ \hat{r}_{1,N} & \hat{r}_{2,N} & \cdots & \hat{r}_{7,N} & e_{5,N} & e_{6,N} & e_{7,N} \end{bmatrix} \tag{3.27}$$

which has dimensions $N \times 10$. Microstructures are represented within the macroscale finite-element model as singular points with homogenised mechanical properties (E_α^H), as depicted in Fig. 3.5. Similarly to the microscale model, the evaluation of the macroscale finite-element model is also computed using FEniCS. (Logg et al., 2012).

This multiscale finite-element approach is representative of the true macroscopic response of the system, so long as a separation of scales exists between the micro and macro scales (as defined by asymptotic homogenisation theory (see Sec. 2.3)). A recent numerical study presented by Cheng et al. (2018), indicated that an order of magnitude is sufficient to achieve a separation of scales.

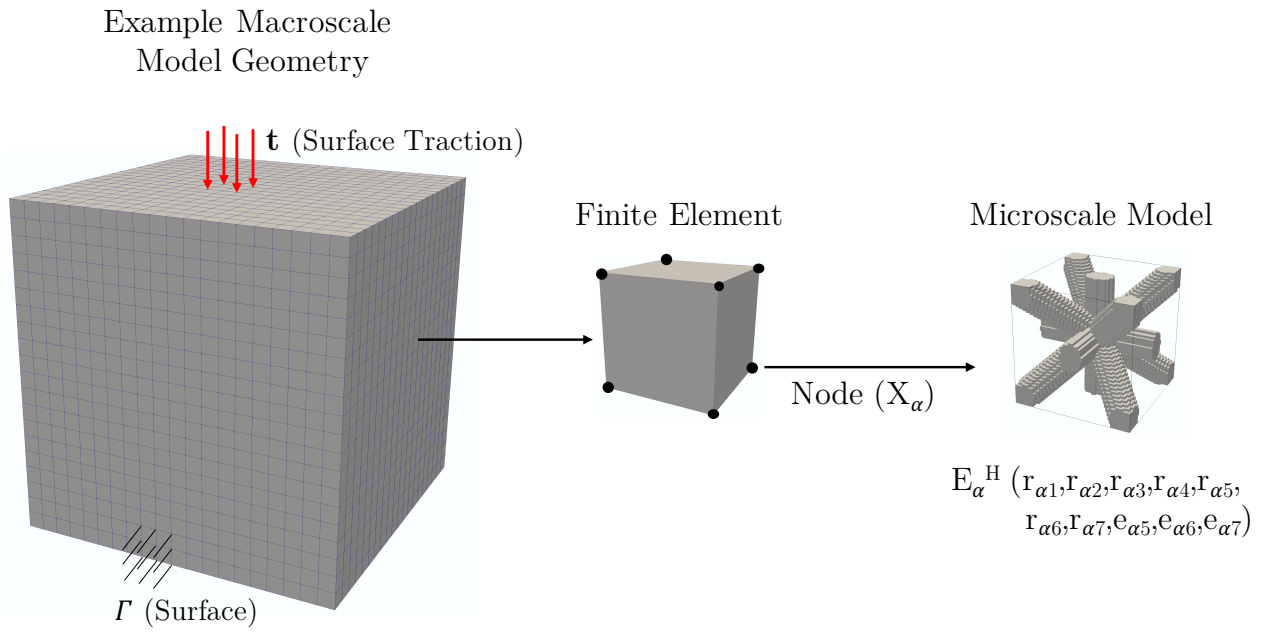


Figure 3.5: The multiscale finite-element model used to evaluate the macroscopic response of a structure

3.4 Concurrent Coupling

The microscale model is coupled to the macroscale model concurrently during optimisation, representing the principal novelty of this framework. This permits only the mechanical property data required to evaluate the macroscale model during each iteration of optimisation to be collected. Additionally, the data collected during optimisation is stored in a re-usable database, for the benefit of successive iterations of optimisation, or entirely new optimisation problems.

To facilitate this approach, in advance of the first optimisation being undertaken, a full factorial Design Of Experiments (DOE)(Antony, 2014) is generated for the microscale parameters,

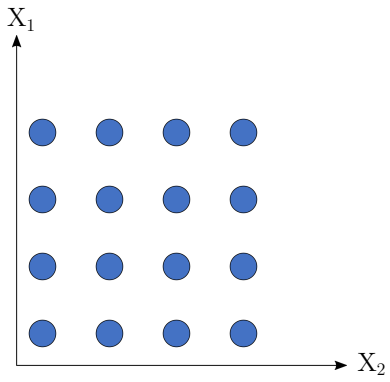
but not populated with simulation data. Nodes within the DOE represent unique permutations of the microscale parameters and in combination span the entire permissible range of design space. The bounds set for the normalised radii ($\hat{r}_{\alpha,1-4} = 0.42, \hat{r}_{\alpha,5-7} = 0.49$) enable a broad range of mechanical properties to be obtained and permit both empty and fully dense unit cells to be configured. A higher normalised radius bound is afforded to each face member, in comparison to the diagonal members, to compensate for the loss in cross-sectional area resulting from the configurable elliptical cross-sections. The maximum aspect-ratio for an elliptical cross-section is set to 5/3 ($AR_{max} = 0.6$), to prevent highly-eccentric members forming in either orthogonal axis. The DOE is constructed using the microscale parameter bounds and level intervals defined in Tab. 3.1. It is important to note, that fewer levels are assigned to the aspect-ratio variables, as smaller variations are noted in the mechanical properties for equivalent changes in the design variables in comparison to the normalised radii. In combination, the DOE contains $153,664,000$ unique permutations of the microscale parameters, which is several orders of magnitude larger than the size of the databases assembled within contemporary sequential frameworks (Imediegwu et al., 2019; Colabella et al., 2019).

Microscale Parameters	No. Levels	Lowest Level	Highest Level	Level Increment
$\hat{r}_{1-4}$	7	0.00	0.42	0.07
$\hat{r}_{5-7}$	8	0.00	0.49	0.07
e_{5-7}	5	0.60	1.40	0.20

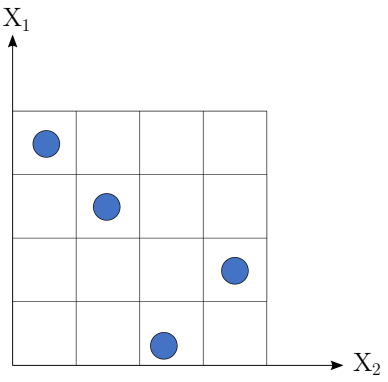
Table 3.1: The individual microscale parameter bounds and level intervals used to generate the full factorial DOE. In combination, this results in a total of $153,664,000$ unique permutations of the microscale parameters.

A full factorial DOE (see Fig. 3.6a) is selected as the foundation of the concurrent coupling methodology, as opposed to other less structured techniques, such as Latin-hypercube sampling (McKay et al., 1979) (see Fig. 3.6b) or orthogonal sampling (Burrage et al., 2015) (see Fig. 3.6c), as the uniformly distributed nodes afford a number of benefits (explained in detail within Secs. 3.4.1 and 3.5.5). It is important to note, that whilst Latin-hypercube sampling and orthogonal sampling typically cover more design space for a fixed number of points in comparison to a full factorial DOE (Alam et al., 2004), the uniformity in the distribution of

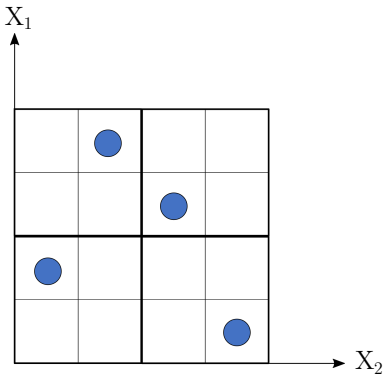
nodes can be leveraged in this instance to vastly reduce the computational cost associated with microscale simulations. This counter-intuitive benefit highlights the atypical application of a DOE within the context of this work.



(a) An example two-dimensional full factorial DOE



(b) An example two-dimensional Latin-hypercube sampling plan



(c) An example two-dimensional orthogonal sampling plan

Figure 3.6: Examples DOEs are presented using full factorial, Latin-hypercube sampling and orthogonal sampling plans. The boxes depicted within Figs. 3.6b and 3.6c denote bins, in which only a single point can allocated within for each row and column.

3.4.1 Material Property Transformations

The regular arrangement of nodes within the DOE permits rotational and reflective symmetries of the microscale geometry to be exploited, enabling the mechanical properties of other geometrically similar microstructures to be derived from a single parent microscale simulation. Sets of geometrically similar microstructures are iteratively derived by applying unique

combinations of three-dimensional rotational matrices

$$R_x(\theta = 90^\circ) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix}, \; R_y(\theta = 90^\circ) = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ -1 & 0 & 0 \end{bmatrix}, \; R_z(\theta = 90^\circ) = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{3.28}$$

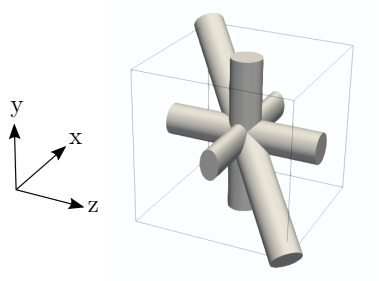
where $R_x(\theta = 90^\circ)$, $R_y(\theta = 90^\circ)$ and $R_z(\theta = 90^\circ)$ denote 90° rotations about the x, y and z axes respectively, and three-dimensional reflection matrices

$$R_{xy} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{bmatrix}, \; R_{xz} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \; R_{yz} = \begin{bmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{3.29}$$

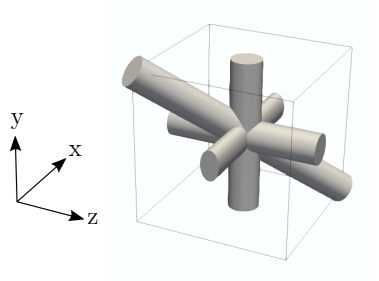
where R_{xy}, R_{xz} and R_{yz} denote reflections about the $X-Y$, $X-Z$ and $Y-Z$ planes respectively. In combination and including the identity matrix, a total of 64 transformations exist of which 48 are unique. The 48 unique transformation matrices are then applied individually to a three-dimensional tensor, ($[20 \times 3 \times 3]$) which represents a generalised skeleton of the microscale geometry using the minimum number of vertices (two per diagonal member and four per face-member to account for aspect-ratio) in order to determine how many of the 48 unique transformation matrices uniquely permute the design variable vector (X_α). A total of 24 unique design variable vector permutations exist, which in some cases enables the mechanical properties of up to 24 geometrically similar microstructures to be derived from a single parent microscale simulation. An example set of geometrically similar microstructures is presented in Fig. 3.7 using the present microscale description.

The mechanical properties of a child microstructure ($E^H_{ij,\text{child}}$) is derived using the mechanical properties of the parent microstructure ($E^H_{ij,\text{parent}}$) and the corresponding transformation matrix

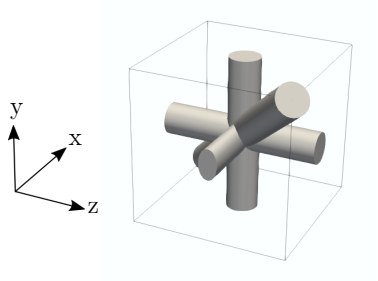
$$T = \begin{bmatrix} t_{11} & t_{12} & t_{13} \\ t_{21} & t_{22} & t_{23} \\ t_{31} & t_{32} & t_{33} \end{bmatrix} \tag{3.30}$$



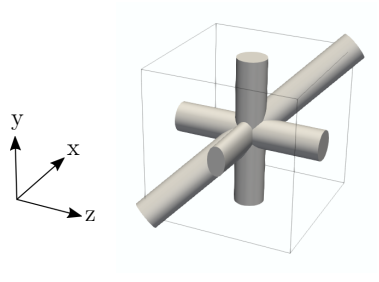
(a) Parent microstructure
$(0.1, 0, 0, 0, 0.1, 0.1, 0.1, 0.6, 1, 1)$



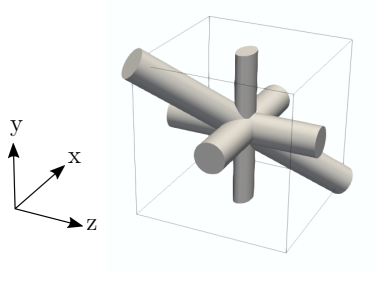
(b) Child microstructure, Transformation R_{yz}
$(0, 0, 0.1, 0, 0.1, 0.1, 0.1, 0.6, 1, 1)$



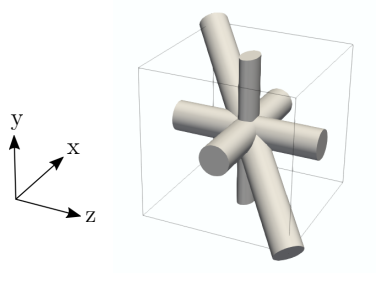
(c) Child microstructure — Transformation R_{xz}
$(0, 0.1, 0, 0, 0.1, 0.1, 0.1, 0.6, 1, 1)$



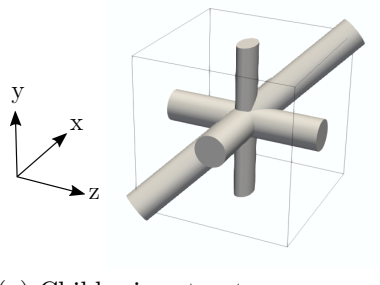
(d) Child microstructure — Transformation $R_{xz}R_{yz}$
$(0, 0, 0, 0.1, 0.1, 0.1, 0.1, 0.6, 1, 1)$



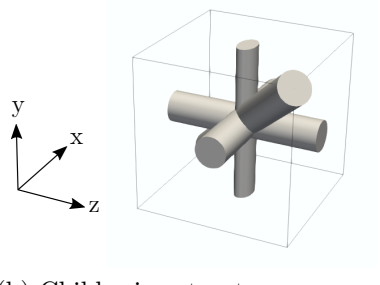
(e) Child microstructure — Transformation R_z
$(0, 0, 0.1, 0, 0.1, 0.1, 0.1, 1, 0.6, 1)$



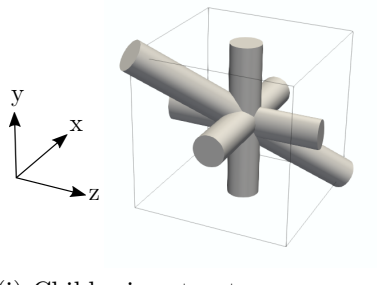
(f) Child microstructure — Transformation R_zR_{yz}
$(0.1, 0, 0, 0, 0.1, 0.1, 0.1, 1, 0.6, 1)$



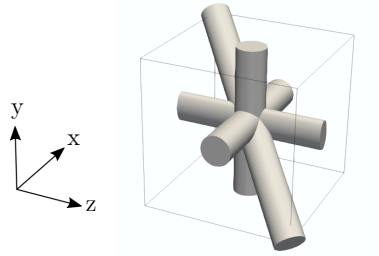
(g) Child microstructure — Transformation R_zR_{xz}
$(0, 0, 0, 0.1, 0.1, 0.1, 0.1, 1, 0.6, 1)$



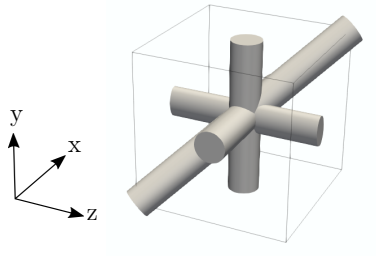
(h) Child microstructure — Transformation $R_zR_{xz}R_{yz}$
$(0, 0.1, 0, 0, 0.1, 0.1, 0.1, 1, 0.6, 1)$



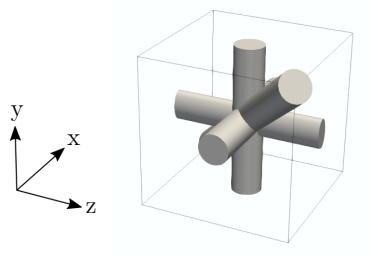
(i) Child microstructure — Transformation R_y
$(0, 0, 0.1, 0, 0.1, 0.1, 0.1, 1, 1, 0.6)$



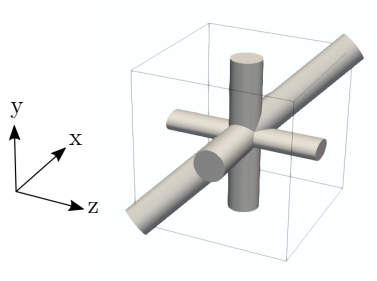
(j) Child microstructure — Transformation R_yR_{yz}
$(0.1, 0, 0, 0, 0.1, 0.1, 0.1, 1, 1, 0.6)$



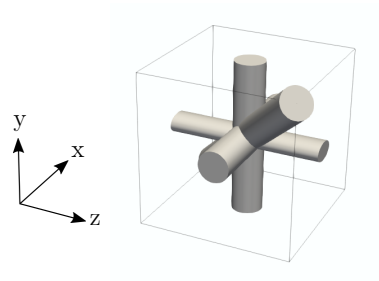
(k) Child microstructure — Transformation R_yR_{xz}
$(0, 0, 0, 0.1, 0.1, 0.1, 0.1, 1, 1, 0.6)$



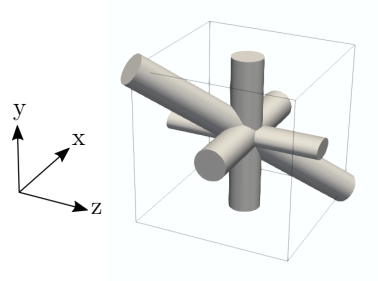
(l) Child microstructure — Transformation $R_yR_{xz}R_{yz}$
$(0, 0.1, 0, 0, 0.1, 0.1, 0.1, 1, 1, 0.6)$



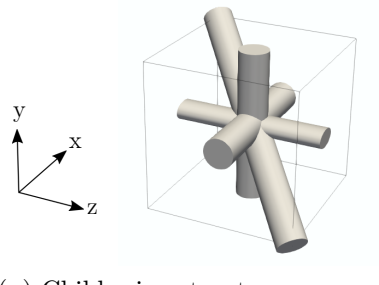
(m) Child microstructure, Transformation $R_y R_z$ $(0, 0, 0, 0.1, 0.1, 0.1, 0.1, 1, 1, 1.4)$



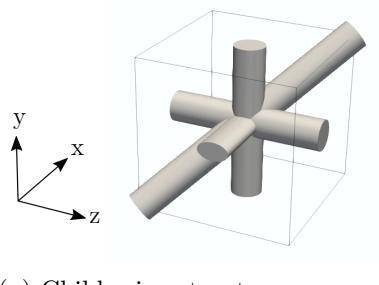
(n) Child microstructure, Transformation $R_y R_z R_{yz}$ $(0, 0.1, 0, 0, 0.1, 0.1, 0.1, 1, 1, 1.4)$



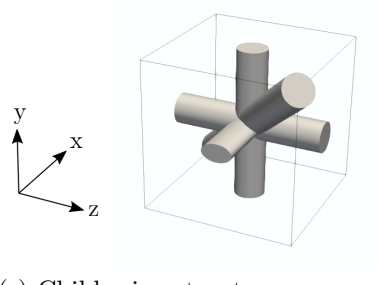
(o) Child microstructure — Transformation $R_y R_z R_{xz}$ $(0, 0, 0.1, 0, 0.1, 0.1, 0.1, 1, 1, 1.4)$



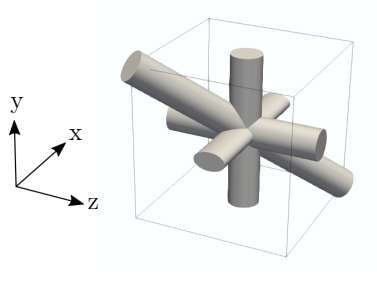
(p) Child microstructure — Transformation $R_y R_z R_{xz} R_{yz}$ $(0.1, 0, 0, 0, 0.1, 0.1, 0.1, 1, 1, 1.4)$



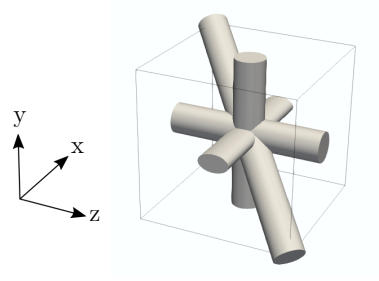
(q) Child microstructure — Transformation R_x $(0, 0, 0, 0.1, 0.1, 0.1, 0.1, 1.4, 1, 1)$



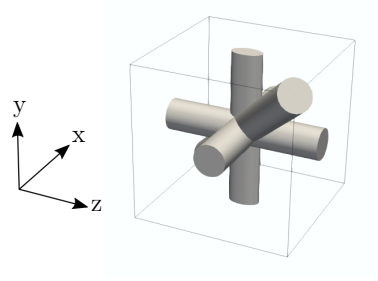
(r) Child microstructure — Transformation $R_x R_{yz}$ $(0, 0.1, 0, 0, 0.1, 0.1, 0.1, 1.4, 1, 1)$



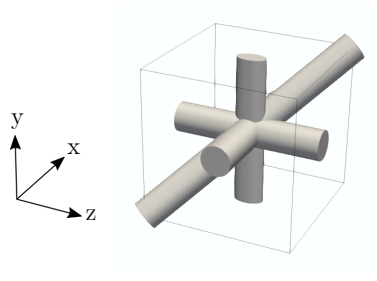
(s) Child microstructure — Transformation $R_x R_{xz}$ $(0, 0, 0.1, 0, 0.1, 0.1, 0.1, 1.4, 1, 1)$



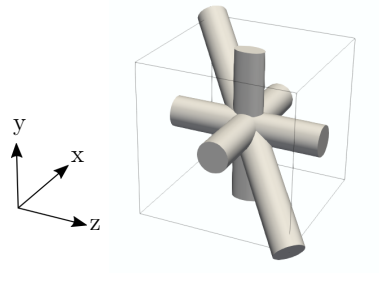
(t) Child microstructure — Transformation $R_x R_{xz} R_{yz}$ $(0.1, 0, 0, 0, 0.1, 0.1, 0.1, 1.4, 1, 1)$



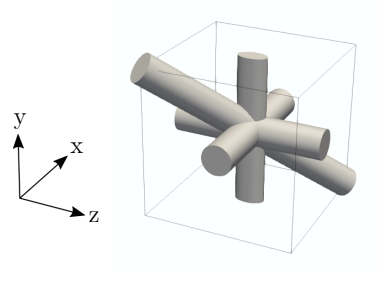
(u) Child microstructure — Transformation $R_x R_z$ $(0, 0.1, 0, 0, 0.1, 0.1, 0.1, 1, 1.4, 1)$



(v) Child microstructure — Transformation $R_x R_z R_{yz}$ $(0, 0, 0, 0.1, 0.1, 0.1, 0.1, 1, 1.4, 1)$



(w) Child microstructure — Transformation $R_x R_z R_{xz}$ $(0.1, 0, 0, 0, 0.1, 0.1, 0.1, 1, 1.4, 1)$



(x) Child microstructure — Transformation $R_x R_z R_{xz} R_{yz}$ $(0, 0, 0.1, 0, 0.1, 0.1, 0.1, 1, 1.4, 1)$

Figure 3.7: An illustration of how combinations of rotation and reflection matrices can be used to derive geometrically similar microstructures (children). In this example, it is possible to derive the mechanical properties of 24 microstructures from a single parent simulation.

using the following relation (Tsai, 1966)

$$E^H_{ij,\text{child}} = M(T) E^H_{ij,\text{parent}} M(T)^{-1} \tag{3.31}$$

where $M(T)$ denotes the block transformation matrix

$$M(T) = \begin{bmatrix} A(T) & B(T) \\ C(T) & D(T) \end{bmatrix} \tag{3.32}$$

where

$$A(T) = \begin{bmatrix} t_{11}^2 & t_{21}^2 & t_{31}^2 \\ t_{12}^2 & t_{22}^2 & t_{32}^2 \\ t_{13}^2 & t_{23}^2 & t_{33}^2 \end{bmatrix} \tag{3.33}$$

$$B(T) = \begin{bmatrix} 2t_{21}t_{31} & 2t_{31}t_{11} & 2t_{11}t_{21} \\ 2t_{22}t_{32} & 2t_{32}t_{12} & 2t_{12}t_{22} \\ 2t_{23}t_{33} & 2t_{33}t_{13} & 2t_{13}t_{23} \end{bmatrix} \tag{3.34}$$

$$C(T) = \begin{bmatrix} t_{12}t_{13} & t_{22}t_{23} & t_{32}t_{33} \\ t_{13}t_{11} & t_{23}t_{21} & t_{33}t_{31} \\ t_{11}t_{12} & t_{21}t_{22} & t_{31}t_{32} \end{bmatrix} \tag{3.35}$$

$$D(T) = \begin{bmatrix} t_{22}t_{33} + t_{23}t_{32} & t_{32}t_{13} + t_{33}t_{12} & t_{22}t_{13} + t_{23}t_{12} \\ t_{23}t_{31} + t_{21}t_{33} & t_{33}t_{11} + t_{31}t_{13} & t_{23}t_{11} + t_{21}t_{13} \\ t_{21}t_{32} + t_{22}t_{31} & t_{31}t_{12} + t_{32}t_{11} & t_{21}t_{12} + t_{22}t_{11} \end{bmatrix} \tag{3.36}$$

Through the application of these rotational and reflective symmetries, it is possible to derive the homogenised mechanical properties for all $153,664,000$ microstructures within the database using only $6,823,320$ calls of the microscale model.

3.4.2 Piecewise Continuous Linear Interpolation

In sequential parametrisation-based frameworks (Imediegwu et al., 2019; Colabella et al., 2019), the homogenised mechanical properties for all $6,823,320$ parent microstructures would be collected prior to optimisation, to facilitate the assembly of a database. In contrast, only the subset of microstructures which enable the evaluation of the macroscale model during each iteration of optimisation is collected within the present framework. To derive this subset, microstructures within the DOE which form enclosing simplices (n-dimensional triangles) around the individual microscale parameter sets (X_α) are computed first

$$\Phi = \bigcup_{\alpha=1}^{N} \phi_\alpha \tag{3.37}$$

where Φ denotes the set of microstructures forming enclosing simplices, ϕ_α denotes the $[n+1]$ microstructures forming a single enclosing simplex around the parameter set X_α. The algorithm used to determine which microstructures form the enclosing simplex is presented in detail within Appendix A. A further subset (Λ) is subsequently derived, which contains the parent microstructures for all microstructures within Φ

$$\Lambda = \bigcup_{m=1}^{N_\Phi} \psi(\Phi_m) \tag{3.38}$$

where N_Φ denotes the number of elements within the set Φ and the function ψ denotes a child to parent mapping function, which is derived in advance of optimisation by leveraging the reflective and rotational symmetries detailed in Sec. 3.4.1. The DOE is then examined to determine which microstructures within Λ have already been collected in previous iterations of optimisation and which microstructures require simulation. After the mechanical properties for all microstructures within Λ have been collected, material property transformation matrices (see Eq. 3.32) are applied to the parent microstructures, to derive the mechanical properties for all of the microstructures within Φ. To reduce the storage requirements for the database, transformations are applied on-the-fly during optimisation for a negligible computational expense, so that only the mechanical properties of parent microstructures are stored.

A linear interpolant is then derived from each simplex for all 21 unique elements of the homogenised elasticity tensor (E^H_α) and also for volume-fraction (V_α), as a function of the continuous microscale parameters using simplex interpolation (Hemingway, 2002)

$$\begin{aligned} E^H_\alpha(X_\alpha) = [&\mathrm{vol}(\phi_{\alpha,2}, \phi_{\alpha,3}, \phi_{\alpha,4}, \ldots, \phi_{\alpha,n+1}, X_\alpha) E^H_{\phi_{\alpha,1}} \\ &+ \mathrm{vol}(\phi_{\alpha,1}, \phi_{\alpha,3}, \phi_{\alpha,4}, \ldots, \phi_{\alpha,n+1}, X_\alpha) E^H_{\phi_{\alpha,2}} \\ &+ \mathrm{vol}(\phi_{\alpha,1}, \phi_{\alpha,2}, \phi_{\alpha,4}, \ldots, \phi_{\alpha,n+1}, X_\alpha) E^H_{\phi_{\alpha,3}} \\ &\vdots \\ &+ \mathrm{vol}(\phi_{\alpha,1}, \phi_{\alpha,2}, \phi_{\alpha,3}, \ldots, \phi_{\alpha,n}, X_\alpha) E^H_{\phi_{\alpha,n+1}}] \\ &/ \mathrm{vol}(\Phi_\alpha) \end{aligned} \tag{3.39}$$

where $\phi_{\alpha,n}$ denotes the n^{th} microstructure used in the construction of the enclosing simplex ϕ_α and 'vol' denotes the function

$$\mathrm{vol}(\Phi_\alpha) = \frac{1}{n!} \left| \det \begin{bmatrix} \phi_{\alpha,2} - \phi_{\alpha,1} \\ \phi_{\alpha,3} - \phi_{\alpha,1} \\ \vdots \\ \phi_{\alpha,n} - \phi_{\alpha,1} \\ \phi_{\alpha,n+1} - \phi_{\alpha,1} \end{bmatrix} \right| \tag{3.40}$$

which is used to calculate the volume enclosed by a simplex (Stein, 1966). This approach is equivalent to modelling the response of the microscale model using a piecewise linear basis, which affords C^0 continuity between adjacent simplicial interpolants. As a result of using a local interpolation based response surface, the evaluation of the microscale parameter space is invariant to the population of additional points within the microscale parameter space, which is in contrast to regression based response surfaces. Consequently, in this instance, local regression based response surfaces could lead to non-convergent solutions as the mechanical properties oscillate.

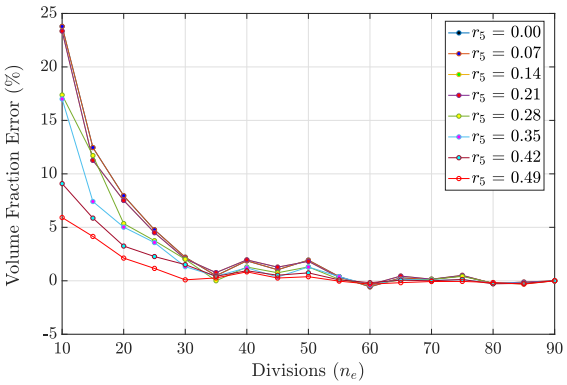
Local interpolation based response surfaces also afford numerous benefits in this context in comparison to global response surface methodologies, such as least-squares polynomial regression (Conte and de Boor, 2017), Kriging (Forrester et al., 2008) and radial basis functions (Buhmann, 2003). The primary benefit is associated with the boundedness of an interpolant, which ensures that the evaluation of the mechanical property space remains physical in all regions of the microscale design space. This is in contrast to the MTO frameworks presented by Imediegwu et al. (2019); Thillaithevan et al. (2021); Nightingale et al. (2021), where it is noted that the sixth-order seven-dimensional polynomials underpredict the extension-extension terms within the elasticity tensor for low volume-fraction microstructures, giving rise to non-physical negative terms. The converse problem is also noted at higher volume-fractions, where the predicted stiffnesses are greater than those exhibited by the base material. It is also possible for least-squares polynomial equations to be over-fitted, which can lead to further errors in the predicted mechanical properties. Linear interpolation based response surfaces can also lead to the under/over prediction of the mechanical property space in between data points. Although, in the limit of an infinitely large microscale database, the predicted values tend to the true values. This is in contrast to least-squares polynomial regression, where the addition of data points does not necessarily improve the accuracy of predictions, or for Kriging where additional points can lead to computationally intractable matrix algebra.

Simplex interpolation is favoured over equivalent rectangular interpolation-based approaches, as fewer simulation nodes are required to derive an interpolant; $[n + 1]$ nodes are required to construct an enclosing simplex, whereas 2^n nodes are required to construct an enclosing hypercube. This is due to the fact that a congruent simplex constructed from the nodes of a unit hypercube, resolves $1/n!$ times the volume of the unit hypercube. However, there is no guarantee that the additional design space enclosed by the hypercube will be used during optimisation. Therefore, as microscale simulations are computationally expensive, it is more efficient to resolve the minimum amount of design space during optimisation.

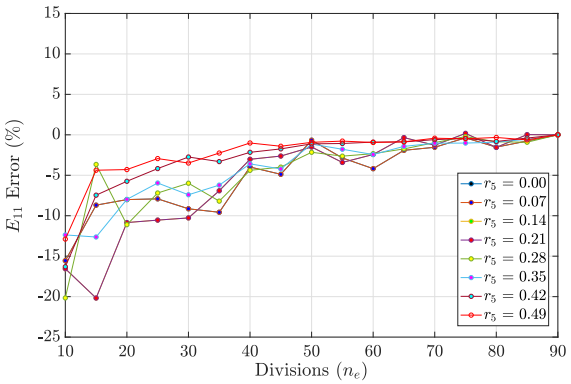
3.4.3 Microscale Model — Mesh Convergence

The computational resources available throughout the duration of this research project placed a limit on both the fidelity of microscale simulations and the number of levels available to use in the definition of each dimension within the full-factorial DOE. It is important to note, that the high performance computing facilities provided by Imperial College London did not facilitate the implementation of concurrent MTO frameworks, as it is not permissible to queue microscale simulations on-the-fly during the optimisation procedure. As a consequence, subsequent calculations presented within this section are based on the unlimited availability of eight Intel Xeon E5-2650 v4 at 2.20GHz CPUs, which represents the largest computational resource available to implement the framework presented within this thesis.

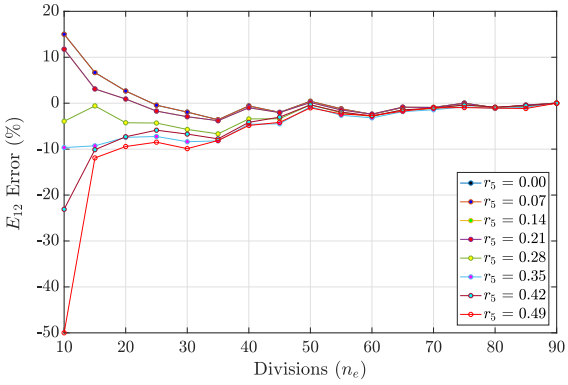
Mesh convergence studies are performed in order to determine an optimal balance between the accuracy of the homogenised mechanical properties derived through the microscale model and the associated computational expense. Mesh refinement in this instance is performed by increasing the number of elements used in the discretisation of the hexahedral cubic mesh, which is parametrised by the number of divisions per unit length (n_e). To sufficiently demonstrate the convergence properties of the microscale model, three cases are presented in Figs. 3.8, 3.9 and 3.10 respectively. In addition, the average computational expenditure incurred for all cases presented within this section is also presented within Fig. 3.11.



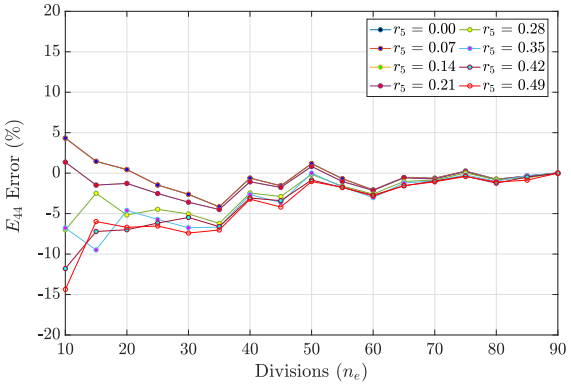
(a) Convergence plot for volume fraction



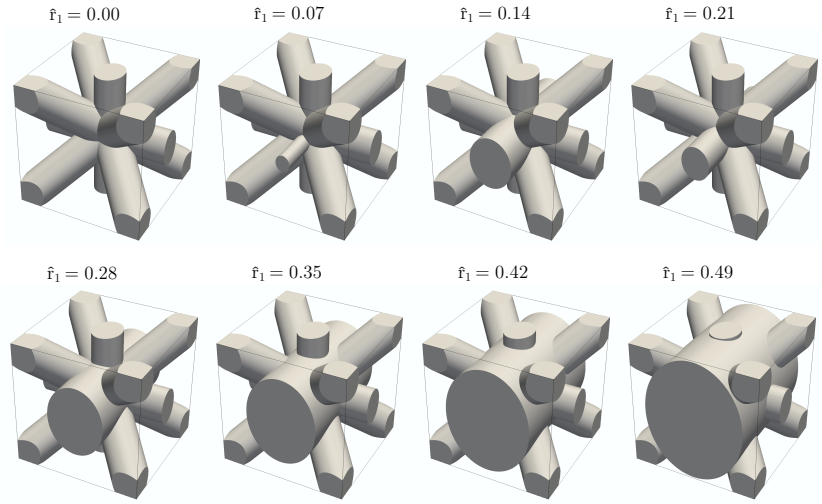
(b) Convergence plot for E_{11}



(c) Convergence plot for E_{12}

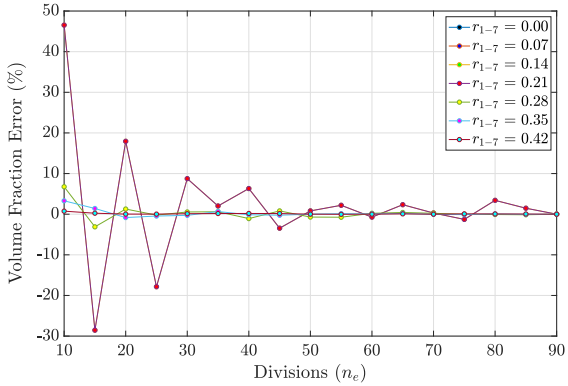


(d) Convergence plot for E_{44}

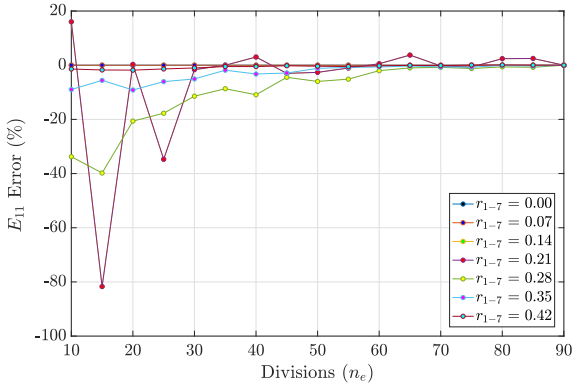


(e) Microstructures investigated

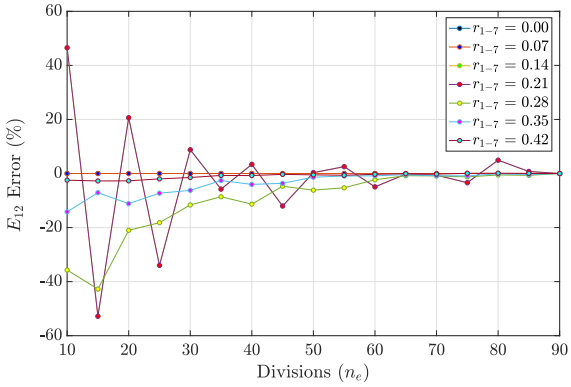
Figure 3.8: Results of a mesh convergence study where structural member $\hat{r}_5$ is incrementally increased from a normalised radius of 0.00 to 0.49 whilst all other members are kept constant with a normalised radius of 0.14. The aspect-ratio variables for all face-members is kept constant at $e_{5,6,7} = 0.8$.



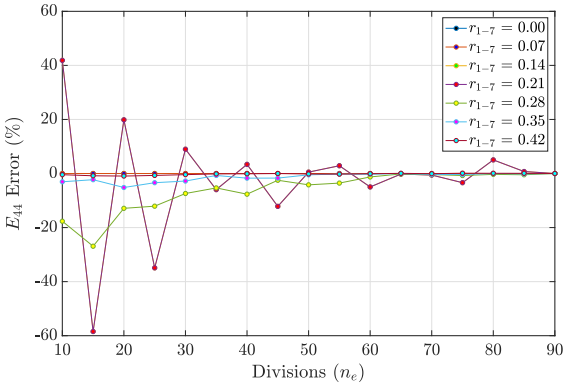
(a) Convergence plot for volume fraction



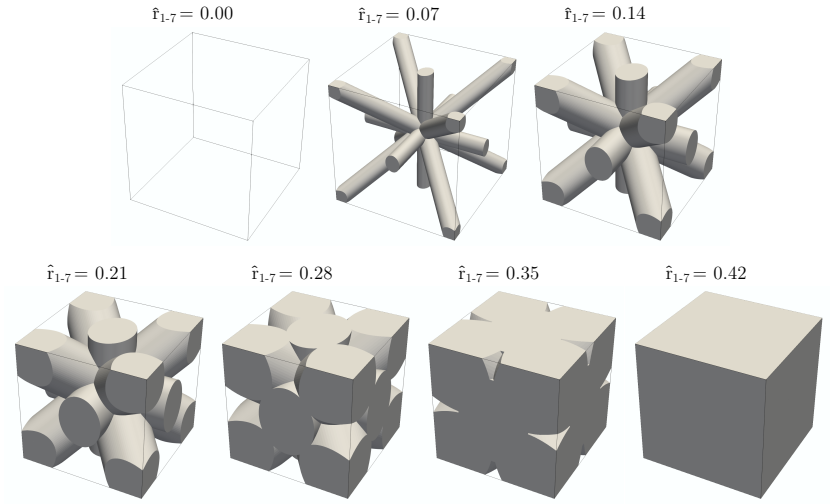
(b) Convergence plot for E_{11}



(c) Convergence plot for E_{12}

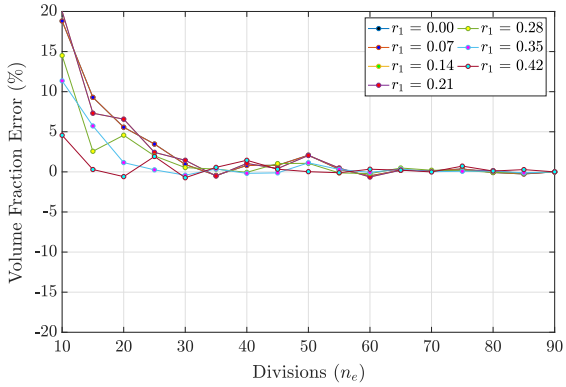


(d) Convergence plot for E_{44}

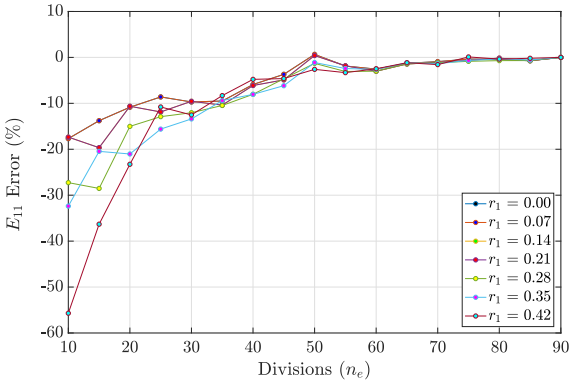


(e) Microstructures investigated

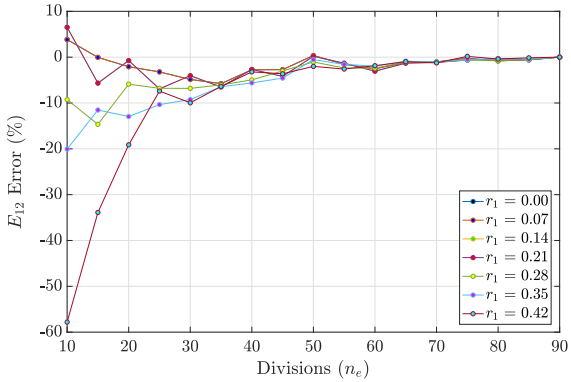
Figure 3.9: Results of a mesh convergence study where structural members $\hat{r}_{1-7}$ are incrementally increased from a normalised radius of 0.00 to 0.42. The aspect-ratio variables for all face-members is kept constant at $e_{5,6,7} = 0.8$.



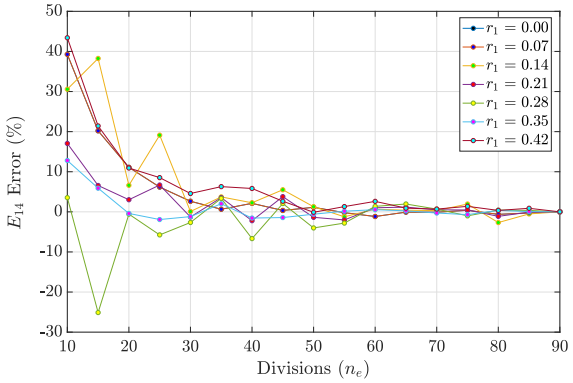
(a) Convergence plot for volume fraction



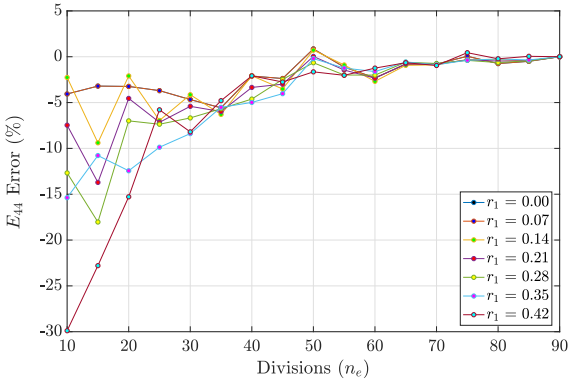
(b) Convergence plot for E_{11}



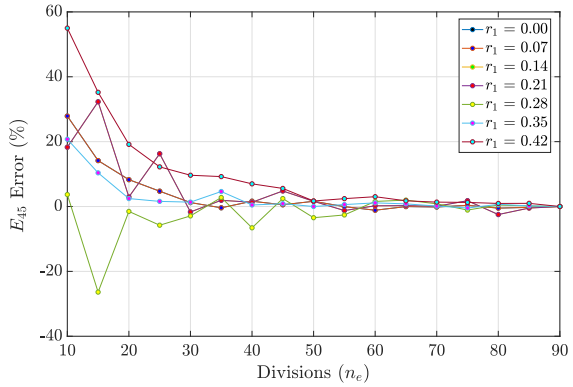
(c) Convergence plot for E_{12}



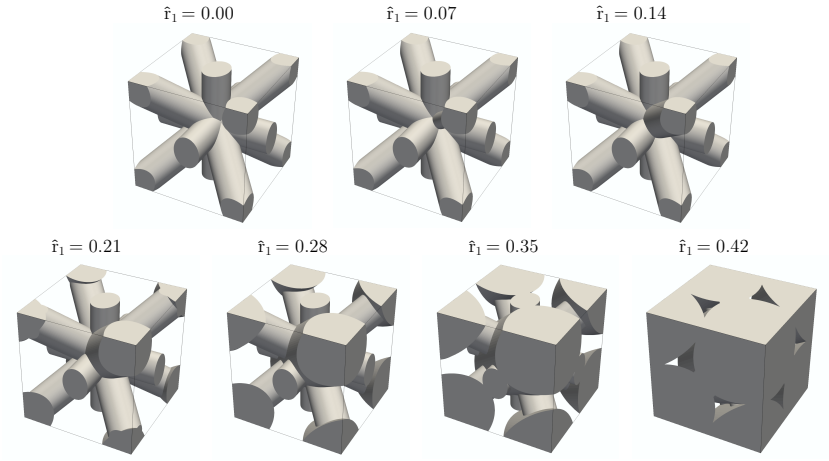
(d) Convergence plot for E_{14}



(e) Convergence plot for E_{44}



(f) Convergence plot for E_{45}



(g) Microstructures investigated

Figure 3.10: Results of a mesh convergence study where structural member $\hat{r}_1$ is incrementally increased from a normalised radius of 0.00 to 0.42 whilst all other members are kept constant with a normalised radius of 0.14. The aspect-ratio variables for all face-members is kept constant at $e_{5,6,7} = 0.8$.

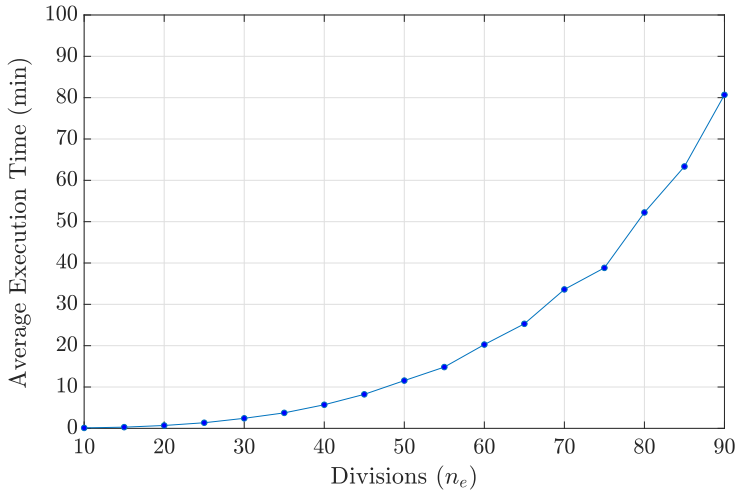


Figure 3.11: The average computation time required to execute the microscale model for each discretisation level presented within Figs. 3.8, 3.9 and 3.10 using a single Intel(R) Xeon(R) CPU E5-2650 v4 at 2.20GHz

It is clearly illustrated within Figs. 3.8, 3.9 and 3.10 that for the 22 microstructures investigated, convergence of the homogenised mechanical properties is obtained at approximately $n_e = 70$ ($< 1\%$ error). The associated computational expense of this discretisation is determined by calculating the computational time required to derive 2% ($\sim 136,000$) of the parent microscale database, which represents a conservative estimate derived from preliminary simulations for the number of microscale simulations required to perform a single optimisation. The required computational time is derived using

$$T_s = \frac{n_s t_s}{n_{cpu}} \tag{3.41}$$

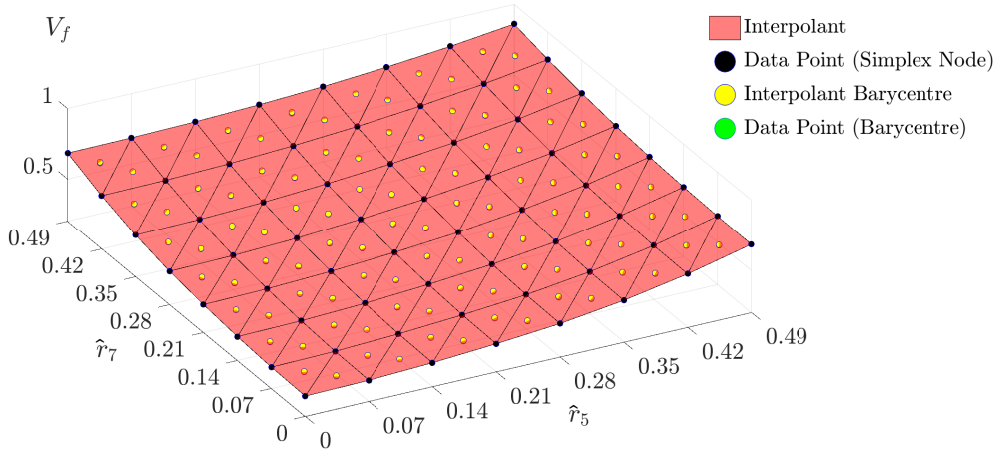
where T_s denotes the total computation time, n_s denotes the number of microscale simulations, t_s denotes the computation time required to perform a single microscale simulation (see Fig. 3.11) and n_{cpu} denotes the number of processing units available to perform computation (eight). This conservative estimate predicts that approximately 400 days of continuous computation are required to derive 2% of the parent microscale database using a discretisation resolution of $n_e = 70$. This represents an intractable length of time given the time-scales associated with this research project. As a result, a discretisation resolution of $n_e = 50$ is chosen instead, as this enables a satisfactory balance between accuracy ($< 5\%$ error) and computation time to be derived ($\sim$128 days).

3.4.4 Mechanical Property Space Visualisation

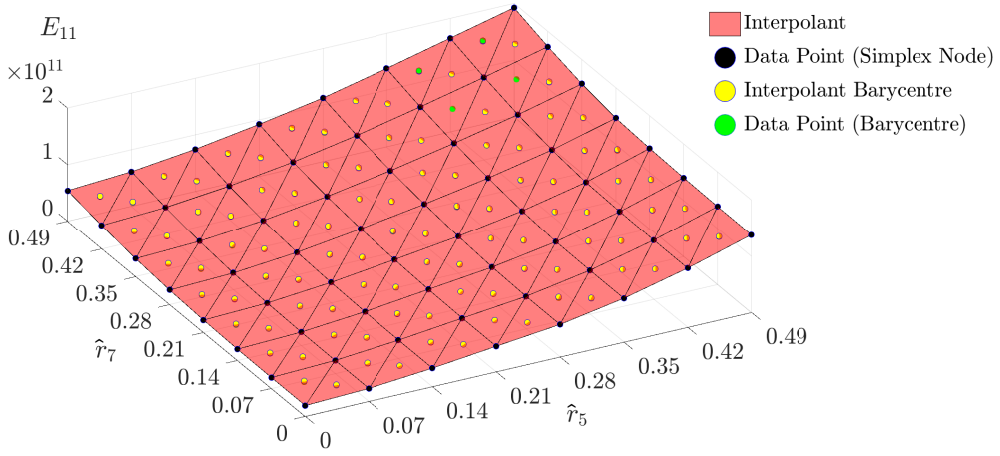
In order to demonstrate the efficacy of the piecewise continuous response surface detailed within Sec. 3.4.2 and the number of levels used for each dimension within the full-factorial DOE (see Tab. 3.1), two-dimensional hyperplanes through the mechanical property space are illustrated for three distinct cases within Figs. 3.12, 3.13 and 3.14. Two-dimensional hyperplanes are used for visualisation, as it is impossible to succinctly communicate this information in higher-dimensional spaces. In order to convey the accuracy of the simplicial interpolants presented within Figs. 3.12, 3.13 and 3.14, the mechanical properties at the barycentre of each simplex are evaluated using both the local interpolant and an additional simulation of the microscale model and are presented alongside the two-dimensional hyperplanes. In addition, the average Root Mean Squared Error (RMSE) for each case is presented within Tab. 3.2 to provide a quantitative measure of error.

It is important to highlight two sources of error associated with the derivation of the RMSE in this context. As a consequence of using a fixed cubic mesh to represent the microscale geometry, it is impermissible to represent features significantly smaller than the characteristic discretisation length (n_e). As a result, a systematic error is introduced, as it is not possible to precisely represent the radius of a structural member given the present discretisation. In addition, as detailed within Sec. 3.4.3, the selected discretisation ($n_e = 50$) of the microscale model enables homogenised mechanical properties within 5% of the converged values to be derived, which therefore introduces a random error in this instance.

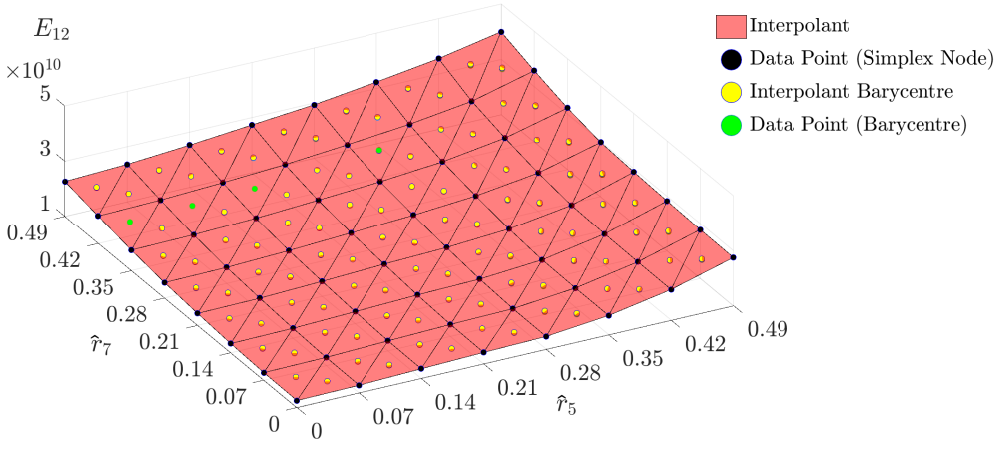
It is clearly illustrated within Figs. 3.12, 3.13 and 3.14 that the piecewise simplicial interpolants combine to form a smooth representation of the mechanical property space. In addition, it is evident within Fig. 3.14 that variations in the aspect-ratio variables lead to smaller changes in the mechanical properties, which enables fewer DOE levels to be used in comparison to the normalised radii. However, large errors up to 20% are noted in Fig. 3.13 between the interpolated and simulated values for large normalised radius values. This error



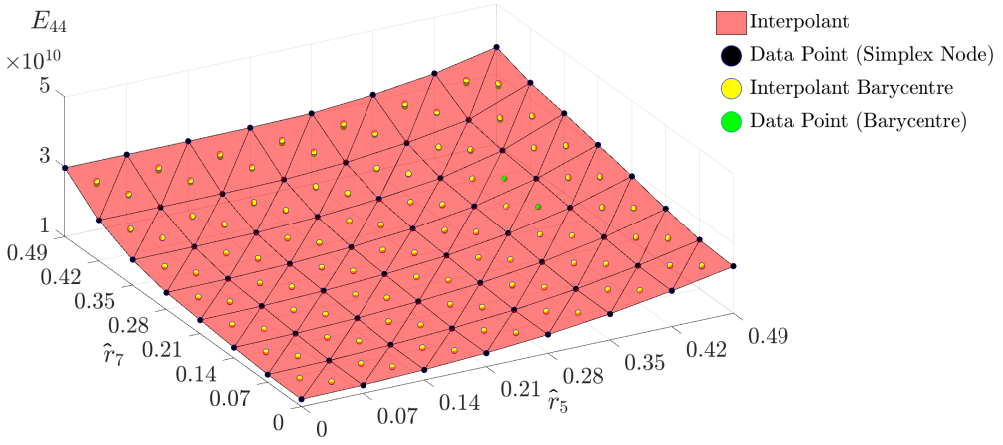
(a) Volume fraction



(b) E_{11}

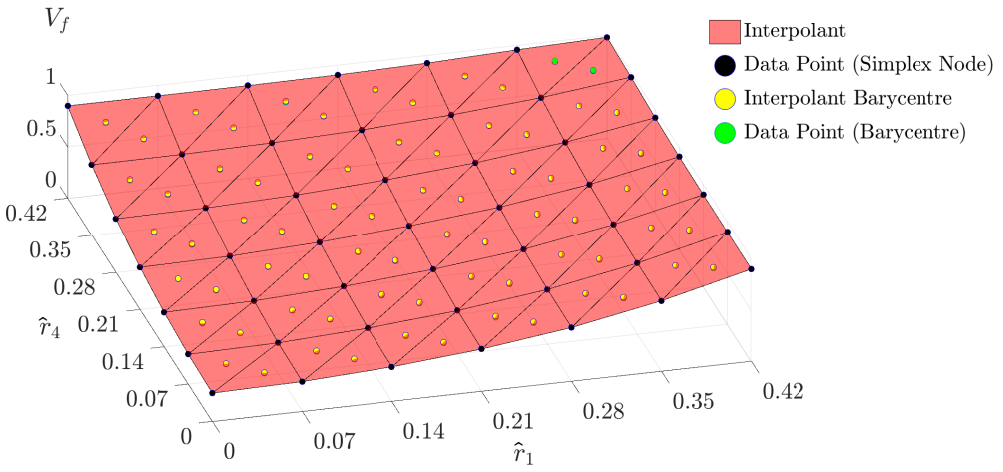


(c) E_{12}

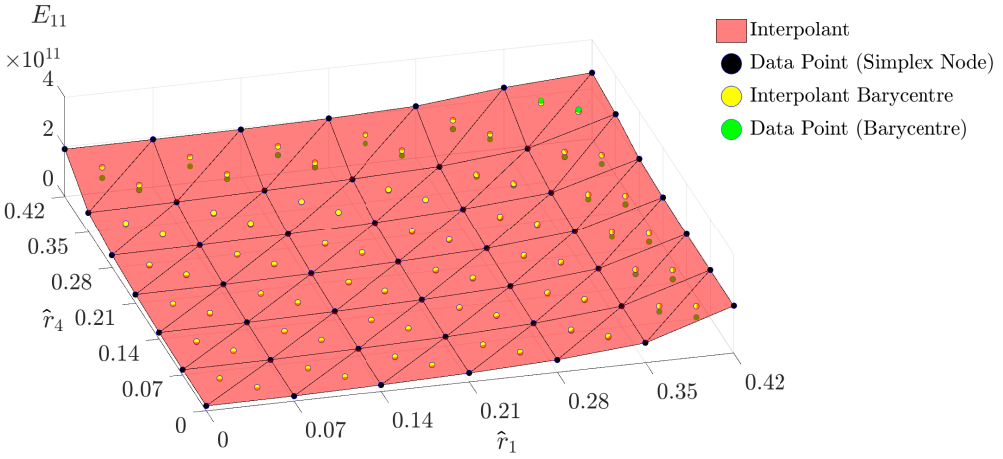


(d) E_{44}

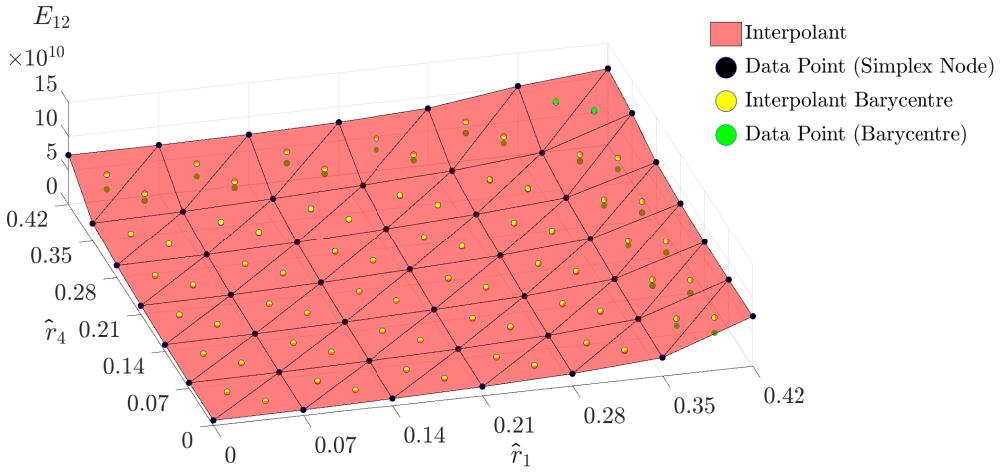
Figure 3.12: The mechanical property space along the hyperplane definition where $\hat{r}_5$ and $\hat{r}_7$ are set to vary, $\hat{r}_{1,2,3,4,6} = 0.14$ and $e_{5,6,7} = 0.8$



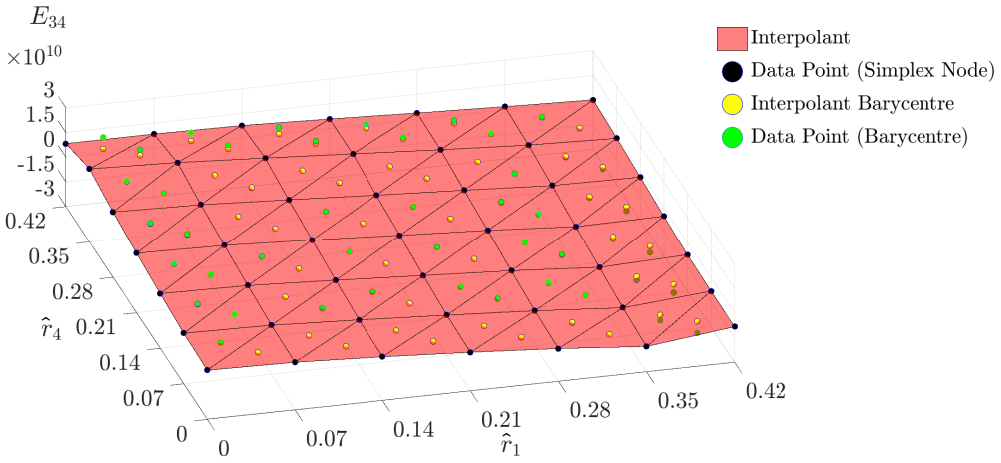
(a) Volume fraction



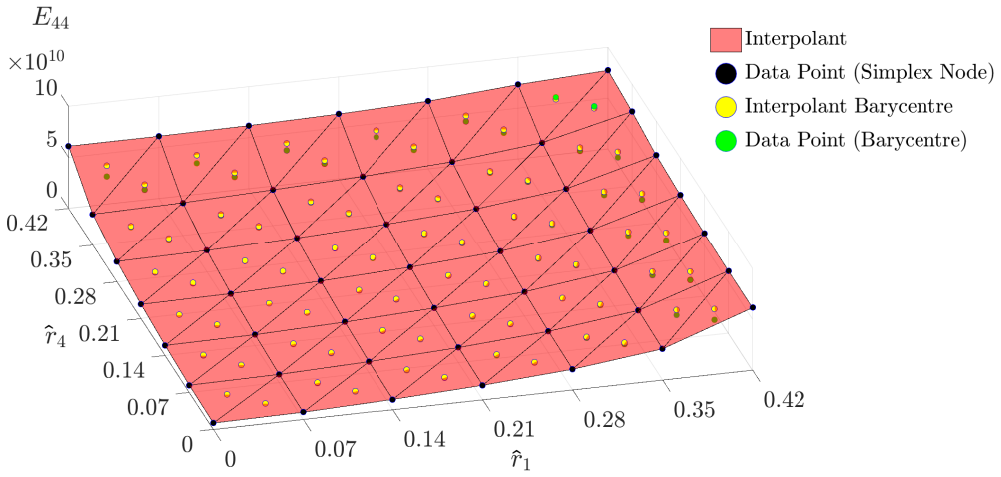
(b) E_{11}



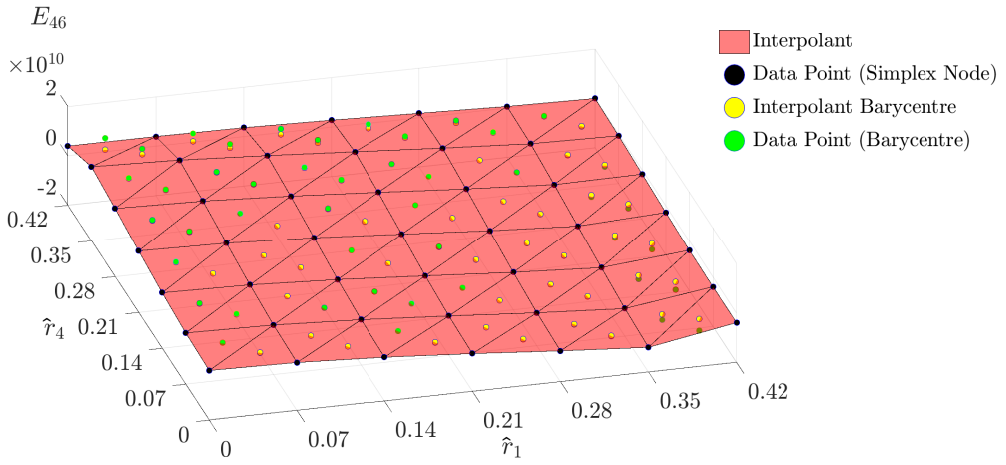
(c) E_{12}



(d) E_{34}

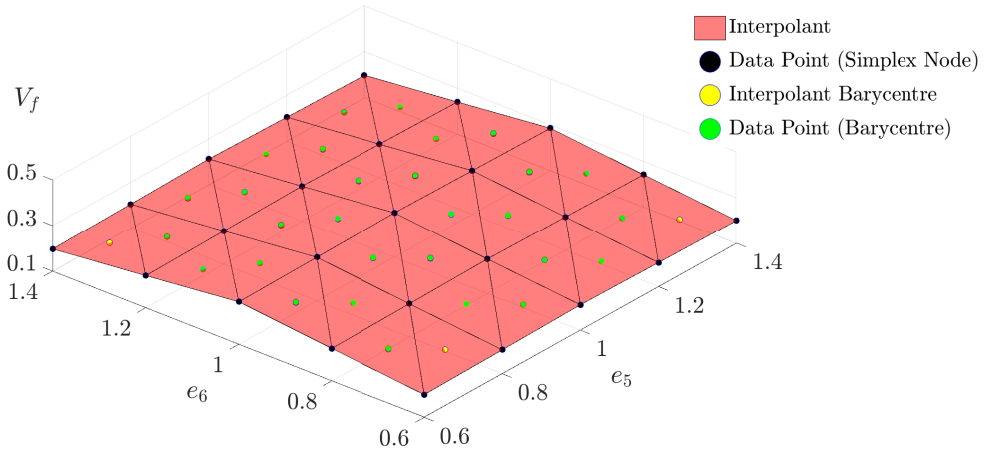


(e) E_{44}

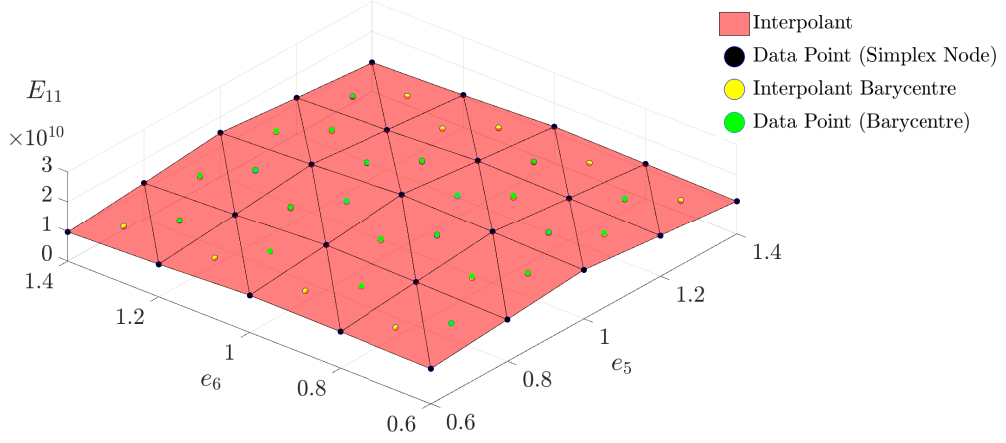


(f) E_{46}

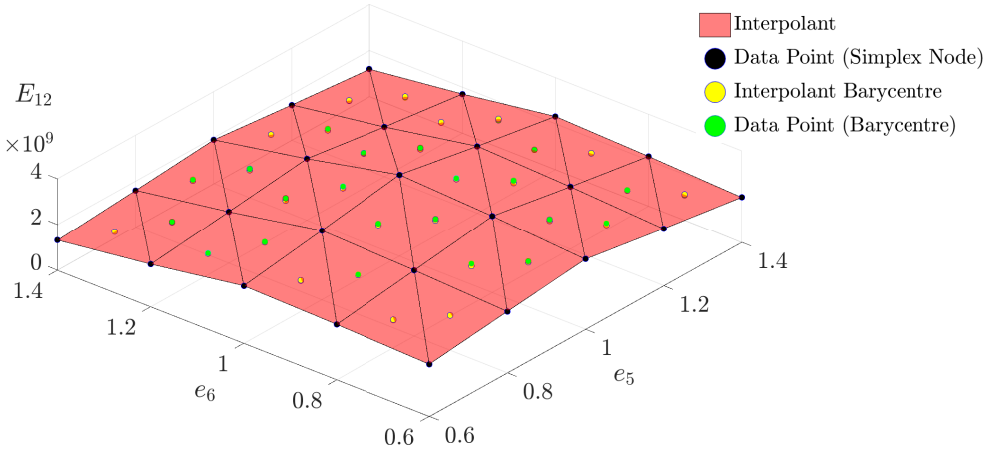
Figure 3.13: The mechanical property space along the hyperplane definition where $\hat{r}_1$ and $\hat{r}_4$ are set to vary, $\hat{r}_{1,4,5,6,7} = 0.14$ and $e_{5,6,7} = 0.8$



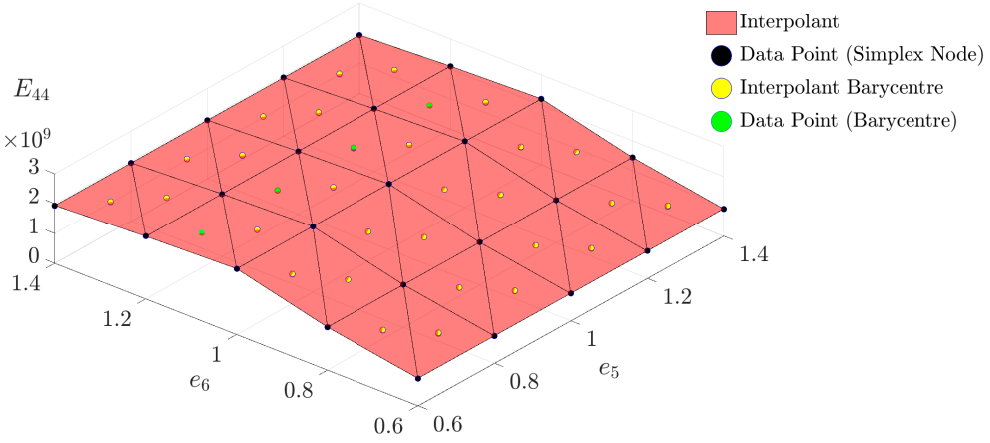
(a) Volume fraction



(b) E_{11}



(c) E_{12}



(d) E_{44}

Figure 3.14: The mechanical property space along the hyperplane definition where e_5 and e_6 are set to vary, $\hat{r}_{1-7} = 0.14$ and $e_7 = 0.8$

Case	Average RMSE
$\hat{r}_5$ vs $\hat{r}_7$ (Fig. 3.12)	1.05%
$\hat{r}_1$ vs $\hat{r}_4$ (Fig. 3.13)	7.43%
e_5 vs e_6 (Fig. 3.14)	1.33%

Table 3.2: The average RMSE associated in the homogenised mechanical properties for the three hyperplanes passing through the mechanical property space illustrated in Figs. 3.12, 3.13 and 3.14

is a consequence of the microstructure tending towards a fully-dense representation as the normalised radius approaches a value of 0.42, which in this instance leads to a superlinear increase in stiffness within the interval $[0.35, 0.42]$, as members coalesce. It would be permissible to reduce this error by adding a further level to the full-factorial DOE, however, this would increase the computational expense to an intractable level given the present computational constraints.

It is noted within Tab. 3.2, that the average RMSE in the mechanical properties for the cases presented in Figs. 3.12 and 3.14 is negligible, given the aforementioned sources of systematic and random error. However, due to the superlinear increase in stiffness noted within the interval $[0.35, 0.42]$ in Fig. 3.13, the average RMSE for this case is larger. In conclusion, the smooth representation of the mechanical property space depicted within Figs. 3.12, 3.13 and 3.14 and the associated RMSEs provides a satisfactory foundation to investigate the benefits afforded by concurrently coupling the micro and macroscale models. Again, it should be reiterated that the errors noted in the mechanical property space for the present framework tends to zero as the available computational resource is increased to infinity.

3.4.5 Concurrent Coupling Example (2D)

To demonstrate visually how a database could be developed over the course of an optimisation, an example is presented and illustrated in Fig. 3.15. To permit clearer visualisation, a two-dimensional microstructure with two configurable thickness parameters (t_1,t_2) is used within this example. Prior to optimisation, a full factorial DOE is generated using 6 uniformly distributed levels for both parameters. In total, the DOE contains 36 unique simulation nodes, for which the material properties can be derived using only 21 nodes, due to the existence of a 90° rotational symmetry.

Again, to permit clearer visualisation, the path of only a single microstructure through the design space is considered during optimisation within this example. At the beginning of the optimisation process, this microstructure is initialised with design variables $[0.2, 0.2]$. An interpolant is then derived to determine the mechanical properties of this microstructure

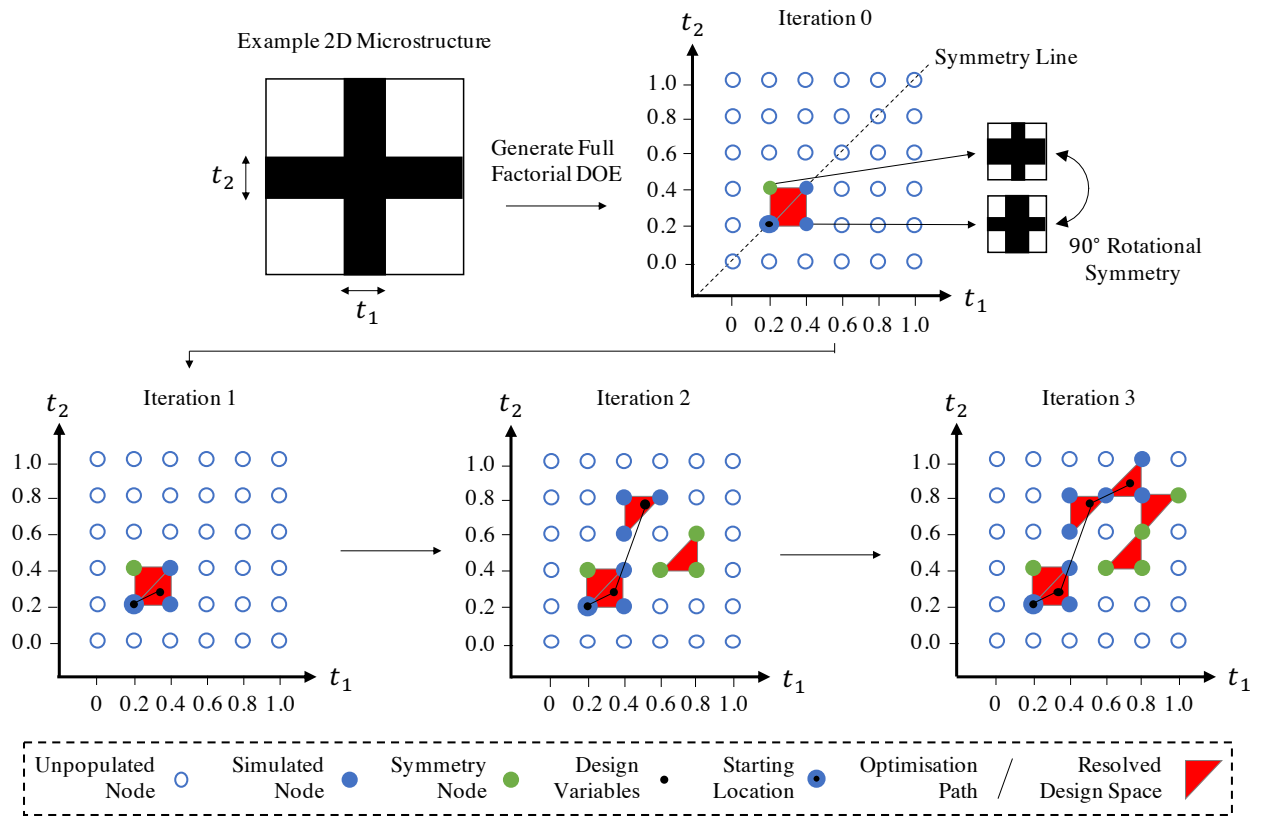


Figure 3.15: A two-dimensional example is provided to demonstrate how a database may evolve over the course of an optimisation, using the concurrent coupling technique detailed within Sec. 3.4. To permit clearer visualisation, the path of only a single microstructure through the design space during optimisation is illustrated.

using the simulation data collected from the nodes $[0.2, 0.2]$, $[0.4, 0.2]$ and $[0.4, 0.4]$, as these microstructures form the enclosing simplex. By consequence, a further region of the design space is simultaneously resolved at no extra cost, as a result of the rotational symmetry which enables the mechanical properties for the microstructure $[0.2, 0.4]$ to be derived from the microstructure $[0.4, 0.2]$. It is important to note that once a region of the design space has been resolved, it can be exhaustively interrogated for a negligible computational expense.

Over the course of the next 3 iterations, the continuous design variables evolve by following sensitivities with respect to functional objectives. To facilitate the evaluation of the macroscale model during each iteration, additional microstructures are simulated to resolve the required regions of design space. By consequence, a re-usable database is assembled during optimisation, which can be leveraged to reduce the computational expense of successive iterations or entirely new optimisation problems.

3.5 Optimisation Formulation

During the optimisation process, the continuous microscale parameters contained with the design variable matrix (X) are varied to fulfil functional objectives at the macroscale. This approach enables structures with precise functionally-graded mechanical properties over two-scales to be derived.

3.5.1 Helmholtz Filter

It is important to note that prior to the evaluation of the macroscale finite-element model during the optimisation procedure, the design variables pass through a Helmholtz filter (Lazarov and Sigmund, 2011) to promote smoothly-varying microscale geometry and to discourage 'checkerboarded' geometries. The Helmholtz filter is implemented by solving the following partial differential equation

$$-R^2\nabla^2\tilde{X}_\beta + \tilde{X}_\beta = X_\beta \tag{3.42}$$

where R denotes the filter length, which is set equal to two times the length of an element for all of the results presented within this thesis. This conservative filter length ensures continuous geometry is derived, which is a fundamental requirement of the heterogeneous multiscale method (Weinan et al., 2007).

3.5.2 Optimisation Formulation

Optimisation problems within this framework are formulated according to

$$\begin{aligned}
\underset{X}{\text{Minimise:}} \quad & J(u(E^H(\tilde{X})), E^H(\tilde{X})) \\
\text{Subject to:} \quad & F(u(E^H(\tilde{X})), E^H(\tilde{X})) = 0 \\
& V_f - \int_\Omega V(\tilde{X}_\alpha) \geq 0 \\
& X_{\alpha\beta} \geq X_{min,\beta},\ \forall\, \beta\ \in\ \{1,2,3,\ldots,n\} \\
& X_{\alpha\beta} \leq X_{max,\beta},\ \forall\, \beta\ \in\ \{1,2,3,\ldots,n\}
\end{aligned} \tag{3.43}$$

where '$J(u(E^H(\tilde{X})), E^H(\tilde{X}))$' denotes the scalar valued objective function and '$F(u(E^H(\tilde{X})), E^H(\tilde{X})) = 0$' denotes the partial differential equation constraining the physics of the macroscale model. For all applications of the present framework presented within this thesis, a linear elastic partial differential equation (represented in weak-variational form)

$$\int_\Omega E_{ijkl}(\boldsymbol{x})\epsilon_{ij}(\boldsymbol{u})\epsilon_{kl}(\boldsymbol{\omega})\ d\Omega = \int_\Omega \boldsymbol{f\omega}\ d\Omega + \int_\Gamma \boldsymbol{t\omega}\ d\Gamma \tag{3.44}$$

is applied to constrain the physics of the macroscale model. A limit is placed on the availability of material within the macroscale domain through a volume-fraction constraint, where $\int_\Omega V(\tilde{X}_\alpha)$ denotes the true volume-fraction and V_f denotes the target volume-fraction. Two further constraints are also applied to bound the design variables between the lowest ($X_{min,\beta}$) and highest ($X_{max,\beta}$) levels outlined within the full factorial DOE (Tab. 3.1).

3.5.3 Sensitivities

Sensitivities of the objective functional with respect to the design variable matrix are calculated through the application of the chain-rule

$$\frac{\partial J}{\partial X} = \left[\sum_{i=1}^{6}\sum_{j=1}^{6}\left(\frac{\partial J}{\partial E_{ij}^H(\tilde{X})} \circ \frac{\partial E_{ij}^H(\tilde{X})}{\partial \tilde{X}}\right)\right]^T \frac{\partial \tilde{X}}{\partial X} \tag{3.45}$$

where $\frac{\partial J}{\partial E^H(\tilde{X})}$ denotes the matrix

$$\frac{\partial J}{\partial E_{ij}^H(\tilde{X})} = \begin{bmatrix} \frac{\partial J}{\partial E_{11}(\tilde{X})} & \frac{\partial J}{\partial E_{12}(\tilde{X})} & \frac{\partial J}{\partial E_{13}(\tilde{X})} & \frac{\partial J}{\partial E_{14}(\tilde{X})} & \frac{\partial J}{\partial E_{15}(\tilde{X})} & \frac{\partial J}{\partial E_{16}(\tilde{X})} \\ \frac{\partial J}{\partial E_{21}(\tilde{X})} & \frac{\partial J}{\partial E_{22}(\tilde{X})} & \frac{\partial J}{\partial E_{23}(\tilde{X})} & \frac{\partial J}{\partial E_{24}(\tilde{X})} & \frac{\partial J}{\partial E_{25}(\tilde{X})} & \frac{\partial J}{\partial E_{26}(\tilde{X})} \\ \frac{\partial J}{\partial E_{31}(\tilde{X})} & \frac{\partial J}{\partial E_{32}(\tilde{X})} & \frac{\partial J}{\partial E_{33}(\tilde{X})} & \frac{\partial J}{\partial E_{34}(\tilde{X})} & \frac{\partial J}{\partial E_{35}(\tilde{X})} & \frac{\partial J}{\partial E_{36}(\tilde{X})} \\ \frac{\partial J}{\partial E_{41}(\tilde{X})} & \frac{\partial J}{\partial E_{42}(\tilde{X})} & \frac{\partial J}{\partial E_{43}(\tilde{X})} & \frac{\partial J}{\partial E_{44}(\tilde{X})} & \frac{\partial J}{\partial E_{45}(\tilde{X})} & \frac{\partial J}{\partial E_{46}(\tilde{X})} \\ \frac{\partial J}{\partial E_{51}(\tilde{X})} & \frac{\partial J}{\partial E_{52}(\tilde{X})} & \frac{\partial J}{\partial E_{53}(\tilde{X})} & \frac{\partial J}{\partial E_{54}(\tilde{X})} & \frac{\partial J}{\partial E_{55}(\tilde{X})} & \frac{\partial J}{\partial E_{56}(\tilde{X})} \\ \frac{\partial J}{\partial E_{61}(\tilde{X})} & \frac{\partial J}{\partial E_{62}(\tilde{X})} & \frac{\partial J}{\partial E_{63}(\tilde{X})} & \frac{\partial J}{\partial E_{64}(\tilde{X})} & \frac{\partial J}{\partial E_{65}(\tilde{X})} & \frac{\partial J}{\partial E_{66}(\tilde{X})} \end{bmatrix} \tag{3.46}$$

which is computed by solving an adjoint equation, derived using the open source python library dolfin-adjoint (Mitusch et al., 2019) which leverages high-level algorithmic differentiation to

derive adjoint models. In this instance, adjoint models provide a computationally inexpensive method to derive gradients (one additional solve of a linear system per microscale parameter), which would otherwise require an intractable computational resource to derive (N additional solves of a linear system per microscale parameter). To demonstrate this, it is prudent to first consider the total derivative of the objective function with respect to the elasticity tensor

$$\frac{DJ}{DE^H(\tilde{X})} = \frac{\partial J}{\partial E^H(\tilde{X})} + \frac{\partial J}{\partial u}\frac{\partial u}{\partial E^H(\tilde{X})} \tag{3.47}$$

where the term $\frac{\partial u}{\partial E^H(\tilde{X})}$ cannot be efficiently calculated, as it requires the solution of $N+1$ linear systems. To avoid this intractable expense, a substitution is performed using the total derivative of the partial differential equation constraining the physics of the macroscale model with respect to the elasticity tensor

$$\frac{DF}{DE^H(\tilde{X})} = \frac{\partial F}{\partial E^H(\tilde{X})} + \frac{\partial F}{\partial u}\frac{\partial u}{\partial E^H(\tilde{X})} = 0 \tag{3.48}$$

which can be rearranged such that $\frac{\partial u}{\partial E^H(\tilde{X})}$ is the subject of the equation

$$\frac{\partial u}{\partial E^H(\tilde{X})} = -\frac{\partial F}{\partial E^H(\tilde{X})}\left[\frac{\partial F}{\partial u}\right]^{-1} \tag{3.49}$$

Substitution of Eq. 3.49 into Eq. 3.47 is then performed

$$\frac{DJ}{DE^H(\tilde{X})} = \frac{\partial J}{\partial E^H(\tilde{X})} - \frac{\partial J}{\partial u}\frac{\partial F}{\partial E^H(\tilde{X})}\left[\frac{\partial F}{\partial u}\right]^{-1} \tag{3.50}$$

It is then permissible to simplify Eq. 3.50 by isolating the terms

$$\frac{\partial J}{\partial u}\left[\frac{\partial F}{\partial u}\right]^{-1}$$

which can be re-arranged into a linear system and solved to derive the adjoint vector (λ)

$$\frac{\partial F}{\partial u}\lambda = \frac{\partial J}{\partial u} \tag{3.51}$$

The adjoint vector can then be substituted into Eq. 3.50 to derive a simplified expression for the total derivative of the objective function

$$\frac{DJ}{DE^H(\tilde{X})} = \frac{\partial J}{\partial E^H(\tilde{X})} - \lambda^T \frac{\partial F}{\partial E^H(\tilde{X})} \tag{3.52}$$

in which all terms are trivial to calculate. As a result, the solution of only one additional linear system (Eq. 3.51) is required to derive the derivative of the objective function with respect to the elasticity tensor. It is important to add that all derivatives presented within this thesis derived through adjoint methods have been verified through finite-differencing.

The partial derivative $\frac{\partial E_{ij}}{\partial \tilde{X}}$ denotes the matrix

$$\frac{\partial E_{ij}(\tilde{X})}{\partial \tilde{X}} = \begin{bmatrix} \frac{\partial E_{11}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{12}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{13}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{14}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{15}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{16}(\tilde{X})}{\partial \tilde{X}} \\ \frac{\partial E_{21}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{22}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{23}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{24}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{25}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{26}(\tilde{X})}{\partial \tilde{X}} \\ \frac{\partial E_{31}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{32}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{33}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{34}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{35}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{36}(\tilde{X})}{\partial \tilde{X}} \\ \frac{\partial E_{41}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{42}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{43}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{44}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{45}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{46}(\tilde{X})}{\partial \tilde{X}} \\ \frac{\partial E_{51}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{52}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{53}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{54}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{55}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{56}(\tilde{X})}{\partial \tilde{X}} \\ \frac{\partial E_{61}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{62}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{63}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{64}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{65}(\tilde{X})}{\partial \tilde{X}} & \frac{\partial E_{66}(\tilde{X})}{\partial \tilde{X}} \end{bmatrix} \tag{3.53}$$

which is computed by finite-differencing between the base nodes of the enclosing simplices and adjacent nodes in each dimension.

As the design variables pass through a Helmholtz filter prior to the evaluation of the objective functional, the sensitivities are calculated with respect to the filtered design variables ($\tilde{X}$). It is therefore necessary to recover the unfiltered representation of the sensitivities to ensure compatibility with the unfiltered design variables (X) during optimisation. As a result, the derivative of the Helmholtz operator, denoted by $\frac{\partial \tilde{X}}{\partial X}$, is applied as an additional step in the sensitivity derivation. Similarly, derivatives of the constraints (Jacobian matrix) are also calculated through the application of a chain rule approach

$$\frac{\partial V(\tilde{X})}{\partial X} = \frac{\partial V(\tilde{X})}{\partial \tilde{X}} \frac{\partial \tilde{X}}{\partial X} \tag{3.54}$$

where the derivative $\frac{\partial V(\tilde{X})}{\partial \tilde{X}}$ is also calculated by finite-differencing between the base nodes of the enclosing simplices and adjacent nodes in each dimension.

3.5.4 Interior-point Optimisation Algorithm

As detailed in Sec. 2.5, the present framework is extended within Sec. 3.5.5 in order to exploit a novel optimisation strategy which seeks to further reduce the computational expense of optimisation by initially optimising a subset of the microscale design parameters before activating additional microscale design parameters as the optimisation progresses. In order to implement this optimisation strategy, it is necessary to have full control over the optimisation solver, which is not possible through the implementation of an open-source optimiser such as IPOPT (Wächter and Biegler, 2006), SNOPT (Gill et al., 2005) or SciPy (L-BFGS-B) (Zhu et al., 1997), as the source code is not user-friendly and/or compiled. As a result, the primal-dual interior point algorithm with line-search outlined within Andrei (2017) is implemented from first principles. This class of optimiser is preferred, as interior point algorithms typically outperform competing methods (sequential quadratic programming, method of moving asymptotes (Svanberg, 1987)) when applied to structural optimisation based problems (Rojas-Labanda and Stolpe, 2015). For the remainder of this thesis, this algorithm is referred to as IPOwDA (Interior Point Optimisation with Dimension Addition).

3.5.4.1 Lagrangian Function

A compact representation of the optimisation problem definition outlined in Eq. 3.43 is used as the subject of the derivation of the interior-point optimisation algorithm

$$\begin{aligned} \underset{x}{\text{Minimise:}} \quad & f(x) \\ \text{Subject to:} \quad & c(x) \geq 0 \\ & x_L \leq x \leq x_U \end{aligned} \tag{3.55}$$

where $x \in \mathbb{R}^n$ denotes the vector of design variables, $f(x) \in \mathbb{R}^1$ denotes the objective function, $c(x) \in \mathbb{R}^m$ denotes the vector of constraints and $x_L \in \mathbb{R}^n$ and $x_U \in \mathbb{R}^n$ denote the lower and upper bounds applied to constrain the design variables respectively.

The first step in the derivation is to reformulate Eq. 3.55 by introducing slack variables (s, v, w) to convert the inequality constraints into equality constraints

$$\begin{aligned} \underset{x}{\text{Minimise:}} \quad & f(x) \\ \text{Subject to:} \quad & c(x) - s = 0 \\ & x - w - x_L = 0 \\ & x + v - x_U = 0 \\ & s \geq 0,\ w \geq 0,\ v \geq 0 \end{aligned} \tag{3.56}$$

where $s \in \mathbb{R}^m$, $v \in \mathbb{R}^n$ and $w \in \mathbb{R}^n$. The inequality constraints associated with the slack variables are subsequently removed by reformulating the objective function to create the log-barrier function

$$\begin{aligned} \underset{x}{\text{Minimise:}} \quad & f(x) - \mu \sum_{i=1}^{m} \ln(s_i) - \mu \sum_{i=1}^{n} \ln(w_i) - \mu \sum_{i=1}^{n} \ln(v_i) \\ \text{Subject to:} \quad & c(x) - s = 0 \\ & x - w - x_L = 0 \\ & x + v - x_U = 0 \end{aligned} \tag{3.57}$$

where μ denotes a scalar constant applied to control the magnitude of the penalisation imposed on the objective to maintain the non-negativity constraints imposed on both the slack variables and Lagrange multipliers. It is then permissible to assemble the Lagrangian function for the constrained system (Eq. 3.57)

$$\begin{aligned} \text{Minimise:} \quad \mathcal{L}(x, s, w, v, z, p, q) = & f(x) - \mu \sum_{i=1}^{m} \ln(s_i) - \mu \sum_{i=1}^{n} \ln(w_i) - \mu \sum_{i=1}^{n} \ln(v_i) \\ & - z^T(c(x) - s) - p^T(x - w - x_L) - q^T(x + v - x_U) \end{aligned} \tag{3.58}$$

where $z \in \mathbb{R}^m$, $p \in \mathbb{R}^n$ and $q \in \mathbb{R}^n$ denote Lagrange multipliers. Reformulating the constrained system outlined in Eq. 3.57 into a Lagrangian function enables an unconstrained system to be derived from a constrained system, where the canonical derivative test of an extremal point

$$\nabla \mathcal{L}_{x,s,w,v,z,p,q} = 0 \tag{3.59}$$

is still a valid method of determining an extremal point of the objective function. First-order optimality conditions can be derived by differentiating Eq. 3.59 with respect to the design variables, slack variables and Lagrange multipliers

$$\nabla \mathcal{L}_x = \nabla f(x) - \nabla c(x)^T z - p + q = 0 \tag{3.60}$$

$$\nabla \mathcal{L}_z = c(x) - s = 0 \tag{3.61}$$

$$\nabla \mathcal{L}_p = x - w - x_L = 0 \tag{3.62}$$

$$\nabla \mathcal{L}_q = x + v - x_U = 0 \tag{3.63}$$

$$\nabla \mathcal{L}_s = SZe - \mu e = 0 \tag{3.64}$$

$$\nabla \mathcal{L}_w = WPe - \mu e = 0 \tag{3.65}$$

$$\nabla \mathcal{L}_v = VQe - \mu e = 0 \tag{3.66}$$

where capitalised variables denote diagonal matrices and e denotes a vector of varying length with all elements set equal to one. In combination, Eqs. 3.60 — 3.66 denote the first-order Karush-Kuhn-Tucker (KKT) conditions, and highlight the existence of an extremal point when all of the derivative tests are satisfied.

3.5.4.2 Reduced KKT Linear System

A relaxed Newton's method is then applied to iteratively determine the optimal set of values for the design variables, slack variables and Lagrange multipliers which minimises the objective function and adheres to all constraints

$$\begin{bmatrix} x_{i+1} \\ z_{i+1} \\ p_{i+1} \\ q_{i+1} \\ s_{i+1} \\ w_{i+1} \\ v_{i+1} \end{bmatrix} = \begin{bmatrix} x_i \\ z_i \\ p_i \\ q_i \\ s_i \\ w_i \\ v_i \end{bmatrix} - \alpha \frac{\nabla \mathcal{L}}{\nabla \otimes (\nabla \mathcal{L})} \tag{3.67}$$

where α denotes the step size. It is then permissible to expand the term

$$\frac{\nabla \mathcal{L}}{\nabla \otimes (\nabla \mathcal{L})}$$

into the KKT linear system

$$\begin{bmatrix} \nabla \mathcal{L}_{xx} & -\nabla c(x) & -I & I & 0 & 0 & 0 \\ \nabla c(x)^T & 0 & 0 & 0 & -I & 0 & 0 \\ I & 0 & 0 & 0 & 0 & -I & 0 \\ I & 0 & 0 & 0 & 0 & 0 & I \\ 0 & S & 0 & 0 & Z & 0 & 0 \\ 0 & 0 & W & 0 & 0 & P & 0 \\ 0 & 0 & 0 & V & 0 & 0 & Q \end{bmatrix} \begin{bmatrix} d_x \\ d_z \\ d_p \\ d_q \\ d_s \\ d_w \\ d_v \end{bmatrix} = - \begin{bmatrix} \nabla f(x) - \nabla c(x)^T z - p + q \\ c(x) - s \\ x - w - x_L \\ x + v - x_U \\ SZe - \mu e \\ WPe - \mu e \\ VQe - \mu e \end{bmatrix} \tag{3.68}$$

where the solution to this linear system provides the search directions for the design variables (d_x), slack variables (d_s, d_w, d_v) and Lagrange multipliers (d_z, d_p, d_q). It is necessary to transform Eq. 3.68 into a symmetric linear system to aid future calculations by multiplying rows 2 and 3 by -1, and rows 4, 5 and 6 by S^{-1}, W^{-1} and V^{-1} respectively

$$\begin{bmatrix} \nabla\mathcal{L}_{xx} & -\nabla c(x) & -I & I & 0 & 0 & 0 \\ -\nabla c(x)^T & 0 & 0 & 0 & I & 0 & 0 \\ -I & 0 & 0 & 0 & 0 & I & 0 \\ I & 0 & 0 & 0 & 0 & 0 & I \\ 0 & I & 0 & 0 & ZS^{-1} & 0 & 0 \\ 0 & 0 & I & 0 & 0 & PW^{-1} & 0 \\ 0 & 0 & 0 & I & 0 & 0 & QV^{-1} \end{bmatrix} \begin{bmatrix} d_x \\ d_z \\ d_p \\ d_q \\ d_s \\ d_w \\ d_v \end{bmatrix} = - \begin{bmatrix} \nabla f(x) - \nabla c(x)^T z - p + q \\ -c(x) + s \\ -x + w + x_L \\ x + v - x_U \\ Ze - \mu S^{-1} e \\ Pe - \mu W^{-1} e \\ Qe - \mu V^{-1} e \end{bmatrix} \tag{3.69}$$

Isolated expressions are then derived for d_s, d_w and d_v

$$d_s = \mu Z^{-1} e - Se - Z^{-1} S d_z \tag{3.70}$$

$$d_s = \mu P^{-1} e - We - P^{-1} W d_p \tag{3.71}$$

$$d_s = \mu Q^{-1} e - Ve - Q^{-1} V d_q \tag{3.72}$$

which enables the linear system defined by Eq. 3.69 to be reduced in dimension to

$$\begin{bmatrix} \nabla\mathcal{L}_{xx} & -\nabla c(x) & -I & I \\ -\nabla c(x)^T & -Z^{-1}S & 0 & 0 \\ -I & 0 & -P^{-1}W & 0 \\ I & 0 & 0 & -Q^{-1}V \end{bmatrix} \begin{bmatrix} d_x \\ d_z \\ d_p \\ d_q \end{bmatrix} = - \begin{bmatrix} \nabla f(x) - \nabla c(x)^T z - p + q \\ -c(x) + Z^{-1} e \\ -x + x_L + P^{-1} e \\ -x_U + x + Q^{-1} e \end{bmatrix} \tag{3.73}$$

Subsequently, isolated expressions are derived for d_p and d_q

$$d_p = \mu W^{-1} e - W^{-1} P(x - x_L) - W^{-1} P d_x \tag{3.74}$$

$$d_q = \mu V^{-1} e - V^{-1} Q(x_U - x) + V^{-1} Q d_x \tag{3.75}$$

which again enables the linear system defined by Eq. 3.73 to be reduced in dimension to

$$\begin{bmatrix} \nabla\mathcal{L}_{xx} + W^{-1}P + V^{-1}Q & -\nabla c(x) \\ -\nabla c(x)^T & -Z^{-1}S \end{bmatrix} \begin{bmatrix} d_x \\ d_z \end{bmatrix} = - \begin{bmatrix} \nabla f(x) - \nabla c(x)^T z - p + q + W^{-1}P(x - x_L) - V^{-1}Q(x_U - x) - \mu W^{-1}e + \mu V^{-1}e \\ -c(x) + \mu Z^{-1}e \end{bmatrix} \tag{3.76}$$

which for convenience is represented as

$$\begin{bmatrix} \nabla\mathcal{L}_{xx} + W^{-1}P + V^{-1}Q & -\nabla c(x) \\ -\nabla c(x)^T & -Z^{-1}S \end{bmatrix} \begin{bmatrix} d_x \\ d_z \end{bmatrix} = - \begin{bmatrix} RHS_1 \\ RHS_2 \end{bmatrix} \tag{3.77}$$

for the remainder of this section. Whilst Eq. 3.77 represents a comparatively smaller linear system with respect to Eq. 3.68, it is not possible to solve Eq. 3.77 without the application of prohibitively expensive direct linear algebra solvers or bespoke factorisation methods which are not readily available. In addition, it is also not possible to either derive an analytical expression for the Hessian matrix '$\nabla\mathcal{L}_{xx}$' or store its dense $[n \times n]$ matrix representation in memory. As a consequence, L-BFGS matrices (Liu and Nocedal, 1989; Byrd et al., 1995) are implemented in compact form to approximate the Hessian matrix

$$\nabla\mathcal{L}_{xx} \approx \sigma I - \begin{bmatrix} -\sigma G & -Y \end{bmatrix} \begin{bmatrix} -\sigma G^T G & -L \\ L^T & -D \end{bmatrix}^{-1} \begin{bmatrix} -\sigma G^T \\ Y^T \end{bmatrix} \tag{3.78}$$

where Y denotes the matrix

$$Y_i = \begin{bmatrix} y_{i-\lambda}, & \cdots, & y_{i-1} \end{bmatrix} \tag{3.79}$$

where the subscripts i and λ denote the iteration number and the L-BFGS memory, and y_i denotes

$$y_i = x_{i+1} - x_k \tag{3.80}$$

and where G denotes the matrix

$$G_i = \begin{bmatrix} g_{i-\lambda}, & \cdots, & g_{i-1} \end{bmatrix} \tag{3.81}$$

where g_i denotes

$$g_i = \nabla_x \mathcal{L}_{i+1} - \nabla_x \mathcal{L}_i \tag{3.82}$$

and where

$$L_{i,jk} = \begin{cases} (g_{i-\lambda-1+j})^T (y_{i-\lambda-1+k}) & \text{if } \; j > k \\ 0 & \text{otherwise} \end{cases} \tag{3.83}$$

and where

$$D_i = \text{diag} \begin{bmatrix} g_{i-\lambda}^T y_{i-\lambda}, & \cdots, & g_{i-1}^T y_{i-1} \end{bmatrix} \tag{3.84}$$

and where σ denotes a positive scaling factor. By substituting this approximate representation of the Hessian matrix into Eq. 3.77, the following linear system is obtained

$$\begin{aligned} \Bigg(& \begin{bmatrix} \sigma I + W^{-1}P + V^{-1}Q & -\nabla c(x) \\ -\nabla c(x)^T & -Z^{-1}S \end{bmatrix} - \\ & \begin{bmatrix} -\sigma G & -Y \\ 0 & 0 \end{bmatrix} \begin{bmatrix} -\sigma G^T G & -L \\ L^T & -D \end{bmatrix}^{-1} \begin{bmatrix} -\sigma G^T & 0 \\ Y^T & 0 \end{bmatrix} \Bigg) \begin{bmatrix} d_x \\ d_z \end{bmatrix} = - \begin{bmatrix} RHS_1 \\ RHS_2 \end{bmatrix} \end{aligned} \tag{3.85}$$

As a consequence of representing the Hessian matrix using L-BFGS matrices in compact form, it is permissible to derive an algebraic expression for the inverse of the LHS of Eq. 3.77 through the application of the Sherman-Morrison-Woodbury identity (Higham, 2002), which provides an analytical expression for the inverse of matrix equations of the form

$$\left[A + UC^{-1}V\right]^{-1} = A^{-1} - A^{-1}U(C + VA^{-1}U)^{-1}VA^{-1} \tag{3.86}$$

which is identical to the linear system defined by Eq. 3.85. The linear system is then substituted into Eq. 3.86 to derive an analytical expression for the search directions d_x and d_z

$$\begin{aligned}
&\begin{bmatrix} d_x \\ d_z \end{bmatrix} = \begin{bmatrix} \sigma I + W^{-1}P + V^{-1}Q & -\nabla c(x) \\ -\nabla c(x)^T & -Z^{-1}S \end{bmatrix}^{-1} \begin{bmatrix} -RHS_1 \\ -RHS_2 \end{bmatrix} - \\
&\begin{bmatrix} \sigma I + W^{-1}P + V^{-1}Q & -\nabla c(x) \\ -\nabla c(x)^T & -Z^{-1}S \end{bmatrix}^{-1} \begin{bmatrix} -\sigma G & -Y \\ 0 & 0 \end{bmatrix} \left(\begin{bmatrix} -\sigma G^T G & -L \\ L^T & -D \end{bmatrix} - \right. \\
&\left. \begin{bmatrix} -\sigma G^T & 0 \\ Y^T & 0 \end{bmatrix} \begin{bmatrix} \sigma I + W^{-1}P + V^{-1}Q & -\nabla c(x) \\ -\nabla c(x)^T & -Z^{-1}S \end{bmatrix}^{-1} \begin{bmatrix} -\sigma G & -Y \\ 0 & 0 \end{bmatrix} \right)^{-1} \begin{bmatrix} -\sigma G^T & 0 \\ Y^T & 0 \end{bmatrix} \\
&\begin{bmatrix} \sigma I + W^{-1}P + V^{-1}Q & -\nabla c(x) \\ -\nabla c(x)^T & -Z^{-1}S \end{bmatrix}^{-1} \begin{bmatrix} -RHS_1 \\ -RHS_2 \end{bmatrix}
\end{aligned} \tag{3.87}$$

which is straightforward to solve except for the inversion of the matrix

$$\begin{bmatrix} \sigma I + W^{-1}P + V^{-1}Q & -\nabla c(x) \\ -\nabla c(x)^T & -Z^{-1}S \end{bmatrix}^{-1}$$

However, it is permissible to trivially invert this matrix using the block matrix inversion identity (Bernstein, 2009)

$$\begin{aligned}
&A^{-1} = \\
&\begin{bmatrix} A_{11}^{-1} + A_{11}^{-1}A_{12}(A_{22} - A_{21}A_{11}^{-1}A_{12})^{-1}A_{21}A_{11}^{-1} & -A_{11}^{-1}A_{12}(A_{22} - A_{21}A_{11}^{-1}A_{12})^{-1} \\ -(A_{22} - A_{21}A_{11}^{-1}A_{12})^{-1}A_{21}A_{11}^{-1} & (A_{22} - A_{21}A_{11}^{-1}A_{12})^{-1} \end{bmatrix}
\end{aligned} \tag{3.88}$$

which takes advantage of the sparse arrowhead pattern (Oĺeary and Stewart, 1996) to simplify computation. As a consequence of the matrix identities leveraged, an extremely computationally efficient method has been derived to calculate the search directions required to iterate Newton's method (see Eq. 3.67).

3.5.4.3 Line Search

After the search-directions are calculated, it is then necessary to determine the optimal step size in reference to minimising the first-order optimality conditions (see Eqs. 3.60 - 3.66). To avoid trying to satisfy multiple conditions independently, a single merit function is formulated

$$F(x_i, z_i, p_i, q_i, s_i, w_i, v_i) = \begin{bmatrix} F_1(x_i, z_i, p_i, q_i, s_i, w_i, v_i) \\ F_2(z_i, p_i, q_i, s_i, w_i, v_i) \end{bmatrix} \begin{bmatrix} F_1(x_i, z_i, p_i, q_i, s_i, w_i, v_i) \\ F_2(z_i, p_i, q_i, s_i, w_i, v_i) \end{bmatrix}^T \tag{3.89}$$

where

$$F_1(x_i, z_i, p_i, q_i, s_i, w_i, v_i) = \begin{bmatrix} \nabla \mathcal{L}_x(x_i, z_i, p_i, q_i, s_i, w_i, v_i) \\ \nabla \mathcal{L}_z(x_i, z_i, p_i, q_i, s_i, w_i, v_i) \\ \nabla \mathcal{L}_p(x_i, z_i, p_i, q_i, s_i, w_i, v_i) \\ \nabla \mathcal{L}_q(x_i, z_i, p_i, q_i, s_i, w_i, v_i) \end{bmatrix} \tag{3.90}$$

and where

$$F_2(x_i, z_i, p_i, q_i, s_i, w_i, v_i) = \begin{bmatrix} \nabla \mathcal{L}_s(z_i, p_i, q_i, s_i, w_i, v_i) \\ \nabla \mathcal{L}_w(z_i, p_i, q_i, s_i, w_i, v_i) \\ \nabla \mathcal{L}_v(z_i, p_i, q_i, s_i, w_i, v_i) \end{bmatrix} \tag{3.91}$$

It is important to note that the first-order optimality conditions are split into two separate sub-functions (F_1, F_2), as this aids with the implementation of the line-search. A back-tracking line-search is implemented to calculate the optimal step-size, which is determined by examining whether a given step-size satisfies a series of conditions in reference to the merit function. If one of the conditions is violated, the step-size is halved, and the process is repeated until an acceptable step-size is found. The first condition

$$x_L \leq x_{i,\alpha} \leq x_U \tag{3.92}$$

where $x_{i,\alpha}$ denotes

$$x_{i,\alpha} = x_i + \alpha_i d_{x,i} \tag{3.93}$$

is applied to ensure that the design variables remain bounded between the lower and upper bounds, and similarly the second condition

$$z_{i,\alpha},\ p_{i,\alpha},\ q_{i,\alpha},\ s_{i,\alpha},\ w_{i,\alpha},\ v_{i,\alpha},\ > 0 \tag{3.94}$$

is applied to ensure that the slack variables and Lagrange multipliers remain positive. A further condition

$$\frac{\min\left\{s_{i,\alpha}z_{i,\alpha},\ w_{i,\alpha}p_{i,\alpha},\ v_{i,\alpha}q_{i,\alpha}\right\}}{\frac{s_{i,\alpha}^T z_{i,\alpha}+w_{i,\alpha}^T p_{i,\alpha}+v_{i,\alpha}^T q_{i,\alpha}}{2n+m}} \geq \frac{\min\left\{s_0 z_0,\ w_0 p_0,\ v_0 q_0\right\}}{\frac{s_0^T z_0+w_0^T p_0+v_0^T q_0}{2n+m}} \tag{3.95}$$

is applied to ensure that the iterate is not prematurely projected into close proximity with the boundary of the feasible region (defined as a relative measure against the initial distance from the boundary of the feasible region). This condition is important to satisfy, as if it is neglected it can lead to a significant increase in the convergence time as the step-size is severely limited whilst in close proximity to the boundary. A further condition is implemented to avoid non-convergence of the algorithm

$$\frac{s_{i,\alpha}^T z_{i,\alpha} + w_{i,\alpha}^T p_{i,\alpha} + v_{i,\alpha}^T q_{i,\alpha}}{\left\|F_1(x_{i,\alpha}, z_{i,\alpha}, p_{i,\alpha}, q_{i,\alpha}, s_{i,\alpha}, w_{i,\alpha}, v_{i,\alpha})\right\|_2} \geq \frac{s_0^T z_0 + w_0^T p_0 + v_0^T q_0}{\left\|F_1(x_0, z_0, p_0, q_0, s_0, w_0, v_0)\right\|_2} \tag{3.96}$$

which prevents the slack variables and Lagrange multipliers from prematurely converging, which can isolate and prevent an iterate from reaching an extremal point. The last condition

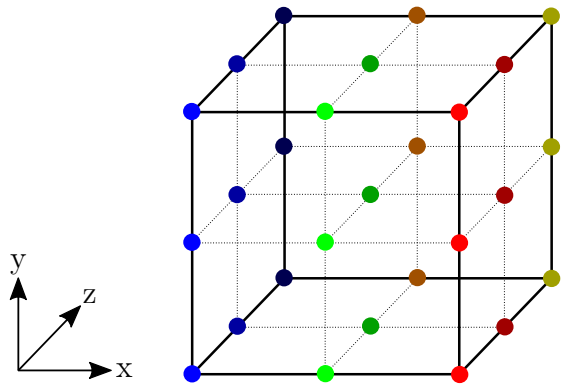
$$F(x_{i,\alpha}, z_{i,\alpha}, p_{i,\alpha}, q_{i,\alpha}, s_{i,\alpha}, w_{i,\alpha}, v_{i,\alpha}) \leq F(x_i, z_i, p_i, q_i, s_i, w_i, v_i) + \alpha_i \beta F'(x_i, z_i, p_i, q_i, s_i, w_i, v_i) \tag{3.97}$$

where β denotes a scaling factor defined in the interval $(0, 1]$, ensures that the step-size provides a sufficient and proportional improvement in the merit function. In combination, Eqs. 3.92 — 3.97 ensure that convergence of the interior-point optimisation algorithm is guaranteed.

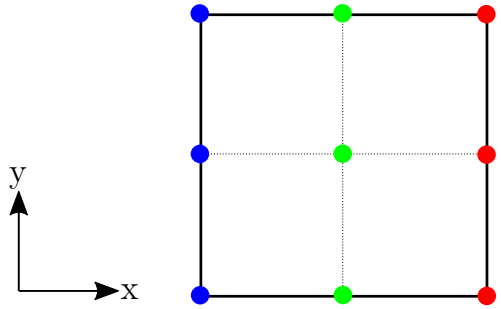
3.5.5 Addition of Microscale Parameters

Presented within this section is an extension to the interior-point optimisation algorithm presented within Sec. 3.5. The primary aim of this extension is to further reduce the computational expense of optimisation by initially optimising a subset of the microscale design parameters, before activating additional microscale design parameters as the optimisation progresses. The premise of this novel optimisation strategy is that the most sensitive microscale parameters have to traverse more microscale design space during the optimisation procedure in comparison to the insensitive microscale parameters. Therefore, by initially optimising only the most sensitive microscale parameters, a significant percentage of the path traversed through the microscale design space is undertaken within a sub-space. Consequently, the computational expense of simplex interpolation is reduced as the number of evaluations of the microscale model required scales linearly with the number of microscale parameters. It is assumed that the number of iterations required to converge the optimisation will increase through the application of this optimisation strategy, however, it is hoped that the reduction in the computational expense associated with microscale simulations outweighs the increased time to convergence.

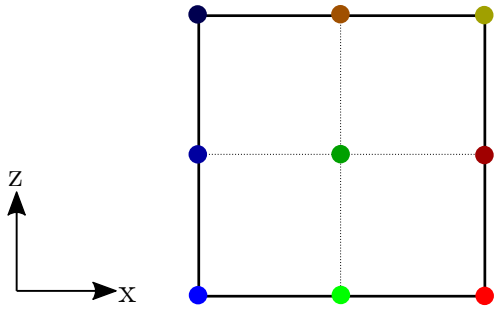
In order to implement this strategy, it is mandatory that C^1 continuity is maintained for each point within the microscale design space with respect to changes in the active microscale parameter subset. It important to note, that without C^1 continuity, it is likely that convergence issues would be encountered. The continuity requirement represents a further reason for the selection of a full-factorial DOE and an interpolation-based response surface, as in combination C^1 continuity is afforded. This is a result of the geometric composition of a full-factorial DOE, which preserves the uniform arrangement of nodes if one or multiple dimensions is removed. As a result, this enables nested full-factorial DOEs to be trivially isolated, as illustrated in Fig. 3.16. Consequently, interpolants derived for both the full set of microscale design parameters and all possible subsets are identical, as the terms associated with the deleted parameters do not vary across the interpolant and are therefore equal to zero. It is important to add that any microscale simulations conducted within a reduced microscale parameter space still populate and contribute towards the full parent microscale database.



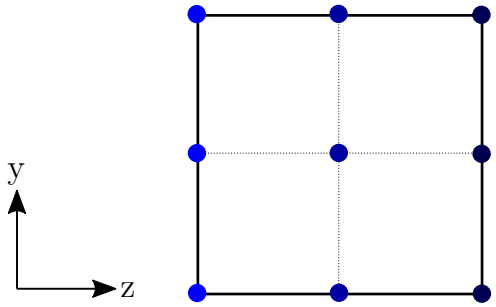
(a) Three-dimensional full-factorial DOE



(b) Isolated two-dimensional full-factorial DOE in the $X-Y$ plane



(c) Isolated two-dimensional full-factorial DOE in the $X-Z$ plane



(d) Isolated two-dimensional full-factorial DOE in the $Y-Z$ plane

Figure 3.16: How multiple nested two-dimensional full-factorial DOEs can be isolated from a three-dimensional full-factorial DOE. Three levels are used in the definition of each dimension, which results in a total of 27 points for the three-dimensional DOE and 9 points for each of the three two-dimensional DOEs.

3.5.5.1 Interior-point Optimisation Algorithm Extension

In order to implement this novel optimisation strategy, it is necessary to derive an aggregate measure for the sensitivity of each microscale parameter to aid relative comparison

$$\theta_\beta = \sum_{\alpha=1}^{N} \left| \nabla J \left(\tilde{X}_{\alpha\beta} \right) \right| \tag{3.98}$$

where θ_β denotes the aggregate sensitivity for the microscale parameter $\tilde{X}_\beta$. For convenience, a normalisation is applied

$$\hat{\theta}_\beta = \frac{\theta_\beta}{\max\{\theta\}} \tag{3.99}$$

to place a limit of 1 on $\hat{\theta}$. The aggregate sensitivities are then used to rank the microscale parameters in order of sensitivity. The optimisation procedure is then initialised using a user-defined number of the most sensitive microscale parameters, denoted for the remainder of this section as the active-set. It is important to add, that the inactive-set of microscale parameters needs to be initialised with a value equal to a DOE level (see Tab. 3.1) to prevent erroneous interpolation. It is important to highlight that an aggregate measure is used, as it is not possible to spatially vary the microscale parameter subset without violating the smoothness assumptions of the heterogeneous multiscale method.

Aggregate sensitivities for all microscale parameters are periodically derived during optimisation to check whether an inactive microscale parameter outranks an active microscale parameter. If this is the case, the inactive microscale parameters which outrank any of the active microscale parameters are introduced into the optimisation procedure as new design variables. Further to this, in situations where an inactive microscale parameter does not outrank an active microscale parameter over the course of an optimisation, a predefined iteration is set at which all microscale parameters are introduced into the active set.

Introducing additional design variables into a partially converged optimisation problem presents a number of challenges, as the KKT linear system (see. Eq. 3.69) transitions from a linear system with dimensions

$$\begin{bmatrix} [n\times n] & [n\times m] & [n\times n] & [n\times n] & 0 & 0 & 0 \\ [m\times n] & 0 & 0 & 0 & [m\times m] & 0 & 0 \\ [n\times n] & 0 & 0 & 0 & 0 & [n\times n] & 0 \\ [n\times n] & 0 & 0 & 0 & 0 & 0 & [n\times n] \\ 0 & [m\times m] & 0 & 0 & [m\times m] & 0 & 0 \\ 0 & 0 & [n\times n] & 0 & 0 & [n\times n] & 0 \\ 0 & 0 & 0 & [n\times n] & 0 & 0 & [n\times n] \end{bmatrix} \begin{bmatrix} [n\times 1] \\ [m\times 1] \\ [n\times 1] \\ [n\times 1] \\ [m\times 1] \\ [n\times 1] \\ [n\times 1] \end{bmatrix} = \begin{bmatrix} [n\times 1] \\ [m\times 1] \\ [n\times 1] \\ [n\times 1] \\ [m\times 1] \\ [n\times 1] \\ [n\times 1] \end{bmatrix} \tag{3.100}$$

to a new linear system with dimensions

$$\begin{bmatrix} [\tilde{n}\times\tilde{n}] & [\tilde{n}\times m] & [\tilde{n}\times\tilde{n}] & [\tilde{n}\times\tilde{n}] & 0 & 0 & 0 \\ [m\times\tilde{n}] & 0 & 0 & 0 & [m\times m] & 0 & 0 \\ [\tilde{n}\times\tilde{n}] & 0 & 0 & 0 & 0 & [\tilde{n}\times\tilde{n}] & 0 \\ [\tilde{n}\times\tilde{n}] & 0 & 0 & 0 & 0 & 0 & [\tilde{n}\times\tilde{n}] \\ 0 & [m\times m] & 0 & 0 & [m\times m] & 0 & 0 \\ 0 & 0 & [\tilde{n}\times\tilde{n}] & 0 & 0 & [\tilde{n}\times\tilde{n}] & 0 \\ 0 & 0 & 0 & [\tilde{n}\times\tilde{n}] & 0 & 0 & [\tilde{n}\times\tilde{n}] \end{bmatrix} \begin{bmatrix} [\tilde{n}\times 1] \\ [m\times 1] \\ [\tilde{n}\times 1] \\ [\tilde{n}\times 1] \\ [m\times 1] \\ [\tilde{n}\times 1] \\ [\tilde{n}\times 1] \end{bmatrix} = \begin{bmatrix} [\tilde{n}\times 1] \\ [m\times 1] \\ [\tilde{n}\times 1] \\ [\tilde{n}\times 1] \\ [m\times 1] \\ [\tilde{n}\times 1] \\ [\tilde{n}\times 1] \end{bmatrix} \tag{3.101}$$

where $\tilde{n}$ denotes the new (larger) number of design variables. Consequently, in order to accommodate an increase in the number of design variables, additional slack variables (p, q) and Lagrange multipliers (w, v) need to be introduced. In contrast, the slack variables (s) and Lagrange multipliers (z) associated with the constraints remain fixed in length, as no additional constraints are imposed on the design variables.

Interior-point optimisation algorithms are sensitive to the initial choice of design variables, slack variables and Lagrange multipliers, as starting points located in close proximity to the boundary of a constraint or bound can lead to a significantly reduced maximum step-size (Gertz et al., 2004). Consequently, in these situations, it is likely that progress with respect to the objective function is ignored, as other first-order optimality conditions (see Eqs. 3.60 — 3.66) temporarily take precedent, resulting in a significant step away from the optimum and towards the interior of the feasible region. It is important to add that significant steps away from the optimal path through the microscale design space are expensive with respect to microscale simulations. As the target of the optimisation strategy presented within this section is to reinitialise a partially converged optimisation problem, it is likely that the aforementioned problems will be encountered.

To alleviate this issue, the partially optimised design variables (x), slack variables (p, q) and Lagrange multipliers (w, v) are used to initialise the corresponding variables in the new optimisation problem, which reduces the impact of initialising these variables in close proximity

to the boundary of the feasible region. It is, however, not possible to carry forward any of the information associated with the L-BFGS matrices, as the variables G, Y, L and D (see Eqs. 3.79 — 3.84) no longer represent square or rectangular matrices and therefore prohibit the linear algebra calculations required to derive the Hessian approximation.

In addition, it is also necessary to carry forward the initialised values for the slack variables (s_0, p_0, q_0) and Lagrange multipliers (w_0, v_0, z_0) in order to enable the line-search conditions defined within Eqs. 3.95 and 3.96 to be successfully applied within the reinitialised optimisation. This is necessary, as the partially converged slack variables and Lagrange multipliers will be in close proximity to the boundary of the feasible region, which in this instance could lead to significant constraints placed on the maximum step-size for the remainder of the optimisation procedure if these variables are re-initialised.

3.6 Macroscale Geometry Reconstruction

Once an optimisation has converged, the optimised normalised radii and aspect-ratio distributions are reconstructed into a single macroscale geometry. The reconstruction is derived by populating the macroscale geometry with volumetric data describing the distance from each structural member surface. The volumetric data uses 25^3 subdivisions for each unit cell. A marching cubes algorithm (Lewiner et al., 2003) is then used to reliably convert the volumetric data into a surface mesh, which enables a single manifold to be derived[2].

3.7 Summary

Presented within this chapter is a detailed breakdown of the methodology applied to construct the concurrently coupled MTO framework and the novel optimisation strategy (which includes the derivation of the foundational interior point optimisation algorithm). These methodologies

[2]Code written by Hewson (Murphy et al., 2020)

are applied to a series of mechanical structural optimisation problems within Chapter 4 to demonstrate their efficacy.

Chapter 4

Multiscale Topology Optimisation Results

The MTO framework presented within Chapter 3 is applied here to solve a range of structural optimisation problems. The results are separated into two sections, one for the base MTO framework and another for the base MTO framework with addition of microscale parameters. Optimisation is terminated for all problems presented within this section when both the objective function decreases by less than 0.01% for five successive iterations and the normalised constraint violation is less than 1×10^{-3}.

4.1 Base Multiscale Topology Optimisation Framework

The base MTO framework is applied within this section to solve two classical compliance minimisation problems. The optimised results and the computational expenditure are presented in detail for both problems. In addition, the optimised results are benchmarked against a SIMP based TO, to understand the relative merit of this approach in comparison to the de facto standard method for structural optimisation. The initialisation conditions are kept constant for both problems presented within this section and are detailed in Tab. 4.1.

Design Variable	Initialisation Value
$\hat{r}_{1-7}$	0.07
e_{5-7}	1.00

Table 4.1: The initialisation values used for all parameters at the beginning of optimisation

4.1.1 Cantilever Beam — Compliance Minimisation

The first geometry considered is a cantilever beam with a uniformly distributed load applied to the top surface, as illustrated in Fig. 4.1. A 30% volume-fraction constraint is applied. Compliance in this instance is quantified by

$$J = \int_{\Gamma_1} \boldsymbol{t} \cdot \boldsymbol{u} \, d\Gamma_1 \tag{4.1}$$

where Γ_1 denotes the top surface of the cantilever geometry, as illustrated in Fig. 4.1. The cantilever geometry is discretised into $24,000$ first-order hexahedral elements using $60 \times 20 \times 20$ divisions in the x, y and z axes respectively. The optimised normalised radii distributions are illustrated in Fig. 4.2, with two thresholds applied to highlight the salient structures. A visualisation threshold $\tilde{\hat{r}}_{\alpha,i} \geq 0.3$ is applied to highlight the dominant structures formed during optimisation and the second (transparent) threshold $0 \leq \tilde{\hat{r}}_{\alpha,1-7} < 0.3$ is applied to highlight where material is allocated within the design domain. The optimised aspect-ratio distributions relating the orthogonal axes of the face members are illustrated in Fig. 4.3. To permit clearer visualisation, these distributions are overlaid onto the primary normalised radii thresholds. Additionally, the convergence plot for this problem is illustrated in Fig. 4.4.

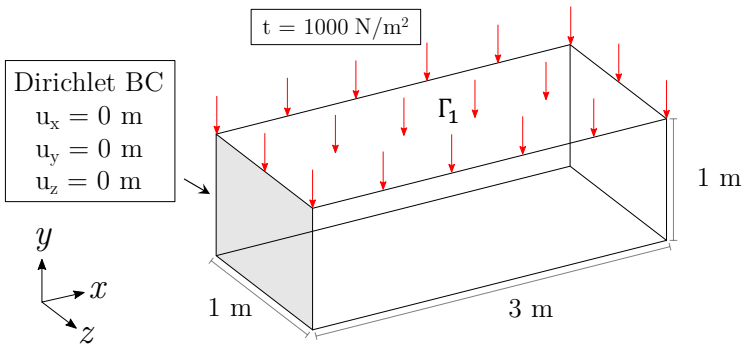
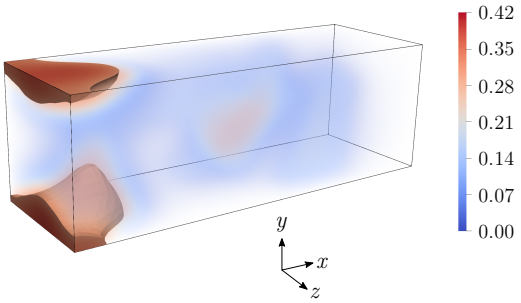
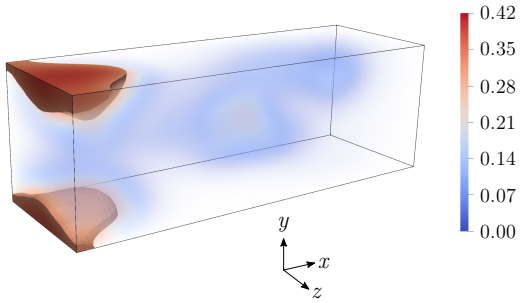


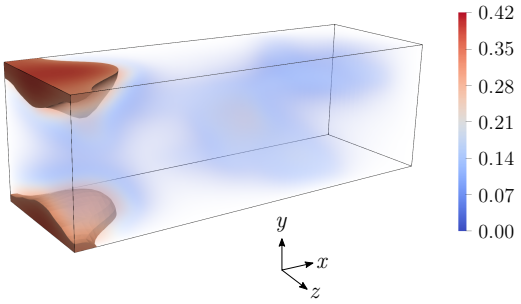
Figure 4.1: The cantilever beam geometry targeted by compliance minimisation, with the loads and boundary conditions highlighted



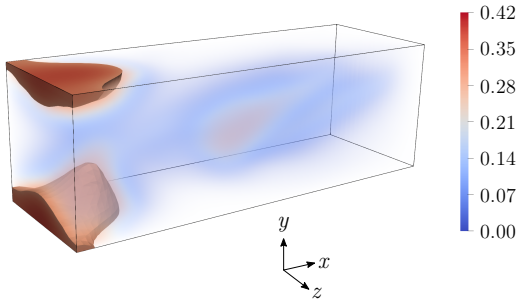
(a) Structural Member 1 ($\tilde{\hat{r}}_{\alpha,1}$)



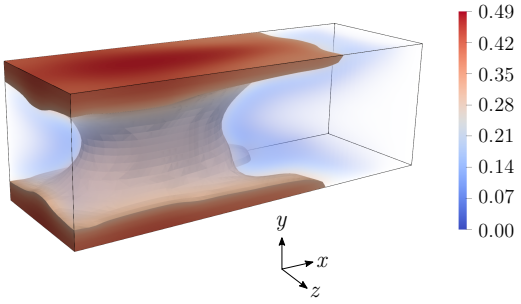
(b) Structural Member 2 ($\tilde{\hat{r}}_{\alpha,2}$)



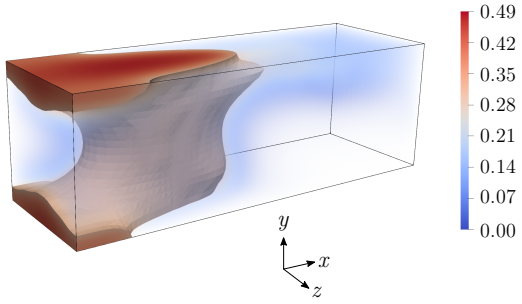
(c) Structural Member 3 ($\tilde{\hat{r}}_{\alpha,3}$)



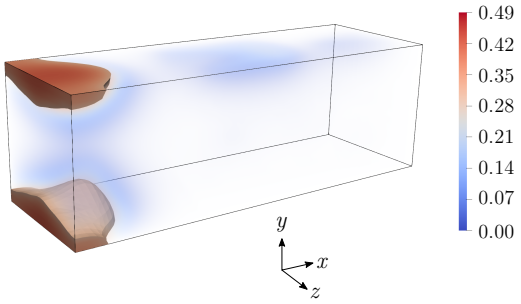
(d) Structural Member 4 ($\tilde{\hat{r}}_{\alpha,4}$)



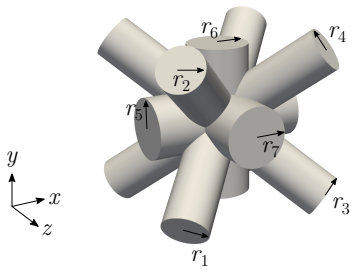
(e) Structural Member 5 ($\tilde{\hat{r}}_{\alpha,5}$)



(f) Structural Member 6 ($\tilde{\hat{r}}_{\alpha,6}$)

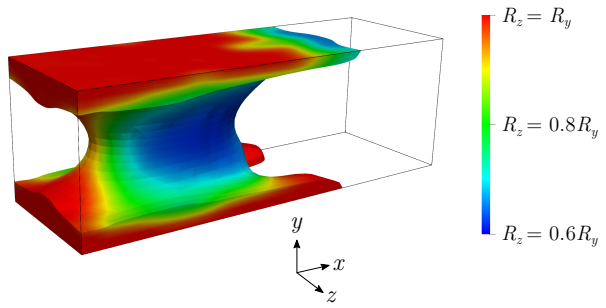


(g) Structural Member 7 ($\tilde{\hat{r}}_{\alpha,7}$)

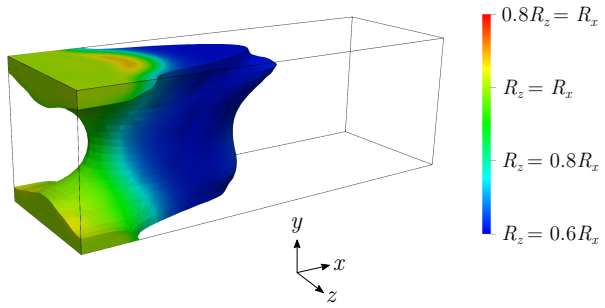


(h) Microstructure for reference

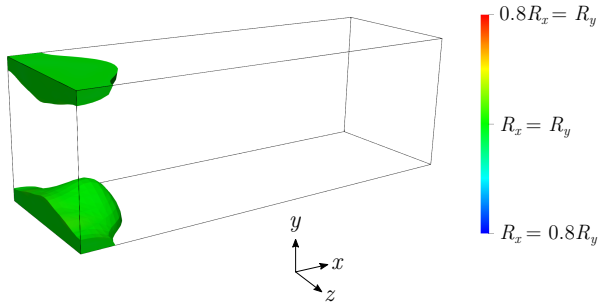
Figure 4.2: The optimised normalised radii distributions for the cantilever compliance minimisation problem. One threshold $\tilde{\hat{r}}_{\alpha,i} \geq 0.3$ is applied to highlight the dominant structures formed during optimisation and the second (transparent) threshold $0 \leq \tilde{\hat{r}}_{\alpha,1-7} < 0.3$ is applied to highlight where material is allocated within the design domain.



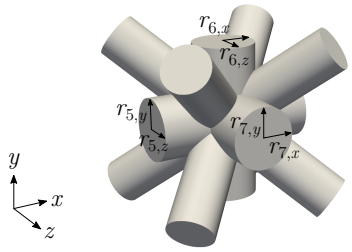
(a) Structural Member 5 — cross-section aspect-ratio ($\tilde{e}_{\alpha,5}$)



(b) Structural Member 6 — cross-section aspect-ratio ($\tilde{e}_{\alpha,6}$)



(c) Structural Member 7 — cross-section aspect-ratio ($\tilde{e}_{\alpha,7}$)



(d) Microstructure for reference

Figure 4.3: The optimised aspect-ratio distributions for the cantilever compliance minimisation problem. To permit clearer visualisation, these distributions are overlaid onto the primary normalised radii thresholds ($\tilde{\hat{r}}_{\alpha,1-7} \geq 0.3$).

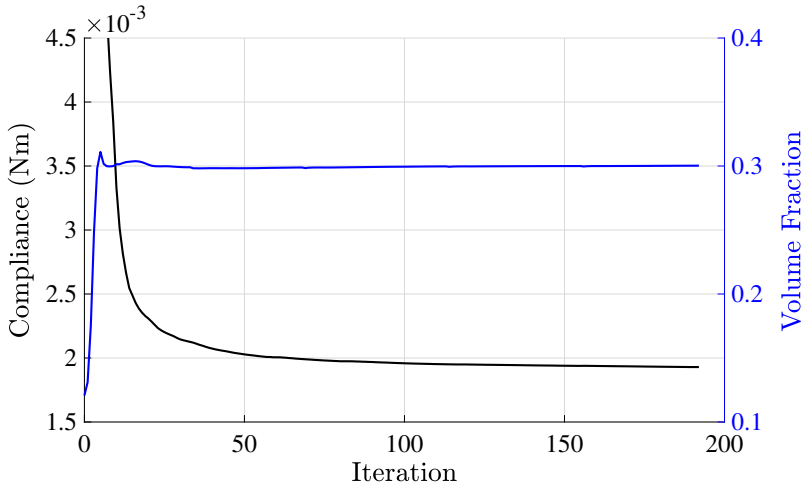
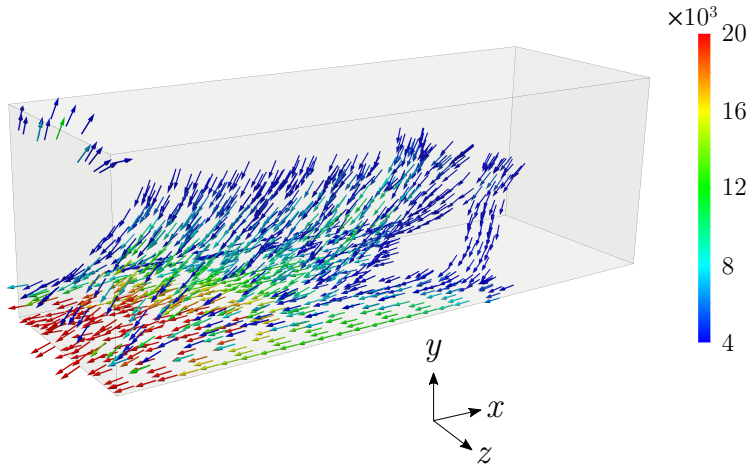


Figure 4.4: The convergence plot for the cantilever beam compliance minimisation

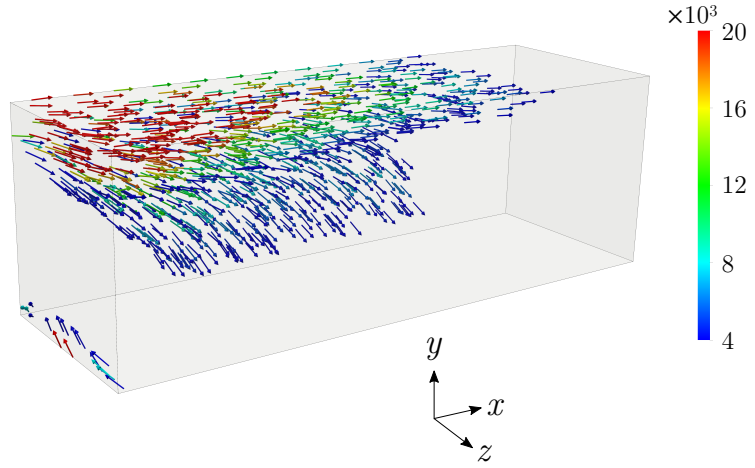
As depicted in Fig. 4.2, material is primarily allocated near to the root of the cantilever beam geometry. It is in this region of the cantilever that the largest principal stresses occur, as illustrated in Fig 4.5. Therefore, by allocating material near to the root, the resistance

to bending is increased. Structural member 5 (Fig. 4.2e) receives the highest allocation of material and is distributed in the shape of an I-beam. The principal stresses in the flanges of the I-beam are orientated parallel to member 5, which enables this member to provide an efficient resistance to the defined load. This is an example of how a microscale parameter can be configured to couple tightly with specific elements of the elasticity tensor (E_{11}), enabling clear accessibility to efficient stiffness characteristics during optimisation. Structural member 6 (Fig. 4.2f), which is orientated parallel to the force vector, forms an approximate I-beam shape also, but with smaller flanges and a thicker web section than member 5. The difference in shape is attributed to the orientation of member 6 aligning more closely to the principal stresses in the web, than at the flanges. The diagonal members (Figs. 4.2a – 4.2d) are primarily distributed at the root of the cantilever beam, as these members are able to provide an efficient resistance in regions where the principal stress directions transition between the web and flanges of the I-beam. Symmetries are noted about the $X - Y$ plane at $z = 0.5$ for members 1 and 4 and members 2 and 3, as these member pairs are identically orientated in the $X - Y$ plane. Structural member 7 (Fig. 4.2g) receives the lowest allocation of material, due to its orientation to the z-axis and is therefore insensitive to the compliance functional. As a result, this member is distributed in similar regions to the diagonal members, to also assist with the transition of the principal stress directions between the web and flanges.

It is apparent that the structural members are primarily distributed in regions where the principal stress directions are parallel to the respective member orientations. However, due to the finite number of members, it is not always possible to have a member directly orientated to the principal stresses. In this situation, a combination of members is used instead. By varying the two orthogonal axes defining each face member, it is possible to make these combinations more efficient. This is demonstrated in Fig. 4.3a for structural member 5, where the aspect-ratio tends to the minimum permissible value (0.6) in the web section of I-beam, where the principal stresses are not directly aligned to any member, but in contrast remained unchanged from the initialisation value (1.0) in regions where there is a direct alignment, for example in the flanges. The normalised radii parameters provide sufficient control over the point properties in regions where the principal stresses are directly aligned and therefore changes



(a) Minimum principal stresses (Pa)



(b) Maximum principal stresses (Pa)

Figure 4.5: The maximum and minimum principal stresses for the optimised cantilever beam geometry, derived from the homogenised stress tensors. To permit clearer visualisation, a threshold is applied to eliminate regions where the principal stresses are less than 4 kPa.

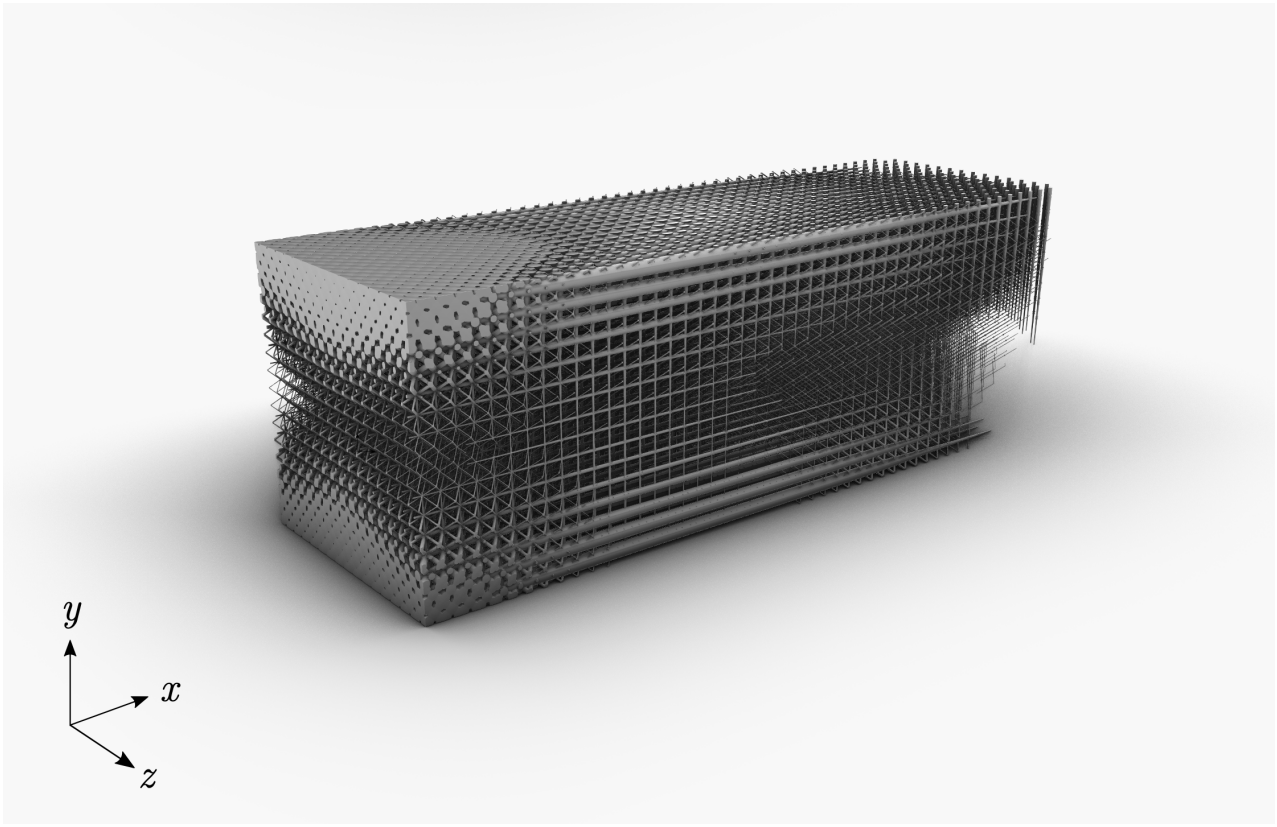
in the aspect-ratios are not required. However, to compensate in regions where there is not a direct alignment, changes to the aspect-ratio enable the secondary component of stiffness to be efficiently increased. For structural member 5, this corresponds to increasing $r_{5,y}$ relative to $r_{5,z}$, enabling the stiffness properties of this member to be enhanced in the y direction at the expense of the insensitive z direction for an equivalent cross-sectional area. A similar pattern is also noted for structural member 6 (Fig. 4.3b), where $r_{6,x}$ is increased relative to $r_{6,z}$ in the web section of the I-beam to increase stiffness in the x direction, whilst sacrificing stiffness in the insensitive z direction. Due to the insensitivity of structural member 7 (Fig. 4.3c) to the compliance functional, no significant changes in aspect-ratio are noted from the initialisation

conditions (see Tab. 4.1).

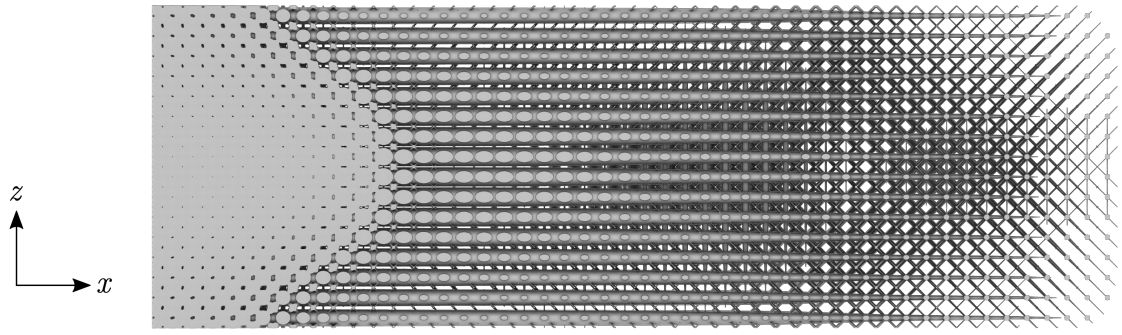
The optimal parameter distributions, illustrated in Figs. 4.2 and 4.3, are combined to construct the physical representation of the macroscale structure, as illustrated in Fig. 4.6. The number of unit cells per dimension is equal to $48 \times 16 \times 16$ in the x, y and z directions respectively and the corresponding manifold mesh is over 2.3 GB in size. The scale and resolution of the reconstructed geometry is selected to aid with visualisation and is therefore coarser than the geometry which would be manufactured.

The computational expense incurred by microscale simulations over the course of this compliance minimisation is illustrated in Fig. 4.7, by plotting the cumulative number of parent microscale simulations required after each iteration. In total, $32,490$ parent microstructures are required to derive the optimised geometry. This correspond to 0.48% of the total parent database, which represents a significant reduction in comparison to sequential frameworks, which require a precomputed fully-populated database. The number of parent microscale simulations required to perform the first 10 iterations of optimisation is modest (364), as all microstructures are initialised with the same values and are therefore initially contained within a single simplex. This requirement then increases linearly at an average of 473 simulations per iteration until iteration 60, at which point the optimisation begins to converge. As a result, a smaller proportion of the design space is explored in successive iterations, therefore increasingly fewer parent microscale simulations are required.

Equivalent compliance minimisations are performed using the cantilever beam geometry, loads, boundary conditions and discretisation for the base 7-dimensional microstructure (Imediegwu et al., 2019) and a SIMP based topology optimisation with a penalisation parameter of 3. A penalisation parameter of 3 is selected for comparison, as it is the value most commonly applied within industrial software and therefore provides a coherent benchmark. An identical discretisation is chosen for the SIMP case to aid comparison, as any increase in the mechanical performance afforded through the MTO framework can be directly attributed to the lower inherent penalisation and the accessibility afforded to anisotropic point properties. As the smoothness requirements for topology optimisations are less stringent than for the present



(a) Isometric projection of reconstructed cantilever beam



(b) Orthogonal projection of reconstructed cantilever beam onto $X - Z$ plane

Figure 4.6: Reconstruction of the optimised parameter distributions (see Figs. 4.2 and 4.3) into a single macroscale geometry

framework, a more typical Helmholtz filter length of 1 cell length (Lazarov and Sigmund, 2011) is used instead.

The compliance functionals for both frameworks are compared in Tab. 4.2. Through the application of the concurrent multiscale framework in combination with the present 10 variable microstructure parameterisation (Fig. 3.1), a 27.4% reduction in the compliance functional is achieved in comparison to SIMP. This is a result of the clear accessibility afforded by the mi-

croscale parameterisation to orthotropic microstructures with efficient stiffness characteristics, in comparison to the fixed isotropic material properties locally available within SIMP.

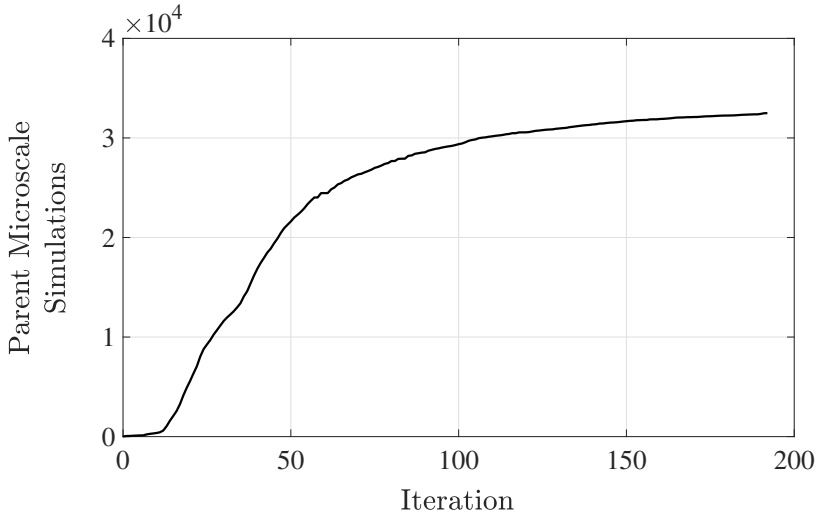


Figure 4.7: The cumulative number of parent microscale simulations required to derive the optimised cantilever beam geometry

A further reduction of 2.1% is also noted in comparison to the base 7-dimensional microstructure, demonstrating the increased accessibility to efficient stiffness characteristics permitted by the three additional design parameters. However, it should be reiterated that the principal objective of this thesis is to demonstrate the computational savings afforded by a concurrent multiscale approach and not the optimality of the present microstructure with respect to the Hashin-Shtrikman upper bounds for bulk and shear moduli (Hashin and Shtrikman, 1963).

Method	Compliance (Nm)
Concurrent Multiscale (10 Parameters)	0.001931
Concurrent Multiscale (7 Parameters)	0.001972
SIMP ($p = 3$)	0.002636

Table 4.2: The optimised compliance functionals derived for the cantilever beam compliance minimisations, for both the present concurrent multiscale framework and a SIMP based topology optimisation. The 10 and 7 parameter optimisations correspond to the present and base (Imediegwu et al., 2019) microstructure respectively.

4.1.2 Unit Cube — Compliance Minimisation

The second geometry targeted for compliance minimisation is a unit-cube with four equal and uniformly distributed loads applied to four unique surfaces, constrained by a single Dirichlet boundary condition, as illustrated in Fig. 4.8. Compliance in this instance is quantified by

$$J = \int_{\Gamma_1} \boldsymbol{t_1} \cdot \boldsymbol{u}\, d\Gamma_1 + \int_{\Gamma_2} \boldsymbol{t_2} \cdot \boldsymbol{u}\, d\Gamma_2 + \int_{\Gamma_3} \boldsymbol{t_3} \cdot \boldsymbol{u}\, d\Gamma_3 + \int_{\Gamma_4} \boldsymbol{t_4} \cdot \boldsymbol{u}\, d\Gamma_4$$

where Γ_i and t_i denote the 4 surfaces and surface tractions illustrated in Fig. 4.8 respectively. The unit-cube geometry is discretised into $27,000$ first-order hexahedral elements using $30 \times 30 \times 30$ divisions in the x, y and z axes respectively.

To demonstrate the reusability of the microscale property data collected during optimisation, a parametric study of volume-fraction is conducted. A total of 7 volume fractions are investigated, starting at 20% and increasing linearly in 5% intervals before stopping at 50%. The optimised normalised radii distributions are illustrated for the 30% volume-fraction case only in Fig. 4.9, using equivalent normalised radii thresholds to those applied to the cantilever beam results, but with a lower bound for the primary threshold of 0.25. Similarly, the optimised aspect-ratio distributions relating the orthogonal axes of the face members are illustrated for the 30% volume-fraction case also in Fig. 4.10. Additionally, the convergence plot for this problem is illustrated in Fig. 4.11.

As depicted in Fig. 4.9, material is allocated for all members at the base of the unit-cube near the displacement constraint, as this is where the largest principal stresses occur, as illustrated in Fig. 4.12. Structural member 6 (Fig. 4.9f) receives the highest allocation of material, as it is orientated parallel to the two loads applied to the top surface of the unit-cube and therefore provides an efficient resistance and a load path to the boundary constraint. Structural member 5 (Fig. 4.9e), is also allocated a significant proportion of the available material, as it is orientated parallel to the two forces applied to the side surfaces ($Y - Z$) of the unit-cube and therefore also provides an efficient resistance. The four loads applied to the unit-cube are aligned to either the x or y axes and act in tension, which results in a compression of

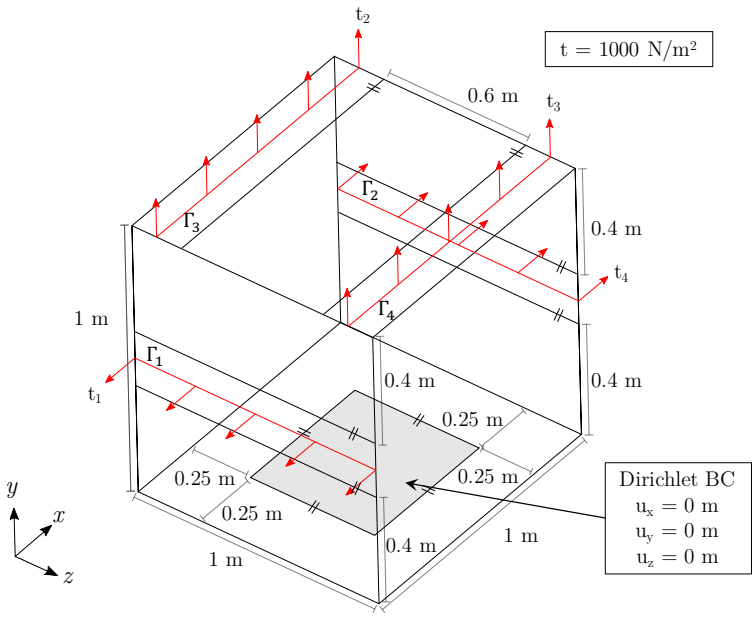
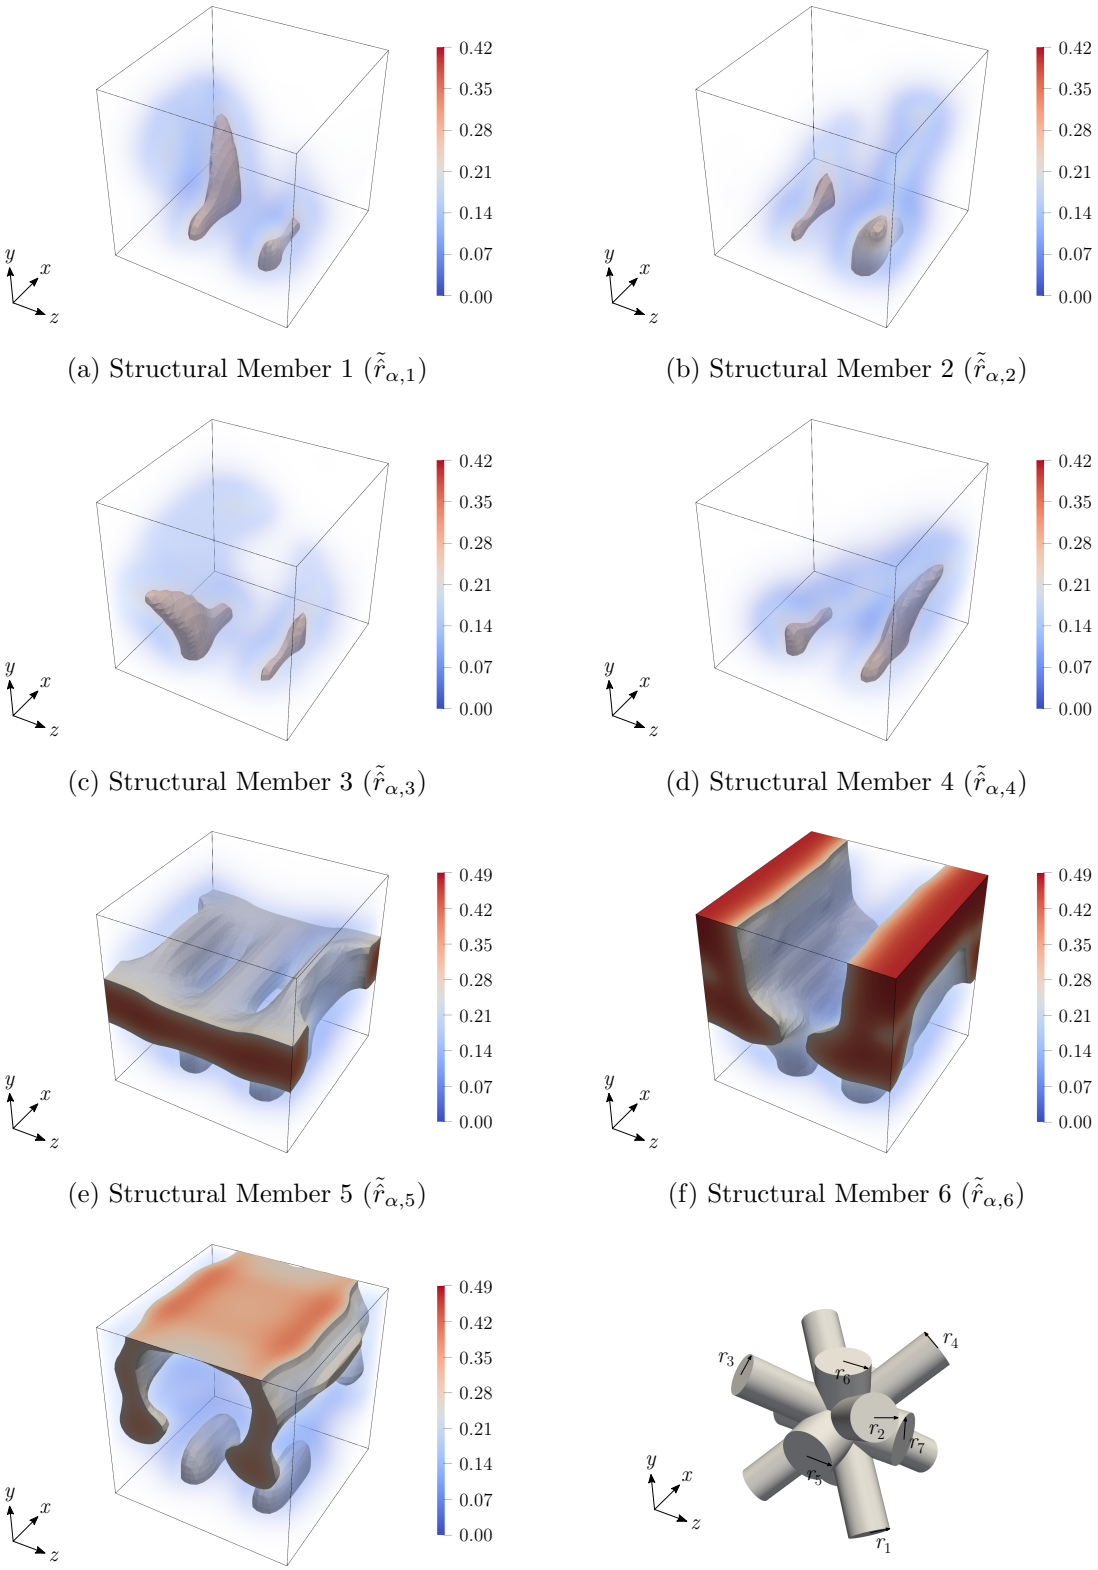


Figure 4.8: The unit-cube geometry targeted by compliance minimisation, with the loads and boundary conditions highlighted

the unit-cube in the z axis, as demonstrated by the minimum principal stresses (Fig. 4.5a). As a result, structural member 7 (Fig. 4.9g), is distributed in regions where the minimum principal stresses are largest, as its orientation to the z-axis enables the compression forces to be efficiently resisted. The diagonal members (Figs. 4.9a – 4.9d) form structures in regions where both force vectors contribute significantly to the displacement field, as the principal stresses in these regions have components in all three coordinate axes. A 90° rotational symmetry in the $X - Z$ plane about the centroid of the base is also noted for the diagonal members, as the orientation of each member is suitable for a different combination of the two force vectors.

The optimised aspect-ratios defining the orthogonal axes of the face members display similar patterns to those noted for the cantilever beam compliance minimisation. In regions where the principal stresses are parallel to a members orientation, the aspect-ratio remains unchanged from the initialisation value of 1 (see Tab. 4.1), but in regions where there is a misalignment, the aspect-ratio is altered to increase stiffness in the secondary direction for an equivalent cross-sectional area. For example, the dominant axis for structural member 6 (Fig. 4.10b) near to the top surface of the unit-cube is $R_{6,z}$, as this is where the strongest compression



(a) Structural Member 1 ($\tilde{\hat{r}}_{\alpha,1}$)

(b) Structural Member 2 ($\tilde{\hat{r}}_{\alpha,2}$)

(c) Structural Member 3 ($\tilde{\hat{r}}_{\alpha,3}$)

(d) Structural Member 4 ($\tilde{\hat{r}}_{\alpha,4}$)

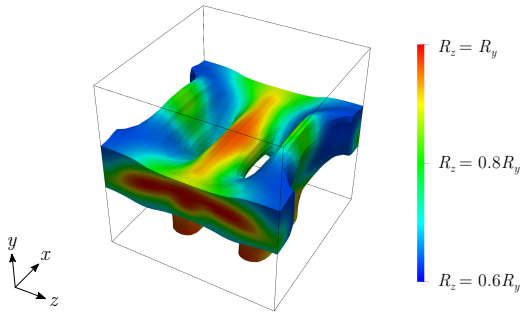
(e) Structural Member 5 ($\tilde{\hat{r}}_{\alpha,5}$)

(f) Structural Member 6 ($\tilde{\hat{r}}_{\alpha,6}$)

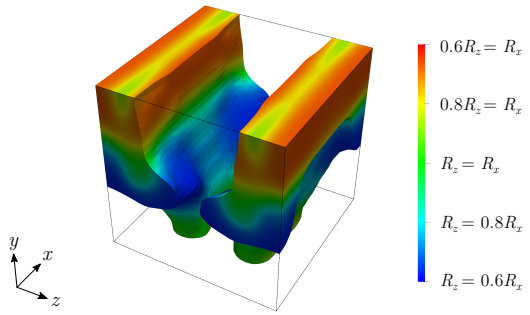
(g) Structural Member 7 ($\tilde{\hat{r}}_{\alpha,7}$)

(h) Microstructure for reference

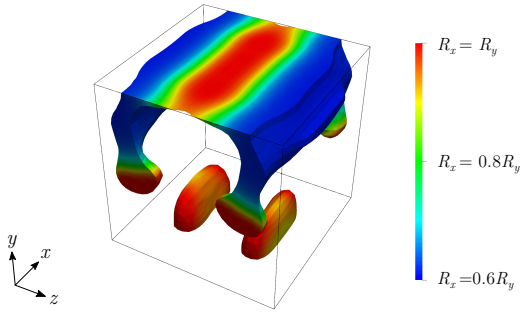
Figure 4.9: The optimised normalised radii distributions for the unit-cube compliance minimisation problem, with a 30% volume-fraction constraint. One threshold $\tilde{\hat{r}}_{\alpha,i} \geq 0.25$ is applied to highlight the dominant structures formed during optimisation and the second (transparent) threshold $0 \leq \tilde{\hat{r}}_{\alpha,1-7} < 0.25$ is applied to highlight where material is allocated within the design domain.



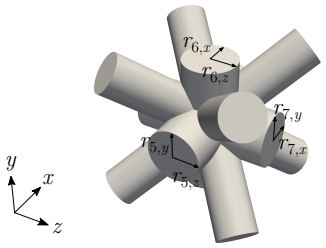
(a) Structural Member 5 — cross-section aspect-ratio ($\tilde{e}_{\alpha,5}$)



(b) Structural Member 6 — cross-section aspect-ratio ($\tilde{e}_{\alpha,6}$)



(c) Structural Member 7 — cross-section aspect-ratio ($\tilde{e}_{\alpha,7}$)



(d) Microstructure for reference

Figure 4.10: The optimised aspect-ratio distributions for the unit-cube compliance minimisation problem, with a 30% volume-fraction constraint. To permit clearer visualisation, these distributions are overlaid onto the primary normalised radii thresholds ($\tilde{\hat{r}}_{\alpha,1-7} \geq 0.25$).

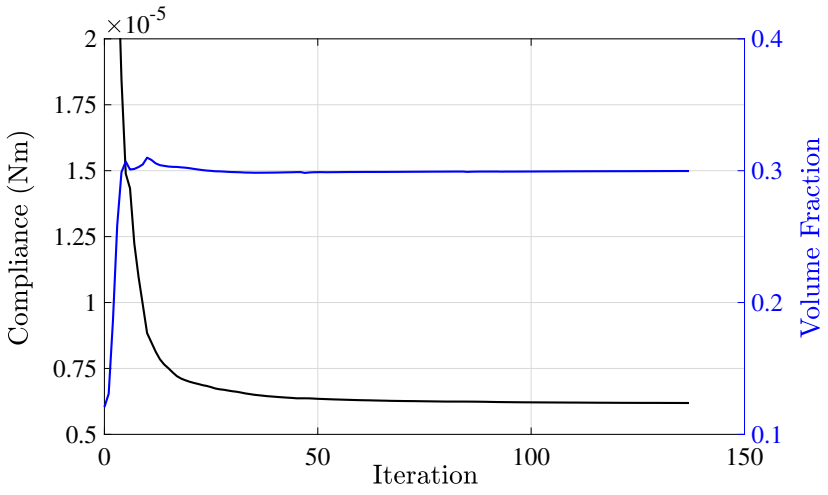
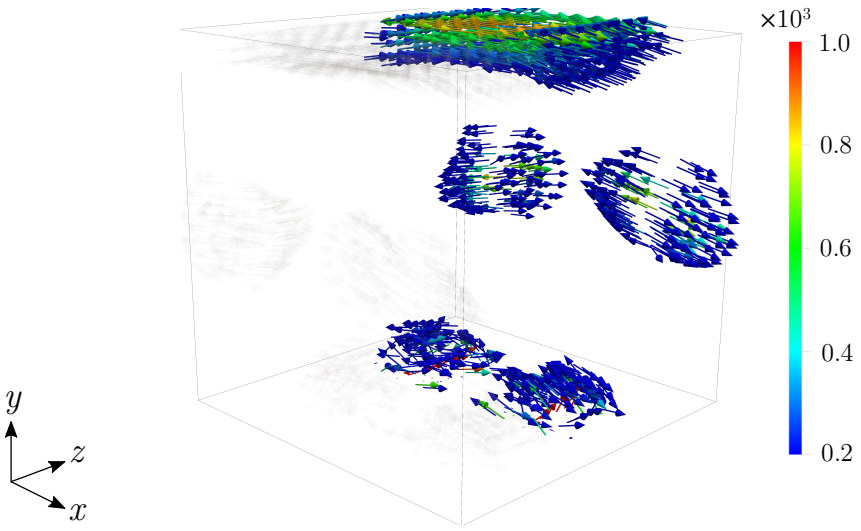


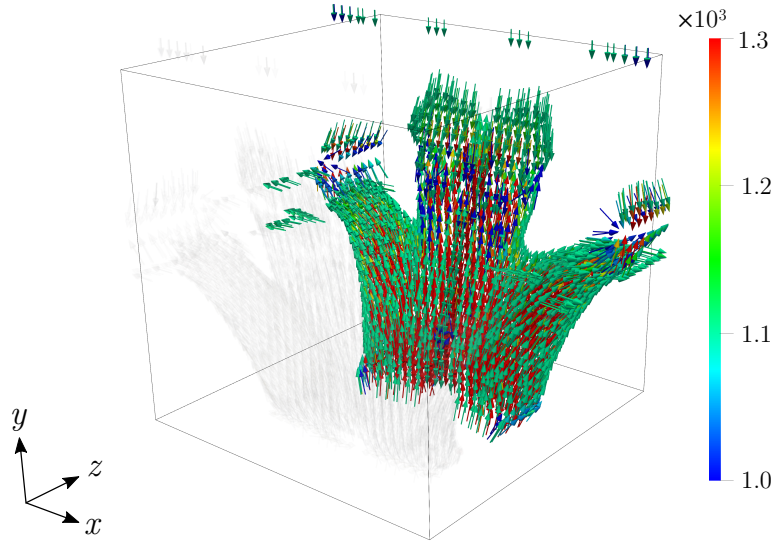
Figure 4.11: The convergence plot for the unit-cube compliance minimisation, with a 30% volume-fraction constraint

forces act. However, the dominant axis transitions to $R_{6,x}$ nearer to the base of the unit-cube, as the loads applied to the side surfaces affect the displacement field more significantly within this region. Similarly, at the corners of the unit-cube, the dominant axis for structural member

5 (Fig. 4.10a) is $R_{5,y}$, as the secondary component of force acts in the y direction, as depicted by the maximum principal stresses (Fig. 4.12b). A similar pattern is also noted for structural member 7 (Fig. 4.10c), where the dominant axis is $R_{7,y}$ at the top of the unit-cube, where the largest maximum principal stresses are aligned to the y axis.



(a) Minimum principal stresses (Pa) — a threshold is applied to eliminate regions where the principal stresses are less than 0.2 kPa.

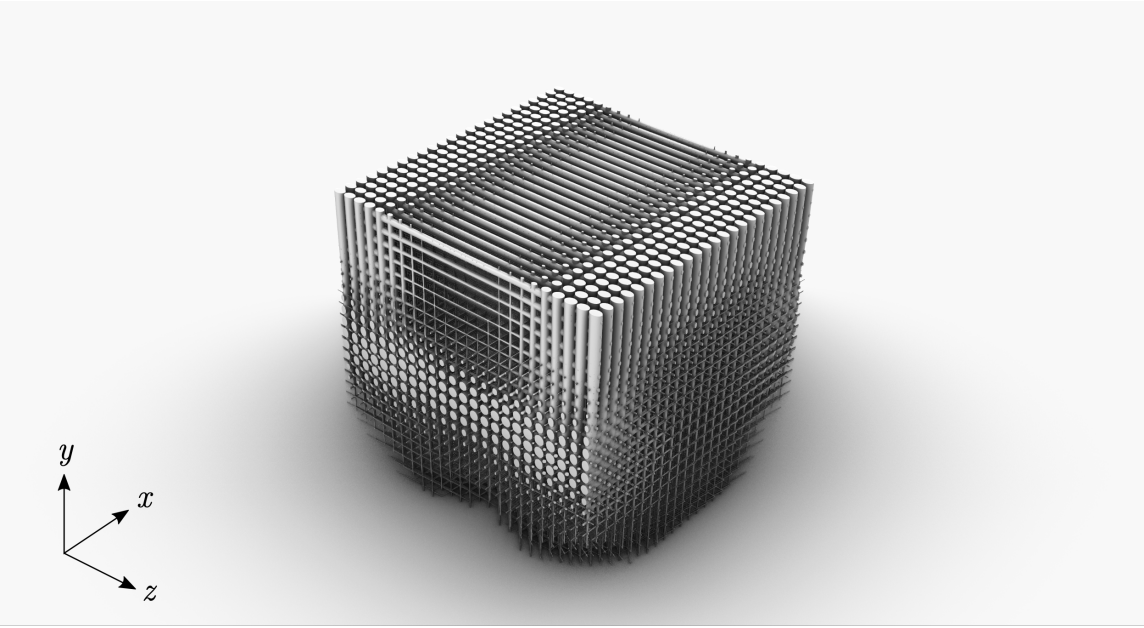


(b) Maximum principal stresses (Pa) — a threshold is applied to eliminate regions where the principal stresses are less than 1 kPa.

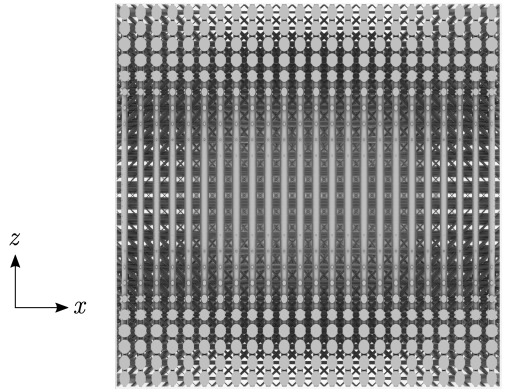
Figure 4.12: The maximum and minimum principal stresses for the optimised unit-cube geometry, derived from the homogenised stress tensors. To permit clearer visualisation, only the principal stresses in one half of the unit-cube domain are illustrated.

The optimised parameter distributions, illustrated in Figs. 4.9 and 4.10, are combined to

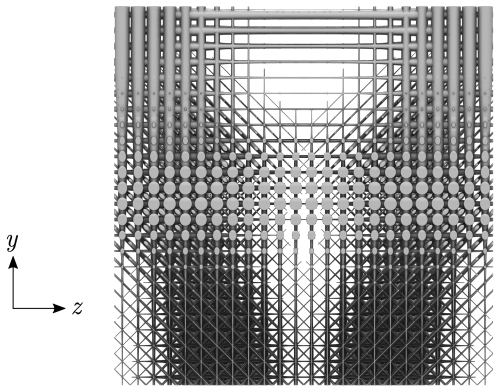
construct the physical representation of the macroscale structure, as illustrated in Fig. 4.13. The number of unit cells per dimension is equal to $24 \times 24 \times 24$ in the x, y and z directions respectively and the corresponding manifold mesh is over 2.6 GB in size. Again, the scale and resolution of the reconstructed geometry is selected to aid with visualisation and is therefore coarser than the geometry which would be manufactured.



(a) Isometric projection of reconstructed cube geometry



(b) Orthogonal projection of reconstructed cube geometry onto $X - Z$ plane



(c) Orthogonal projection of reconstructed cube geometry onto $Y - Z$ plane

Figure 4.13: Reconstruction of the optimised parameter distributions (see Figs. 4.9 and 4.10) into a single macroscale geometry

The individual computational expense incurred by microscale simulations for all 7 optimisations, each with a unique volume-fraction constraint, is illustrated in Fig. 4.14. The optimisations were conducted in ascending order of volume-fraction, starting at 20% and finishing at

50%. The microscale property data collected during the cantilever beam compliance minimisation serves as the initial database for this parametric study.

The 20% volume-fraction case requires a total of 20, 434 parent microscale simulations to derive the optimised geometry, however 5, 049 of those were already derived during the cantilever beam compliance minimisation. As a result, the reusability of the database permitted a reduction in the computational expense incurred by microscale simulations of 24.7%. In general, excluding the 25% volume-fraction case, the number of parent microscale simulations required increased with volume-fraction, as all simulations used the same initialisation conditions and therefore started at the same volume-fraction of 12%. Therefore, the optimisations with higher permitted volume-fractions have to cover more design space to reach the optimised solution. Over the course of the 7 optimisations, on average 63.1% of the parent microscale data required is collected in a previous optimisation. This represents a significant reduction in the computational expense of microscale simulations incurred during optimisation and demonstrates the re-usability of the collected data.

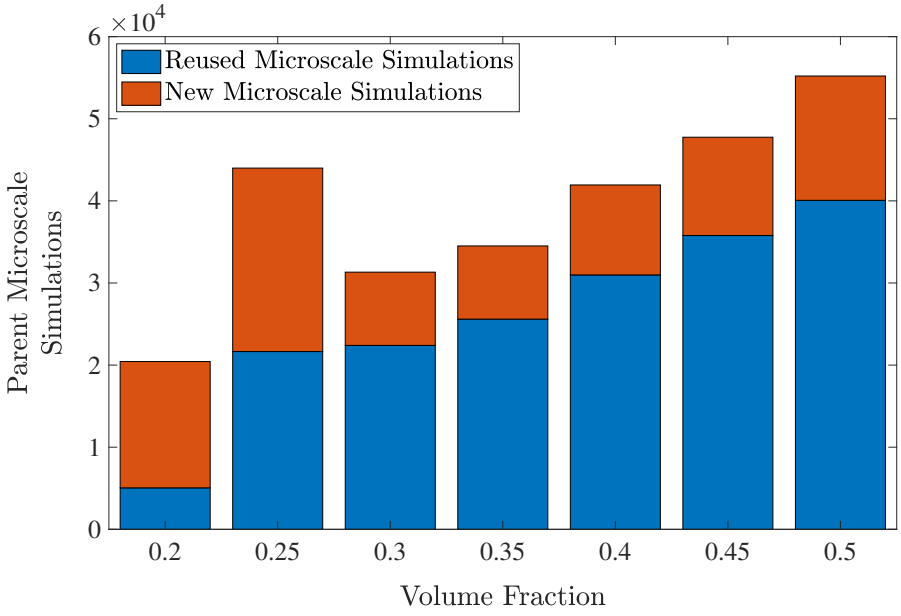


Figure 4.14: The number of parent microscale simulations required to derive the optimised geometries for each volume-fraction case within the parametric study. The optimisations were performed in ascending order of volume-fraction.

The results collected from the parametric study are used to derive a more coherent benchmark against a SIMP based TO. Equivalent optimisations are performed for all 7 volume-fractions,

again using a SIMP algorithm with a penalisation parameter of 3, identical discretisation and a Helmholtz filter length of 1 cell length. The compliance functionals derived by both the present framework and SIMP, are compared and illustrated in Fig. 4.15. It is clearly demonstrated, that the present concurrent multiscale optimisation framework outperforms SIMP for all 7 volume-fractions, with a compliance functional on average 38.3% lower. Additionally, an approximate mass saving is also calculated for the present framework, by determining the volume-fraction required by SIMP to achieve an equivalent compliance for the 30% volume-fraction case. This is computed by intersecting a compliance isoline (see Fig. 4.15) with the compliance curve generated by the results from the SIMP parametric study. Approximately 25% more material is required for SIMP in comparison to the present framework. Similarly, to the cantilever beam compliance minimisation, this advantage is a result of the clear accessibility to efficient orthotropic microstructures afforded by the microscale parameterisation and the ability to perform optimisation without penalising the intermediate values.

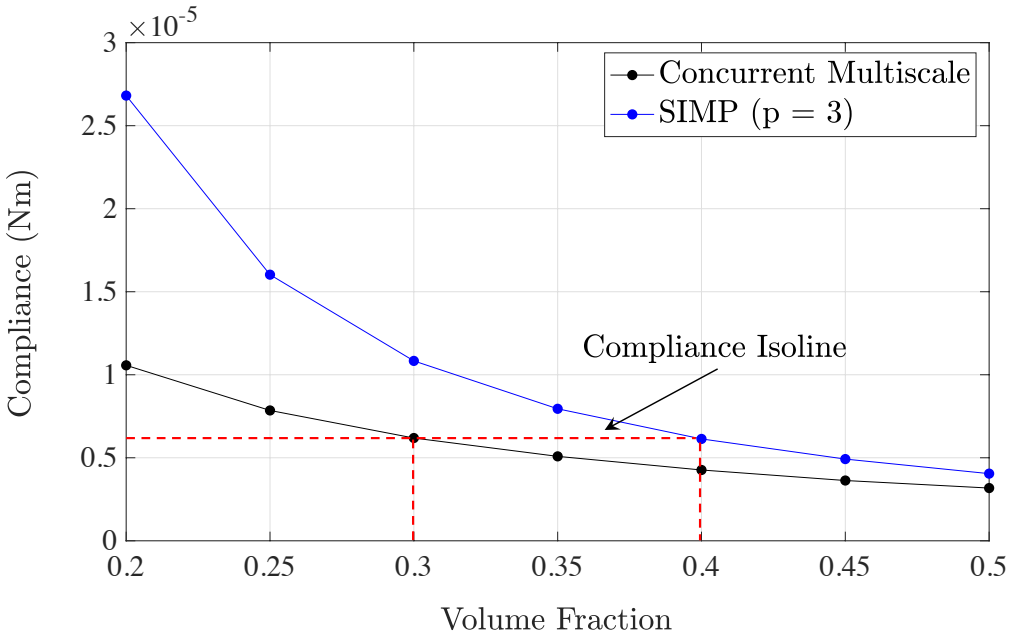


Figure 4.15: The compliance values derived from a parametric study of volume-fraction for the unit-cube, for both the present framework and a SIMP based TO

4.1.3 Computational Expense Incurred by Microscale Simulations

Over the course of the 7 unit-cube compliance minimisations and the single cantilever beam compliance minimisation, a total of 126,122 parent microscale simulations were required to derive the optimised results. This corresponds to resolving 1.85% of the parent microstructure database and 0.09% of the total microstructure database. This represents a significant reduction in the computational expense in comparison to contemporary sequential methods, which require fully-populated databases to be assembled prior to the first optimisation. In addition, the average number of parent microscale simulations required per optimisation for the present framework is only 18,017, which is lower than the requirement to assemble the databases in contemporary sequential frameworks which have 3 fewer design variables, for example Imediegwu et al. (2019) and Colabella et al. (2019) require 40,817 and 41,990 simulations respectively.

4.1.4 Comparison to IPOPT

In order to demonstrate the performance of the interior-point algorithm (IPOwDA) presented within Sec. 3.5.4, identical compliance minimisation problems are solved to those presented within Secs. 4.1.1 and 4.1.2 (30% volume-fraction cases only) using the present MTO framework, however, the established open-source interior point optimisation algorithm IPOPT (Wächter and Biegler, 2006) is applied instead. The converged objective functionals and the number of iterations required to achieve convergence are detailed for the cantilever beam case within Tab. 4.3 and for the unit-cube case within Tab. 4.4.

Method	Compliance (Nm)	Convergence Iteration
IPOwDA	0.001931	192
IPOPT	0.001928	179

Table 4.3: The optimised compliance functionals and the number of iterations required to achieve convergence for both IPOwDA and IPOPT for the cantilever beam case

Method	Compliance (Nm)	Convergence Iteration
IPOwDA	6.189×10^{-6}	139
IPOPT	6.167×10^{-6}	132

Table 4.4: The optimised compliance functionals and the number of iterations required to achieve convergence for both IPOwDA and IPOPT for the unit-cube case, subject to a 30% volume fraction constraint

On average, IPOPT takes 10 fewer iterations to achieve convergence and derives an objective functional 0.25% lower. The reason for implementing an interior-point optimisation algorithm from first principles is to enable the novel optimisation strategy presented within Sec. 3.5.5 to be tested. Therefore, the performance demonstrated by IPOwDA is deemed to be sufficiently close to IPOPT, as it is not intended to achieve equivalent performance to the 'best-in-class' open-source interior point optimisation algorithm. Consequently, IPOwDA is considered to be a suitable platform to explore the addition of microscale parameters optimisation strategy.

4.2 Addition of Microscale Parameters

The optimisation strategy detailed within Sec. 3.5.5 is applied within this section in combination with the present MTO framework to determine whether a reduction in the computational expense of optimisation is afforded through this approach. The initialisation conditions are kept constant for both cases presented within this section and are detailed in Tab. 4.5.

Design Variable	Initialisation Value
$\hat{r}_{1-7}$	0.14
e_{5-7}	1.00

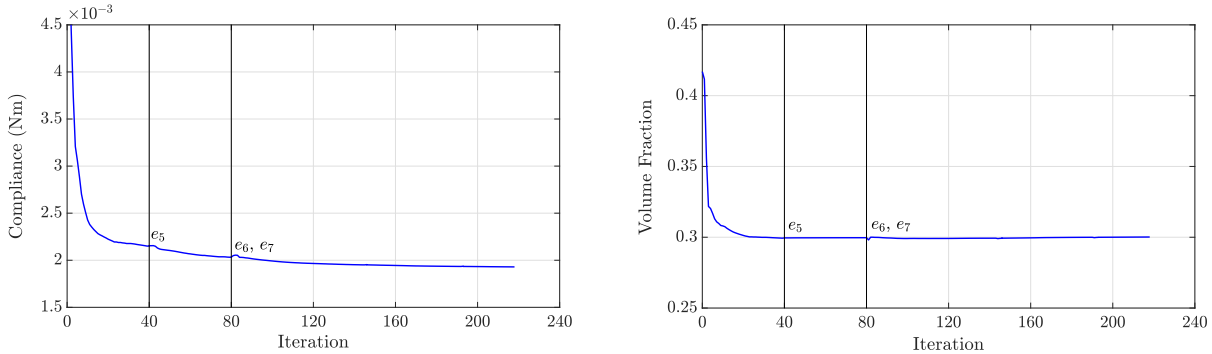
Table 4.5: The initialisation values used for all parameters at the beginning of optimisation

For both cases presented within this section, the initial number of active dimensions is set equal to seven. This number is selected to ensure all seven structural members are activated at the beginning of the optimisation procedure, as the aspect-ratio variables are relatively insensitive to the objective in comparison. This is important, as insensitive inactive structural members allocated a non-zero normalised radius account for a significant proportion of the total volume-

fraction, therefore, once these structural members are activated and the optimisation procedure is restarted, large steps would be required through the microscale design space to reach the new optimal solution. These large steps are counter-productive, as they will require a large number of parent microscale simulations.

4.2.1 Cantilever Beam — Compliance Minimisation

The first case considered is identical to the cantilever beam compliance minimisation problem presented within Sec. 4.1.1 and illustrated in Fig. 4.1. In this instance, the normalised aggregative sensitivity (see. Eq. 3.98) of each dimension is periodically evaluated every 40 iterations. In addition, it is stipulated that all microscale parameters must be introduced on or before iteration 80, to ensure that the optimisation problem is not fully converged before all microscale parameters are introduced. As detailed within Sec. 3.5.5, the criteria for introducing a new dimension during the optimisation procedure is if the aggregate sensitivity of an inactive dimension is larger than the smallest aggregate sensitivity derived for an active dimension. Convergence of the compliance functional and the volume fraction constraint is illustrated in Fig. 4.16 for this case. In addition, the normalised aggregate sensitivities derived during each sensitivity evaluation are detailed in Tab. 4.6.



(a) Compliance functional

(b) Volume fraction

Figure 4.16: Convergence plots for the cantilever beam compliance minimisation case, where optimisation is initially performed using only the seven normalised radii ($\hat{r}_{1-7}$). At iteration 40, the optimisation procedure is restarted and the microscale parameter e_5 is added to the active set. Then, at iteration 80, the optimisation procedure is again restarted and the microscale parameters e_6, e_7 are forcibly added to the active set.

	Normalised Aggregate Sensitivity									
	$\hat{\theta}_{\hat{r}_1}$	$\hat{\theta}_{\hat{r}_2}$	$\hat{\theta}_{\hat{r}_3}$	$\hat{\theta}_{\hat{r}_4}$	$\hat{\theta}_{\hat{r}_5}$	$\hat{\theta}_{\hat{r}_6}$	$\hat{\theta}_{\hat{r}_7}$	$\hat{\theta}_{e_5}$	$\hat{\theta}_{e_6}$	$\hat{\theta}_{e_7}$
Iteration 0	0.554	0.547	0.547	0.554	1.000	0.262	0.218	0.059	0.010	0.011
Iteration 40	0.658	0.529	0.529	0.658	1.000	0.433	0.076	0.137	0.026	0.002
Iteration 80	0.738	0.599	0.599	0.738	1.000	0.535	0.120	0.088	0.032	0.004

Table 4.6: The normalised aggregate sensitivities derived during each sensitivity evaluation for the cantilever beam case. The iteration in which a dimension is activated is highlighted using a red font.

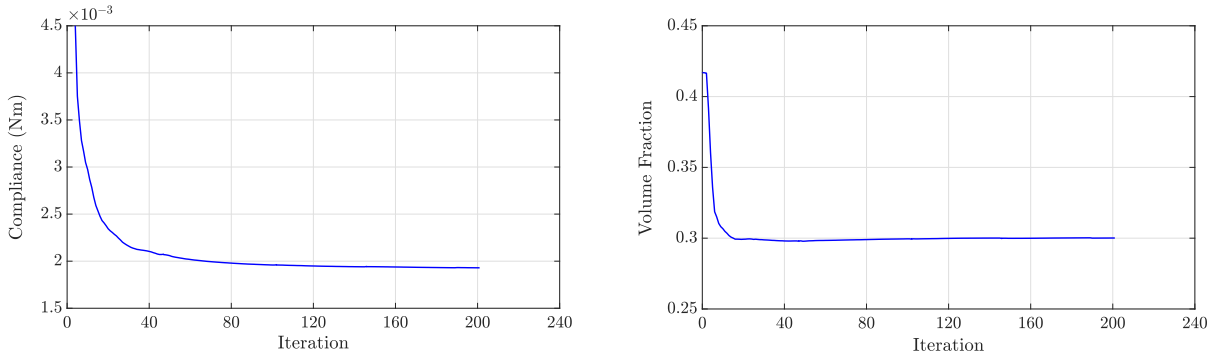
As detailed in Tab. 4.6, the optimisation procedure is initially started using only the seven structural members, as they constitute the seven microscale parameters with the largest normalised aggregate sensitivities in this instance. At iteration 40, the optimisation procedure is restarted, as the aggregate sensitivity of the aspect-ratio variable e_5 is larger than the aggregate sensitivity of structural member r_7. It is clearly illustrated in Fig. 4.16, that the restart procedure has a negligible influence on the convergence of both the objective functional and the volume fraction constraint. This demonstrates the efficacy of the restart procedure detailed within Sec. 3.5.5. The optimisation procedure is restarted again during iteration 80, as this constitutes the last iteration before all microscale parameters must be introduced. Similarly, it is noted that the impact of the restart procedure on the convergence of both the compliance functional and the volume fraction constraint is again negligible.

In order to determine the merits of this approach, an equivalent compliance minimisation is performed using all ten microscale parameters from the beginning of the optimisation procedure. Convergence of the compliance functional and the volume fraction constraint is illustrated in Fig. 4.17 for this case. In addition, the compliance functionals derived from both cases are detailed in Tab. 4.7.

Dimension Addition	Compliance (Nm)	Convergence Iteration
Activated	0.001929	218
Disabled	0.001930	201

Table 4.7: The optimised compliance functionals and the number of iterations required to achieve convergence for both cases presented within this section

As detailed within Tab. 4.7, equivalent values for the compliance functionals are obtained



(a) Compliance functional (b) Volume fraction

Figure 4.17: Convergence plots for the cantilever beam compliance minimisation case, where optimisation is performed using all ten microscale parameters, and the dimension activation functionality is disabled

at convergence, which further demonstrates the efficient convergence properties of IPOwDA after a partially converged optimisation is restarted. However, as expected, 17 more iterations of optimisation are required in order to reach convergence for the case where the dimension addition functionality is activated in comparison to when it is disabled.

The computational expense incurred by parent microscale simulations for the two cases presented within this section are compared in Fig. 4.18. In addition, the average distance traversed through the microscale design space by each microstructure (X_α) during the optimisation procedure is calculated using a piecewise summation of the Euclidean norm

$$S = \sum_{\alpha=1}^{N} \left(\frac{\sum_{i=2}^{N} \|X_{\alpha,i+1} - X_{\alpha,i}\|_2}{N} \right) \tag{4.2}$$

where i denotes the optimisation iteration, and is illustrated for both cases presented within this section in Fig. 4.19. A 5.5% reduction in the computational expense incurred by parent microscale simulations is noted for the case where the dimension addition functionality is activated, in comparison to when it is disabled. It is prudent to add, that the large increases noted within Fig. 4.18 during each sensitivity evaluation are as a result of the additional nodes required to calculate the derivative $\frac{\partial E_{ij}(\tilde{X})}{\partial \tilde{X}}$ for all ten microscale parameters. In contrast, however, the average distance traversed through the microscale design space is 14.4% larger for the case where the dimension addition functionality is activated, in comparison to when

it is disabled. Consequently, the reduction in the number of parent microscale simulations is a result of initially traversing through sub-spaces of the ten-dimensional hyperspace, as the hyper-volume resolved by each simplex is inversely proportional to its dimensionality (see Sec. 3.4.2). Therefore, fewer simplices are required to resolve a unit volume within a sub-space of the ten-dimensional hyperspace, in comparison to the number required within the ten-dimensional hyperspace. This is aptly demonstrated for this case at iteration 20, where the average arc lengths are approximately equal for both cases (see Fig. 4.19), however, approximately 84% fewer simulations are required to traverse this distance within the seven-dimensional hyperspace (see Fig. 4.18), when the dimension addition functionality is activated, in comparison to when it is disabled.

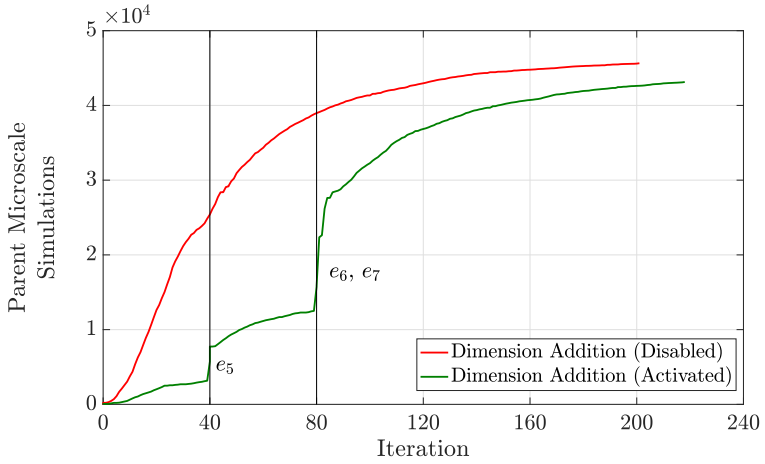


Figure 4.18: A comparison for the number of parent microscale simulations required to derive the optimised solutions for both cases presented within this section

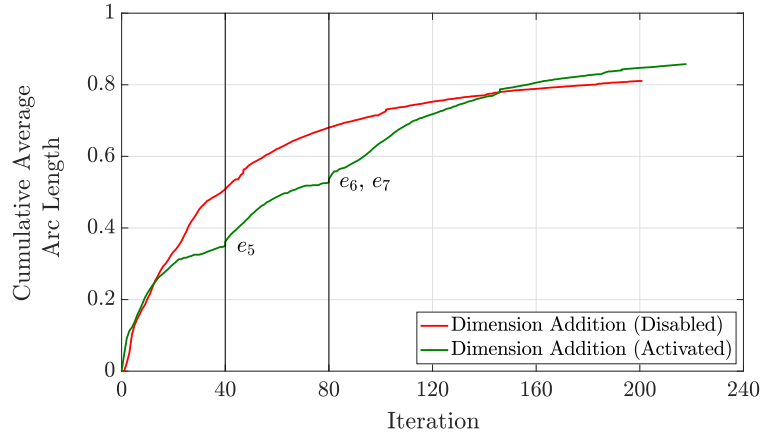
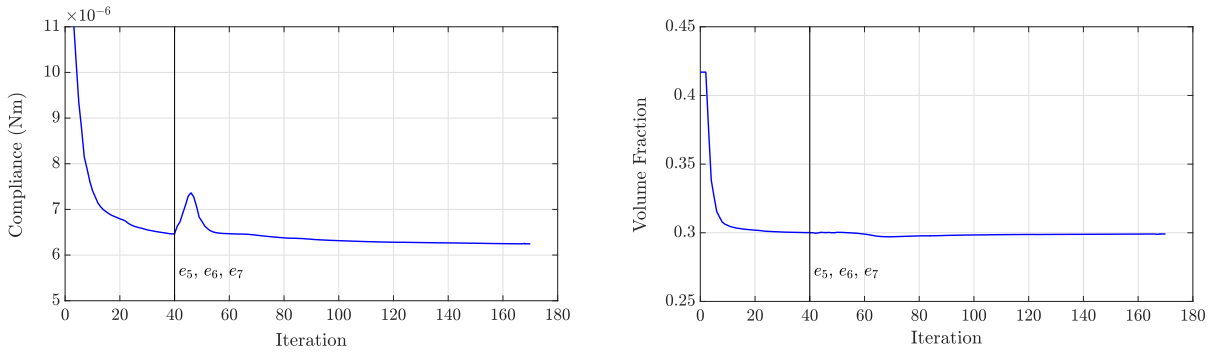


Figure 4.19: A comparison of the cumulative average arc length for both cases presented within this section

4.2.2 Unit Cube — Compliance Minimisation

The second case considered is identical to the unit-cube compliance minimisation problem presented within Sec. 4.1.2 and illustrated in Fig. 4.8. In this instance, it is stipulated that all microscale parameters must be introduced by iteration 40, as an equivalent preliminary optimisation of this case demonstrated that the aspect-ratio variables do not satisfy the requirements to be introduced during any sensitivity evaluation. Convergence of the compliance functional and the volume fraction constraint is illustrated in Fig. 4.20 for this case. In addition, the normalised aggregate sensitivities derived during each sensitivity evaluation are detailed in Tab. 4.8.



(a) Compliance functional (b) Volume fraction

Figure 4.20: Convergence plots for the unit-cube compliance minimisation case, where optimisation is initially performed using only the seven normalised radii ($\hat{r}_{1-7}$). At iteration 40, the optimisation procedure is restarted and the microscale parameters e_5, e_6, e_7 are forcibly added to the active set.

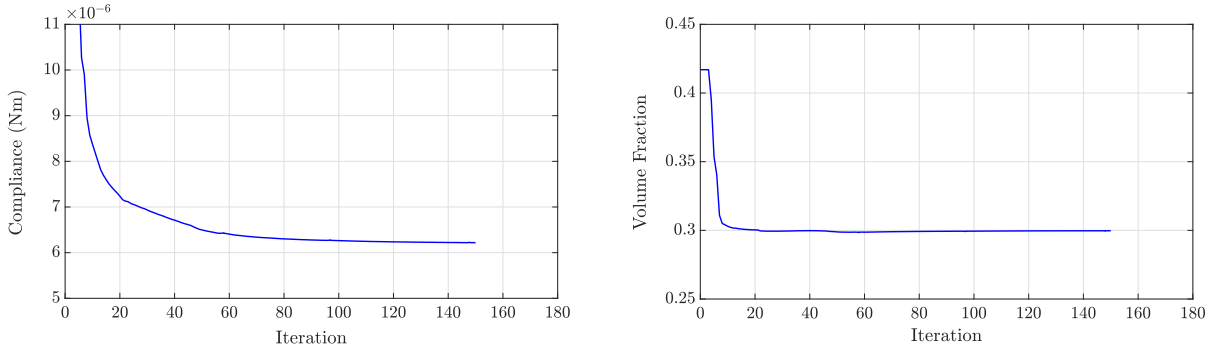
	Normalised Aggregate Sensitivity									
	$\hat{\theta}_{\hat{r}_1}$	$\hat{\theta}_{\hat{r}_2}$	$\hat{\theta}_{\hat{r}_3}$	$\hat{\theta}_{\hat{r}_4}$	$\hat{\theta}_{\hat{r}_5}$	$\hat{\theta}_{\hat{r}_6}$	$\hat{\theta}_{\hat{r}_7}$	$\hat{\theta}_{e_5}$	$\hat{\theta}_{e_6}$	$\hat{\theta}_{e_7}$
Iteration 0	0.771	0.771	0.771	0.771	0.386	1.000	0.424	0.059	0.019	0.018
Iteration 40	0.312	0.312	0.312	0.312	0.521	1.000	0.421	0.051	0.123	0.024

Table 4.8: The normalised aggregate sensitivities derived during each sensitivity evaluation for the unit-cube case. The iteration in which a dimension is activated is highlighted using a red font.

As detailed in Tab. 4.8, the optimisation procedure is initially started using only the seven structural members, as they constitute the seven microscale parameters with the largest normalised aggregate sensitivities in this instance. At iteration 40, the optimisation procedure is

restarted, as the microscale parameters e_5, e_6, e_7 are forcibly introduced. In this case, a significant perturbation to the compliance functional is noted after the optimisation procedure is restarted. It is assumed that the cause of this perturbation is a result of not being able to re-use the L-BFGS matrices from the initial optimisation procedure, as all other information is carried forwards.

Similarly, an equivalent compliance minimisation is performed for this case using all ten microscale parameters from the beginning of the optimisation procedure. Convergence of the compliance functional and the volume fraction constraint is illustrated in Fig. 4.21 for this case. In addition, the compliance functionals derived from both cases presented within this section are detailed in Tab. 4.9.



(a) Compliance functional (b) Volume fraction

Figure 4.21: Convergence plots for the unit-cube compliance minimisation case, where optimisation is performed using all ten microscale parameters, and the dimension activation functionality is disabled

Dimension Addition	Compliance (Nm)	Convergence Iteration
Activated	6.24×10^{-6}	170
Disabled	6.22×10^{-6}	150

Table 4.9: The optimised compliance functionals and the number of iterations required to achieve convergence for both cases presented within this section

As detailed within Tab. 4.9, equivalent values for the compliance functionals are obtained at convergence, which further demonstrates the efficient convergence properties of IPOwDA after a partially converged optimisation is restarted. However, as expected, 20 more iterations of optimisation are required in order to reach convergence for the case where the dimension addition functionality is activated in comparison to when it is disabled.

The computational expense incurred by parent microscale simulations for the two cases presented within this section are compared in Fig. 4.22. In addition, the average distance traversed through the microscale design space by each microstructure (X_α) during the optimisation procedure is also illustrated for both cases in Fig. 4.22. In this instance, only a 0.4% reduction in the number of parent microscale simulations is afforded. This is a result of the problematic restart to the optimisation procedure, which led to a significant increase in the number of parent microscale simulations required per iteration. An identical relationship between the number of parent microscale simulations, the dimensionality of the sub-space and the cumulative average arc length is also noted for this case.

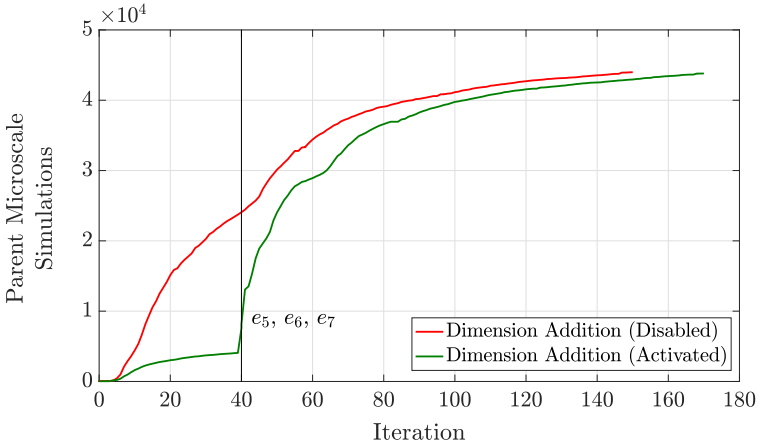


Figure 4.22: A comparison for the number of parent microscale simulations required to derive the optimised solutions for both cases presented within this section

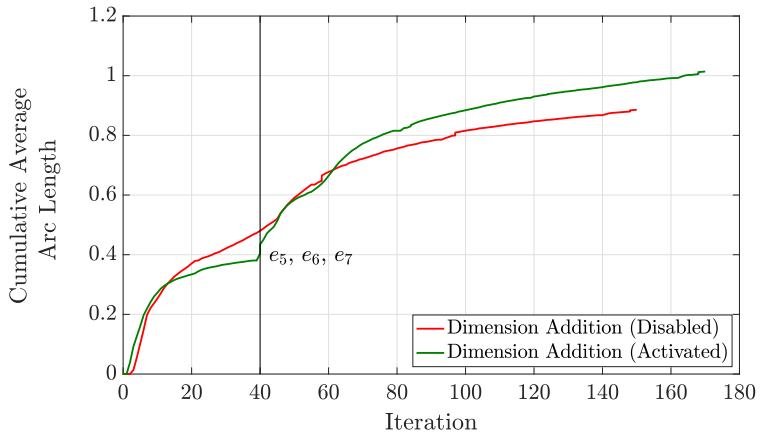


Figure 4.23: A comparison of the cumulative average arc length for both cases presented within this section

Chapter 5

Design for Additive Manufacturing

As detailed in Sec. 2.4.4, it is common practice within open-walled (truss-like) MTO frameworks to place a non-zero lower bound on the radii of the constituent structural members during optimisation, in order to constrain the minimum length scale of an optimised geometry and to prohibit the formation of unsupported members or isolated assemblies of overhanging members (Imediegwu et al., 2019; Wang et al., 2020a; Thillaithevan et al., 2021; Nightingale et al., 2021). Unsupported members and isolated assemblies of overhanging members are explicitly prohibited by a non-zero lower bound constraint, as a non-zero lower bound fixes the microscale topology and therefore guarantees the existence of supporting structures for all structural members. However, further structural performance is traded for manufacturability as a result of this prohibitive lower bound constraint, as accessibility is denied to microstructures with low volume-fractions.

The work presented within this chapter seeks to circumvent the necessary application of prohibitive lower bound constraints to guarantee manufacturability, by integrating novel AM constraints within the present MTO framework. To simplify the derivation of these AM constraints and to enable a generalised framework for open-walled microstructures to be presented, the parametrisation is simplified by removing the aspect-ratio variables (see Fig. 5.1). As a consequence, it is no longer necessary to include an eighth DOE level ($\hat{r}_{5-7} = 0.49$) for the normalised radius of a face member (see Tab. 5.1).

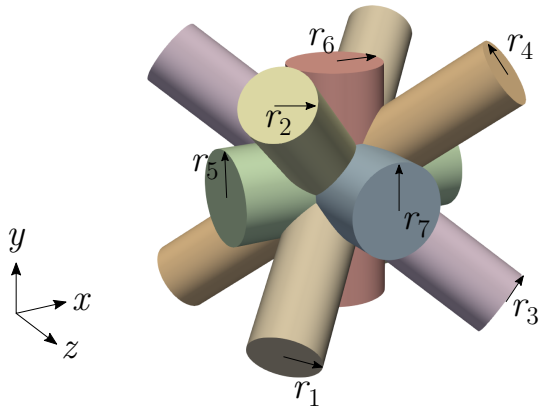


Figure 5.1: The parametrised microscale geometry used for all application of the present MTO framework with integrated AM constraints

Microscale Parameters	No. Levels	Lowest Level	Highest Level	Level Increment
$\hat{r}_{1-7}$	7	0.00	0.42	0.07

Table 5.1: The individual microscale parameter bounds and level intervals used for AM constrained cases.

5.1 Manufacturability Requirements

In order to formulate constraints without the necessity of lower bound constraints on member radius to guarantee manufacturability, it is prudent to first consider the design restrictions AM processes place on open-walled microstructures within open-walled MTO frameworks. As is the case with all AM processes, the minimum length scale of a geometric feature is restricted by the printer's resolution. If features exist with smaller length scales than those permitted by the printer, significant differences are expected between the desired and fabricated geometries, as well as unpredictable mechanical performance. For open-walled microstructures, this places a limit on the minimum manufacturable radius (r_{min}) of the constituent structural members; defined as half of the length of the printer resolution (r_p), as illustrated in Fig. 5.2. Therefore, for an arbitrary open-walled microstructure manufactured through AM processes, the constituent structural members must be an element of the set $\{0, [r_{min}, r_{max}]\}$ to be considered manufacturable, where r_{max} denotes the maximum allowable radius. Importantly, structural members with a radius of zero are included within this set, as trivially, they do not require

fabrication.

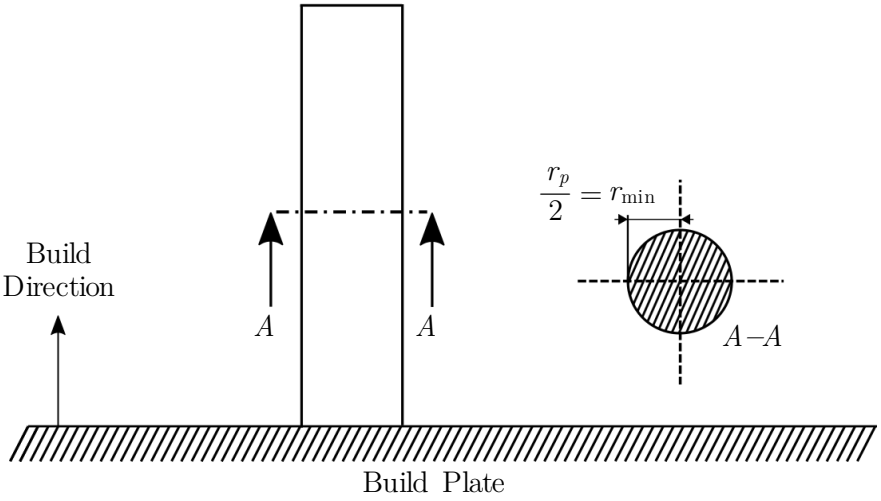


Figure 5.2: Illustration of how the minimum manufacturable radius (r_{min}) of a structural member is derived as function of the printer resolution (r_p).

For structural compliance, it is desirable to enable the radius of a structural member to equal zero during optimisation, as this permits accessibility to microstructures with increased specific directional stiffnesses and more importantly those devoid of material. As a consequence, however, this enables the microstructural topology to vary spatially during optimisation, which can lead to the formation of unsupported and isolated assemblies of overhanging members, as illustrated in Fig. 5.3 using a two-dimensional example microstructure.

In the generalised case of open-walled microstructures, overhanging members are often identified by those with an angle of greater than 45° relative to the build direction-vector (Jiang et al., 2018a), however, it is printer dependent. Conversely, members with an angle of less than 45° relative to the build direction-vector are often denoted as self-supporting members. In this instance, diagonal members are orientated at 35.3° relative to the build direction-vector. As a consequence, self-supporting members are identified by those with an angle of less than 35.3° relative to the build direction-vector for the remainder of this thesis, as this enables multiple members to provide support during the optimisation procedure.

Most printing architectures are able to cope with relatively small horizontal overhangs before fabrication issues are encountered, which enables overhanging members to be manufactured, so long as self-supporting members are in close proximity (Flores et al., 2020). Consequently, in

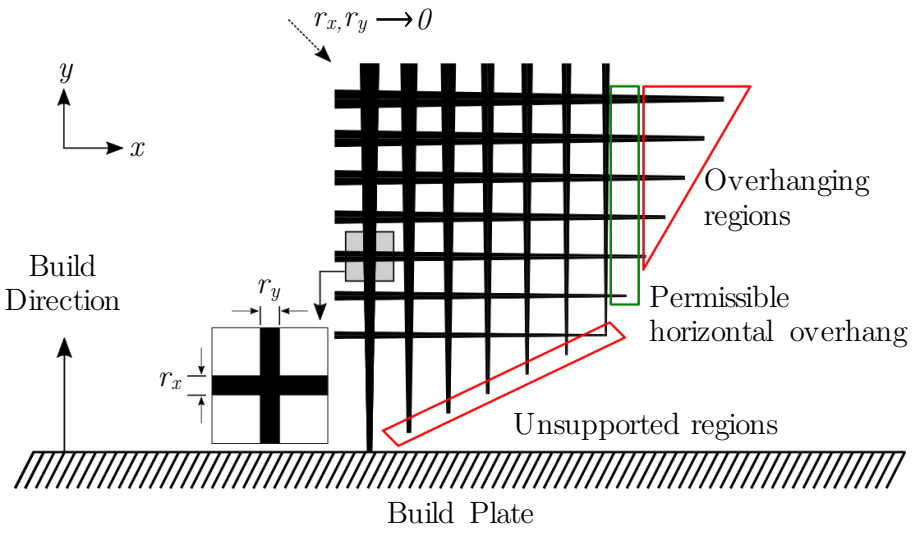


Figure 5.3: An assembly of an example two-dimensional parametrised microstructure, which is purposely functionally-graded to demonstrate how unsupported members and isolated assemblies of overhanging members can form as a result of permitting the microstructural topology to vary spatially during optimisation.

the absence of self-supporting members, an assembly of overhanging members is problematic, as insufficient support is provided between printing layers. As a result, thermal stresses can accumulate, resulting in isolated assemblies of overhanging members warping or distorting, which in extreme cases can lead to the inability to continue the fabrication process (Muthusamy et al., 2018).

A common approach to alleviate the problems associated with unsupported members or isolated assemblies of overhanging members is to include temporary support structures to reinforce problematic features during fabrication (Jiang et al., 2018b, 2019). However, this approach is not cost-effective, as a significant percentage of the build material is sacrificial. In addition, the acute complexity of optimised multiscale geometries further prohibits the already difficult procedure of removing support structures. As a consequence, it is economically advantageous to discourage these features by reformulating the optimisation problem, as this enables a suitable balance between structural optimality, manufacturability and build cost to be obtained.

A further consequence of permitting the radii of structural members to equal zero during optimisation is the potential formation of slender members at the macroscale, which are prone

to buckling and at risk of being knocked over during the fabrication process. Slender members in this context are characterised by an assembly of identically orientated members in a direction parallel to their respective orientation and with all other members absent, as illustrated in Fig. 5.4 using a two-dimensional example microstructure. As a consequence of the linear elastic assumptions applied within the present MTO framework and the numerical homogenisation procedure, recognition of slender members at the macroscale is prohibited and leads to the over-prediction of a slender member's axial stiffness. Without the application of prohibitively expensive nonlinear finite-element analyses, geometric constraints are required to prohibit the formation of slender members during optimisation.

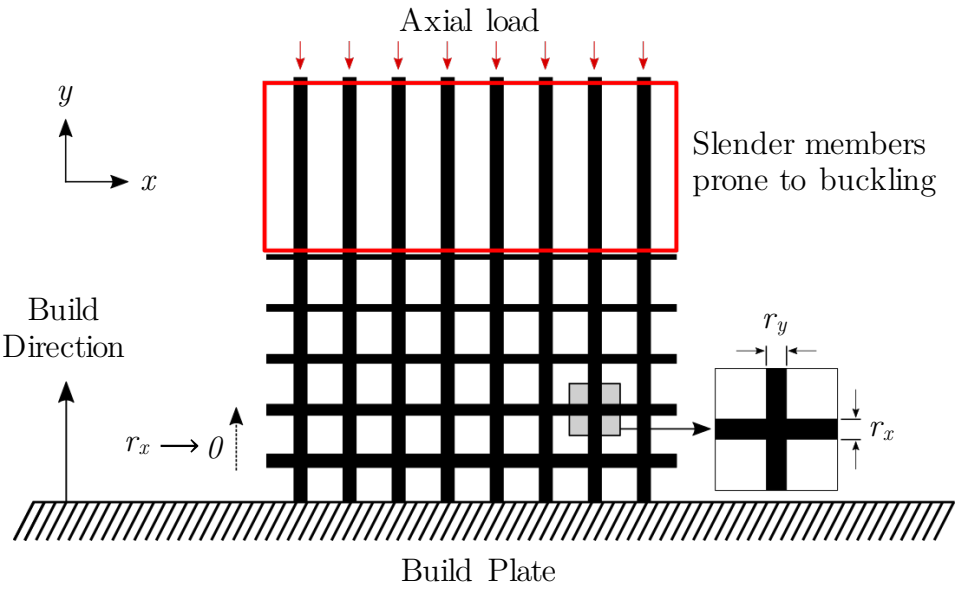


Figure 5.4: An assembly of an example two-dimensional parametrised microstructure, which has been purposely functionally-graded to demonstrate how slender members can form as a result of permitting the microstructural topology to vary spatially during optimisation.

In summary, optimised structures, derived through open-walled MTO frameworks, can be considered manufacturable via AM processes if the following constraints are satisfied:

- The radii of all constituent structural members are an element of the set $\{0, [r_{min}, r_{max}]\}$.
- There are no unsupported members, isolated assemblies of overhanging members or slender members.

It should be noted that the formation of locally-enclosed volumes is excluded from this list, as it is an assumed property of open-walled microstructures.

5.2 Minimum Length Scale Control

To control the minimum length scale of a constituent structural member, whilst affording maximum design freedom during optimisation, an augmented projection filter is applied to the normalised radii

$$\hat{\tilde{X}}_{\alpha\beta} = \tilde{X}_{\alpha\beta}\frac{\tanh(\eta\hat{r}_{min}) + \tanh(\eta(\tilde{X}_{\alpha\beta} - \hat{r}_{min}))}{\tanh(\eta\hat{r}_{min}) + \tanh(\eta(1 - \hat{r}_{min}))} \tag{5.1}$$

where $\hat{\tilde{X}}_{\alpha\beta}$ denotes a projected and filtered normalised radius, $\hat{r}_{min}$ denotes the minimum manufacturable normalised radius and η denotes the Heaviside thresholding parameter (Lazarov et al., 2016). The augmented projection filter applies a non-linear mapping, which penalises the normalised radii below the minimum manufacturable normalised radius, but tends back to a linear relationship above this value, as depicted in Fig. 5.5. The severity of the penalisation is a function of the Heaviside thresholding parameter (η); the higher this value the more nonlinear the function becomes. Importantly, in the limit

$$\lim_{\eta\to\infty} \hat{\tilde{X}}_{\alpha\beta} = \begin{cases} \tilde{X}_{\alpha\beta}, & \text{if } \tilde{X}_{\alpha\beta} \geq \hat{r}_{min} \\ 0, & \text{otherwise} \end{cases} \tag{5.2}$$

the manufacturable set of radii is recovered. As a result, the application of this filter enables optimisation of the normalised radii to be performed within the interval $[0, \hat{r}_{max}]$, but subsequently post-processed into the set $\{0, [\hat{r}_{min}, \hat{r}_{max}]\}$ with relative ease. In addition, it is permissible to use different values of $\hat{r}_{min}$ for each constituent structural member of the microscale parametrisation, which is of use for printing architectures where the minimum manufacturable normalised radius varies relative to the build direction. However, in this instance, the minimum manufacturable normalised radius is assumed to be constant, irrespective

of member orientation relative to the build direction vector.

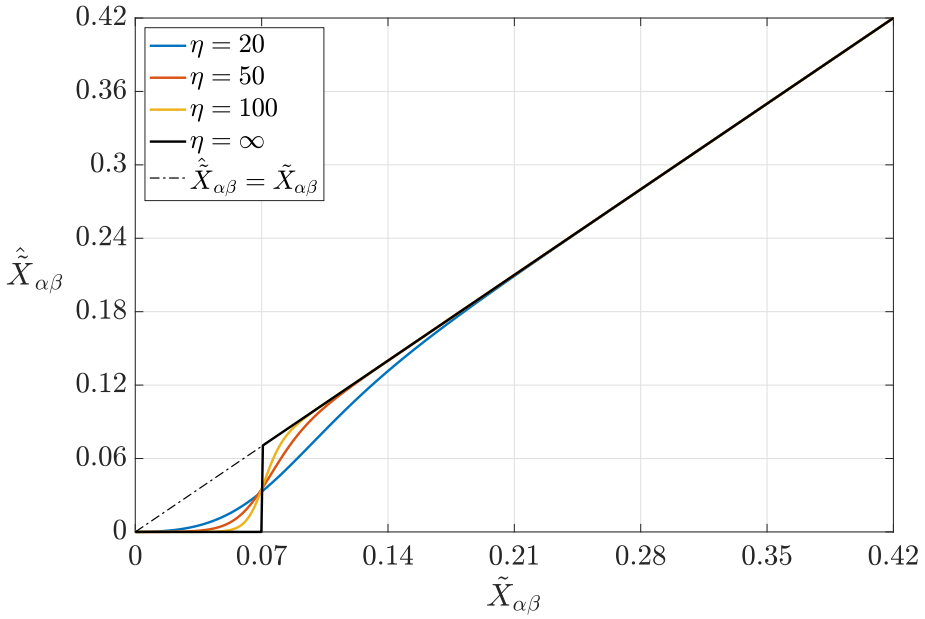


Figure 5.5: How the normalised radii are projected as a result of passing through the augmented projection filter (5.1). A total of four different Heaviside thresholding parameters (η) are illustrated, all using an identical minimum manufacturable normalised radius of $\hat{r}_{min} = 0.07$.

5.3 Unsupported and Overhanging Members

As detailed in Sec. 5.1, by enabling the normalised radius of a structural member to equal zero during optimisation, it is possible for unsupported members, isolated assemblies of overhanging members and slender members to form during optimisation. To prohibit the formation of unsupported and isolated assemblies of overhanging members, a novel geometric constraint is introduced to ensure that the optimised geometry is self-supporting.

A structural member is considered to be supported if it intersects with a self-supporting member with a normalised radius greater than or equal to the minimum manufacturable normalised radius, which is positioned in-between itself and the build plate. Therefore, for a structure to be considered self-supporting, an assembly of interconnected self-supporting members should exist between each member and the build plate. In the case of a diagonal

member, support can be provided from beneath by any one of the four diagonal members which intersect with this member at the base of the unit-cell, as illustrated in Fig. 5.6a. Furthermore, in the case of a face member, support can only be provided from beneath by an identically orientated face member, as illustrated in Fig. 5.6b.

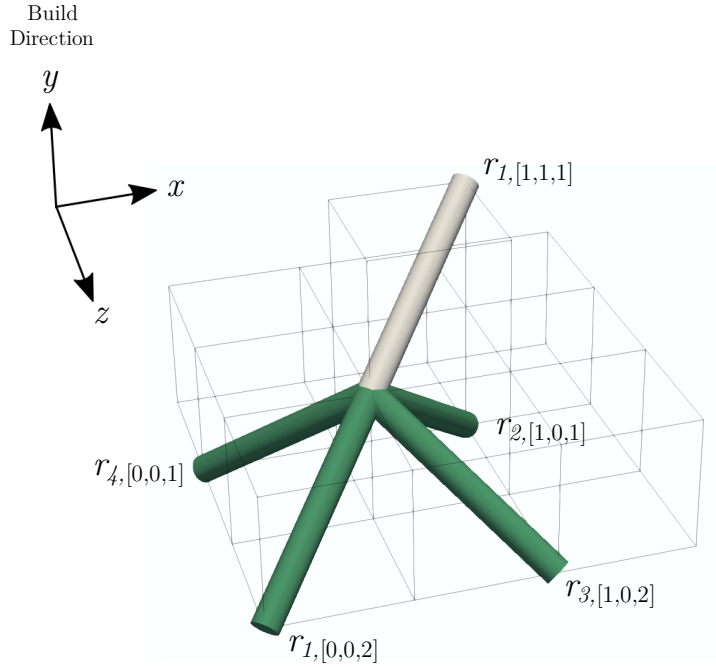
Additionally, as all members intersect at the centroid of the unit-cell by definition, it is also permissible for a structural member to be supported by a self-supporting member from within the same microstructure, as illustrated in Fig. 5.7. This support mechanism is referred to as intra-cell support for the remainder of this section. Permitting intra-cell support is crucial, as it enables overhanging members to be supported and allows interconnected assemblies of different members to provide support, enabling face members for example to be supported by diagonal members. It should be noted, however, that intra-cell support permits one half of a structural member to be unsupported, as illustrated in Fig. 5.8. This issue arises whenever a diagonal member is only supported intra-cell by a face member and in the converse situation where a face member is only supported intra-cell by a diagonal member. However, as these features do not contribute significantly to the stiffness of the macroscale structure, they are permitted during optimisation and subsequently removed during post-processing.

To discourage isolated assemblies of overhanging members and unsupported members from within the optimised solution, a relaxed binary integral constraint is introduced to enforce the support requirements outlined above for each member. Initially, a standard projection-filter is applied to the normalised radii

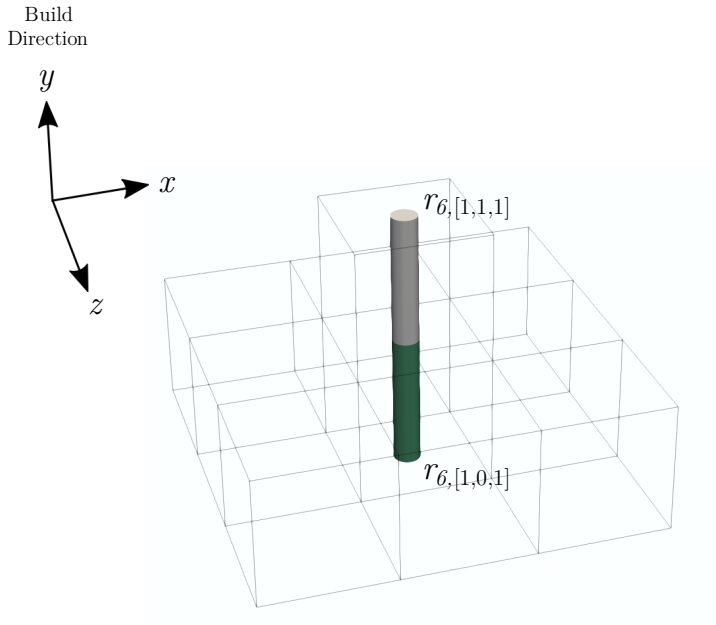
$$\phi_{\alpha\beta} = \frac{\tanh(\eta\hat{r}_{min}) + \tanh(\eta(\tilde{r}_{\alpha\beta} - \hat{r}_{min}))}{\tanh(\eta\hat{r}_{min}) + \tanh(\eta(1 - \hat{r}_{min}))} \tag{5.3}$$

where ϕ denotes a quasi-Boolean function, which is applied to determine whether a normalised radius is above or below the minimum manufacturable radius, as depicted in Fig. 5.9. In the limit

$$\lim_{\eta\to\infty} \phi_{\alpha\beta} = \begin{cases} 1, & \text{if } \tilde{X}_{\alpha\beta} \geq \hat{r}_{min} \\ 0, & \text{otherwise} \end{cases} \tag{5.4}$$

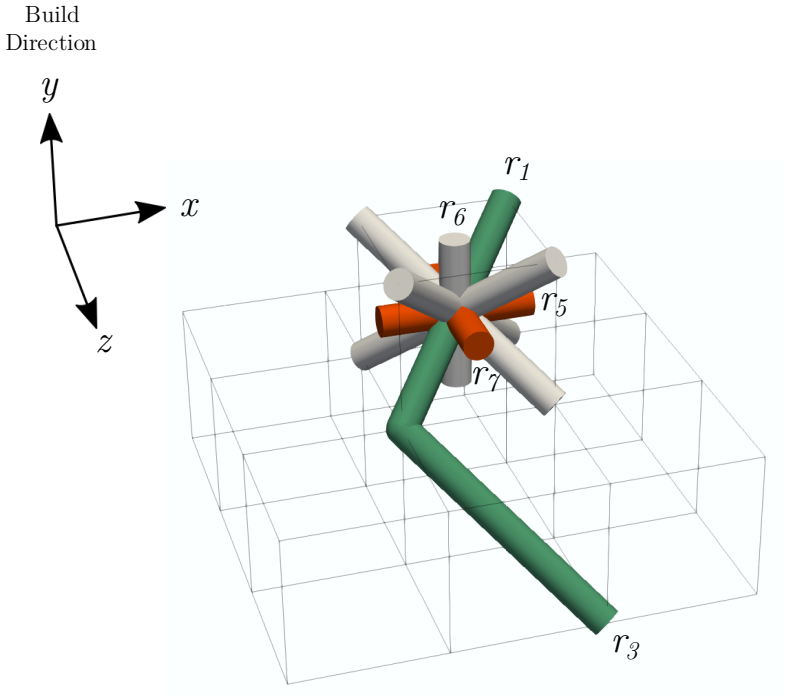


(a) The four structural members capable of providing support to structural member r_1 from beneath. If at least one of these supporting members has a normalised radius greater than or equal to the minimum manufacturable normalised radius ($\hat{r}_{min}$), structural member r_1 is considered to be supported. In addition, it is important to note that the supporting structures for diagonal members r_{2-4} are located in rotationally symmetric positions.

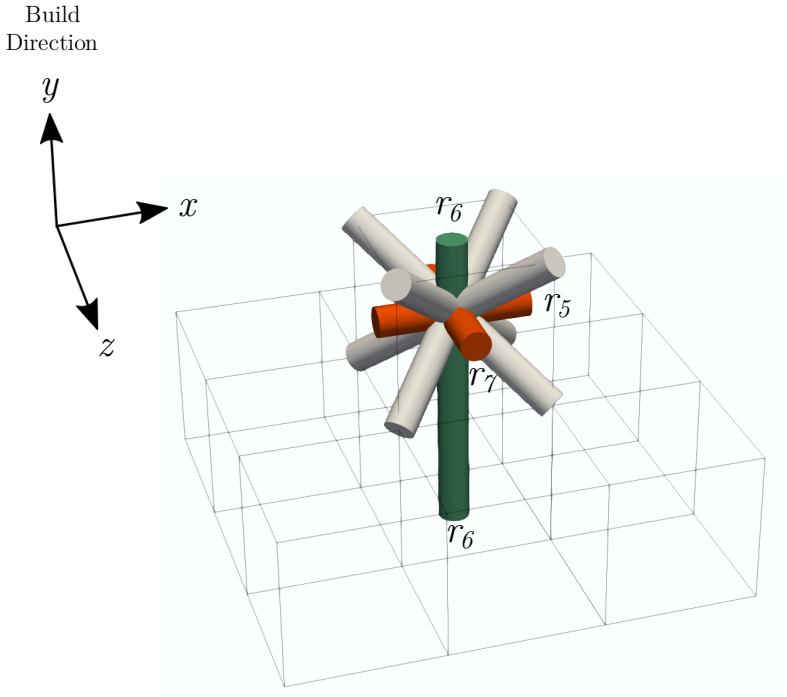


(b) The only structural member capable of providing support to structural member r_6 from beneath. If this supporting member has a normalised radius greater than or equal to the minimum manufacturable normalised radius ($\hat{r}_{min}$), structural member r_6 is considered to be supported.

Figure 5.6: The members capable of providing support to diagonal and face members from beneath are depicted using structural members r_1 and r_6 respectively. The structural members capable of providing support are illustrated in green, and the supported members in grey. Unit vector notation is used to help identify the location of each structural member.



(a) Intra-cell support is provided by structural member r_1 to structural members r_{2-7}, as it is supported from below by structural member r_3.



(b) Intra-cell support is provided by structural member r_6 to structural members $r_{1-5,7}$, as it is supported from below by structural member r_6.

Figure 5.7: How intra-cell support is provided by both diagonal and face members respectively. The structural members providing support are illustrated in green, overhanging members are illustrated in orange and self-supporting members are illustrated in grey.

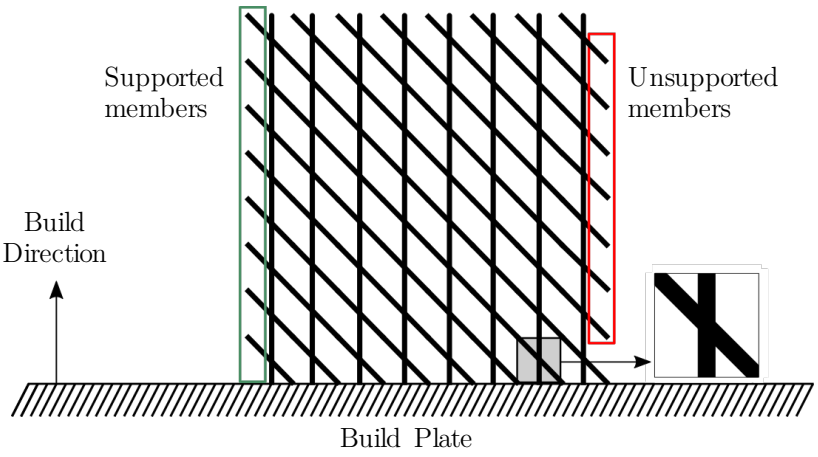


Figure 5.8: An assembly of an example two-dimensional microstructure, which demonstrates the formation of unsupported members as a consequence of permitting intra-cell support. As these features do not contribute significantly to the stiffness of the macroscale structure, they are permitted during optimisation and then subsequently removed during post-processing.

a truly Boolean function is obtained. The quasi-Boolean nature of this function is leveraged to derive constraints analogous to logic gates.

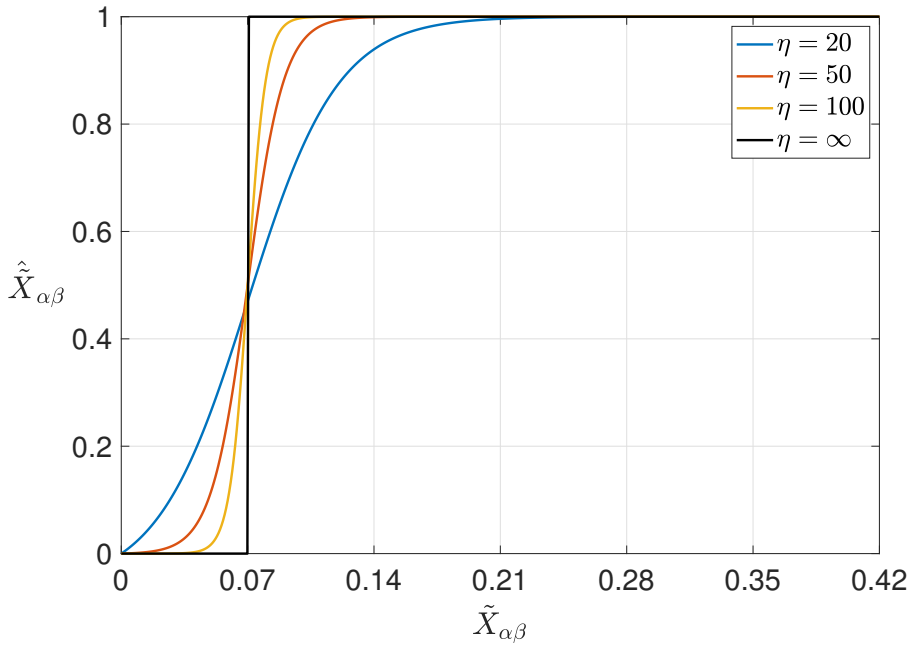


Figure 5.9: A graph depicting the quasi-Boolean nature of the projection filter (Eq. 5.3), which is applied to determine whether a normalised radius is above or below the minimum manufacturable normalised radius. A total of four different Heaviside thresholding parameters (η) are illustrated, all using an identical minimum manufacturable normalised radius of $\hat{r}_{min} = 0.07$.

The relaxed binary integral constraint is composed of a series of logical expressions, which

in combination determine whether or not unsupported members or isolated assemblies of overhanging members exist within the optimised solution. To begin, a logical expression is formulated for each constituent structural member ($X_{\alpha\beta}$) to determine whether or not it is supported from beneath by another member

$$\zeta_{I,\alpha\beta} = \prod_{i=1}^{n(\Psi_{\alpha\beta})} \left(1 - \Psi_{\alpha\beta,i}\right) \tag{5.5}$$

where $\Psi_{\alpha\beta}$ denotes a set containing the projected values (ϕ) for all self-supporting members which can potentially provide support to structural member $X_{\alpha\beta}$, and $\zeta_{I,\alpha\beta}$ denotes a quasi-Boolean function which in the limit $\eta \to \infty$ equals 0 if member $X_{\alpha\beta}$ is supported from beneath and conversely 1 if it is not supported. The location of the normalised radii composing the set $\Psi_{\alpha\beta}$ is identified in advance of optimisation for each constituent member, as a function of the build direction and the location of the build plate.

This logical expression is then embedded within another logical expression to determine whether a member is supported from beneath and/or receives intra-cell support. In combination, this identifies members within microstructure X_α which have a radius greater than or equal to the minimum manufacturable normalised radius, are self-supporting and supported from beneath

$$\zeta_{II,\alpha} = \prod_{\beta=1}^{n(\phi_{\alpha\beta})} \left(1 - \phi_{\alpha\beta}\left(1 - \zeta_{I,\alpha\beta}\right)\right) \tag{5.6}$$

where $\zeta_{II,\alpha}$ denotes a quasi-Boolean function which in the limit $\eta \to \infty$ equals 0 if a self-supporting member which is also supported from beneath is present, or 1 if either of these two statements are false. As all members intersect at the centroid of a unit-cell, this value is identical for all constituent members of microstructure X_α.

In its present form, this logical expression (Eq. 5.6) determines whether or not supporting structures exist within each microstructure. If the integral of this logical expression for all constituent structural members within the macroscale domain is constrained to sum to zero during optimisation, then at least one self-supporting member would be required within each constituent microstructure to satisfy this constraint. This would result in a sub-optimal design,

as the derivation of microstructures devoid of material is not possible. As a result, this logical expression is extended to

$$\zeta_{III,\alpha\beta} = \phi_{\alpha\beta}^{P} \zeta_{II,\alpha} \tag{5.7}$$

which accounts for the situation where the normalised radius of a structural member equals zero and therefore, trivially, does not require support, and where $\zeta_{III,\alpha}$ denotes a quasi-Boolean function which in the limit $\eta \to \infty$ equals 0 if a member which requires support is supported or if a member does not require support because its radius is equal to zero, and conversely equals 1 for all other eventualities. It should be noted that $\phi_{\alpha\beta}$ is raised to the power 'P', as this enables the constraint violation incurred for a structural member traversing the non-manufacturable interval $(0, \hat{r}_{min})$ to be reduced, which enables more members to traverse the non-manufacturable set concurrently for a fixed global constraint violation. In addition, it importantly does not alter the logic.

Consequently, the integral of this logical expression now provides a measure for whether the macroscale structure is self-supporting or not

$$Z = \sum_{\alpha=1}^{N} \sum_{\beta=1}^{7} \zeta_{III,\alpha\beta} \leq \epsilon \tag{5.8}$$

where Z denotes the relaxed binary integral constraint for unsupported and overhanging members, and ϵ denotes a small positive non-zero value. If this constraint is satisfied during optimisation, the macroscale structure will be self-supporting. Importantly, this constraint is continuous, implicitly discourages isolated assemblies of overhanging members and permits assemblies of interconnected self-supporting members of varying orientation to support the macroscale structure. Through this formulation, the maximum permissible flexibility in the macroscale geometry is afforded, whilst ensuring it is self-supporting. In addition, as this constraint is applied during optimisation, an optimal compromise between structural optimality and manufacturability is obtained.

Sensitivities of the relaxed binary integral constraint are calculated through analytical differentiation. Differentiation is performed through the application of the product rule, which

enables linear expressions to be separated from the binary integral constraint with only a single instance of the dependent variable, which enables these expressions to be trivially differentiated. In advance of optimisation, vectorised expressions are derived for each component of the quasi-Boolean function $\zeta_{III,\alpha\beta}$. As a result, both the evaluation of the relaxed binary integral constraint and the sensitivities can be computed for a negligible computational expense.

5.4 Slender Members

Similarly, slender members are discouraged from within the optimised solution through the application of a further relaxed binary integral constraint, which also leverages the standard projection-filter (Eq. 5.3) to precondition the normalised radii prior to input. The foundational assumption of this constraint is that the problematic buckling modes associated with slender members (see Sec. 5.1) can be alleviated through bracing. Bracing in this context corresponds to an assembly of identically orientated members that intersect with a slender member and therefore increase the lateral stiffness of this assembly and mitigate the bucking modes, as depicted in Fig. 5.10.

To begin, assemblies of identically orientated members which have the potential to be slender are identified at the macroscale through the application of a logical expression

$$\upsilon_{I,\alpha} = \prod_{j=1}^{l_s} \Xi_{\alpha\beta,j} \tag{5.9}$$

where $\Xi_{\alpha\beta}$ denotes a set containing the projected values (ϕ) for the assembly of members passing through microstructure X_α, identically orientated to structural member $X_{\alpha\beta}$ and within a predefined integer distance of $\frac{l_s}{2}$ unit-cell lengths (see Fig. 5.10), and $\upsilon_{I,\alpha}$ denotes a quasi-Boolean function, which in the limit $\eta \to \infty$ equals 1 if an assembly is present and conversely 0 if one is not. The location of the normalised radii composing the set $\Xi_{\alpha\beta}$ are identified in advance of optimisation, as a function of the minimum slender member length (l_s).

Equivalently, a further logical expression is introduced to identify assemblies of members

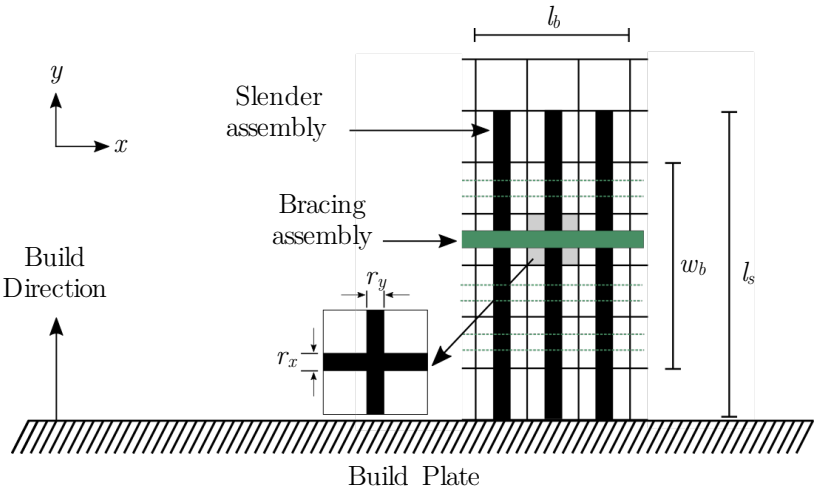


Figure 5.10: An assembly of three slender members (black) using an example two-dimensional parametrised microstructure. The problems associated with slender members can be abated by a bracing assembly of members (green); only one brace is required. The minimum length of an assembly before it is considered to be slender or bracing is denoted by l_s and l_b unit-cell lengths respectively. In addition, an assembly is only considered to be bracing if it intersects the slender member within $\frac{w_b}{2}$ unit-cell lengths of the centroid.

capable of bracing the assembly identified by Eq. 5.9

$$\upsilon_{II,\alpha\beta} = \prod_{i=1}^{w_b} \prod_{j=1}^{7} \left(1 - \prod_{k=1}^{l_b} \delta_{\beta k} \Psi_{\alpha\beta,ijk} \right) \tag{5.10}$$

where $\Psi_{\alpha\beta,ij}$ denotes a set containing the projected values (ϕ) for each assembly of identically orientated members capable of forming a brace of length l_b unit-cell lengths, within $\frac{w_b}{2}$ unit-cell lengths of microstructure X_α (see Fig. 5.10), and $\upsilon_{II,\alpha\beta}$ denotes a quasi-Boolean function, which in the limit $\eta \to \infty$ equals 0 if a brace exists and conversely 1 if a brace does not exist. In addition, the Kronecker delta function ($\delta_{\beta k}$) is applied to remove the mathematical possibility of slender members bracing themselves within this logical expression. In order to discount bracing structures intersecting at the tips of a slender member the inequality constraint

$$w_b + 2 <= l_s \tag{5.11}$$

should be obeyed, as bracing structures located at the tips do not significantly increase the orthogonal stiffness of a slender member.

Consequently, by finding the product of the two logical expressions presented in Eqs. 5.9 and5.10, it is possible to determine whether a slender member composed of structural member $X_{\alpha\beta}$ is present and whether or not it is braced according to

$$\upsilon_{III,\alpha\beta} = \upsilon_{II,\alpha\beta}\upsilon_{I,\alpha\beta} \tag{5.12}$$

where $\upsilon_{III,\alpha\beta}$ denotes a quasi-Boolean function which in the limit $\eta \to \infty$ equals 1 if a slender member is present and not braced, and 0 if either of one these two statements are false.

Consequently, the integral of this logical expression now provides a measure for whether slender members are present within the macroscale geometry

$$U = \sum_{\alpha=1}^{N}\sum_{\beta=1}^{7} \upsilon_{III,\alpha\beta} \leq \epsilon \tag{5.13}$$

where U denotes the relaxed binary integral constraint for slender members. Again, it is important to note that this constraint is continuous and permits the maximum geometric flexibility, whilst discouraging slender members from within the optimised solution. Furthermore, akin to the unsupported and overhanging members constraint, sensitivities can be derived for a negligible computational expense.

5.5 Post-processing Algorithm

The AM constraints presented within Secs. 5.2 — 5.4 are purposely formulated using continuous functions to enable their integration within a gradient based optimisation algorithm. As a result, the projected variables ($\hat{\tilde{X}}$, ϕ) are not truly binary, as illustrated in Figs. 5.1, 5.3 where $\eta \neq \infty$. Consequently, the optimised normalised radii distributions may contain normalised radii within the non-manufacturable interval $(0, \hat{r}_{min})$. Similarly, the satisfaction of the relaxed binary integral constraints may not correspond to a perfectly self-supporting geometry or the absence of slender members. It is therefore necessary to introduce a post-optimisation processing stage to clean-up the minor numerical errors resulting from the continuous functions

used for projection.

Prior to the execution of the post-processing algorithm, a threshold is applied to the optimised design variables ($\hat{\tilde{X}}^*_{\alpha\beta}$) to ensure all normalised radii are an element of the manufacturable set $\{0, [\hat{r}_{min}, \hat{r}_{max}]\}$

$$\hat{\tilde{X}}^*_{\alpha\beta,PP} = \begin{cases} X^*_{\alpha\beta}, & \text{if } \hat{\tilde{X}}^*_{\alpha\beta} \geq \hat{r}_{min} \\ \hat{r}_{min}, & \text{if } A\hat{r}_{min} \leq \hat{\tilde{X}}^*_{\alpha\beta} < \hat{r}_{min} \\ 0, & \text{otherwise} \end{cases} \tag{5.14}$$

where $\hat{\tilde{X}}^*_{\alpha\beta,PP}$ denotes the post-processed optimised design variables and A denotes a scaling factor defined in the interval $[\epsilon, 1]$, which is applied to ensure the normalised radii just below the minimum manufacturable normalised radius are not rounded down to zero. Similarly, the projected variables applied within the relaxed binary integral constraints are also thresholded

$$\phi^*_{\alpha\beta,PP} = \begin{cases} 1, & \text{if } \hat{\tilde{X}}^*_{\alpha\beta} \geq \hat{r}_{min} \\ 1, & \text{if } A\hat{r}_{min} \leq \hat{\tilde{X}}^*_{\alpha\beta} < \hat{r}_{min} \\ 0, & \text{otherwise} \end{cases} \tag{5.15}$$

where $\phi^*_{\alpha\beta,PP}$ denotes the post-processed projected variables. Both sets of post-processed variables are required within the post-processing algorithm. The pseudo-code for this post-processing algorithm is presented in Algorithm 1.

In summary, each constituent member of the macroscale geometry is inspected to determine whether it is unsupported, forms part of an isolated overhanging assembly or whether it forms part of a slender member. If any of these statements are true, the normalised radius is set equal to zero. As a result, unsupported members, isolated overhanging assemblies of members and slender members are guaranteed to be removed during the post-processing procedure. Importantly, this inspection procedure traverses the macroscale domain in order of ascending normal distance from the build plate. Therefore, the support structures investigated for each member already form part of a verified support structure. It is important to state that the

Algorithm 1: Post-processing algorithm

```
for α = 1:N do
    for β = 1:7 do
        if υ*_{III,αβ,PP} = 1 then
            X̂*_{αβ,PP} = 0                    // Member X̂*_{αβ,PP} forms
                                               // part of a slender assembly
        else
            pass                               // Member X̂*_{αβ,PP} does not form
                                               // part of a slender assembly
    if ζ*_{II,α,PP} = 1 then
        X̂*_{α,PP} = 0                         // Members X̂*_{α,PP} are not
                                               // supported
    else
        pass                                   // Members X̂*_{α,PP} are supported
```

post-processing stage should be executed using the mesh intended for fabrication, as both the supporting structures and slender members are defined as a function of the mesh's topology. Consequently, if fabrication is performed using a different mesh resolution, an intermediate interpolation stage between meshes should be performed before the post-processing stage to ensure that the fabricated geometry is manufacturable.

Naturally, by post-processing the optimised solution, the volume-fraction of the macroscale geometry is perturbed away from the desired value. In order to recover the desired volume-fraction, a one-dimensional optimisation is undertaken to search for the optimal scaling factor (A), which ensures that the post-processed normalised radii distributions return a volume fraction equal to the desired value.

5.6 Updated Optimisation Formulation

As detailed in Sec. 3.5, the microscale parameters serve as continuous design variables during optimisation enabling functional objectives at the macroscale to be targeted. Through the application of this methodology, it is permissible to derive self-supporting, hierarchical structures with varying topology, suitable for fabrication through AM processes.

Optimisation problems are formulated in accordance with Eq. 3.43 with the addition of the two relaxed binary integral constraints, leading to

$$\begin{aligned}
\underset{X}{\text{Minimise:}} \quad & J(u(E^H(\hat{\tilde{X}})), E^H(\hat{\tilde{X}})) \\
\text{Subject to:} \quad & F(u(E^H(\hat{\tilde{X}})), E^H(\hat{\tilde{X}})) = 0 \\
& V_f - \int_\Omega V(\hat{\tilde{X}}_\alpha) \geq 0 \\
& 1 - Z(\phi(\tilde{X})) \geq 0 \\
& 1 \times 10^{-4} - U(\phi(\tilde{X})) \geq 0 \\
& \hat{\tilde{X}}_{\alpha\beta} \geq X_{min,\beta}, \ \forall\, \beta \ \in \ \{1, 2, 3, \ldots, n\} \\
& \hat{\tilde{X}}_{\alpha\beta} \leq X_{max,\beta}, \ \forall\, \beta \ \in \ \{1, 2, 3, \ldots, n\}
\end{aligned} \tag{5.16}$$

where $X_{min,\beta}$ and $X_{max,\beta}$ refer to the values presented within Tab. 5.1. It is prudent to add that the constraint violation permitted for each respective binary integral constraint is chosen such that it is permissible for a structural member to traverse the non-manufacturable interval $(0, \hat{r}_{min})$ during optimisation, but not so large such that the satisfaction of these constraints does not correspond to the absence of problematic fabrication features.

5.6.1 Sensitivities

Sensitivities are derived in accordance with Eq. 3.45, however, an additional step is required to account for the application of projection filters applied within this section. For the objective function, this corresponds to

$$\frac{\partial J}{\partial X} = \left[\sum_{i=1}^{6} \sum_{j=1}^{6} \left(\frac{\partial J}{\partial E_{ij}^H(\tilde{X})} \circ \frac{\partial E_{ij}^H(\tilde{X})}{\partial \tilde{X}} \right) \right]^T \frac{\partial \hat{\tilde{X}}}{\partial \tilde{X}} \frac{\partial \tilde{X}}{\partial X} \tag{5.17}$$

where

$$\frac{\partial \hat{\tilde{X}}_{\alpha\beta}}{\partial \tilde{X}_{\alpha\beta}} = \frac{\eta \tilde{X}_{\alpha\beta} \left(1 - \tanh\left(\eta \left(\tilde{X}_{\alpha\beta} - \hat{r}_{min} \right) \right)^2 \right) + \tanh\left(\eta \hat{r}_{min} \right) + \tanh\left(\eta \left(\tilde{X}_{\alpha\beta} - \hat{r}_{min} \right) \right)}{\tanh\left(\eta \hat{r}_{min} \right) + \tanh\left(\eta \left(1 - \hat{r}_{min} \right) \right)} \tag{5.18}$$

and for the constraints

$$\frac{\partial V(\hat{\tilde{X}})}{\partial X} = \frac{\partial V(\hat{\tilde{X}})}{\partial \hat{\tilde{X}}} \frac{\partial \hat{\tilde{X}}}{\partial \tilde{X}} \frac{\partial \tilde{X}}{\partial X} \tag{5.19}$$

$$\frac{\partial Z(\phi(\tilde{X}))}{\partial X} = \frac{\partial Z(\phi(\tilde{X}))}{\partial \phi} \frac{\partial \phi}{\partial \tilde{X}} \frac{\partial \tilde{X}}{\partial X} \tag{5.20}$$

$$\frac{\partial U(\phi(\tilde{X}))}{\partial X} = \frac{\partial U(\phi(\tilde{X}))}{\partial \phi} \frac{\partial \phi}{\partial \tilde{X}} \frac{\partial \tilde{X}}{\partial X} \tag{5.21}$$

where

$$\frac{\partial \phi_{\alpha\beta}}{\partial \tilde{X}_{\alpha\beta}} = \frac{\eta \left(1 - \tanh\left(\eta \left(-\hat{r}_{min} + \tilde{X}_{\alpha\beta}\right)\right)^2\right)}{\tanh\left(\eta \hat{r}_{min}\right) + \tanh\left(\eta \left(1 - \hat{r}_{min}\right)\right)} \tag{5.22}$$

5.7 Summary

Presented within this chapter is a detailed breakdown of the methodology applied to construct the novel AM constraints. This methodology is applied to a series of mechanical structural optimisation problems within Chapter 6 to demonstrate their efficacy.

Chapter 6

Design for Additive Manufacturing Results

The AM constraints presented within Chapter 5 are applied in combination with the base MTO here to solve a range of structural optimisation problems. Similarly, optimisation is terminated for all problems presented within this section when both the objective function decreases by less than 0.01% for five successive iterations and the normalised constraint violation is less than 1×10^{-3}.

6.1 Additive Manufacturing Constraints

The framework presented within Sec. 5 is applied within this section to solve two classical compliance minimisation problems. The initialisation conditions are kept constant for both problems presented within this section and are detailed in Tab. 6.1. The variables used in the definition of the AM constraints are kept constant for both problems presented within this section and are listed in Tab. 6.2. Consequently, assemblies of identically orientated members are considered to be slender or bracing if they are greater than or equal to 5 or 3 unit-cell lengths respectively. In order to discount bracing assemblies located at the tips of a slender member, Eq. 5.11 is satisfied by setting $w_b = 3$. In addition, the Heaviside thresholding parameter

(η) is defined in all cases as a function of the minimum manufacturable normalised radius, ensuring that ϕ and $\hat{\tilde{X}}$ equal 0.95 and $0.95\hat{r}_{min}$ respectively when $\tilde{X} = 1.2\hat{r}_{min}$. This ensures that equivalent penalisation is applied for different values of the minimum manufacturable normalised radius. It is important to note that the target of this work is to demonstrate the increased mechanical performance permissible through the integration of the novel AM constraints presented within this thesis, in comparison to the application of prohibitive lower bound constraints. As a result, whilst the AM parameters (Tab. 6.2) applied are representative, they are not specific to a single printing architecture.

Design Variable	Initialisation Value
$\hat{r}_{1-7}$	$\hat{r}_{min}$

Table 6.1: The initialisation values used for all parameters at the beginning of optimisation

AM Constraint Parameter	Value
P (Eq. 5.8)	10
l_s (Eq. 5.9)	5
l_b (Eq. 5.10)	3
w_b (Eq. 5.10)	3

Table 6.2: The parameters used in the definition of the two relaxed binary integral constraints (see Eqs. 5.8 and 5.13)

To simplify the interpretation of the optimised results, only build directions aligned to the three orthogonal axes are considered. As a result, independent of the build direction, all four diagonal members are always self-supporting. Further to this, the face member aligned to the build direction is also always self-supporting, however, the remaining two face members are not and are therefore considered to be overhanging members. Consequently, the presence of these two face members should only be permitted within the optimised solution if intra-cell support is provided. Furthermore, the build plate is defined for all cases by the plane normal to the build direction vector passing through the coordinate $(0, 0, 0)$. In addition, the discussion of the optimised results is limited within this section to manufacturability.

6.1.1 Unit Cube — Compliance Minimisation

The first geometry targeted by compliance minimisation is a unit-cube with two equally distributed loads applied in the $X - Z$ plane at $y = 0.1$ m and a single Dirichlet boundary constraint, as illustrated in Fig. 6.1. A limit is placed on the availability of material within the macroscale domain by a 30% volume-fraction constraint. Compliance in this instance is quantified by

$$J = \int_{\Gamma_1} \boldsymbol{t} \cdot \boldsymbol{u} \, d\Gamma_1 + \int_{\Gamma_2} \boldsymbol{t} \cdot \boldsymbol{u} \, d\Gamma_2 \tag{6.1}$$

where Γ_1 and Γ_2 denotes the two surfaces where the equally distributed loads are applied (see Fig. 6.1). The unit-cube geometry is discretised into $8,000$ first-order hexahedral elements using $20 \times 20 \times 20$ divisions in the x, y and z axes respectively.

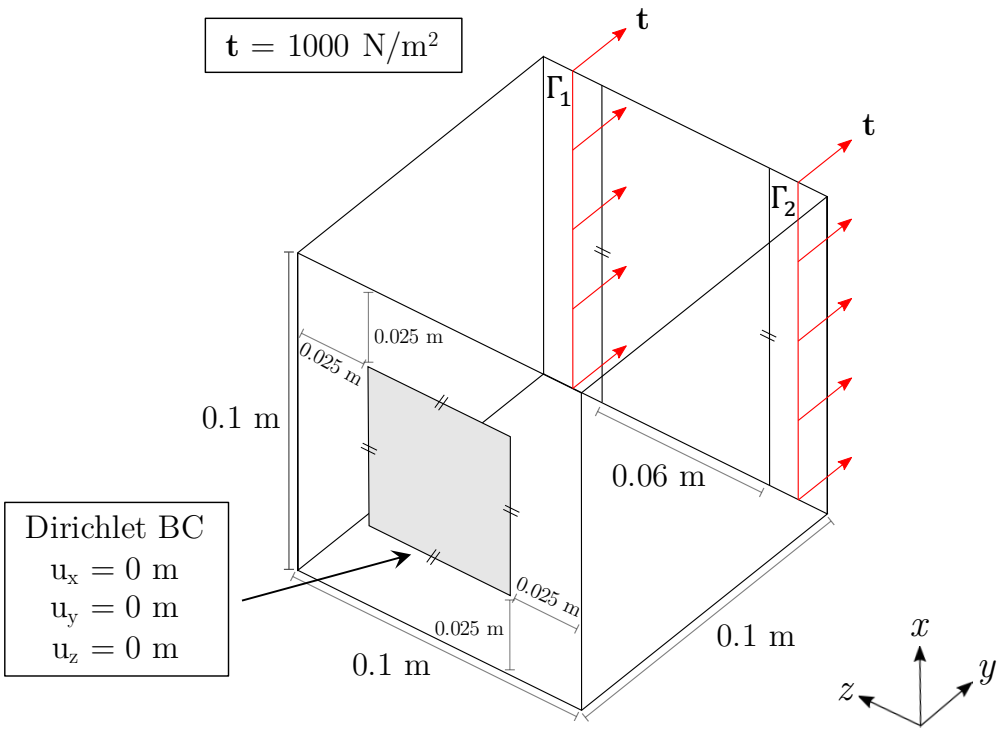


Figure 6.1: The unit-cube geometry subjected to compliance minimisation, with the loads and boundary conditions highlighted

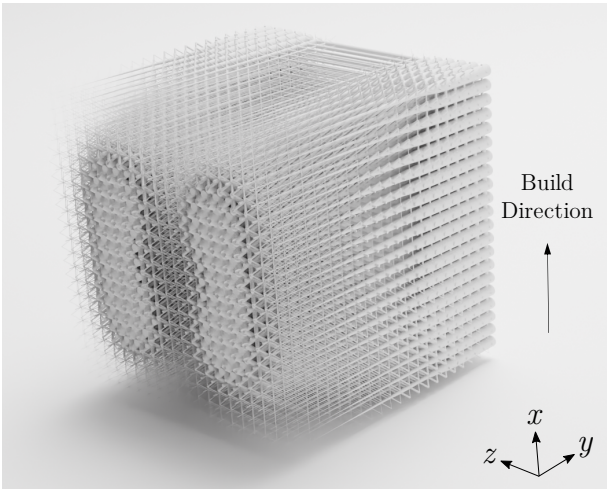
Initially, this problem is solved using the base multiscale optimisation framework and a lower bound constraint on the normalised radii of 0, in order to demonstrate the formation of the problematic fabrication features detailed in Sec. 5.1. To be consistent with subsequent cases

and to aid comparison, problematic fabrication features are identified for this case using a minimum manufacturable normalised radius of 0.07 and the parameters defined in Tab. 6.2. In addition, it is also assumed that the build direction vector is aligned to the x-axis and that the build plate is defined by the plane normal to the x-axis passing through the coordinate $(0, 0, 0)$. The optimised reconstructed macroscale geometry for this case is illustrated in Fig. 6.2, including annotations highlighting the problematic fabrication features.

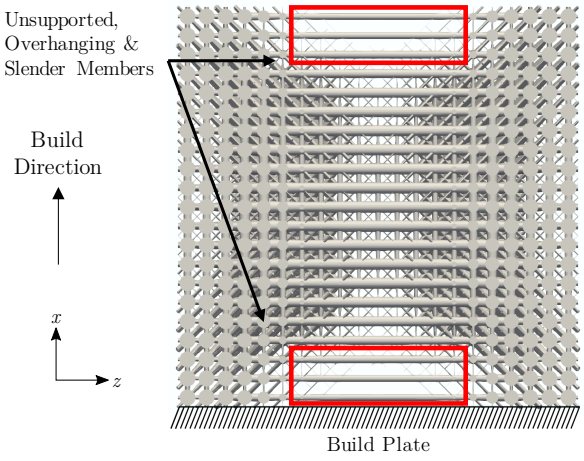
Spatial variations in the microstructural topology are clearly illustrated in Fig. 6.2, which is a result of permitting the normalised radii to equal zero during optimisation. Consequently, this leads to the formation of unsupported members, isolated assemblies of overhanging members and slender members, as illustrated in Figs. 6.2b and 6.2c. In summary, by permitting unconstrained optimisation of the normalised radii within the non-manufacturable interval $(0, 0.07)$, a series of problematic fabrication features have arisen.

This problem is solved once more using the base framework and a lower bound constraint on the normalised radii of 0. However in this instance, the integrated AM constraints presented within Sec. 5 are included. Similarly to the unconstrained case, the minimum manufacturable normalised radius is set to 0.07, the build direction vector is aligned to the x-axis and the build plate is defined by the plane normal to the x-axis passing through the coordinate $(0, 0, 0)$. It is important to note, that this build direction is specifically chosen, as structural member r_5 (aligned to the x-axis) is the member allocated the least amount of material within the unconstrained case (see Fig. 6.2). Therefore, this build direction should present the hardest challenge to the optimiser, as allocating further material to structural member 5 is clearly sub-optimal with respect to the compliance functional. The pre-processed and post-processed optimised reconstructed macroscale geometries for this case are illustrated in Fig. 6.3.

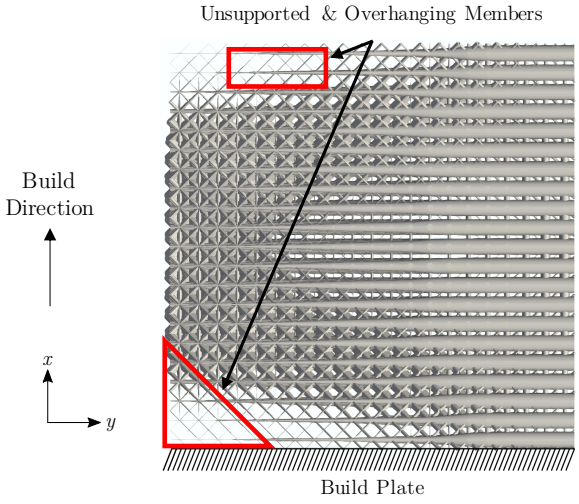
It is clearly illustrated in Fig. 6.3 that the problematic fabrication features highlighted within the unconstrained case have been successfully eliminated by the integrated AM constraints. The biggest change noted between the constrained and unconstrained cases is the addition of supporting structures located near to the build plate, most notably using structural member r_5 (aligned to the build direction). In addition, structural member r_5 is allocated further



(a) Isometric projection of the reconstructed unit-cube geometry



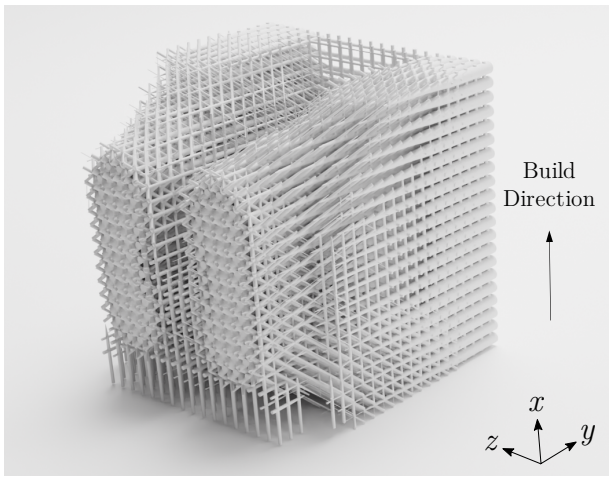
(b) Orthogonal projection of the reconstructed unit-cube geometry onto the $X - Z$ plane



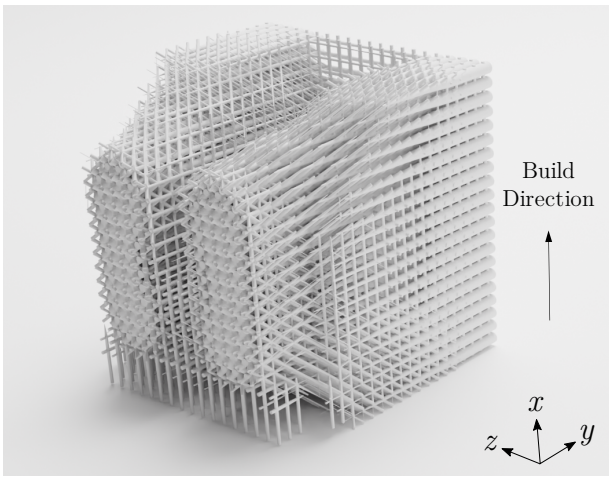
(c) Orthogonal projection of the reconstructed unit-cube geometry onto the $X - Y$ plane

Figure 6.2: The optimised reconstructed macroscale geometry derived after performing a compliance minimisation on the unit-cube geometry (see Fig. 6.1) using the base multiscale optimisation framework and a lower bound constraint on the normalised radii of 0. Examples of the problematic fabrication features outlined in Sec. 5.1 are highlighted using red boxes. It is assumed that the minimum manufacturable normalised radius is equal to 0.07, the build direction vector is aligned to the x-axis and the build plate is defined by the plane normal to the x-axis passing through the coordinate $(0, 0, 0)$, to aid subsequent comparison.

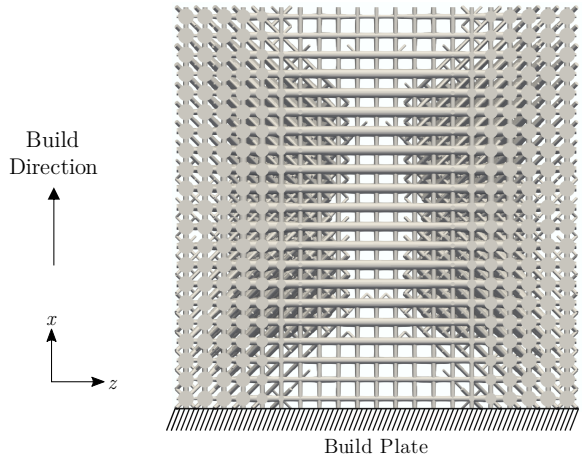
material in the $X - Z$ plane at $y = 0.1$ m (see Fig. 6.3c) to brace the slender members noted in the unconstrained case. Whilst additional supporting and bracing members are now included, the salient structures are identical in both the unconstrained and constrained cases. It is clear to see how effective the AM constraints are at discouraging problematic fabrication features, as negligible changes are noted between the pre-processed (Fig. 6.3a) and



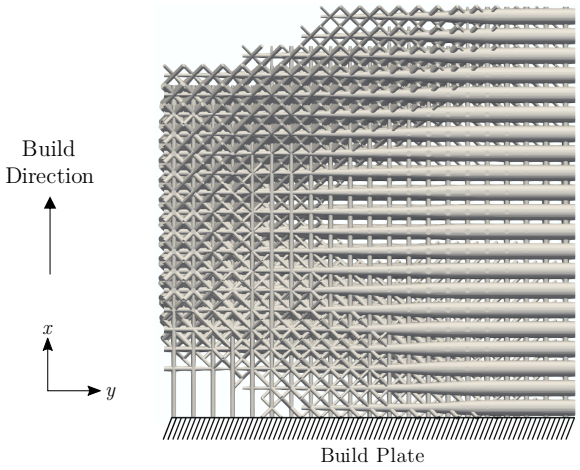
(a) Isometric projection of the pre-processed reconstructed unit-cube geometry



(b) Isometric projection of the post-processed reconstructed unit-cube geometry



(c) Orthogonal projection of the post-processed reconstructed unit-cube geometry onto the $X - Z$ plane



(d) Orthogonal projection of the post-processed reconstructed unit-cube geometry onto the $X-Y$ plane

Figure 6.3: The pre-processed and post-processed, optimised reconstructed macroscale geometry for the unit-cube compliance minimisation, performed using the base framework with integrated AM constraints. In this instance, the minimum manufacturable normalised is set equal to 0.07, the build direction vector is aligned to the x-axis and the build plate is defined by the plane normal to the x-axis passing through the coordinate $(0, 0, 0)$.

post-processed (Fig. 6.3b) geometries.

In order to demonstrate the efficiency of the augmented projection filter, a histogram comparing the optimised normalised radii distributions for both the constrained and unconstrained cases is illustrated in Fig. 6.4. It is clearly noted in Fig. 6.4 that the majority of the normalised radii in the unconstrained case are located within the non-manufacturable interval $(0, 0.07)$.

In contrast, the non-manufacturable interval is more sparsely populated in the constrained case, as more structural members are allocated either a normalised radius of zero or one above the minimum manufacturable normalised radius. It is also noted that as the normalised radius increases in Fig. 6.4, the constrained distribution reverts back to the unconstrained distribution, demonstrating that above the minimum manufacturable normalised radius the augmented projection filter does not significantly influence the optimisation.

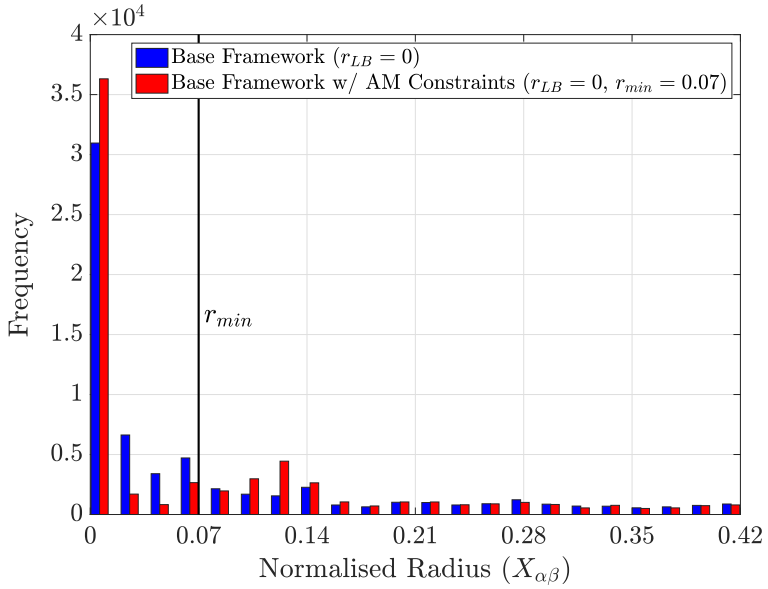


Figure 6.4: Comparing the optimised normalised radius distributions in both the unconstrained and constrained cases for the unit-cube compliance minimisation. A total of 21 uniformly spaced intervals are used in the construction of this histogram.

The optimised compliance functionals for the two cases presented within this section, alongside three further cases to aid comparison, are listed in Tab. 6.3. The additional cases include, a further application of the base multiscale optimisation framework with a prohibitive lower bound constraint set equal to the minimum manufacturable normalised radius (0.07) and two further applications of the base framework with integrated AM constraints for build directions aligned to the y and z axes respectively.

Out of the five cases investigated, the case with the lowest compliance functional is where the base multiscale optimisation framework is applied with a lower bound constraint on the normalised radii of 0, as this set-up permits the maximum flexibility in the macroscale geometry during optimisation. By increasing the lower bound constraint to the assumed minimum

manufacturable normalised radius of 0.07, the objective function increases by 22.8%. This is not surprising, as the majority of the normalised radii are located below this lower bound in the unconstrained case (see Fig. 6.4). Therefore, a significant amount of material is allocated in sub-optimal regions of the design space in this instance. On average, the compliance functionals for the post-processed constrained cases are 3.4% higher and 20.1% lower than the unconstrained cases with lower bound constraints of 0 and 0.07 applied respectively. Therefore, through the inclusion of the AM constraints presented within this Sec. 5 a significant percentage of the mechanical performance is recovered whilst maintaining manufacturability. In addition, it is also noted that post-processing the constrained optimisations results in a negligible impact on the objective functionals, as they only increase by 0.42% on average, further demonstrating the efficacy of this approach.

Framework Compliance (Nm)	Compliance (Nm)	Post-processed
Base ($\hat{r}_{LB} = 0$)	4.280×10^{-9}	N/A
Base ($\hat{r}_{LB} = 0.07$)	5.542×10^{-9}	N/A
Base w/ AM Constraints ($\hat{r}_{LB} = 0$, $\hat{r}_{min} = 0.07$ Build Direction: x-axis)	4.424×10^{-9}	4.438×10^{-9}
Base w/ AM Constraints ($\hat{r}_{LB} = 0$, $\hat{r}_{min} = 0.07$, Build Direction: y-axis)	4.412×10^{-9}	4.430×10^{-9}
Base w/ AM Constraints ($\hat{r}_{LB} = 0$, $\hat{r}_{min} = 0.07$, Build Direction: z-axis)	4.391×10^{-9}	4.416×10^{-9}

Table 6.3: The optimised compliance functionals for the unit-cube problem derived using both the base framework and the base framework with integrated AM constraints

6.1.2 Cantliever Beam — Compliance Minimisation

The second geometry subjected to compliance minimisation is a cantilever beam with a single distributed load applied to the bottom surface and a single Dirichlet boundary constraint, as illustrated in Fig. 6.5. A limit is again placed on the availability of material within the macroscale domain by a 30% volume-fraction constraint. Compliance in this instance is quantified by the homogenised mechanical properties

$$J = \int_{\Gamma_1} \boldsymbol{t} \cdot \boldsymbol{u} \, d\Gamma_1 \tag{6.2}$$

where Γ_1 denotes the section of the bottom surface where the distributed load is applied (see Fig. 6.5). The cantilever beam geometry is discretised into $12,288$ first-order hexahedral elements using $48 \times 16 \times 16$ divisions in the x, y and z axes respectively.

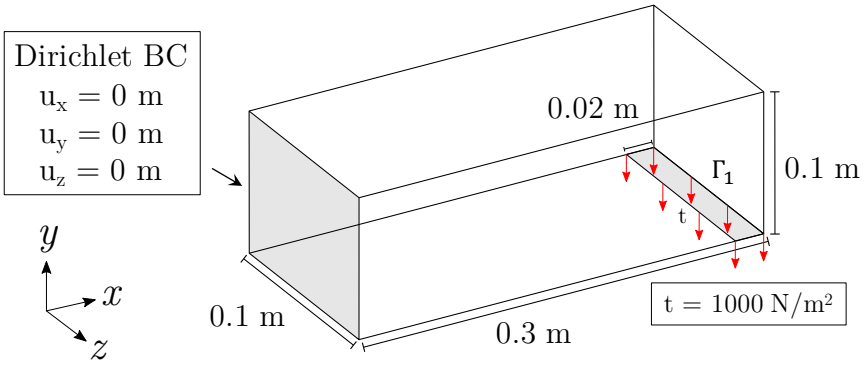
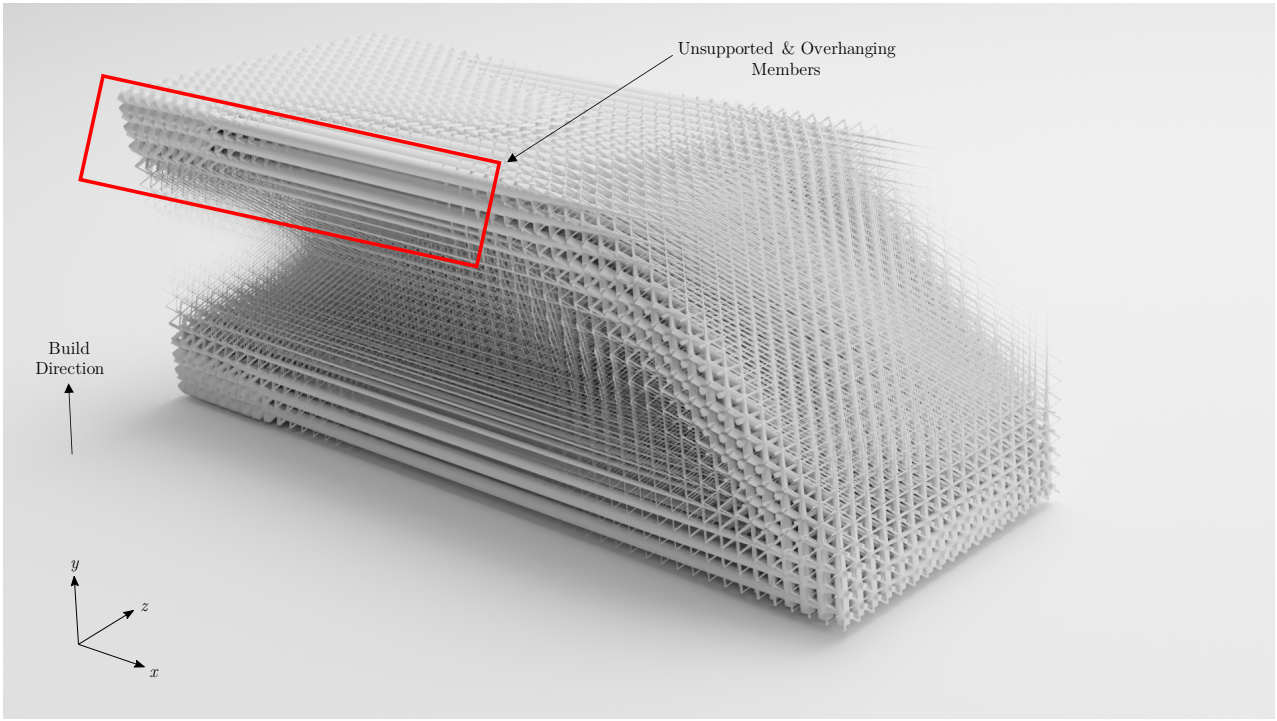


Figure 6.5: The cantilever beam geometry subjected to compliance minimisation, with the loads and boundary conditions highlighted

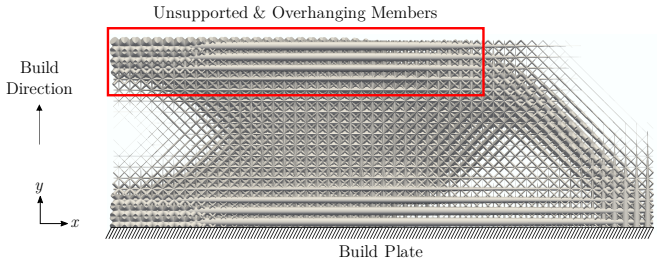
In order to demonstrate the formation of problematic fabrication features, this problem is solved at first using the base multiscale optimisation framework and a lower bound constraint on the normalised radii of 0. Similarly, to aid subsequent comparison, problematic fabrication features are identified in this instance using a minimum manufacturable normalised radius of 0.07, a build direction aligned to the y-axis and a build plate defined by the plane normal to the y-axis passing through the coordinate $(0, 0, 0)$. The optimised reconstructed macroscale geometry for this case is illustrated in Fig. 6.6, including annotations highlighting the problematic fabrication features.

As depicted in Fig. 6.6, it is again noted that as a consequence of permitting the normalised radii to equal zero during optimisation, the microstructural topology varies spatially. Consequently, this again leads to the formation of unsupported members, isolated assemblies of overhanging members and slender members, as illustrated in Fig 6.6. In this instance, a large section of the macroscale geometry is unsupported. This is due to the I-beam shape of the macroscale geometry, which whilst structurally efficient results in unsupported and overhanging flange sections of the I-beam.

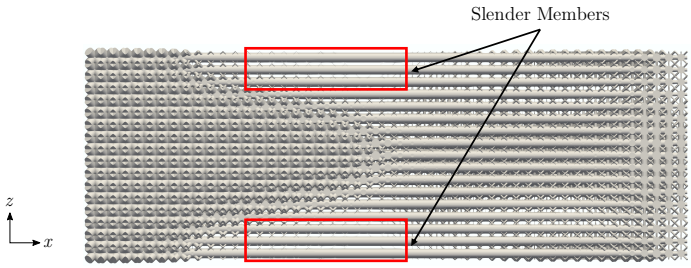
This problem is solved once more using the base framework and a lower bound constraint on the normalised radii of 0. However, in this instance, the integrated AM constraints presented



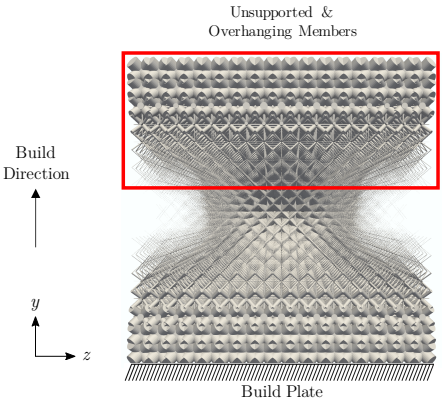
(a) Isometric projection of the reconstructed cantilever beam geometry



(b) Orthogonal projection of the reconstructed cantilever beam geometry onto the $X - Y$ plane



(c) Orthogonal projection of the reconstructed cantilever beam geometry onto the $X - Z$ plane



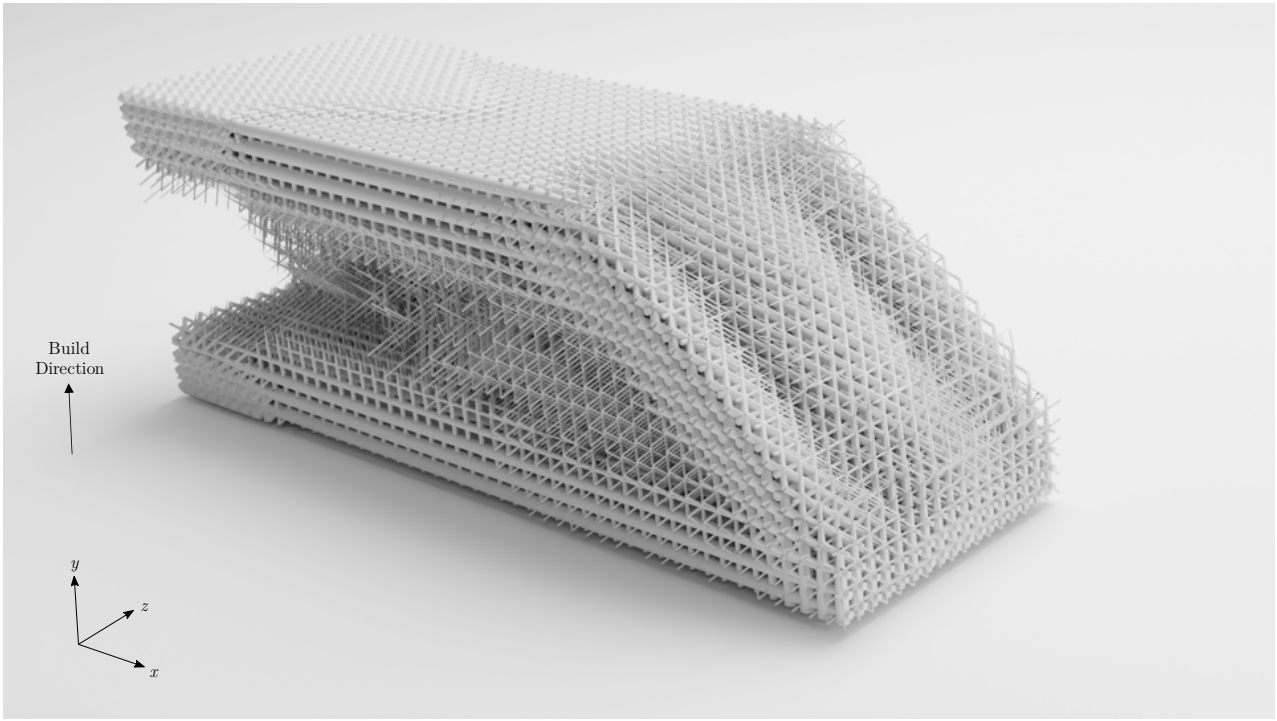
(d) Projection of the reconstructed cantilever beam geometry onto the $Y - Z$ plane

Figure 6.6: The optimised reconstructed macroscale geometry derived after performing a compliance minimisation on the cantilever beam geometry (see 6.5 using the base multiscale optimisation framework and a lower bound constraint on the normalised radii of 0. Examples of the problematic fabrication features outlined in Sec. 5.1 are highlighted using red boxes. It is assumed that the minimum manufacturable normalised radius is equal to 0.07, the build direction vector is aligned to the y-axis and the build plate is defined by the plane normal to the y-axis passing through the coordinate $(0, 0, 0)$, to aid subsequent comparison.

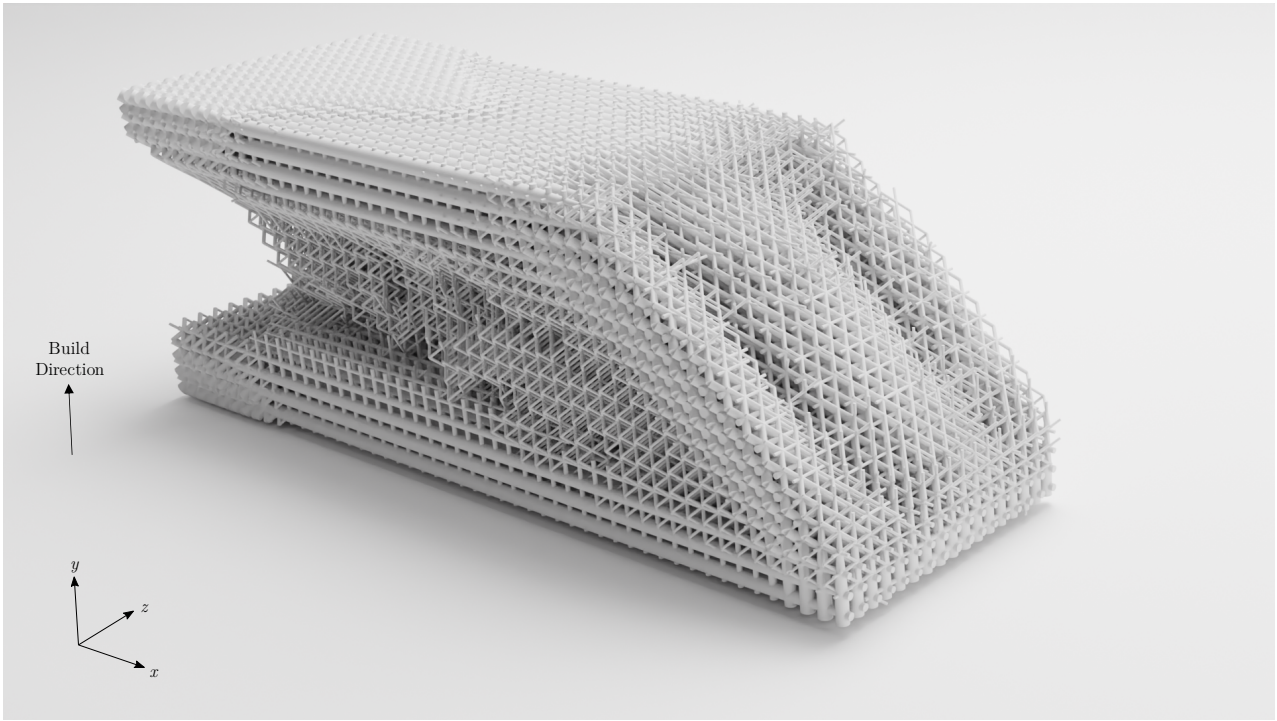
within this Sec. 5 are included. Similarly to the unconstrained case, the minimum manufacturable normalised radius is set to 0.07, the build direction vector is aligned to the y-axis and the build plate is defined by the plane normal to the y-axis passing through the coordinate $(0, 0, 0)$. A build direction vector aligned to the y-axis is chosen in this instance as it presents a significant challenge to the optimiser given how many structural members are unsupported and overhanging in the unconstrained case. The pre-processed and post-processed optimised macroscale reconstructions for this case are illustrated in Fig. 6.6.

It is clearly illustrated in Fig. 6.6, that through the inclusion of the integrated AM constraints, the problematic fabrication features identified in the unconstrained case are no longer present. Most notably, the upper flange section of the I-beam is now supported by diagonal members which originate from the web section, as illustrated in Fig. 6.6e. Furthermore, the slender members located in the flanges are now braced by structural member 6 (aligned to the build direction). In both instances, it is clear to see how the additional supporting structures contribute significantly to minimising the compliance functional whilst enabling manufacturability. In addition, the salient structures are again identical to those noted in the unconstrained case. It it clear to see that minimal differences are again noted between the pre-processed (Fig. 6.7a) and post-processed (Fig. 6.7b) geometries, with only a few diagonal members removed near to the root of the cantilever beam, and near to the exterior where the normalised radius of multiple diagonal members has been increased to the minimum manufacturable normalised radius.

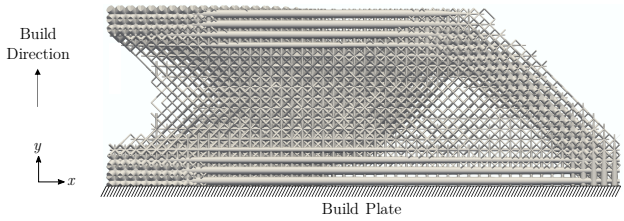
To further demonstrate the efficacy of the augmented projection filter, a histogram comparing the optimised normalised radii distributions for both the unconstrained and constrained cases is illustrated in Fig. 6.7. A similar pattern is noted to the unit-cube compliance minimisations, where it is observed that the majority of the normalised radii are located within the non-manufacturable interval $(0, 0.07)$ in the unconstrained case. Similarly, the normalised radii are shifted either towards a normalised radius of zero or to one above the minimum manufacturable normalised radius in the constrained case. Again, it is noted that the constrained normalised radii distribution reverts back to the unconstrained distribution above the minimum manufacturable normalised radius.



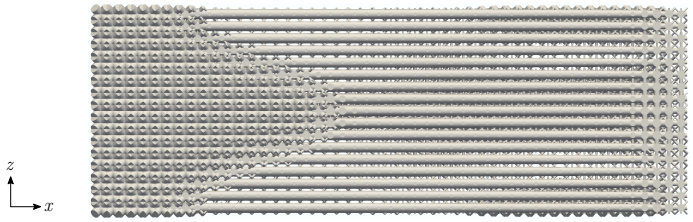
(a) Isometric projection of the pre-processed reconstructed cantilever beam geometry



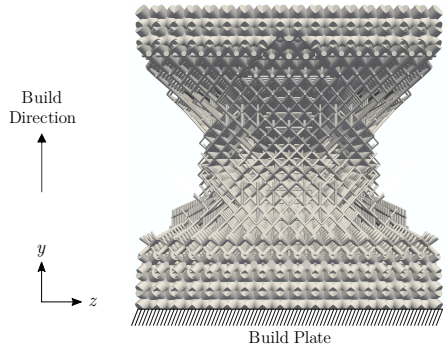
(b) Isometric projection of the post-processed reconstructed cantilever beam geometry



(c) Orthogonal projection of the post-processed reconstructed cantilever beam geometry onto the $X-Y$ plane



(d) Orthogonal projection of the post-processed reconstructed cantilever beam geometry onto the $X-Z$ plane



(e) Projection of the post-processed reconstructed cantilever beam geometry onto the $Y-Z$ plane

Figure 6.6: The post-processed and post-processed, optimised reconstructed macroscale geometry for the cantilever beam compliance minimisation, performed using the base framework with integrated AM constraints. In this instance, the minimum manufacturable normalised is set equal to 0.07, the build direction vector is aligned to the y-axis and the build plate is defined by the plane normal to the y-axis passing through the coordinate $(0, 0, 0)$.

To further demonstrate the increased mechanical performance permissible through the integrated AM constraints presented within this section, parametric studies for both the lower bound constraint and the minimum manufacturable normalised radius are conducted for the cantilever beam problem. A total of six lower bound constraints are investigated, starting at 0 and increasing linearly in steps of 0.02 before finishing at 0.1. In addition, a total of five minimum manufacturable normalised radii are investigated, starting at 0.02 and increasing linearly in steps of 0.02 before finishing at 0.1. A minimum manufacturable normalised radius of 0 is not considered as it represents an unconstrained case. The results from both parametric studies are illustrated in Fig. 6.8.

It is clearly illustrated in Fig. 6.8 that as the lower bound constraint and the minimum manufacturable normalised radius increases, the difference between the optimised compliance functionals increases. On average, the compliance functional is 16.5% lower when comparing the base framework with integrated AM constraints to the base framework. In addition, the compliance functionals for the present framework is only on average 2.95% larger when compared

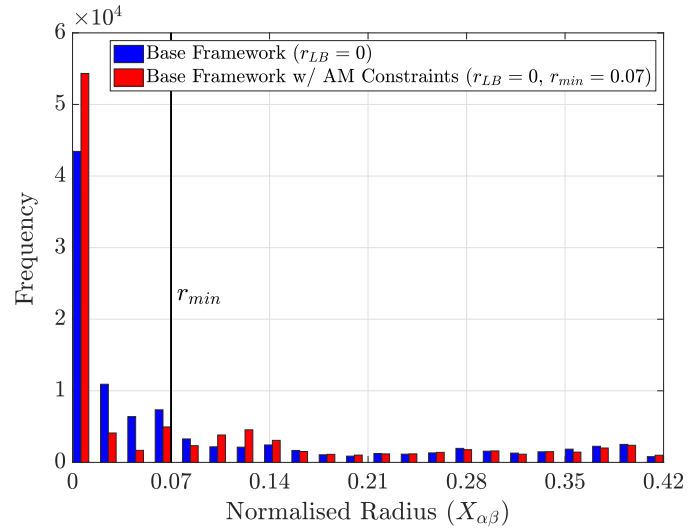


Figure 6.7: Comparing the optimised normalised radius distributions in both the unconstrained and constrained cases for the cantilever beam compliance minimisation. A total of 21 uniformly spaced intervals are used in the construction of this histogram.

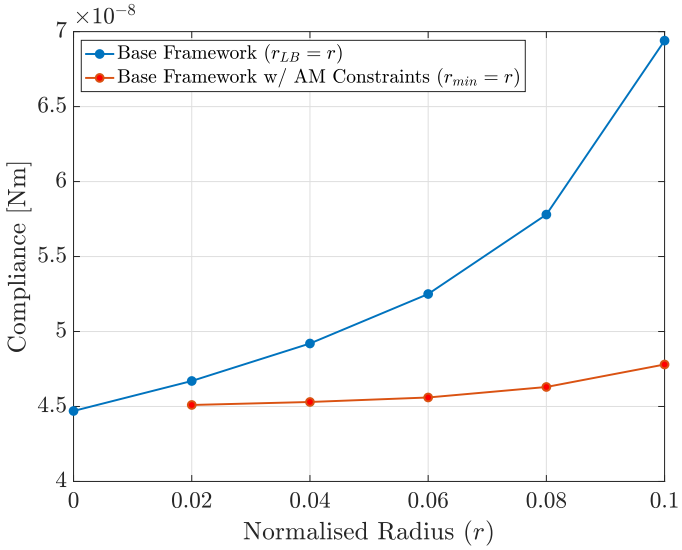


Figure 6.8: The compliance functionals derived from the parametric studies performed for both the lower bound constraint on the normalised radii and the minimum manufacturable normalised radii

to the base framework with a lower bound constraint on the normalised radii of 0. Therefore, it is further demonstrated that through the inclusion of the novel AM constraints presented within this Sec. 5, it is permissible to derive manufacturable structures with substantially increased mechanical performance in comparison to equivalent frameworks with prohibitive lower bound constraints applied.

Chapter 7

Conclusion

A novel multiscale topology optimisation framework with concurrent coupling between scales is successfully implemented and presented within this thesis. By embedding a concurrent coupling strategy, only the microscale data required to evaluate the macroscale model during each iteration of optimisation is collected. The microscale data collected during optimisation is stored in a reusable database to reduce the computational expense of successive iterations or entirely new optimisation problems. This permits a previously intractable number of design variables (10) to be used in the parametrisation of the microscale geometry, which enables accessibility to a greater range of extremal point properties during optimisation.

Two classical compliance minimisation problems are solved and benchmarked against a Solid Isotropic Material with Penalisation (SIMP) based topology optimisation, using a penalisation parameter of 3 and an identical discretisation. The first problem is a cantilever beam, subjected to a 30% volume-fraction constraint. A reduction in the compliance functional of 27.4% is noted in comparison to the equivalent SIMP case and only 0.48% of the total parent database is required to derive this result. For the second compliance minimisation problem, a parametric study of volume-fraction is undertaken using a unit-cube geometry, to demonstrate the reusability of the microscale data collected during the optimisation procedure. A total of 7 volume-fractions are investigated between 20% and 50%. On average, 63.1% of the microscale data required to derive the optimised solution is collected in a previous optimisation.

In addition, the present framework outperformed SIMP for all 7 volume-fraction cases, with a compliance functional on average of 38.3% lower. The advantage of this framework over SIMP based topology optimisations, is the clear accessibility afforded by the microscale parameterisation to orthotropic microstructures with efficient stiffness characteristics, in comparison to the fixed isotropic material properties locally available within a SIMP based topology optimisation. In total, only 1.85% of the total parent database is required to derive all 8 optimised solutions presented within Sec. 4.1. This represents a significant improvement in comparison to contemporary sequential methods, which require fully-populated databases prior to undertaking optimisation.

In addition, a novel optimisation strategy is implemented, which seeks to further reduce the computational expense associated with parent microscale simulations, by initially optimising a subset of the microscale design parameters before activating additional microscale design parameters as the optimisation progresses. Two applications of this methodology are presented using both the cantilever beam case and the unit-cube case presented within Sec. 4.1. A 5.5% and a 0.4% reduction in the number of parent microscale simulations is noted for the cantilever beam and unit-cube cases respectively, in comparison to equivalent simulations which do not leverage the dimension addition capability. The reduction in the number of parent microscale simulations is a result of initially traversing through sub-spaces of the ten-dimensional hyperspace, as the hyper-volume resolved by each simplex is inversely proportional to its dimensionality (see Sec. 3.4.2). Therefore, fewer simplices are required to resolve a unit volume within a sub-space of the ten-dimensional hyperspace, in comparison to the number required within the ten-dimensional hyperspace.

In addition, novel additive manufacturing constraints for open-walled microstructures are also successfully implemented and presented within this thesis. Through the application of this methodology, it is permissible to derive self-supporting, hierarchical structures with varying topology, suitable for fabrication through additive manufacturing processes. This is achieved through the application of an augmented projection filter, which permits control over the minimum length scale and two relaxed binary integral constraints which prohibit the formation of unsupported members, isolated assemblies of overhanging members and slender

members during optimisation. To demonstrate the efficacy of this approach, two classical compliance minimisation problems are solved using the present multiscale topology optimisation framework with and without the integrated additive manufacturing constraints, to enable comparison. The first problem considered is the compliance minimisation of a unit-cube, with a minimum manufacturable normalised radius set equal to 0.07. This problem is solved using the base framework with a lower bound constraint on the normalised radii of zero. As a result of the flexibility permitted, a series of problematic fabrication features arise in the optimised solution. However, through the integration of the additive manufacturing constraints, these problematic features are eliminated with a only a 3.4% increase in the compliance functional noted on average.

A further application of the base framework with a lower bound constraint set equal to the minimum manufacturable normalised radius is solved. On average, the compliance functional is 22.8% higher in this case in comparison to the constrained cases, demonstrating the increased mechanical performance permissible through this approach. The second problem considered is a cantilever beam geometry. Similarly, through the application of the base framework with a lower bound constraint on the normalised radii of zero, a series of problematic fabrication features again arise. Upon application of the additive manufacturing constraints, these features are again eliminated. To further demonstrate the increased mechanical performance permissible through this approach, parametric studies for both the lower bound constraint and the minimum manufacturable normalised radius are performed for this case. On average, the compliance functional decreases by 16.5% through the inclusion of the additive manufacturing constraints in comparison to equivalent prohibitive lower bound constraints. In summary, it has been shown that through the application of the additive manufacturing constraints presented within this thesis, it is possible to derive manufacturable geometries with significantly increased mechanical performance relative to contemporary open-walled multiscale optimisation frameworks which apply prohibitive lower bound constraints to guarantee manufacturability.

The methodology presented within this thesis enables the derivation of manufacturable optimised multiscale structures for a significantly reduced computational expense in comparison to contemporary multivariate MTO frameworks. The computational savings afforded by concurrently coupling the micro and macroscale models during optimisation is invariant to the defined parametrisation. Consequently, the present methodology enables accessibility to a space of manufacturable designs, unobtainable through sequential methods. This enables a step towards the geometric freedom afforded by hierarchical MTO frameworks for a significantly reduced computational cost, which has broad repercussions within the field of multiscale topology optimisation.

7.1 Summary of Thesis Achievements

The specific contributions of this thesis are summarised as follows:

- The development of a concurrently coupled parametrisation based multiscale topology optimisation framework, which enables multiscale topology optimisation to be performed at a significantly reduced computational cost in comparison to equivalent sequentially coupled frameworks.

- The development of a bespoke optimisation strategy, which further reduces the computational cost of optimisation, by initially optimising only the most sensitive microscale parameters, before introducing the insensitive parameters at a latter stage of the optimisation procedure.

- The development of novel additive manufacturing constraints for open-walled microstructures, which permits self-supporting, hierarchical structures with varying topology, suitable for fabrication through additive manufacturing processes to be derived.

7.2 Future Work

To conclude this thesis, areas for future work are suggested. As detailed within Sec. 3.4.3, a maximum is placed on the discretisation resolution of the microscale model, as a function of the computational resources available throughout the duration of this research project. In turn, this constraint also places a maximum on both the number of microscale parameters available to parametrise the microscale geometry and the number of levels available to discretise each dimension within the full-factorial DOE. Consequently, the computation time required to execute the microscale model severely restricts the fidelity of the present methodology and denotes the greatest challenge faced to extend multivariate multiscale topology optimisation frameworks to more complex problems.

It is therefore recommended that both software and hardware based acceleration strategies are explored to reduce the time required to execute the microscale model. One suggestion is to leverage a GPU compatible solver at the core of the microscale model, as significant speed-ups in the linear algebra computation are afforded by GPUs in comparison to optimised CPUs (Mozaffar et al., 2019). Another suggestion is to investigate the benefits afforded by fast-Fourier-transform based microscale simulations, which have been demonstrated to offer significant computational savings for trivial microstructural geometries (Gierden et al., 2021). This will enable a further step towards the geometric freedom afforded by hierarchical multiscale topology optimisation frameworks for a significantly reduced computational cost.

It is also suggested that a further AM constraint is developed and integrated within the present MTO framework, which aims to reduce the difficulty of powder removal from within optimised open-walled multiscale structures. This could be achieved by reducing the maximum manufacturable normalised radius, such that sufficiently large paths are created for powder removal. Alternatively, computational fluid dynamic analysis could be coupled to the optimiser to directly discourage constrictive powder removal paths.

In the long term, it is also recommended that an algorithm is developed which is capable of extracting succinct parametrisations from three-dimensional hierarchical multiscale topology op-

timisation frameworks. As detailed in Sec. 2.4.1, it is at present computationally intractable to execute three-dimensional hierarchical multiscale topology optimisation frameworks, however, these frameworks are capable of deriving optimal multiscale topologies without problem specific information. Assuming the performance of computers continues to increase in accordance with Moore's law, there will be a point in time where it is possible to execute three-dimensional hierarchical multiscale topology optimisation frameworks using modest high performance computing facilities. This will enable the following three step process to be implemented to derive optimal multiscale topologies at a significantly reduced computational cost

1. Assemble datasets for singular or hybrid objectives detailing optimal microscale topologies using a hierarchical multiscale topology optimisation framework
2. Apply model order reduction techniques to extract objective specific parametrisations from assembled datasets
3. Integrate derived parametrisations within the present multiscale topology optimisation framework to derive (quasi)-optimal multiscale topologies for a significantly reduced computational expense

Alternatively, inspiration could be taken from free material optimisation based approaches (Kocvara et al., 2008) to derive parametrisations where each parameter is able to directly manipulate a single element of the elasticity tensor. Through the application of inverse homogenisation techniques, this would enable a generalised parametrisation to be derived which is capable of efficiently targeting all forms of structural optimisation based problems. Specifically, this would have application within the field of mechanical compliance, where only nine parameters would be required to derive optimal orthotropic microscale topologies (Sigmund et al., 2016; Daynes et al., 2017; Wu et al., 2021a,b).

Appendix A

Derivation of Enclosing Simplex

To derive a simplex interpolant for the continuous microscale parameter set X_α, the microstructures within the full factorial DOE which form an enclosing simplex around X_α need to be located first. This set is a subset of the microstructures which form the enclosing hypercube, which are trivial to locate as a full factorial DOE represents a collection of n-dimensional hypercubes. Initially, only two microstructures from the hypercube set are required, the microstructure closest to the origin (ξ_{min}) and conversely the microstructure furthest away from the origin (ξ_{max}). To derive the microstructure closest to the origin, all elements within the microscale parameter set

$$X_\alpha = [0.22, 0.3, 0.11, 0.41, 0.03, 0.42, 0.19, 0.72, 1.31, 0.87]$$

are rounded down to the nearest DOE level (see Tab. 3.1)

$$\xi_{min} = [0.21, 0.28, 0.07, 0.35, 0.00, 0.42, 0.14, 0.60, 1.20, 0.80]$$

and conversely the microstructure furthest away from the origin is derived by rounding up all elements to the nearest DOE level

$$\xi_{max} = [0.28, 0.35, 0.14, 0.42, 0.07, 0.49, 0.21, 0.80, 1.40, 1.00]$$

Then by performing an exhaustive binary search, all 2^n microstructures within the hypercube set can then be found by permuting the minimum and maximum values derived for each microscale parameter.

The enclosing hypercube can be dissected into $n!$ congruent simplices using $n!/2!(n-2)!$ bisecting hyperplanes, of which all pass through ξ_{min} and ξ_{max} by definition. As a consequence, all $n!$ simplices within this hypercube are constructed using both ξ_{min} and ξ_{max}. To determine the remaining $n-1$ microstructures defining the enclosing simplex, a hyperplane bisection algorithm is used. To simplify the bisection process, the relative position of each microscale parameter within the hypercube is first computed

$$X'_{\alpha} = (X_{\alpha} - \xi_{min}) \oslash (\xi_{max} - \xi_{min}) \tag{A.1}$$

where X'_{α} denotes the transformed microscale parameters and for the present example is equal to

$$[0.14, 0.29, 0.57, 0.86, 0.43, 0.00, 0.71, 0.60, 0.55, 0.35]$$

A bisecting hyperplane is defined by the equality of a unique combination of two microscale parameters. Depending on which side of the hyperplane the transformed parameters fall, a different set of microstructures can be eliminated from the hypercube set. An example is illustrated in Fig. A.1, which depicts a hypercube, bisected by a hyperplane defined by the equality of the two transformed microscale parameters, $X'_{\alpha,1}$ and $X'_{\alpha,2}$.

The transformed vector of microscale parameters illustrated in Fig. A.1, possesses an $X'_{\alpha,2}$ value larger than its $X'_{\alpha,1}$ value. As a result, the coordinate lies above the line $X'_{\alpha,1} = X'_{\alpha,2}$ and therefore the microstructures with transformed coordinates $(1, 0, ..., X'_{\alpha,n})$ are eliminated from the hypercube set. This process is repeated for every hyperplane combination until only $n+1$ microstructures remain in the hypercube set. The reverse transformation is then applied

$$X_{\alpha} = (X'_{\alpha} + \xi_{min}) \circ (\xi_{max} - \xi_{min}) \tag{A.2}$$

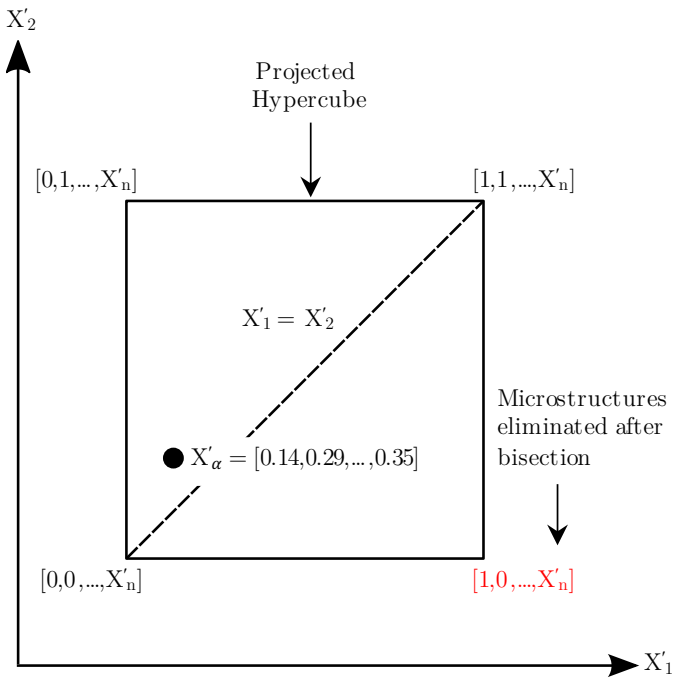


Figure A.1: Hypercube bisection example, projected onto the $X'_{\alpha,1} - X'_{\alpha,2}$ plane for clearer visualisation.

to recover the microscale parameters which form the enclosing simplex (Φ_α). Although this implementation enables the microstructures forming the enclosing simplex to be derived, it is computationally expensive in a high-dimensional space, as a large number of bisections are required.

To reduce the computational expense of this process, a novel method is presented which indirectly considers multiple bisecting hyperplanes simultaneously. The purpose of a bisecting hyperplane is to determine whether the transformed value of one parameter is larger than another. Instead of considering each hyperplane individually, by sorting the transformed vector of design variables into order of descending magnitude

$$[0.86, 0.71, 0.60, 0.57, 0.55, 0.43, 0.35, 0.29, 0.14, 0.00]$$

and recording the order of the sorted array

$$(X'_{\alpha,4}, X'_{\alpha,7}, X'_{\alpha,8}, X'_{\alpha,3}, X'_{\alpha,9}, X'_{\alpha,5}, X'_{\alpha,10}, X'_{\alpha,2}, X'_{\alpha,1}, X'_{\alpha,6})$$

all of the bisecting hyperplanes can be considered simultaneously. As the first design variable in the sorted array is larger than all other design variables by definition, it will always be incremented from the base point first, as all other permutations will be eliminated after bisection. This process is then repeated for all elements within the sorted array. In summary, to find the enclosing simplex, the base point of the enclosing hypercube (ξ_{min}) is incremented to the next DOE level in the order of the design variables within the sorted array, with each intermediate node serving as a microstructure within the enclosing simplex. For the microscale parameter set X_α, the microstructures forming the enclosing simplex are

$$\begin{aligned}
\phi_\alpha = \{&[0.21, 0.28, 0.07, 0.35, 0.00, 0.42, 0.14, 0.60, 1.20, 0.80],\\
&[0.21, 0.28, 0.07, 0.42, 0.00, 0.42, 0.14, 0.60, 1.20, 0.80],\\
&[0.21, 0.28, 0.07, 0.42, 0.00, 0.42, 0.21, 0.60, 1.20, 0.80],\\
&[0.21, 0.28, 0.07, 0.42, 0.00, 0.42, 0.21, 0.80, 1.20, 0.80],\\
&[0.21, 0.28, 0.14, 0.42, 0.00, 0.42, 0.21, 0.80, 1.20, 0.80],\\
&[0.21, 0.28, 0.14, 0.42, 0.00, 0.42, 0.21, 0.80, 1.40, 0.80],\\
&[0.21, 0.28, 0.14, 0.42, 0.07, 0.42, 0.21, 0.80, 1.40, 0.80],\\
&[0.21, 0.28, 0.14, 0.42, 0.07, 0.42, 0.21, 0.80, 1.40, 1.00],\\
&[0.21, 0.35, 0.14, 0.42, 0.07, 0.42, 0.21, 0.80, 1.40, 1.00],\\
&[0.28, 0.35, 0.14, 0.42, 0.07, 0.42, 0.21, 0.80, 1.40, 1.00],\\
&[0.28, 0.35, 0.14, 0.42, 0.07, 0.49, 0.21, 0.80, 1.40, 1.00]\}
\end{aligned}$$

This method is extremely computationally efficient, as the only intensive process is the sorting of the array, which is in contrast to the large number of bisecting hyperplane calculations needed otherwise. During optimisation, this calculation is required over one million times, so the reduction in execution time is significant.

Bibliography

Aage, N., Andreassen, E., Lazarov, B., and Sigmund, O. (2017). Giga-voxel computational morphogenesis for structural design. *Nature*, 550:84–86.

Alam, F. M., McNaught, K. R., and Ringrose, T. J. (2004). A comparison of experimental designs in the development of a neural network simulation metamodel. *Simulation Modelling Practice and Theory*, 12(7):559–578.

Allaire, G., Jouve, F., and Toader, A.-M. (2002). A level-set method for shape optimization. *Comptes Rendus Mathematique*, 334(12):1125–1130.

Allaire, G., Jouve, F., and Toader, A.-M. (2004). Structural optimization using sensitivity analysis and a level-set method. *Journal of Computational Physics*, 194(1):363–393.

Andrei, N. (2017). *Continuous Nonlinear Optimization for Engineering Applications in GAMS Technology*. Springer.

Antony, J. (2014). 6 - full factorial designs. In Antony, J., editor, *Design of Experiments for Engineers and Scientists (Second Edition)*, pages 63 – 85. Elsevier, Oxford, second edition edition.

Arabnejad, S. and Pasini, D. (2013). Mechanical properties of lattice materials via asymptotic homogenization and comparison with alternative homogenization methods. *International Journal of Mechanical Sciences*, 77:249–262.

Attaran, M. (2017). The rise of 3-d printing: The advantages of additive manufacturing over traditional manufacturing. *Business Horizons*, 60(5):677–688.

Bellman, R., Bellman, R., and Corporation, R. (1957). *Dynamic Programming*. Rand Corporation research study. Princeton University Press.

Bendsøe, M. P. (1989). Optimal shape design as a material distribution problem. *Structural optimization*, 1:193–202.

Bendsøe, M. P. and Kikuchi, N. (1988). Generating optimal topologies in structural design using a homogenization method. *Computer Methods in Applied Mechanics and Engineering*, 71(2):197 – 224.

Bendsøe, M. P. and Sigmund, O. (2004). *Topology Optimization: Theory, Methods and Applications*. Springer.

Bernstein, D. S. (2009). *Matrix Mathematics: Theory, Facts, and Formulas (Second Edition)*. Princeton University Press.

Bici, M., Broggiato, G. B., and Campana, F. (2017). *Topological Optimization in Concept Design: starting approach and a validation case study*, pages 289–299. Springer International Publishing.

Bjørk, R., Bahl, C., and Insinga, A. (2017). Topology optimized permanent magnet systems. *Journal of Magnetism and Magnetic Materials*, 437:78 – 85.

Bourdin, B. (2001). Filters in topology optimization. *International Journal for Numerical Methods in Engineering*, 50(9):2143–2158.

Brackett, D., Ashcroft, I., and Hague, R. (2011). Topology optimization for additive manufacturing. *22nd Annual International Solid Freeform Fabrication Symposium - An Additive Manufacturing Conference, SFF 2011*, pages 348–362.

Bruns, T. E. and Tortorelli, D. A. (2003). An element removal and reintroduction strategy for the topology optimization of structures and compliant mechanisms. *International Journal for Numerical Methods in Engineering*, 57(10).

Buhmann, M. D. (2003). *Radial Basis Functions: Theory and Implementations*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press.

Burger, M., Hackl, B., and Ring, W. (2004). Incorporating topological derivatives into level set methods. *Journal of Computational Physics*, 194(1):344–362.

Burrage, K., Burrage, P., Donovan, D., and Thompson, B. (2015). Populations of models, experimental designs and coverage of parameter space by latin hypercube and orthogonal sampling. *Procedia Computer Science*, 51:1762–1771.

Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C. (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208.

Chen, D., Skouras, M., Zhu, B., and Matusik, W. (2018). Computational discovery of extremal microstructure families. *Science Advances*, 4(1).

Cheng, L., Bai, J., and To, A. C. (2018). Functionally graded lattice structure topology optimization for the design of additive manufactured components with stress constraints. *Computer Methods in Applied Mechanics and Engineering*, 344:334 – 359.

Cheng, L., Bai, J., and To, A. C. (2019). Functionally graded lattice structure topology optimization for the design of additive manufactured components with stress constraints. *Computer Methods in Applied Mechanics and Engineering*, 344:334–359.

Cheng, L., Zhang, P., Biyikli, E., Bai, J., Robbins, J., and To, A. (2017). Efficient design optimization of variable-density cellular structures for additive manufacturing: Theory and experimental validation. *Rapid Prototyping Journal*, 23.

Christensen, P. and Klarbring, A. (2008). *An Introduction to Structural Optimization*, volume 153.

Colabella, L., Cisilino, A. P., Fachinotti, V., and Kowalczyk, P. (2019). Multiscale design of elastic solids with biomimetic cancellous bone cellular microstructures. *Struct. Multidiscip. Optim.*, 60(2):639–661.

Conte, S. and de Boor, C. (2017). *Chapter Six: Approximation*, pages 235–293. Society for Industrial and Applied Mathematics.

Daynes, S., Feih, S., Lu, W. F., and Wei, J. (2017). Optimisation of functionally graded lattice structures using isostatic lines. *Materials & Design*, 127:215–223.

Dbouk, T. (2017). A review about the engineering design of optimal heat transfer systems using topology optimization. *Applied Thermal Engineering*, 112:841 – 854.

Dühring, M. B., Jensen, J. S., and Sigmund, O. (2008). Acoustic design by topology optimization. *Journal of Sound and Vibration*, 317(3):557 – 575.

Edwards, C. S., Kim, H. A., and Budd, C. J. (2007). An evaluative study on ESO and SIMP for optimising a cantilever tie—beam. *Structural and Multidisciplinary Optimization*, 34(5).

Flores, I., Kretzschmar, N., Azman, A. H., Chekurov, S., Pedersen, D. B., and Chaudhuri, A. (2020). Implications of lattice structures on economics and productivity of metal powder bed fusion. *Additive Manufacturing*, 31:100947.

Forrester, A., Sobester, A., and Keane, A. J. (2008). *Constructing a Surrogate*, chapter 2, pages 33–76. John Wiley & Sons, Ltd.

Fulmanski, P., Laurain, A., Scheid, J.-F., and Sokoowski, J. (2008). Level set method with topological derivatives in shape optimization. *International Journal of Computer Mathematics*, 85(10):1491–1514.

Garner, E., Kolken, H. M. A., Wang, C. C. L., Zadpoor, A. A., and Wu, J. (2019). Compatibility in microstructural optimization for additive manufacturing. *Additive Manufacturing*, 26:65–75.

Gaynor, A. T. and Guest, J. K. (2016). Topology optimization considering overhang constraints: Eliminating sacrificial support material in additive manufacturing through design. *Structural and Multidisciplinary Optimization*, 54(5):1157–1172.

Gebhardt, A. and Htter, J.-S. (2016). 1 - basics, definitions, and application levels. In Gebhardt, A. and Htter, J.-S., editors, *Additive Manufacturing*, pages 1–19. Hanser.

Gertz, M., Nocedal, J., and Sartenar, A. (2004). A starting point strategy for nonlinear interior methods. *Applied Mathematics Letters*, 17(8):945–952.

Gierden, C., Waimann, J., Svendsen, B., and Reese, S. (2021). A geometrically adapted reduced set of frequencies for a fft-based microstructure simulation. *Computer Methods in Applied Mechanics and Engineering*, 386:114131.

Gill, P. E., Murray, W., and Saunders, M. A. (2005). Snopt: An SQP algorithm for large-scale constrained optimization. *SIAM Review*, 47(1):99–131.

Guo, H., Takezawa, A., Honda, M., Kawamura, C., and Kitamura, M. (2020). Finite element simulation of the compressive response of additively manufactured lattice structures with large diameters. *Computational Materials Science*, 175.

Hashin, Z. and Shtrikman, S. (1963). A variational approach to the theory of the elastic behaviour of multiphase materials. *Journal of the Mechanics and Physics of Solids*, 11(2):127 – 140.

Hassani, B. and Hinton, E. (1998). A review of homogenization and topology optimization ihomogenization theory for media with periodic structure. *Computers & Structures*, 69(6).

Hemingway, P. (2002). *n-Simplex Interpolation*.

Higham, N. J. (2002). *Accuracy and Stability of Numerical Algorithms*. Society for Industrial and Applied Mathematics.

Hollister, S., Fyhrie, D., Jepsen, K., and Goldstein, S. (1991). Application of homogenization theory to the study of trabecular bone mechanics. *Journal of Biomechanics*, 24(9):825–839.

Hollister, S. J. and Kikuchi, N. (1992). A comparison of homogenization and standard mechanics analyses for periodic porous composites. *Computational Mechanics*, 10(2):73–95.

HyperWorks (2017). *OptiStruct*. Altair Engineering, Inc., United States.

Imediegwu, C. (2020). Multiscale structural, thermal and thermo-structural optimization towards three-dimensional printable structures.

Imediegwu, C., Murphy, R., Hewson, R., and Santer, M. (2019). Multiscale structural optimization towards three-dimensional printable structures. *Structural and Multidisciplinary Optimization*, 60(2):513–525.

Jiang, J., Stringer, J., and Xu, X. (2019). Support optimization for flat features via path planning in additive manufacturing. *3D Printing and Additive Manufacturing*, 6(3):171–179.

Jiang, J., Stringer, J., Xu, X., and Zhong, R. Y. (2018a). Investigation of printable threshold overhang angle in extrusion-based additive manufacturing for reducing support waste. *International Journal of Computer Integrated Manufacturing*, 31(10):961–969.

Jiang, J., Xu, X., and Stringer, J. (2018b). Support structures for additive manufacturing: A review. *Journal of Manufacturing and Materials Processing*, 2(4).

Johnson, T. E. and Gaynor, A. T. (2018). Three-dimensional projection-based topology optimization for prescribed-angle self-supporting additively manufactured structures. *Additive Manufacturing*, 24:667–686.

Kadic, M., Milton, G. W., van Hecke, M., and Wegener, M. (2019). 3d metamaterials. *Nature Reviews Physics*, 1(3):198–210.

Keane, A. and Nair, P. (2005). *Architectures for Multidisciplinary Optimization*, chapter 9, pages 359–385. John Wiley & Sons, Ltd.

Kocvara, M., Stingl, M., and Zowe, J. (2008). Free material optimization: recent progress. *Optimization*, 57(1):79–100.

Lakes, R. (1987). Foam structures with a negative poisson's ratio. *Science*, 235(4792).

Langelaar, M. (2016). Topology optimization of 3d self-supporting structures for additive manufacturing. *Additive Manufacturing*, 12:60–70.

Larson, M. G. and Bengzon, F. (2013). *The Finite Element Method: Theory, Implementation, and Applications*. Springer Publishing Company, Incorporated.

Lazarov, B. S. and Sigmund, O. (2011). Filters in topology optimization based on helmholtz-type differential equations. *International Journal for Numerical Methods in Engineering*, 86(6):765–781.

Lazarov, B. S., Wang, F., and Sigmund, O. (2016). Length scale and manufacturability in density-based topology optimization. *Archive of Applied Mechanics*, 86(1):189–218.

Lewiner, T., Lopes, H., Vieira, A. W., and Tavares, G. (2003). Efficient implementation of marching cubes' cases with topological guarantees. *Journal of Graphics Tools*, 8(2):1–15.

Li, H., Luo, Z., Gao, L., and Qin, Q. (2018). Topology optimization for concurrent design of structures with multi-patch microstructures by level sets. *Computer Methods in Applied Mechanics and Engineering*, 331:536 – 561.

Li, Q., Chen, W., Liu, S., and Tong, L. (2016). Structural topology optimization considering connectivity constraint. *Structural and Multidisciplinary Optimization*, 54(4):971–984.

Liu, D. C. and Nocedal, J. (1989). On the limited memory bfgs method for large scale optimization. *Mathematical Programming*, 45(1):503–528.

Liu, H., Zong, H., Tian, Y., Ma, Q., and Wang, M. Y. (2019). A novel subdomain level set method for structural topology optimization and its application in graded cellular structure design. *Structural and Multidisciplinary Optimization*, 60(6):2221–2247.

Liu, L., Yan, J., and Cheng, G. (2008). Optimum structure with homogeneous optimum truss-like material. *Computers & Structures*, 86(13):1417–1425.

Liu, P., Kang, Z., and Luo, Y. (2020a). Two-scale concurrent topology optimization of lattice structures with connectable microstructures. *Additive Manufacturing*, 36:101–427.

Liu, P., Kang, Z., and Luo, Y. (2020b). Two-scale concurrent topology optimization of lattice structures with connectable microstructures. *Additive Manufacturing*, 36.

Liu, S., Li, Q., Chen, W., Tong, L., and Cheng, G. (2015). An identification method for enclosed voids restriction in manufacturability design for additive manufacturing structures. *Frontiers of Mechanical Engineering*, 10(2):126–137.

Liu, S., Li, Q., Liu, J., Chen, W., and Zhang, Y. (2018). A realization method for transforming a topology optimization design into additive manufacturing structures. *Engineering*, 4(2):277 – 285.

Liu, Y., Zheng, G., Letov, N., and Zhao, Y. (2020c). A survey of modeling and optimization methods for multi-scale heterogeneous lattice structures. *Journal of Mechanical Design*, 143:1–59.

Logg, A., Mardal, K.-A., Wells, G. N., et al. (2012). *Automated Solution of Differential Equations by the Finite Element Method*. Springer.

Luo, J., Luo, Z., Chen, L., Tong, L., and Wang, M. (2008). A semi-implicit level set method for structural shape and topology optimization. *J. Comput. Phys.*, 227:5561–5581.

Mattheck, C. and Burkhardt, S. (1990). A new method of structural shape optimization based on biological growth. *International Journal of Fatigue*, 12(3).

McKay, M., Beckman, R., and Conover, W. (1979). Comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239–245.

Mirzendehdel, A. M., Behandish, M., and Nelaturi, S. (2020). Topology optimization with accessibility constraint for multi-axis machining. *Computer-Aided Design*, 122.

Mitusch, S., Funke, S., and Dokken, J. (2019). dolfin-adjoint 2018.1: automated adjoints for fenics and firedrake. *Journal of Open Source Software*, 4(38):1292.

Mozaffar, M., Ndip-Agbor, E., Lin, S., Wagner, G. J., Ehmann, K., and Cao, J. (2019). Acceleration strategies for explicit finite element analysis of metal powder-based additive manufacturing processes using graphical processing units. *Computational Mechanics*, 64(3):879–894.

Multiphysics, C. (1998). Introduction to comsol multiphysics®. *COMSOL Multiphysics, Burlington*, 9:2018.

Murat, F. and Tartar, L. (1997). *Calculus of Variations and Homogenization*, pages 139–173. Birkhäuser Boston.

Murphy, R., Imediegwu, C., Hewson, R., and Santer, M. (2021). Multiscale structural optimization with concurrent coupling between scales. *Structural and Multidisciplinary Optimization*, 63(4):1721–1741.

Murphy, R. D., Imediegwu, C., Hewson, R., Santer, M. J., and Muir, M. (2020). *Multiscale Concurrent Multi-Objective Structural Optimization of a Goose Neck Hinge*.

Muthusamy, M., Safaee, S., and Chen, R. K. (2018). Additive manufacturing of overhang structures using moisture-cured silicone with support material. *Journal of Manufacturing and Materials Processing*, 2(2).

Nightingale, M., Hewson, R., and Santer, M. (2021). Multiscale optimisation of resonant frequencies for lattice-based additive manufactured structures. *Structural and Multidisciplinary Optimization*, 63:1–15.

Norato, J., Bell, B., and Tortorelli, D. (2015). A geometry projection method for continuum-based topology optimization with discrete elements. *Computer Methods in Applied Mechanics and Engineering*, 293:306–327.

Oĺeary, D. P. and Stewart, G. W. (1996). Computing the eigenvalues and eigenvectors of arrowhead matrices.

Otero, F., Oller, S., and Martinez, X. (2018). Multiscale computational homogenization: Review and proposal of a new enhanced-first-order method. *Archives of Computational Methods in Engineering*, 25(2):479–505.

Panesar, A., Abdi, M., Hickman, D., and Ashcroft, I. (2018). Strategies for functionally graded lattice structures derived using topology optimisation for additive manufacturing. *Additive Manufacturing*, 19:81–94.

Penumakala, P. K., Santo, J., and Thomas, A. (2020). A critical review on the fused deposition modeling of thermoplastic polymer composites. *Composites Part B: Engineering*, 201.

Plocher, J. and Panesar, A. (2019). Review on design and structural optimisation in additive manufacturing: Towards next-generation lightweight structures. *Materials & Design*, 183:108164.

Querin, O. M., Steven, G. P., and Xie, Y. M. (1998). Evolutionary structural optimisation (ESO) using a bidirectional algorithm. 15(8).

Rodrigues, H., Guedes, J. M., and Bendsoe, M. P. (2002). Hierarchical optimization of material and structure. *Structural and Multidisciplinary Optimization*, 24(1):1–10.

Rojas-Labanda, S. and Stolpe, M. (2015). Benchmarking optimization solvers for structural topology optimization. *Structural and Multidisciplinary Optimization*, 52(3):527–547.

Sethian, J. and Wiegmann, A. (2000). Structural boundary design via level set and immersed interface methods. *Journal of Computational Physics*, 163(2):489–528.

Sigmund, O. (1994). Materials with prescribed constitutive parameters: An inverse homogenization problem. *International Journal of Solids and Structures*, 31(17).

Sigmund, O. (2000). A new class of extremal composites. *Journal of the Mechanics and Physics of Solids*, 48(2).

Sigmund, O. (2007). Morphology-based black and white filters for topology optimization. *Structural and Multidisciplinary Optimization*, 33(4):401–424.

Sigmund, O., Aage, N., and Andreassen, E. (2016). On the (non-)optimality of michell structures. *Structural and Multidisciplinary Optimization*, 54(2).

Sigmund, O. and Maute, K. (2013). Topology optimization approaches. *Structural and Multidisciplinary Optimization*, 48(6):1031–1055.

Sivapuram, R., Dunning, P. D., and Kim, H. A. (2016). Simultaneous material and structural optimization by multiscale topology optimization. *Structural and Multidisciplinary Optimization*, 54(5):1267–1281.

Smith, M. (2009). *ABAQUS/Standard User's Manual, Version 6.9*. Dassault Systèmes Simulia Corp, United States.

Stein, P. (1966). A note on the volume of a simplex. *The American Mathematical Monthly*, 73(3):299–301.

Stolpe, M. and Svanberg, K. (2001). On the trajectories of penalization methods for topology optimization. *Structural and Multidisciplinary Optimization*, 21(2):128–139.

Svanberg, K. (1987). The method of moving asymptotesa new method for structural optimization. *International Journal for Numerical Methods in Engineering*, 24(2):359–373.

Thillaithevan, D., Bruce, P., and Santer, M. (2021). Stress-constrained optimization using graded lattice microstructures. *Structural and Multidisciplinary Optimization*, 63(2):721–740.

Tsai, S. (1966). Mechanics of composite materials. Technical report, Air Force Materials Laboratory Technical Report, AFML-TR-66-149.

van Dijk, N. P., Langelaar, M., and van Keulen, F. (2012). Explicit level-set-based topology optimization using an exact heaviside function and consistent sensitivity analysis. *International Journal for Numerical Methods in Engineering*, 91(1).

van Dijk, N. P., Maute, K., Langelaar, M., and van Keulen, F. (2013). Level-set methods for structural topology optimization: a review. *Structural and Multidisciplinary Optimization*, 48(3):437–472.

Veselago, V. G. (1968). The electrodynamics of substances with simultaneously negative values of ϵ and μ. *Soviet Physics Uspekhi*, 10(4):509–514.

Wächter, A. and Biegler, L. T. (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57.

Wang, C., Gu, X., Zhu, J., Zhou, H., Li, S., and Zhang, W. (2020a). Concurrent design of hierarchical structures with three-dimensional parameterized lattice microstructures for additive manufacturing. *Structural and Multidisciplinary Optimization*, 61(3):869–894.

Wang, F., Lazarov, B. S., and Sigmund, O. (2011). On projection methods, convergence and robust formulations in topology optimization. *Structural and Multidisciplinary Optimization*, 43(6):767–784.

Wang, F.-S. and Chen, L.-H. (2013). *Heuristic Optimization*, pages 885–885. Springer New York.

Wang, J., Dai, G., and Huang, J. (2020b). Thermal metamaterial: Fundamental, application, and outlook. *iScience*, 23(10).

Wang, M. Y., Wang, X., and Guo, D. (2003). A level set method for structural topology optimization. *Computer Methods in Applied Mechanics and Engineering*, 192(1):227 – 246.

Weinan, E., Engquist, B., Li, X., Ren, W., and Vanden-Eijnden, E. (2007). Heterogeneous multiscale methods: A review. *Communications in Computational Physics*, 2(3):367–450.

Wu, J., Sigmund, O., and Groen, J. (2021a). Topology optimization of multi-scale structures: a review. *Structural and Multidisciplinary Optimization*, 63:14551480.

Wu, J., Wang, W., and Gao, X. (2021b). Design and optimization of conforming lattice structures. *IEEE Transactions on Visualization and Computer Graphics*, 27(1):4356.

Xie, Y. and Steven, G. (1993). A simple evolutionary procedure for structural optimization. *Computers & Structures*, 49(5):885–896.

Xu, B., Han, Y. S., Zhao, L., and Xie, Y. M. (2019). Topology optimization of continuum structures for natural frequencies considering casting constraints. *Engineering Optimization*, 51(6):941–960.

Yang, R. J. and Chahande, A. I. (1995). Automotive applications of topology optimization. *Structural optimization*, 9(3):245–249.

Yap, C. Y., Chua, C. K., Dong, Z. L., Liu, Z. H., Zhang, D. Q., Loh, L. E., and Sing, S. L. (2015). Review of selective laser melting: Materials and applications. *Applied Physics Reviews*, 2(4):041101.

Young, V., Querin, O. M., Steven, G. P., and Xie, Y. M. (1999). 3d and multiple load case bi-directional evolutionary structural optimization (beso). *Structural optimization*, 18(2).

Yu, H., Huang, J., Zou, B., Shao, W., and Liu, J. (2020). Stress-constrained shell-lattice infill structural optimisation for additive manufacturing. *Virtual and Physical Prototyping*, 15(1):35–48.

Zhang, S., Norato, J. A., Gain, A. L., and Lyu, N. (2016). A geometry projection method for the topology optimization of plate structures. *Structural and Multidisciplinary Optimization*, 54(5):1173–1190.

Zhang, S., Yin, L., and Fang, N. (2009). Focusing ultrasound with an acoustic metamaterial network. *Phys. Rev. Lett.*, 102.

Zhou, M. and Rozvany, G. (1991). The coc algorithm, part ii: Topological, geometrical and generalized shape optimization. *Computer Methods in Applied Mechanics and Engineering*, 89(1):309–336.

Zhu, B., Skouras, M., Chen, D., and Matusik, W. (2017). Two-scale topology optimization with microstructures. *ACM Trans. Graph.*, 36(4).

Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J. (1997). Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Trans. Math. Softw.*, 23(4):550560.

Zhu, J., Shou, H., Wang, C., Zhou, L., Yuan, S., and Zhang, W. (2021). A review of topology optimization for additive manufacturing: Status and challenges. *Chinese Journal of Aeronautics*, 34(1):91–110.

Zhu, J.-H., Zhang, W.-H., and Xia, L. (2016). Topology optimization in aircraft and aerospace structures design. *Archives of Computational Methods in Engineering*, 23(4):595–622.